University of Southampton

Faculty of Engineering, Science and Mathematics

School of Electronics and Computer Science

Towards a payment-free mechanism for promoting cooperation through scoring rules and reciprocity

by

Teodor Ghetiu

September 2008

A dissertation submitted in partial fulfilment of the degree of

MSc in Complexity Science

by examination and dissertation

Abstract

This work presents a model of payment-free cooperation addressing the problem of information elicitation within a distributed information system. Making use of the quadratic scoring rule, a mathematical instrument developed in the forecasting domain and proved to motivate experts in reporting truthfully their probabilistic estimates, the model moves away from previous approaches that relied on financial incentives for their effectiveness. Instead of payments, the concept of transferable utility is tested. Agents are allowed to exchange estimates, process aimed at increasing the precision of their private model and thus their expected utility. We show that this leads to a "prisoner's dilemma" scenario, situation that we address through a mechanism that equips agents with trust modelling abilities and adds a "stable roommates" pair-allocation service. Although the pairs generated are stable, the mechanism is shown to be effective mostly when cooperators are faced with very dishonest defectors, while softer forms of defection remain more profitable. We show that the main factors behind this situation are the agent's method of decoding their trust models, coupled with the fixation of pairs in time sustained by the pair allocation service. A critical element is the agents' exploration-exploitation trade-off.

Acknowledgements

The list of people to thank for stretches a long way, up to my home town and Romania. So all my appreciation and respect to my family and friends, for supporting me throughout this wonderful ECS journey. Also Dr. Seth Bullock - our MSc Coordinator, Dr. Richard Watson - our tutor, and all the members of the SENSe group are special landmarks and I thank them to. To all the colleagues from MSc Complexity Science and all the friends brought forward by "random processes", I thank them for their "selective pressure" but also for helping me keep the "mutation rate" in acceptable ranges.

The work presented in this paper is part of a different phase of the MSc, and for this I carry great respect and gratitude to Dr. Alex Rogers, my supervisor, who should be Sir Alex Rogers for the way he guided and understood me, and for his calm when critical situations were met.

All in all, the greatest thanks to $\sum_{\alpha}^{\Omega}$, who gives a meaning to everything.

Contents

List of Figures

Chapter 1

Introduction

Information elicition is the process by which an agent aims at obtaining information about the private valuations of other information provider agents. The classical example is that of decision makers querying "experts" for their private opinion on subjects of interest: what will be the next day's weather, how good is a certain product or will the stock market see a positive evolution? These are all questions that require a probabilistic answer, coming from an information provider that is able to make such predictions.

For information elicitation to be effective, two challenges must be addressed. The first is called the elicitation problem and stands for the fact that forming and reporting one's opinion implies certain costs, but the information only benefits others, not direclty the provider. The second is the evaluation problem and refers to the information provider's honesty. In some cases, giving negative feedback is inhibited by the fear of retaliation or it is just more profitable to comply to a general opinion rather then expressing one's true position. Subjectivity plays thus a decisive role: it can steer decision makers towards obtaining an accurate view on the information required or a totally inaccurate one.

Due to the pervasiveness of information elicitation, several scientifical fields have studied and developed means of dealing with these two challenges. Forecasting, economics, computer science have all contributed to the design of systems that induce honest reporting and motivate effort. The main concept behind them is that of providing rewards, be them in the form of payments or other financial-like means. While this certainly addresses the economic part of the problem, it increases the requirements for transforming these systems from a theoretical state to a practical one. The nature of payments, they way they will handled, the currency used, etc are all issues that prevent reward systems from being implemented at a large scale.

1.1 Payment-free cooperation

This paper proposes a model for eliciting information without using financial incentives, but relying on the concepts of *transferable utility* and *reciprocity*. Moving away from the study of financial rewards and payment-based incentives, we focus here on the implications brought forward by viewing an agent as part of a community of information providers. Within such a framing, there might be more than one agent modelling the same event and they could benefit from sharing their predictions. The more data they obtain, the higher the precision of their personal estimate can be. As rewards are in direct relation with the quality of information agents provide, holding estimates of higher precisions induces an implicit increase in utility and thus in the probable rewards they will receive. Accordingly, cooperation in the form of truthful information exchanges, is a dominant strategy.

However, if we view the setup as a rational one, with agents having different stakeholders and thus different interests, cooperation is not dominant anymore. The incentive to free-ride or defect becomes more important than that for being truthful and we find agents being involved in a "prisoner's dilemma". Here is where we propose and study a payment-free mechanism (PFM) for promoting cooperation. We base this mechanism on the use of scoring rules for dealing with the evaluation problem while reciprocity, implemented through trust modelling techniques and a "stable roommates" allocation (SRA) service that matches agents according to their preferences, addresses the elicitation problem.

In more detail, each agent is equipped with trust modelling abilities through which it can evaluate the degree of honesty of those with which it interacted. Based on these evaluations, agents form preference lists that express their option towards interacting with the available agents. These are submitted to the pair-allocation service that in turn outputs a set of optimal agent-pairs. Interactions take place according to this output, agents update their trust models and the cycle repeats.

Our experiments prove that cooperation can be promoted through payment-free incentives, but under certain constraints. Defection, in its strong forms, is easier to detect and separate in a minimal number of interactions. Cooperators are then efficiently separated from this type of defectors and the pair allocator service allows them to have stable interactions. Softer forms of defection are harder to handle. When the agent's trust model becomes accurate it may realise it had been interacting with a defector, but just changing its preferences is not guaranteed to lead to a different pair allocation. In fact the bigger the agent population is and the higher the variety of defection levels, the harder it is for pure cooperators to be matched and thus obtain a dominance of cooperation. Truthfulness, in its pure form, requires then more flexibility from the agent's part and the mechanism.

1.2 Brief background

The two challenges of information elicitation have been previously addressed in a large body of scientifical work. The elicitation problem has been dealt with through the use of financial incentives. A relevant example is Cremer and McLean's [9, 10] study proving that budget-balancing payments can be used for extracting agent's private information. A different approach was taken for solving the evaluation problem. Custom mathematical instruments were developed, especially in the forecasting domain. An important class is that of proper scoring rules [60, 23] that are proven to motivate agents into revealing correctly their private estimates, this implying the full expression of their uncertainties.

Scoring rules and financial incentives were first combined by Johnson et al [31] who provided a method for computing budget-balancing transfer payments based on proper scoring rules. Further studies extended the applicability of such mechanisms to continuous information scenarios [56], and more recently, to different variation of mechanisms that relax the assumptions made by fully rational scenarios [41, 50].

The common denominator for previous work is the use of financial incentives for dealing with the elicitation problem. Also, the solutions provided are still based on centralised grounds. Mechanisms rely on the existence of a special agent, the "centre", that receives agent estimates and then returns rewards according to the scores obtained.

1.3 Summary and outline

In this paper, we contribute to the state of the art in the following ways:

1. We describe a new payment-free model for eliciting information that enables agents to freely exchange their forecasts in order to increase their utility. The model is formally evaluated and is proved to be leading towards a "prisoner's dilemma" scenario,thus making reciprocity a relevant candidate for promoting cooperation.

2. We propose then a mechanism that uses trust modelling techniques and a "stable roommates" algorithm for matching agents according to their preferrences, and show in what circumstances this mechanism is able to effectively promote pure cooperation.

3. We finally prove, through normal and evolutionary simulations, that pure cooperation can only be profitable in certain restrained contexts and that when using more complex setups, "soft" defectors are able to exploit cooperators. We conclude that using the trust models needs a higher efficiency, while the way agents address the exploration-exploitation trade-off is also causal. Matching agents through the "stable roommates" algorithm brings stability to the teams of cooperators, but may also prevent new similar pairs from forming.

The paper will proceed as follows. Chapter 2 presents the background literature for this research. Chapter 3 introduces the payment-free model and the mechanism developed for using it. In chapter 4 we describe our findings followed by a brief evaluation of the degree to which our research objectives have been met, in chapter 5. Finally, in chapter 6 we present our conclusions, proposed future work and the implications relative to complexity science.

Chapter 2

Literature review

This chapter presents the main theoretical domains and tools onto which our research work is based. Section 2.1 introduces mechanism design theory, followed by a synthesis of the scoring rules domain in section 2.3 and a brief detailing of the stable roommates problem in section 2.2.

2.1 Mechanism design theory

Mechanism design theory is an important body of scientifical work that provides a framework for analysing and developing systems which display *system-wide properties* while being composed of multiple self-interested agents, each with their private states and objectives. The theory relates to microeconomics and game theory while its use in modern sciences qualifies it as a relevant transdisciplinary domain.

Objectives of mechanism design (MD) are usually system properties like: *truthfulness*, *budget balance*, *maximised social welfare* or *Pareto-optimality*[12]. It is important to note that, due to the self-interested nature of agents, the above mentioned properties may be achieved only if they are in the best interest of each agent individually. In other words, the desired expected outcome must be a result of each agent playing its best strategy.

2.1.1 Important concepts

The concept that marked the creation of the field was that of *mechanism*, defined by Hurwicz in a 1960 article [25]. According to it, a mechanism was a communication system that enabled its participants to send messages to one another or to a "message center", and that mapped every set of received messages to an outcome (e.g. allocation of goods), through a pre-defined rule. This definition stands as the framework for comparing markets and similar institutions to a large range of alternative systems of various natures [48]. It also gives insights into the two main types of mechanisms, centralised and decentralised ones, the difference being made by the presence or absence of the "message center".

Further on, an agent is viewed as having a *type* expressed as private information, and a set of available *strategies* that specify what *actions* it should take according to the state of the world, in order to maximise its *utility*. At any time, an agent has a preferred (dominant) strategy that can be predicted on the basis of a *solution concept* - a rule that states how the game will be played i.e. agents using dominant strategy equilibrium, Nash equilibrium or Bayes-Nash equilibrium.

In normal circumstances and most real-case scenarios, agents may have different stakeholders with potentially conflicting goals. This means that the designer will not be able to control directly an agent's strategy, thus, in order to obtain system-wide properties it has to create a mechanism that incentifies agents in the desired direction. This is where the *revelation principle* steps in. Stating that under weaker conditions any mechanism can be redefined as an equivalent *incentive-compatible* (IC) *direct-revelation mechanism* (DRM) [44, 51], the revelation principle provides the conceptual link between problems and solutions found in representative and heterogeneous agent systems (systems having identical or different types of agents).

Direct-revelation and incentive-compatibility are the main means of overcoming the selfish nature of agents. If direct-revelation relates to setups where making direct claims about ones preferences to the mechanism is the only action available to an agent, incentive-compatible are those direct-revelation mechanisms where information on these preferences is *reported truthfuly*, at equilibrium.

The last important concept is that of *social choice function* (SCF) which is a function that selects an outcome (the desired system-wide property), given the set of types from agents taking part in the game. Social choice functions are the methods through which mechanism designers implement system wide properties. Referring to the recommender system example, if the desired property would be that of truthfullness, the SCF could measure how close the received valuations are to a value that is believed to be truthful. If such a value is not available, other mechanisms could be employed (see section 2.3.2).

For a brief mathematical exposition of mechanism design theory, please consult appendix A.1.

2.1.2 State of the art

The main field of mechanism design has seen important developments over the past decades. Catalised by new results and new problems posed in microeconomics and game theory, mechanism desing plays also an important part in achieving the important aims assumed by computer science.

In order to address algorithmic problems, the sub-field of *algorithmic mechanism design* (AMD) was created. The seminal work presented in 1999 by Nisan and Ronen [46] and then published in 2001 [47] defined the term of AMD, while introducing a centralised model that combined incentive-compatibility and tractability. By providing a rigurous formal definition for the model, their work received received multiple extensions and consequently the creation of new sub-fields.

Branching off from AMD, the new field of *distributed algorithmic mechanism design* (DAMD) was introduced by Feigenbaum et al [14]. DAMD had the same goals as AMD, that is incentive-compatibility and computational tractability, but having them framed in a distributed environment. In other words, the agents, the information and the model were all distributed. For a more detailed exposition on decentralised mechanisms and DAMD, Mu'alem [42] and Fiegenbaum et al [15] provide good references. [1]

Another important subfield of MD is *computational mechanism design* (CMD). This is one of the important vectors of bringing mechanisms and results of MD closer to practical implementation, as it relaxes a central set of assumptions present in MD [8]. Thus, issues like agent rationality, communication infrastructure robustness or agent population composition that are viewed by MD as theoretical concepts and thus fail safe or unbounded, are perceived by MD through the computational requirements they imply. As a consequence, agent rationality is not assumed anymore unbounded, communication has its risks of failure while the agent population might change in numbers or type distribution [12]. CMD seeks to address all these limitations, applying MD techniques to computational problems. Current literature sees CMD techniques being applied in fields involving practical applications of computer science e.g. information fusion within sensor networks [57].

2.2 Stable roommates problem

The stable roommates problem (SRP) is an extension to the stable marriage problem (SMP), introduced by Gale and Shapley in a 1962 article [19]. If the challenge for the SMP is to find stable matchings between two groups of men and women, each of them ranking the members of the opposite sex according to their preferences, the SRP deals with only one group of individuals.

Both problems are very important as they stable solutions to individual's preferences. In the context of SMP, a pair is stable if both the man and woman have no other available partner of opposite sex that they prefer more than they current match. The only difference for the SRP is that there is no opposite sex, all the pairs being done within the same population.

Following Gane and Shapley's seminal work, McVitie and Wilson [40] proposed a recursive algorithm for obtaining all stable matchings given a SMP instance. Tests revealed that their algorithm, and the others previously published had a complexity of $\mathcal{O}(n^2)$, this motivating Knuth into undertaking an extensive study on the SMP and its implications [34]. He analysed previous algorithms and made connections to other combinatorial problems, finally detailing twelve new research problems, one of which being SRP. While previously known that the difference between SMP and SRP was that for SRP, roommate instances existed for which stable matching could not be obtained, Knuth proved that the SRP could also have multiple solutions. He exemplified this by using an instance of size 8 (table 2.1) that yields 3 stable matchings: {1/2, 3/4, 5/8, 6/7}, {1/5, 2/6, 3/7, 4/8} and {1/4, 2/3, 5/6, 7/8}. Finally, he challenged researchers into defining an efficient algorithm for solving the SRP, suggesting that the problem might be NP-complete.

Following Knuth's challenge, numerous improvements were done to the "fundamental" algorithm of McVitie and Wilson, in addition to the proposals of new ones. Thus, Irving [27, 28] solved the SRP through an efficient algorithm that either found a stable matching for all participants or generated a negative result.

More recent work saw the assumptions made for the SRP being relaxed, while its area of applicability was visibly extended. To exemplify, Irvin et al propose in 2002 the SRP with ties [29], agents

Person	Preference list						
1	2	5	4	6	7	8	3
2	3	6	1	7	8	5	4
3	4	7	2	8	5	6	1
4	1	8	3	5	6	7	2
5	6	1	8	2	3	4	7
6	7	2	5	3	4	1	8
7	8	3	6	4	1	2	5
8	5	4	7	1	2	3	6

Table 2.1: Knuth's stable roommates instance of size 8

being permitted here to express ties in their preference lists. Also Fleiner [17] introduces the SRP with choice functions, that allows an agent to be involved in multiple relations. Another study by Fleiner et al [16] sees the assumption of linear preferences being relaxed, agents having here partial preferences. While not iterating through all the SRP's extensions, we leave as references the work of Toka et al [64] and the way Levin et al [36] motivate the use of a SRP based mechanism within the PeerWise Internet routing overlay. Finally, for a discussion on the complexity of SMP and its variants or on the existence of stable roommate matchings we cite Cechlarova [6] and Chung [7], followed by a recent survey on the SRP from Iwama et al [30].

2.3 Scoring rules

Scoring rules have been introduced during the early 1970s', authors like Hendrickson and Buehler [23] or Savage [60] pioneering their use in the domain of forecasting. The necessity for developing such rules has been motivated by the fact that decision makers often relied on experts for making predictions but they lacked the means of making sure these predictions were accurately generated and reported. Scoring rules are especially designed for assessing the quality of a probabilistic estimate, while not penalising their degree of uncertainty.

Two main problems are faced by a decision maker requesting a forecast from an expert: the *elicitation* problem and the *evaluation* problem [18]. If the elicitation problem is described as the decision maker's need for assuring that it is in the expert's best interest to produce the forecast in a way that accurately reflects its best judgement, the evaluation problem is defined as the need to assess the quality of a forecast generated by the expert. As scoring rules assign numerical scores to forecasts based on their closeness to the actual values of the predicted parameters, they deal implicitely with the evaluation problem. In case the experts understand the way they are evaluated and want to maximise their expected score, the rules also address the elicitation problem (Murphy and Winkler present a more detailed discussion on this issue [43]).

To sum up, mechanisms designed for promoting truthfulness can benefit from the use of scoring rules, both from having their structure simplified (as they can address both the evaluation and elicitation problems simultaneously) but also from having a better mathematical grounding.

2.3.1 Strictly proper scoring rules

The most important subclass of scoring rules is that of *strictly proper scoring rules*. These are directly related to incentive-compatible direct-revelation mechanisms and to social choice functions as truthfulness or precision, thus having the potential of being used in a wide range of domains.

Being *proper* means that a scoring rule will enable a forecaster to maximise its expected score by revealing its true estimate, while *strictly proper* means that this maximum score is unique (and attainable only through reporting the estimate accurately). These features correspond to observations stating that the the value of forecasts can be much higher if a measure of their uncertainty is also made available [49, 65].

A set of strictly proper scoring rules has been developed, out of which the quadratic, spherical and logarithmic rules have received special attention. Studies like [18], [20] or [21] have all extended the theory of scoring rules, especially the strictly proper ones, and their area of applicability. The way scores are assigned to probabilistic estimates is a modern approach to incentivising cooperation within rational environments and is well reflected within current research on information elicitation.

For a brief mathematical exposition of strictly proper scoring rules, please see appendix A.2.

2.3.2 State of the art

The main use of scoring rules is that of addressing the difficulties of information elicitation. Having a heterogeneous agent system means that *defection*, manifested through agents supplying erroneous or imprecise information, can be a dominant strategy (as proved by through the "prisoner's dilemma" game [3, 2]). In order to promote cooperation, research has been undertaken in the direction of developing computational trust models, understanding the evolution of trust and designing mechanisms that may efficiently implement it.

Results span from the use of more "ad hoc" trust metrics [26, 58, 62] to the framing of trust within probability theory [45, 54, 55]. Research has addressed both formal aspects of modelling trust [35, 59] and its applicability to well known problems e.g. the prisoner's dilemma [38]. Finally, reputation systems extend the features of trust models by adding means of eliciting reputation reports from other agents [52, 26].

Current efforts in trust modelling and reputation system development are yet to be sufficient in order to obtain efficient IC DRM's that are readily to be transformed into real-world applications. Niches are still open for mechanisms that may better prevent, detect and discourage defection. This is where scoring rules come into play.

Past years have brought forward relevant examples of scoring rules based mechanisms for promoting cooperation. Kilgour and Gerchak [33] elaborate a competitive payment-system based on centering each forecaster's payment at the average value for the generated forecasts.

More recently, Miller et al. [41] motivate and rigourously define a system using scoring rules that functions with or without the presence of objective results. They prove the economic properties of the mecanism, the way score-based payments correctly incentify expert agents in a) participating in the mecanism, b) investing the right ammount of effort necessary for obtaining a probabilistic estimate of a required precision and c) reporting it truthfully. In absence of an objective value for the modelled parameter, sequential interaction is proposed.

Sequential interaction mechanisms are studied further on by Smorodinsky and Tennenholtz [61] or Teacy et al. [63], their works proving that it is more efficient to use this type of ordered interactions than requesting information from an entire set of agents simultaneously. The list of reseach directions is expanded by studies on the robustness of information elicitation mechanisms [66] or the means of balancing forecast risk against returns so that risk averse forecasters have also the incentive of reporting honestly [32].

In 2008, Papakonstantinou et al. [50] propose a novel, two-stage mechanism for eliciting private information without knowing the exact cost functions of participating agents, this representing a relaxation of assumptions present in [41]. The mechanism involves running an auction for selecting the agent with the lowest cost and then rewarding its service through a score-based payment. The economical properties of the mecanism are proved and the study also compares the three main scoring rules, demonstrating that the logarithmic one has the most desirable properties, but as it is improper in case of very small probabilities, the spherical scoring rule is to be preferred.

2.3.3 Critical view

The common denominator of this body of research is the presence of *central agents* that run the interaction protocol and the use of *payments* for obtaining incentive-compatibility. While relaxing the centralised condition is an objective of distributed mechanism design (DMD), removing the payment system is a well motivated step towards achieving decentralisation.

Taking into consideration all the theoretical benefits of scoring rules, using them for computing payments implies the availability of a payment system with a proper currency definition and a proper consideration towards the financial mechanisms involved. Moreover, a centre is necessary for executing the currency transfers, thus implying that a decentralised setting should move from using payments to using a different mechanism. This work aims at studying the use of *transferable utility* as a replacement for payments, moving away from the implications made by such incentive mechanisms.

In addition, a fair amount of presented mechanisms (see [33, 41]) use scoring rules for dichotomous events (binary variables). Recent trends are moving away from this constraining setups, focusing on the use of continuous distributions as a standard for generating forecasts.

To sum up, there is a need for studying other forms of incentives rather than the consacrated payments, to simplify the elicitation mechanisms and bring them closer to decentralisation.

Chapter 3

Payment-free cooperation

The following chapter details the model of payment-free information elicitation proposed, followed by a mathematical and experimental analysis of the properties yielded by agent interactions. We show that the model leads to a "prisoner's dilemma" and consequently define a mechanism aimed at promoting cooperation through this type of information elicitation.

3.1 Information elicitation scenario

The starting point for this model of payment-free coopeation is the work undertaken by Papakonstantinou et al. [50], work that used score-based payments as incentives for cooperation within a rational environment. In the context of information elicitation, cooperation is understood as truthful reporting of one's private information. From this perspective, the mechanisms that aim at promoting cooperation are the incentive-compatible direct-revelation ones.

In order to transform the above mentioned model into a "payment-free" solution two changes were applied: the concept of a center interested in acquiring probabilistic estimates was removed, and with it the payment-based method of incentivising agents. Accordingly, we are considering a system composed of $Q \geq 2$ rational, risk neutral agents that are able to generate unbiased but noisy estimates $x \in \mathbb{R}$ of precision $\theta \in \mathbb{R}^+$. The estimates are modelled as Gaussian random variables, defined by $x \sim N(x_0, 1/\theta)$, where x_0 is the true value of the estimated parameter, which is considered to be made public at a later moment in time. As an example, in a weather prediction setting the estimates could be temperature forecasts while the true value would be the actual temperature registered at the chosen future date.

Producing estimates is a costly operation, this being captured by a cost function $c(\theta)$. For reasons of simplicity, we consider this function to be linear and identical for each individual. This means that agents will actually incurr costs in the form of $c\,\theta$, where $c > 0$ is a global system parameter. In addition, agents have a utility function $U : \mathbb{R}^2 \times \mathbb{R}^+ \rightarrow \mathbb{R}$ that is composed of the cost incurred by an agent for generating an estimate and the score it could receive by reporting it (eq. 3.1). This function represents a model of utility agents associate to estimates, a prediction of the utility value they could obtain based on what they report. The motivation for such a predictive approach, and not the standard way of subtracting expenses from incomes, is that the mechanism is designed to be domain neutral, so in a real-case scenarion, the utility function can always be adapted to it.

$$U(x_0; \hat{x}, \hat{\theta}) = S(x_0; N(\hat{x}, 1/\hat{\theta})) - c(\theta) \tag{3.1}$$

Here $S(x_0; N(\hat{x}, 1/\hat{\theta}))$ represents the *scoring function* - a scoring rule that associates a score to the Gaussian estimate $N(\hat{x}, 1/\hat{\theta})$ reported by the agent, while θ is the precision used for generating its own measurement.

To conclude, agents can choose to interact with others and exchange their probabilistic estimates or not to interact and use only data from their own measurements. Estimates have the role of *transferable utility* as by fusing them, agents have the opportunity of obtaining higher scores at the same price, due to the increase precision (θ's are additive with respect to fusion).

3.2 Model analysis

Out of the three important scoring rules, the quadratic one (eq. 3.2) was selected for conducting the analysis and studying the model, due to its mathematical simplicity and stability with respect to the range of scores it outputs. As a reminder, the quadratic scoring rule, in our setup, has the formula:

$$S(x_0; N(x, 1/\theta)) = 2N(x_0; x, 1/\theta) - \frac{1}{2}\sqrt{\frac{\theta}{\pi}} \tag{3.2}$$

Considering a system composed of two agents, there can be five interaction scenarios: a) no interaction, b) reciprocal cooperation, c) defecting or d) being defected upon and e) mutual defection. Before proceeding to their analysis, additional formulas and concepts that will be used throughout the remainder of the paper need to be presented.

Firstly, after the interaction of two agents providing two truthful estimates, $N(x_1, 1/\theta_1)$ and $N(x_2, 1/\theta_2)$, each will be in the possesion of an identical and more precise fused estimate $N(x_f, 1/\theta_f)$, having $x_f\theta_f = (x_1\theta_1 + x_2\theta_2)$ and $\theta_f = \theta_1 + \theta_2$. This implies that, even if truthful, agents will not benefit equally from their interactions, and this is part of motivation for developing a mechanism that rectifies this situation.

Secondly, the expected utility of an agent by:

$$\overline{U}(\theta) = \overline{S}(\hat{\theta}) - c\,\theta \tag{3.3}$$

where $\hat{\theta}$ is the precision of the reported estimate that normally differs from θ in the case of fused estimates or defecting agents. The expected score $\overline{S}(\theta)$ can be obtained as in [50], by integrating the scoring function over the expected outcome. For the quadratic scoring rule, its formula is:

$$\overline{S}(\theta) = \frac{1}{2}\sqrt{\frac{\theta}{\pi}} \tag{3.4}$$

Finally, we can compute the optimal precision θ^* that an agent should use, by maximising its expected utility i.e. computing $\overline{U'}(\theta^*) = 0$.

Having laid these foundations, we can continue by detailing the five interaction scenarios.

3.2.1 No interaction (N)

By not interacting, an agent will hold only its estimate of precision θ and the associated utility will follow the formula presented in equation 3.1. In order to maximise it, the optimal precision θ^*_N that the agent should use is:

$$\theta^*_N = \frac{1}{16\pi c^2} \tag{3.5}$$

In this scenario, the maximum expected utility $\overline{U}(\theta^*_N)$ is equal to:

$$\overline{U}(\theta^*_N) = \overline{S}(\theta^*_N) - c\,\theta^*_N = \frac{1}{16\pi c} \tag{3.6}$$

Figure 3.1 presents the relation between an agent's utility and the precision used, when considering costs $c = 10^{-3}$ and $c = 0.3$ respectively. As a default, experimental results are averaged out over 10^5 iterations and plots contain both analytical and experimental values, if not stated otherwise. Also we will use $c = 0.3$ as the standard for the following plots.

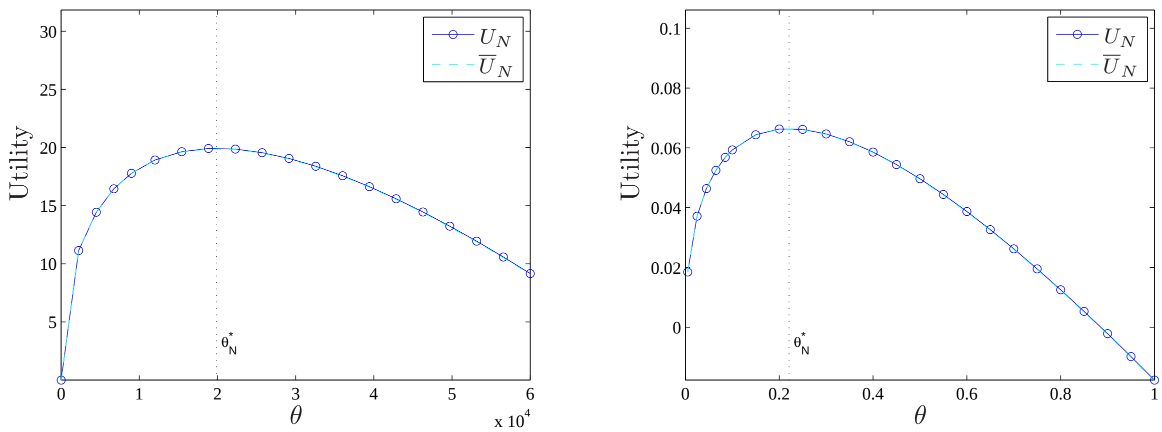

(a) Utility when not interacting ($c = 10^{-3}$)

(b) Utility when not interacting ($c = 0.3$)

Figure 3.1: Utility when not interacting

It can be observed that the cost function influences the optimal precision an agent should use but not the fact that the utility function has a unique maximum, $U(\theta^*_N)$. In addition, the utility falls to zero quicker when using precisions lower than the optimal one, rather than higher.

3.2.2 Mutual cooperation (CC)

In this scenario a cooperating agent generates estimates of precision θ, that are then exchanged with other similar cooperators. Through these interactions, agents obtain fused estimates of double the initial precision $\theta_f = 2\,\theta$. The utility function will be then:

$$U_{CC}(\theta) = S(2\theta) - c\,\theta \tag{3.7}$$

Based on the previous formula, the optimal precision θ^*_{CC} in case of mutual cooperation is:

$$\theta^*_{CC} = \frac{1}{8\pi c^2} \tag{3.8}$$

Comparing equations 3.5 and 3.8 we can observe that $\theta^*_{CC} = 2\,\theta^*_N$, which is an interesting property of the mechanism: **cooperation provides incentives for using double the precision of individualism**. The maximised expected utility will also be double the one obtained by not cooperating:

$$\overline{U}(\theta^*_{CC}) = 2\,\overline{U}(\theta^*_N) = \frac{1}{8\pi c}$$

Figure 3.2 presents the expected and experimental utility values obtained by agents with respect to their measurment precisions, in the case of mutual cooperation. The same trend is observed as for U_N: using lower precisions is potentially more harmful than slightly higher ones.

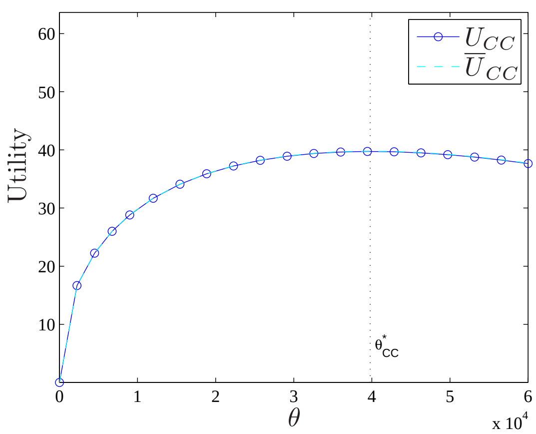

Figure 3.2: Utility of mutual cooperators

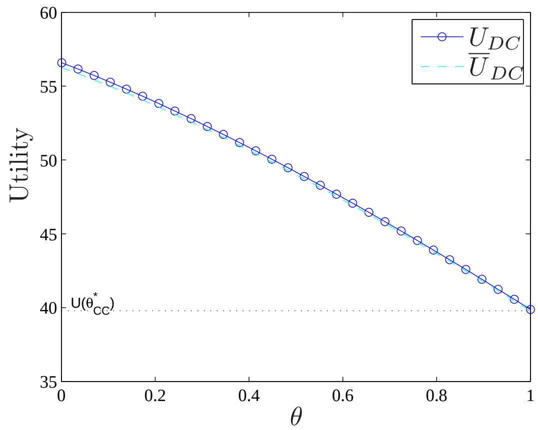

Figure 3.3: Utility of defecting a cooperator

3.2.3 Defecting against a cooperator (DC)

Defection is generally understood as a situation where an agent missreports its estimate by under or overreporting the precision it used. As an example, because agents exchange their estimates before finding out the true value x_0, defectors could choose to generate random Gaussian distributions instead of investing resources for doing measurements, but report them as being of high precision.

To better describe defection, a new parameter has been introduced: $\alpha \in [0,1]$. We consider it the "degree of defection" or the defection coeficient. Consequently, each agent's precision can be expressed as a percentage of θ_{CC}. In the DC case, $\theta_{DC} = \alpha\theta_{CC}$.

Figure 3.3 presents expected and actual utility values obtained by defecting agents, when they are interacting with cooperating ones. It is clear that they obtain utilities higher than cooperators, utilities that gradually converge towards $U_{CC}(\theta^*_{CC})$ as α becomes closer to 1.

We conclude this scenario by mentioning that the defector's expected utility has the following experimentally confirmed formula:

$$\begin{aligned}\overline{U}_{DC}(\alpha) &= \frac{1}{2}\sqrt{\frac{(\alpha+1)\theta_{CC}}{\pi}} - c\,\alpha\theta_{CC} \\ &= \frac{\sqrt{2(\alpha+1)} - \alpha}{8\pi c}\end{aligned} \tag{3.9}$$

3.2.4 Cooperating with a defector (CD)

In this scenario, the cooperating agent believes it is receiving an estimate of precision θ_{CC} and fuses it accordingly. However, its counterpart may just pretend it is using θ_{CC} but instead adopt a different precision, $\theta_{DC} = \alpha\theta_{CC}$. Such a situation may have a strong impact on the cooperator's utility, as shown in figure 3.4.

The precision of the cooperator's fused estimate is dependent on the degree of defection α of its interaction partner, dependency presented in figure 3.5 from both an analytical and experimental point of view ($\theta_{CC} = 1$). Derivation of the fused estimate's precision θ_f can be obtained in the following way:

$$\begin{aligned} x_1 &= N(x_0, \theta) & \sigma_f^2 &= \frac{1}{4}(\frac{1}{\theta} + \frac{1}{\alpha\theta}) \\ x_2 &= N(x_0, \alpha\theta) & \theta_f &= \frac{4\alpha}{\alpha+1}\theta \\ x_f &= \frac{x_1+x_2}{2} \end{aligned} \tag{3.10}$$

As undetected defection alters the precision of a cooperator's fused estimate, the expected score is not reflected anymore by the initial formula (eq. 3.4). The correct form is derived by reintegrating the scoring function over the distribution of the actual outcome. If the agent believes it holds an estimate of precision θ_b and reports it, while in reality the precision was θ_f, then the expected score will be:

$$\overline{S}(\theta_b, \theta_f) = 2\sqrt{\frac{\theta_b\theta_f}{2\pi(\theta_b+\theta_f)}} - \frac{1}{2}\sqrt{\frac{\theta_b}{\pi}} \tag{3.11}$$

In the CD scenario, a cooperator's expected score and utility will have the following formulas:

$$\overline{S}(\theta_{CC}) = 2\sqrt{\frac{2\alpha\theta_{CC}}{(3\alpha+1)\pi}} - \frac{1}{2}\sqrt{\frac{2\theta_{CC}}{\pi}} \tag{3.12}$$

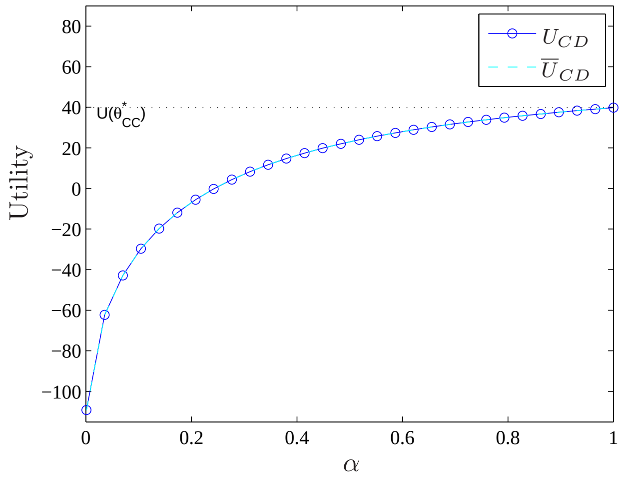

Figure 3.4: Utility when being defected upon

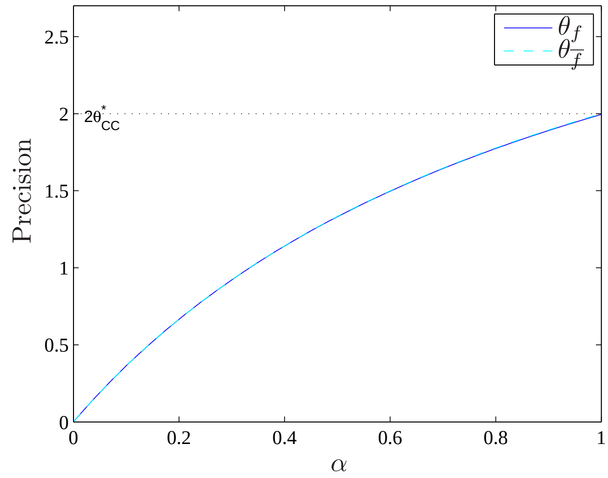

Figure 3.5: Fused estimate's precision

$$\overline{U}_{CD}(\alpha, c) = 2\sqrt{\frac{2\alpha\theta_{CC}}{(3\alpha+1)\pi}} - \frac{1}{2}\sqrt{\frac{2\theta_{CC}}{\pi}} - c\,\theta_{CC} \tag{3.13}$$

3.2.5 Mutual defection (DD)

Reciprocal defection is encountered when both agents use precisions in the form of $\alpha\theta_{CC}$. If we were to consider the simple case where α values are identical, the precision of a fused estimate will not be $\theta_f = (\alpha+1)\theta_{CC}$ as believed by each of them, but:

$$\theta_f = \frac{\alpha(\alpha+1)^2}{\alpha^2+1}\theta_{CC} \tag{3.14}$$

Respecting equation 3.11, the expected score and utility in the DD scenario will be:

$$\overline{S}(\theta_{CC}) = 2(\alpha+1)\sqrt{\frac{\alpha\theta_{CC}}{2\pi(2\alpha^2+\alpha+1)}} - \frac{1}{2}\sqrt{\frac{(\alpha+1)\theta_{CC}}{\pi}} \tag{3.15}$$

$$\overline{U}_{DD}(\alpha, c) = 2(\alpha+1)\sqrt{\frac{\alpha\theta_{CC}}{2\pi(2\alpha^2+\alpha+1)}} - \frac{1}{2}\sqrt{\frac{(\alpha+1)\theta_{CC}}{\pi}} - c\,\alpha\theta_{CC} \tag{3.16}$$

Figure 3.6 presents analytical and experimental results for the utility obtained in case of mutual defection case, while figure 3.7 depicts the relation between the defection coeficient α and the precision θ_f of the fused estimate ($\theta_{CC} = 1$).

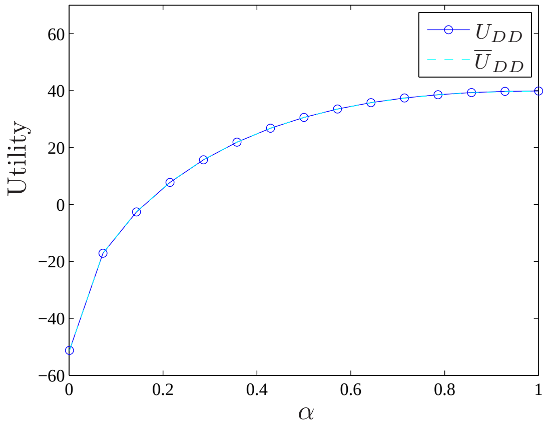

Figure 3.6: Utility of mutual defection

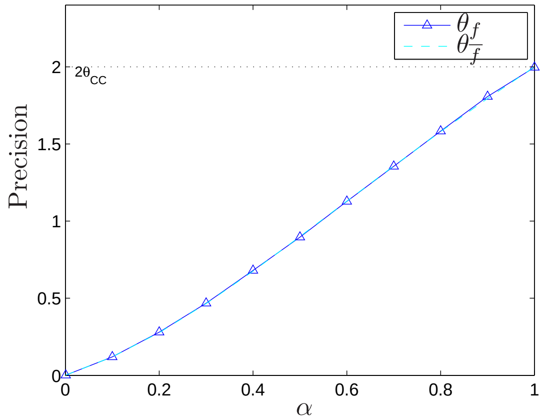

Figure 3.7: Fused estimate's precision

3.2.6 Prisoner's dilemma scenario

Comparing the utilities obtained in each of the five interaction types, we observe that **the model leads to a classical "prisoner's dilemma" scenario** (PD), defection yielding higher utilities than cooperation. In other words, $U_{DC} \geq U_{CC} \geq U_{DD} \geq U_{CD}$. To exemplify this, we plot agent utilities with respect to the defection coeficient α and present them in figure 3.8.

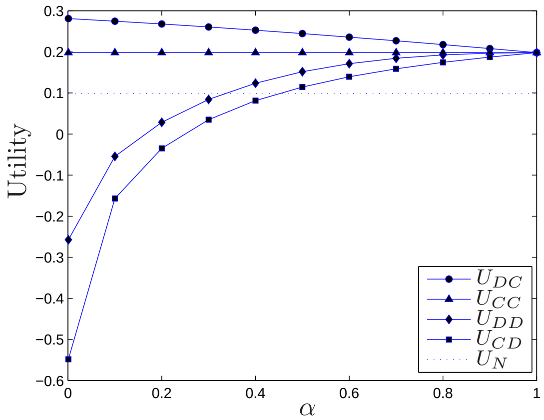

Figure 3.8: "Prisoner's dilemma" setting

Reaching a "prisoner's dilemma" situation has important implications over the way the current study will progress. Firstly, it objectifies the need for developing a mechanism that may effectively promote cooperation based on this model of payment-free cooperation. Secondly, it connects this work to the PD and IPD literature, providing thus a relevant candidate for solving the MD problem: *reciprocity*. If Axelrod's tournament [3, 2] proved that simple strategies like TIT FOR TAT can overcome more complex ones, although they are not "invincible" [5, 13], it shows that reciprocity is both closer to reality and advantageous in handling the IPD problem. For more information on PD, Rapoport and Chammah's 1965 book [53] is a good reference, while studies on the IPD problem [4, 37, 22, 11] together with a review from Hoffman [24] provide further details on more recent results.

In our setting however, defection is not as easy to detect as in the "prisoner's dilemma". If in that case, defection was a dichotomous event having a direct public consequence, here its effects come in degrees and shades. As we rely on continuous distributions for representing an agent's information, we also base the evaluation of an estimate on the continuous score outputted by the scoring rule. The score always reflects quality or precision of the estimate, but also the unavoidable noisiness that affects the agent's measurements. Consequently, approximating α is the proper way of detecting defection and this leads to it being binary classified into: pure cooperation for $\alpha = 1$ and defection for $\alpha < 1$. To better express the degree of defection, we are going to use throught the rest of the paper the terms: *strong* ($\alpha < 0.3$), *intermediate* ($\alpha \in [0.3, 0.6)$) and *mild* or *soft* ($\alpha \geq 0.6$) defection.

Finally, we note that from a specific value of α, cooperating is more profitable than being isolated, even in case of being defected against. Consequently, interacting (indiferent of cooperation or defection) is a dominant strategy for values of α grater or equal to the one where $\overline{U}_{CD}$ and $\overline{U}_N$ meet. This is independent of any other variables and it has the value $\alpha \approx 0.4$.

3.3 Payment-Free Mechanism

Following our analysis, we define a new mechanism that aims at promoting cooperation without using payments. As our problem is of the "prisoner's dilemma" type, the solution proposed by the mechanism is that of implementing *reciprocity*. In this study, reciprocity stems from adding *trust models* at the agent level and a *pair-matching service* (PMS) at the system level. While trust models allow agents to better appreciate the truthfulness of their counterparts, the PMS has the role of increasing the probability of cooperator-cooperator interactions and decreasing that of defectors interacting with cooperators.

Trust modelling relies on estimate exchanges between agents. Each estimate speaks about the owner's truthfulness in the sense that it can be compared to the true value (obtained later) and its probable α calculated. In our case, we base our evaluation on the distance D between an agent's own estimate and the ones it receives. As it knows the α used for its measurements, the distance D allows it to compute the probability distribution of the α used by its counterpart. Trust updates are done through the classic Bayesian rule:

$$P(\alpha|D) = \frac{P(D|\alpha)P(\alpha)}{\sum_k P(\alpha_k)} \tag{3.17}$$

In support of calculating $P(\alpha|D)$, if we consider two agents i and j exchanging estimates $N(x_i, 1/\alpha_i\theta)$ and $N(x_j, 1/\alpha_j\theta)$, then:

$$P(D|\alpha) = N(|x_i - x_j|; x_0, \sigma_D^2) \qquad \text{and} \qquad \sigma_D^2 = \frac{1}{\alpha_i\theta} + \frac{1}{\alpha_\Omega\theta}$$

where $\alpha_\Omega = (0, 1]$ represents the range of α value.

Figure 3.9 shows how the model adapts according to the number of encounters between a cooperator and a defector having $\alpha = 0.3$. In order to obtain the value of α, an agent computes the aggregate $\dot{\alpha} = \alpha_\Omega P(D|\alpha)$ as it is more risk neutral than just using the maximum a posteri.

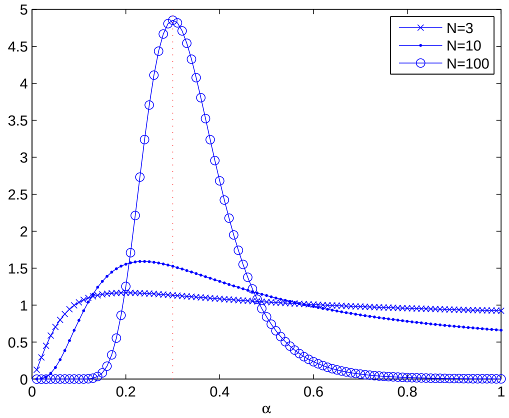

Figure 3.9: Trust convergence

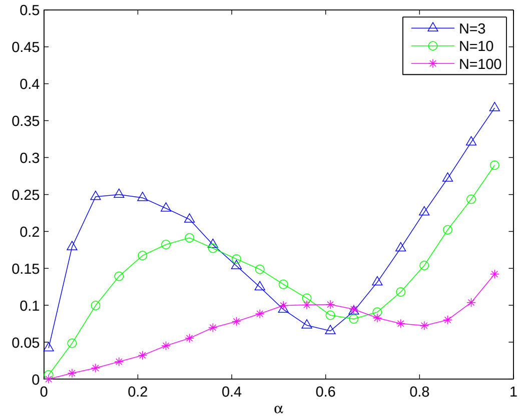

Figure 3.10: Difference between α and $\dot{\alpha}$

This being said, and although the model shows a good approximation, we find that it doesn't converge with the same speed for all values of α. Figure 3.10 underlines this issue by plotting the difference between the true values of α and the approximated ones, over different number of rounds. As shown, for $\alpha < 0.1$ and $\alpha \in (0.5, 0.7)$ the model is quick to adapt, while for the other values it needs over 100 rounds to become more accurate. While choosing this estimation technique, we state the need for further studies using different methods.

Based on the agent encounters, they will now be able to develop preferences stating their interest in interacting with the available agents. Each agent holds a set of $\dot{\alpha}$ values associated to its potential partners and by sorting them descendingly, preference lists may be obtained. These are submitted then to the pair-matching service, implemented here according to Irvin's 1985 [27] method of solving the "stable roommates" problem. As a result, a set of stable pairs is obtained, this guaranteeing that each agent gets to interact with the first available agent on its preference list.

A consequence of modelling the α value of the various interaction partners is that it may use now the computed $\dot{\alpha}\theta_{CC}$ instead of the standard $\theta_C C$, when fusing estimates. In other words, we have two types of fusing strategies: the naïve one, when agents will always fuse as they would be interacting with a cooperator and the adaptive one, where $\dot{\alpha}$ will be used.

The mechanism runs as a synchronous network, interactions taking place in series of games divided in $N \geq 1$ rounds. During a round, each agent may access the mechanism once or choose not to interact. The sequence of actions taken during a round is as follows:

a) each agent reports its set of preferences to the PMS
b) the PMS outputs a list containing the agent pairs, which is guaranteed to be stable. If necessary, the two least preferred agent will be iteratively removed from the pair-allocation process, until stable pairs are obtained
c) agents generate their estimates

d) agents that formed pairs exchange their estimates while the ones that weren't matched use only their personal estimate

e) trust models are updated and the preferences sorted accordingly

3.4 Methodology

Following the removal of financial incentives, the main research question that is to be addressed is:

To what extent can cooperation be promoted in a decentralised information system, through the use of a scoring rules and reciprocity?

We aim at identifying if the mechanism is *effective* and if so, under what other conditions. That is, delimitating the set of assumptions and parameter values that allow cooperation to be promoted and then concentrating on the possible causes for this not to be achieved.

In undertaking this research we propose using a simulation based study starting from very simple setups, where parameters are completely predefined, and advancing towards more realistic scenarios, allowing a higher degree of randomness. The central metric that we require is the average agent utility, metric that will help compare pure cooperation with other degrees of defection.

As the mechanism is domain neutral, concepts like utility or number of interactions are not mapped to a range or acceptable values and are thus harder to label. If we are referring to the convergence time, $N < 20$ might be a requirement in one field while in others $N > 1000$ is acceptable. Taking into consideration the trust model used, we adopt $N = 100$ as our reference, and throughout our experiments we will evaluate the effectiveness of the mechanism in this given "time". If defection will still dominate, we will investigate the average N value that allows cooperation to be more profitable.

Finally, we group our experiments into two phases. The first phase will use homogeneous populations (identical α's) and random interactions, thus allowing us to calibrate our metrics. Following it, we will use "polarised" populations, composed by pure cooperators and defectors of various types and the SRA service. This setup will provide further results as to how well cooperators can cope, when they are in contact with defectors. The second phase will use agents of random types, providing a more realistic ground for assessing the mechanism. Results will be cross-analysed via evolutionary simulations.

Chapter 4

Findings

In this chapter we present the results obtained from testing the mechanism. As previously stated, we followed a two-stage plan in running our experiments. The first stage aimed at calibrating the metrics and providing a general view over the way the mechanism works under close to ideal setups. Following this phase, simulations involving more realistic conditions were taken, in order to validate the degree to which cooperation can be promoted.

As constant parameters we used $c = 0.1995$ which then implied $\theta^*_{CC} = 1$ and $\theta^*_N = 0.5$. Variable were the number of agents Q or rounds per game N, the distribution of α values within the population and the type of estimate fusion used (naïve or adaptive). Each simulation was repeated 2000 times so as to have a minimal statistical error (less than 10^{-4}).

4.1 Simulations under strong assumptions

The first set of simulations is based on the use of simple settings that should facilitate the achievement of the mechanism's objective: making cooperation the dominant strategy. We introduce then a group of pure cooperators into a population of homogeneous defectors and observe how well the mechanism functions. As defection coeficients, α values ranging from 0.1 to 0.9 are used, as these two extremes are sufficient for allowing us to make claims.

Prior to these tests, we run a set of simulations with homogeneous populations that will allow us to calibrate the values yielded by using naïve and adaptive fusion. Finally, as a constant throughout this first phase of simulations, the population size was $Q = 10$ while the number of rounds per game was $N = 100$.

4.1.1 Homogeneous populations

Considering our population formed of agents with identical types (α), we evaluate how their utilities are influenced by the way estimates are fused. "Stable roommates" allocation is not employed here, random pairing being used. As agents provide overall the same quality of estimates, there is no need for generating preferences. Figure 4.1 shows the average results obtained after 100 rounds.

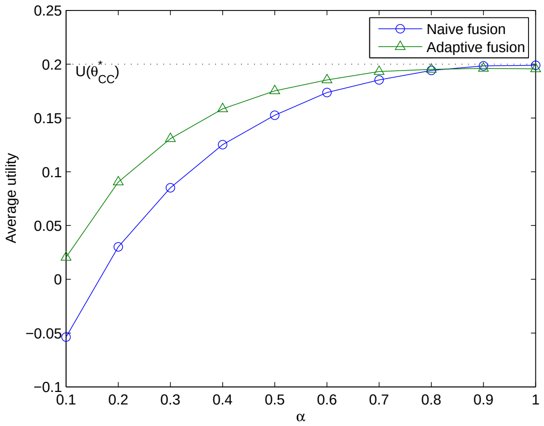

Figure 4.1: Influence of estimate fusion technique

It can be observed that adaptive fusion has a beneficial impact, especially on defectors. If previously, low value α's would have brought them negative utilities, now they can obtain profit from using even the highest levels of defection. This situation is motivated by the deminished penalty they incurr. If previously their reports would have presented overinflated precisions $\hat{\theta}_i = (\alpha_i + 1)\theta^*_{CC}$, now they report less so they lose less. The results can be extrapolated to cooperators, receiving lower penalties from their inevitable interactions with defectors, but this remains to be tested in a more complex setup.

This parameter clarified, we continue by evaluating the way pair allocation influences the overall results.

4.1.2 Pair allocation influence

This series of tests uses a polarised population composed of pure cooperators and strong defectors ($\alpha = 0.1$). The agent numbers are biased so that cooperators represent $P_C = 60\%$ of the whole population, this parameter being in turn studied in later experiments. By default, agents use naïve fusion.

We run the mechanism using random interactions and separately SRA, and contrast their results. Through random interactions, a defector has approximately P_C chances of encountering a cooperator, so overall, defectors should obtain higher utilities. We prove this by simulating and show the results in figure 4.2. As expected, defectors easily exploit cooperators, while these are worst off than if they wouldn't have interacted at all. The plotted $U(\theta^*_N)$ facilitates this comparison. In addition, the defector's average utility is approximately 40% lower than the optimal one in the CD case (section 3.2.3), due to the inter-defector interactions.

Figure 4.3 shows the results from the SRA tests. In contrast to the random allocation, cooperators are better off than the defectors, and there is a clear separation between them. The mechanism is effective in this conditions, in spite cooperators don't reach the expected utility value for θ^*_{CC}.

The fact that defectors obtain negative utility stems from the both the mechanism separating them from cooperators and from their naïve way of fusing estimates.

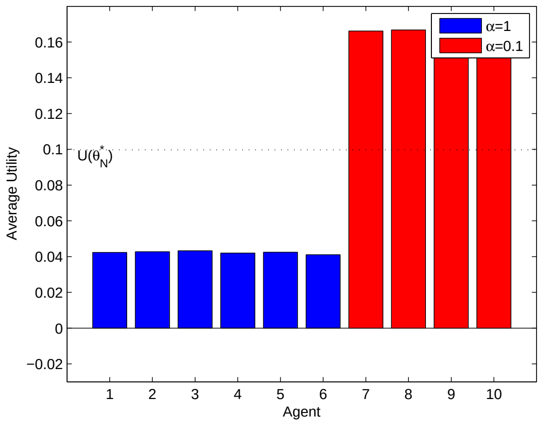

Figure 4.2: Average utilities through random pair allocation

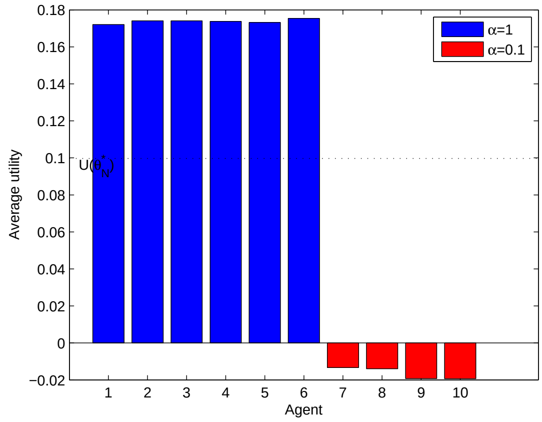

Figure 4.3: Average utilities through stable roommate pair allocation

Up to this stage we can say that results are overall motivating. There are issues that need to be raised though. First, we used a very low defection coeficient ($\alpha = 0.1$). It would be relevant to see under what circumstances cooperation is still more profitable. Secondly, the agents applied naïve fusion which is not a very realistic assumption and also implies a loss of utility from both parts. Finally, the type of the exploration-exploitation trade-off plays a role too in the obtained results. The strategy used so far was purely greedy, but a method like ϵ-Greedy should be more efficient for setups of higher complexity.

We are going to continue then by extending the SRA tests for the use of other α values and see up to which point can the defectors be isolated and cooperation be promoted.

4.1.3 Isolation of defectors

The main objective addressed by this sequence of tests is to identify the maximum α up to which pure cooperation is more profitable than any degree of defection. To do so, we test the possible combinations of cooperators and defectors, with α's ranging from 0.1 to 0.9. Figure 4.4 shows the results, contrasting cooperators with defectors and the results from the homogeneous population tests.

It can be observed that for $\alpha < 0.4$ cooperators obtain higher utilities while for $\alpha \geq 0.4$ defectors have the advantage. Accordingly, the value of α we were looking for is approximately 0.4. This reminds us of the $\alpha \approx 0.4$ from section 3.2.6, where cooperation was preferred to not interacting at all, even when being defected upon. We could translate this into: *from the moment interacting becomes profitable, defection becomes dominant* or *interaction leads to a "prisoner's dilemma" situation.*

We can then say that, under this setup, the mechanism is efficient against strong defection, then gradually becomes permeable to more intermediate forms until it fails after $\alpha \geq 0.4$.

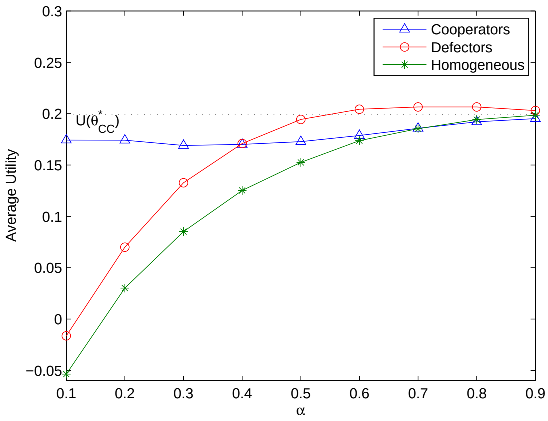

Figure 4.4: Average utilities using naïve fusion

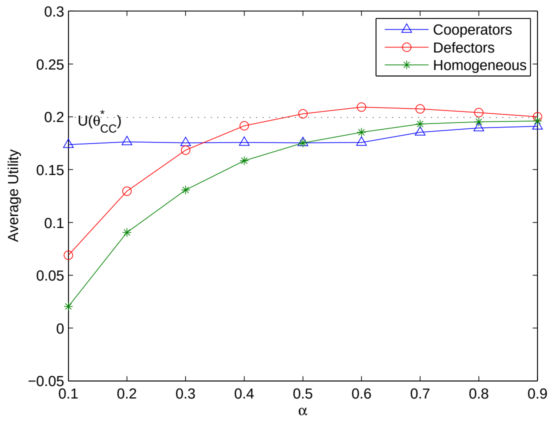

Figure 4.5: Average utilities using adaptive fusion

The difference between the defectors' values and those from the homogeneous population indicate how much they are able to exploit cooperators. More properly speaking, how much utility they gain from encountering cooperators. In turns of number of interactions between pure cooperators and other types of defectors, we find that the value increases from 4% in case of $\alpha = 0.1$ up to 20% for $\alpha = 0.7$, after which fluctuations are minimal. Figure 4.4 shows actually that from $\alpha = 0.7$ cooperators obtain the same utilities as in the homogeneous case, formed only by defectors.

Having clarified the issue of incentive compatibility under this (naïve) setup, we also note the positive facts, before moving on to the next set of experiments. Thus, the utility of the cooperators never falls under 85% of the maximum value and the mean utility over the range of α's is approximately 90% of this maximum. So naïve cooperators risk of obtaining utilities on average 15% lower, situation still better than the analytical one presented in section 3.2.4.

4.1.4 Defector's dominance

The above dominance of defection underlines the effect cooperator-defector interactions have. Such an outcome seems not to be caused by a single factor. Since the analysis of the mechanism, there were signs of nonlinear approximation of α's through the use of the aggregate method. This influenced the number of interactions necessary for obtaining accurate estimates and when considering that agents had $N = 100$ rounds at their disposal, we have more question marks raising. Accordingly, we hypothesise that defection becomes more profitable due to one or a combination of:

a) agents' naïve way of fusing estimates
b) the number of rounds within a game being too low
c) trust modelling inefficiencies
d) exploration-exploitation trade-off suboptimal

Figure 4.5 shows the results obtained by testing the first hypothesis. Respecting the findings from the homogeneous population tests, we can make two observations: firstly, the defectors' average utility matches the one of the cooperators' earlier, from $\alpha > 0.3$ and secondly, from $\alpha = 0.6$, cooperators optain lower utilities than if they would have defected among equally defecting agents. Fusing adaptively gives a slight advantage to defectors, so naïve fusion is not a valid hypothesis.

For our second hypothesis, we used games of 100 up to 10000 rounds. We find that, while for defection coeficients of 0.4 and 0.5, cooperation becomes profitable after 200 respectively 500 rounds, for higher α's, $\mathcal{O}(N) = 10^3$ is a minimum value. In addition, even when using 10000 rounds, pure cooperators could not obtain an average utility higher than $\alpha = 0.9$ adaptive defectors. This means that this hypothesis cannot be fully validated either. It gives us though more insights into what actually happens.

Because agents start with no measure of uncertainty about their potential counterparts, their trust models are defined by a flat distribution with mean 1 and variance 0, while the associated $\dot{\alpha}$ value is also 1. No matter what kind of agent they encounter, the first interaction will lead to an updated value $\dot{\alpha} < 1$, that will cause them to change priorities until they would have interacted at least once with all the other agents in the system. So after the first $Q - 1$ rounds, all the agents would have met once. This is where the hazardous situation starts. Because the trust models are not accurate after one interaction, preferences may see cooperators opting for interacting with medium or even strong defectors. The game continues and trust models update until new preferences are formed, at least from the part of cooperators who "realised" they were being defected. However, by this time, the other agents will have stabilised their $\dot{alpha}$ values so the formation of new and stable pairs of cooperators is highly improbable.

To have a better understanding of the degree to which the use of trust models influences the observed outcome, we return to the third hypothesis and test it. We do this by making the α values public. Agents fuse correctly the received estimates, so in this context the role of the trust models can be properly differentiated. For this, we compare the average utilities obtained by agents who generate their preferences directly on the public α's (C_p and D_p) with those resulted from the use of trust estimates (C_t and D_t). Results are shown in figure 4.6.

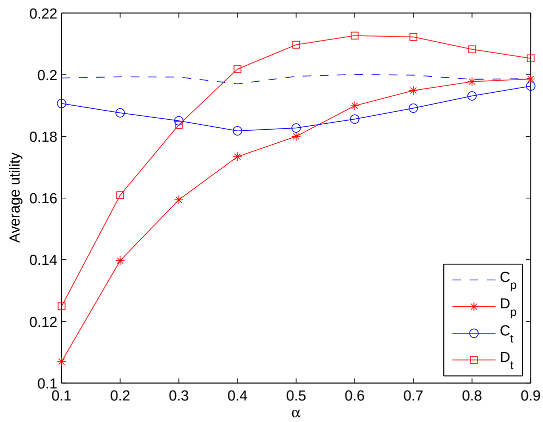

Figure 4.6: Influence of trust model over the average utilities

It can be observed that, with perfect awareness of α values, cooperators have always a higher utility, while when using the trust models, the trend observed in previous tests is repeated. In other words, the SRA cannot do more than the agents themselfs trust it would be good for them. Even more, when an agent changes its preferences due to the decrease of $\dot{\alpha}$ values (defection), its other preferred partners are already undergoing stable interactions with other agents.

To sum up, considering a population formed of two types of agents - cooperators and defectors, the mechanism is able to promote cooperation after a small number of encounters ($N \leq 100$), if $\alpha \leq 0.3$, while for $\alpha > 0.3$ this outcome is achievable up to a certain point, with the price of N being one order of magnitude higher.

The main findings with respect to this payment-free mechanism can be considered reached and the chapter concluded. We provide though a listing of the additional experiments made, considering they will stimulate further research or even provide insights into possible solutions making the mechanism incentive-compatible.

4.2 Additional experiments

In this section, we will only summarise the additional experiments we undertook, in order to obtain more insights into the possible solutions for promoting cooperation and also to evaluate the mechanism within a more complex setting.

The subject of exploration-exploitation trade-off was first to be addressed. As there was no special strategy implemented in the agents, in order to have them testing other agents in the later stages of the game too, not only in the first Q rounds, we added a ϵ-Greedy feature to the agent's specifications. In more detail, with a rate $r = 0.5\,Q$ we had agents update their $\dot{\alpha}$ values to 1. The result was that unilateral changes in preference were made, so the SRA prevented new beneficial pairs from efficiently forming. More than that, defectors had more chances of interacting with cooperators, due to the sudden "change of mind" from their part. A different version of ϵ-Greedy was used, now introducing an empiric protocol for making two agents change their preferences towards forming a pair. The result was that overall, all agents obtained higher utilities and there was a higher variation in the formation of pairs, but the main objective of promoting cooperation was not achieved.

Another series of tests concerned the use of populations with more agents and a higher variation of α values. Specifically, we used $Q = 30$ and we generated α's from a normal distribution. Adaptive fusion was again contrasted to the naïve one but overall results showed that a mild degree of defection lead to higher utilities than pure cooperation. We lowered the maximum value of α from 1 to 0.8. This lead to a similar situation where agents using $\alpha \approx 0.7$ to be most profitable. Figure 4.7 exemplifies the obtained results.

In order to observe the general evolution of such a setup, we undertook an evolutionary approach. Preliminary tests showed that naïve fusion induced incentive-compatiblity under the strong assumption of $\alpha < 0.2'$, whereas adaptive fusion and populations with higher defection coeficients lead to a constant decrease of the mean α. Thus, defection is again confirmed as dominant. Figure 4.8 provides a view over the evolution of the mean α when using such a random population, obtained through the use of two standard selection methods, roulette wheel selection (RWS) and stochastic universal sampling (SUS).

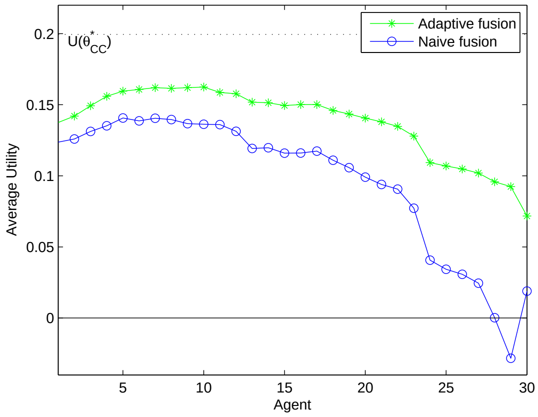

Figure 4.7: Average utilities in a random population

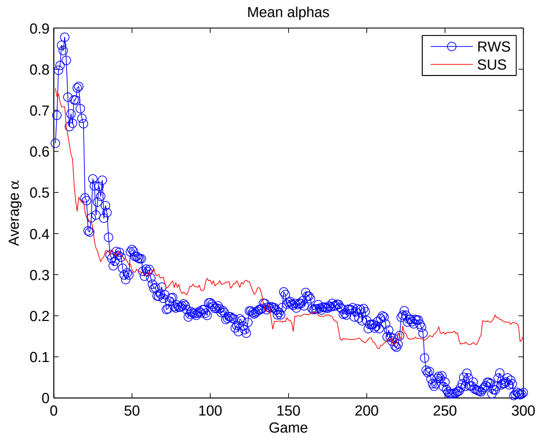

Figure 4.8: Evolution using RWS and SUS selection methods

4.3 Practical considerations

There are a couple of issues that should be discussed here. The first is related to the SRA's efficiency while the second considers the parameters of the implemented simulation.

The SRA is a centralised method for generating stable pairs of agents. Thus, it is bound to have scalability and efficiency issues. From the scalability point of view, we look at the number of preference lists based on which the SRA could not generate a stable solution. Our first tests using $Q = 10$ show that, when using a polarised population, the closer we get to the critical α where defection becomes dominant, the higher the number of unstable solutions generated. Figure 4.9 shows the detailed dependancy of these parameters, scaled at the value of $N = 100$.

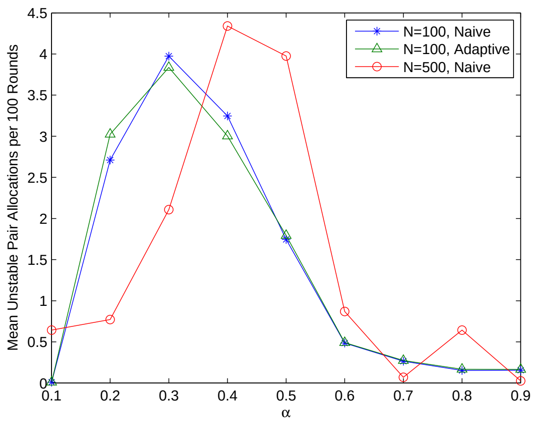

Figure 4.9: Unstable pair allocations

According to these findings, naïve or adaptive fusion lead to similar numbers of unstable allocations, while in the case of $N = 500$, where the values are 5 times higher than the plotted ones, $\alpha = 0.5$ is where the maximum is reached. When considering the second setting with $Q = 30$ and $\alpha \in (0, 1)$, we find that the mean number of unstable allocations is significantly higher, the rule of thumb saying "complexity is problematic" being respected.

With respect to efficiency, a recursive implementation of SRA leeds to a overwhelmingly high number of function calls that become problematic from $Q = 100$. The algorithm is composed of two phases and the second one is where the problem arises. In case a non-recursive strategy is adopted, the algorithm should be more scalable.

Concerning simulation times, in our case, using Matlab on a AMD X2 Turion machine with 1.8 GHz processor frequency and Windows Vista as operating system, our tests required:

Q	$N = 100$	$N = 200$
10	20s	40s
20	45s	1m 43s
30	1m 25s	2m 35

Table 4.1: Average simulation times depending on N and Q

Chapter 5

Evaluation

In this chapter we analyse the degree to which our research goals were answered, taking into consideration both the model of payment-free cooperation and the mechanism proposed for promoting it.

5.1 Model for payment-free cooperation

The main factor that is behind the development of this model is the existence of a informed community, who's members can benefit from sharing information amongst them in a truthful way. By not addressing the subject of information acquisition, we focused mainly on the holders of information as part of a distributed information system. The gained advantage is that within such a system, a natural way of obtaining information is by trading for it one's personal data, both parts benefiting from this.

The model associates agents an utility function U (eq. 3.1) that describes how an individual models the profit it would obtain by delivering information to a client (not present in the model). While U is composed of two terms of different natures, a score and a cost, the score can always be scaled and transformed into a payment [41, 50], thus U is just a simplified and generic version of a utility function that can be adapted to any domain of implementation.

An analysis of the types of agents and their interactions is provided. Important outcomes are the values for the expected utilities and optimal precisions, depending on the type of interaction taking place. The analysis shows that using the quadratic scoring rule is an important motivation for engaging in interactions: the optimal precision is double the one for acting solitary, the utility being double too.

The incentive for free-riding or defecting is also addressed. The main result is that, when using the quadratic scoring rule the model leads to a "prisoner's dilemma" situation (section 3.2.6). The cause for this situation relies on the system itself, as it brings agents of different types, deceitful or honest, together. This is why a mechanism for running the model is required.

To conclude, the model's analysis provides the necessary data for building a mechanism capable of using it and promoting cooperation. As improvements, adressing in greater detail the mutual defection case, where agents have different α's (see B.1) and extending the analysis of the mathematical equations in order to identify more critical points, will provide further research ground.

5.2 Mechanism for promoting payment-free cooperation

Overall, the proposed mechanism manages to promote cooperation but under strong assumptions. We layed efforts in order to make pure cooperation become the dominant strategy, and in certain theoretical setups we achieved this goal. Section 4.1 details these settings and shows that, through the use of this mechanism, pure cooperators are still better off than strong defectors ($\alpha < 0.3$), even if not obtaining their optimal utility.

The PFM struggles however to handle intermediate and mild defection. We can address this situation by letting agents undergo more interactions, allowing pure cooperators to obtain better results this way. There are values of defection though that are still unsolvable i.e. $\alpha \geq 0.9$. Noise levels coupled with the trust modelling technique used, play a great role here. In addition, the "stable roommates" pair allocation method provides a measure of stability, which in some cases becomes too rigid for allowing a cooperator to revert to more profitable interactions, after being paired with a defector.

The situation is of course more complex and there are relevant reasons for believing that payment-free cooperation is not far from reality. This first body of work leaves an open door for expanding the notion of adaptiveness within the mechanism. Here for example we tested variants of ϵ-Greedy exploration and a preliminary protocol that gave more flexibility to the pairing process. Defection was still dominant, but this flexibility could lead effectiveness.

Finally, tests using randomly distributed α values showed that agents that are most truthful, be it pure cooperation or not, obtain lower utilities the the relative mild and intermediate defectors. This leads to a collapse of the mean α value in the evolutionary simulations. So the mechanism is not evolutionary stable, at least under the assumption of $N \leq 500$.

To conclude, the work presented provides a good model for obtaining payment-free cooperation, but for this to come in pure forms, a more refined mechanism that benefits from adaptiveness and decentralisation, should be developed.

Chapter 6

Conclusion

In this paper we tried to set the direction for a new wave of researches in the domain of information elicitation. In an era of decentralisation, pervasiveness and ubiquity, we consider centralised mechanisms for eliciting information both fundamental and out of phase with the current high requirements. By moving away from the classical economic way of using financial incentives, we take a new step towards decentralisation.

The proposed model relies on the existence of a community of information providers that are able to generate probabilistic estimates or forecasts of upcoming events. As the reward an agent might obtain depends on the quality or precision of such estimates, we focus on how to increase theis feature and not on the way to transact them. The precision is what dictates the potential value of information. Accordingly, agents are allowed to exchange estimates between themselfs in order to improve their precision, without increasing their costs. An estimate once generated, it can be used for obtaining n other estimates without incurring further costs. The model is said to be payment-free as it motivates the elicitation of private information without the provision of financial incentives, the actual estimates being the incentives for interaction.

Our analysis shows that the model features important properties. Firstly, it motivates agents who choose to cooperate, in generating estimates of double the precision of those who opt for not interacting. This leads to cooperators gaining double the utility of individualistic agents. Secondly, defection, taking the form of missreporting one's precision, is more profitable than cooperation, this leading to a "prisoner's dilemma" situation. In order to prevent this and to promote cooperation, we develop a simple mechanism that does not rely on payments.

We enable agents to approximate the other's honesty through the use of Bayesian trust models and then to generate preference lists that are used by a pair-allocation service for computing optimal pairs. After starting with random preferences, agents begin converging on a set of preferred partners, this causing the pairs outputted by the pair-allocation service to also stabilise.

The mechanism is tested in different conditions, using populations ranging from small and strictly defined to bigger and more complex ones. Simulations show that the mechanism is effective in the first case, while when dealing with soft defection, it is not able to prevent cooperators from interacting with defectors. Approximating the agents' honesty requires a number of interactions taking place, number that is not linear over all the defection types. Consequently, pure cooperation is harder to detect than strong defection and this is why coalitions of pure cooperators have low chances of forming in populations with a high variety of defectors.

The final result is that soft defectors exploit cooperators, the chain continuing downstream. Separation still exists between pure cooperators and strong defectors, but the overall dominant strategy is defection.

6.1 Future work

Both the proposed model and mechanism for payment-free cooperation enable and require further developments in order to be fully effective. Our work relied on a set of strong assumptions, thus leaving the door open for new studies that will relax these constraints or fully replace concepts we used.

As a first step, extending the study through the use of other scoring rules seems highly fertile. An initial analysis already shows that the logarithmic scoring rule induces different properties in the model, one of which being the fact that cooperation no longer motivates the use of higher precisions. If the quadratic scoring rule made optimal the use of higher precisions for those who chose to interact, the logarithmic scoring rule makes no such difference (see appendix B.2). This can also be traduced as "cooperators incurring lower penalties", feature that might prove important.

Secondly, the PFM proposed here has at least three components that require further study. We mention the way an agent evaluates its trust models, the exploration-exploitation trade-off and the way pairs are allocated. A method that would combine the strengths of aggregation and maximum a posteri could certainly allow agents to approximate $\dot{\alpha}$ values quicker and this would have a great influence over the average utilities. The balance between exploration and exploitation is also essential and this should be coupled with the confidence an agent has in its $\dot{\alpha}$ approximation. Here the concept of confidence was not implemented and the effects were visible.

"Stable roommates" allocation is an efficient solution for matching agents that have good confidence in their preferences. In a more complex setup though, having a bigger popuation implies obtaining stable pairs with a greater difficulty. Not only is this allocation method prone to scalability issues but in some cases it may prevent new beneficial pairs from being formed. Using a clustered version of pair allocation, through a mechanism that provides more than one such service, may improve the overall results. Groups of agents that start having more interactions could decide to use a private pair allocator and thus avoid being "crowded" with many defectors.

6.2 Philosophical aspects

By looking at the results obtained throughout this study, we can extrapolate a series of judgements that might be considered for personal use.

First of all, the decisive factor in benefiting from eliciting information, the one that was partially addressed through this mechanism is that of knowing well the other "agents" in one's community. Tests showed that when using public α values, the number of unstable matches was minimal, in addition to pure cooperation being favoured. Thus, one element that favours the formation of compact groups or societies is facilitating individuals to know well one another. In terms of management, when this issue is not an objective it is not surprising that the "unstable matches" come in higher numbers.

Another interesting fact is provided by the agents' trust modelling abilities and the way evaluation of trust models was undertaken. That is, our Bayesian model proved effective and its empirical analysis showed it mapped well to the defection levels of other agents (figure 3.9). However, decoding the trust model into an optimal $\dot{\alpha}$ value can still be improved. Extrapolating this fact over a psychological framing, we can reflect on the way people do their own approximations. Are they sometimes gullible but trained to see the potential of others (maximum a posteri)? Do they always think of both ends of the spectrum (aggregation) or just view everything with caution (minimum posteri)? Finding the optimal solution in a computational setting should be validated in a real one, and vice-versa.

Naïvity deserves its share of attention too. Its main effect was that of agents loosing utility due to higher incurred penalties. The positive outcome was that cooperators were easier to separate from defectors and strong defectors obtained negative utility, whereas through adaptive fusion defectors were the ones who gained more. This is then a third piece of a puzzle, the first being the way trust models are used and the second, the exploration-exploitation trade-off. The less naïve is one's data fusion method, the higher its utility, irrespective of one's "type". In the real world, adaptiveness may be associated with wisdom. Consequently, focusing the educational programmes in developing wisdom more than delivering information, will have the main impact of people loosing less "utility".

Bibliography

[1] Nir Andelman and Yishay Mansour. A sufficient condition for truthfulness with single parameter agents. In *EC '06: Proceedings of the 7th ACM conference on Electronic commerce*, pages 8–17, New York, NY, USA, 2006. ACM.

[2] Robert Axelrod. *The evolution of cooperation*. Basic Books, 1984.

[3] Robert Axelrod and William D. Hamilton. The evolution of cooperation. *Science, New Series*, 221(4489):1390–1396, March 1981.

[4] B. Robert Boyd. Mistakes allow evolutionary stability in the repeated prisoners dilemma game. *Journal of Theoretical Biology*, 136:47–56, 1989.

[5] Robert Boyd and Jeffrey P. Lorberbaum. No pure strategy is evolutionarily stable in the repeated prisoners dilemma game. *Nature*, 327:58–59, 1987.

[6] Katarna Cechlrov. On the complexity of exchange-stable roommates. *Discrete Appl. Math.*, 116(3):279–287, 2002.

[7] Kim-Sau Chung. On the existence of stable roommate matchings. *Games and Economic Behavior*, 33:206–230, 2000.

[8] V. Conitzer and T. Sandholm. Complexity of mechanism design. In *Uncertainty in Artificial Intelligence Conf*, pages 103–110. Morgan Kaufmann, 2002.

[9] J. Crmer and R. McLean. Optimal selling strategies under uncertainty for a discriminating monopolist when demands are interdependent. *Econometrica*, 53:345–361, 1985.

[10] J. Crmer and R. McLean. Full extraction of surplus in bayesian and dominant strategy auctions. *Econometrica*, 56:1247–1257, 1988.

[11] Paul Darwen and Xin Yao. Coevolution in iterated prisoners dilemma with intermediate levels of cooperation: Application to missile defense. *Int. J. of Computational Intelligence and Applications*, 2(1):83–107, March 2002.

[12] Rajdeep K Dash, David C Parkes, and N. R. Jennings. Computational mechanism design: A call to arms. *IEEE Intelligent Systems*, 18(6):40–47, 2003.

[13] Joseph Farrell and Roger Ware. Evolutionary stability in the repeated prisoners dilemma. *Theoretical Population Biology*, 36:161–166, 1989.

[14] J. Feigenbaum, C. Papadimitriou, and S. Shenker. Sharing the cost of multicast transmissions. *Journal of Computer and System Sciences*, 63:21–41, 2001.

[15] Joan Feigenbaum and Scott Shenker. Distributed algorithmic mechanism design: Recent results and future directions. In *In Proceedings of the 6th International Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications*, pages 1–13. ACM Press, 2002.

[16] Tamas Fleiner, Robert W. Irving, and David F. Manlove. An algorithm for a super-stable roommates problem. Technical report, EGRES, 2008.

[17] Tams Fleiner. The stable roommates problem with choice functions. Technical report, EGRES, 2008.

[18] Daniel Friedman. Effective scoring rules for probabilistic forecasts. *Management Science*, 29(4):447–454, April 1983.

[19] D. Gale and L. S. Shapley. College admissions and the stability of marriage. *American Mathematical Monthly*, 69:9–15, 1962.

[20] A. Garratt, K. Lee, Pesaran, M. H., and Y Shin. Forecast uncertainties in macroeconomic modelling: An application to the uk economy. *Journal of the American Statistical Association*, 98:829–838, 2003.

[21] Tilmann Gneiting and Adrian E. Raftery. Strictly proper scoring rules, prediction and estimation. *Journal of the American Statistical Association*, 102(477):359–378, 2007.

[22] P. Grim. The undecidability of the spatialized prisoners dilemma. *Theory and Decision*, 42(1):53–80, 1997.

[23] Arlo D. Hendrickson and Robert J. Buehler. Proper scores for probability forecasters. *Annals of Mathematical Statistics*, 42(6):1916–1921, 1971.

[24] Robert Hoffmann. Twenty years on: The evolution of cooperation revisited. *Journal of Artificial Societies and Social Simulation*, 3(2):–, 2000.

[25] Leonard Hurwicz. Optimality and informational efficiency in resource allocation processes. In K.J. Arrow, S. Karlin, and P. Suppes, editors, *Mathematical Methods in the Social Sciences*, pages 27–46. Stanford University Press, Stanford, CA, 1960.

[26] T. D. Huynh. *Trust and Reputation in Open Multi-Agent Systems*. PhD thesis, University of Southampton, June 2006.

[27] Robert W. Irving. An efficient algorithm for the ”stable roommates” problem. *Journal of Algorithms*, 6:577–595, 1985.

[28] Robert W. Irving. On the ”stable roommates” problem. Technical report, University of Glasgow, Computing Science Department, 1986.

[29] Robert W. Irving and David F. Manlove. The stable roommates problem with ties. *Journal of Algorithms*, 43(1):85–105, 2002.

[30] Kazuo Iwama and Shuichi Miyazaki. A survey of the stable marriage problem and its variants. In *ICKS '08: Proceedings of the International Conference on Informatics Education and Research for Knowledge-Circulating Society (icks 2008)*, pages 131–136, Washington, DC, USA, 2008. IEEE Computer Society.

[31] Scott Johnson, John Pratt, and Richard Zeckhauser. Efficiency despite mutually payoff-relevant private information: the finite case. *Econometrica*, 58:873–900, 1990.

[32] D. J. Johnstone. The value of a probability forecast from portfolio theory. *Theory and Decision*, 63(2):153–203, 2007.

[33] D. Marc Kilgour and Yigal Gerchak. Elicitation of probabilities using competitive scoring rules. *Decision Analysis*, 1(2):108–113, 2004.

[34] D.E. Knuth. *Mariages stables et leurs relations avec dautres problmes combinatoires.* Presses de l'Universite de Montreal, 1976.

[35] K. Krukow, M. Nielsen, and V. Sassone. A logical framework for history-based access control and reputation systems. *Journal of Computer Security*, 16(1):63–101, 2008.

[36] Dave Levin, Randolph Baden, Cristian Lumezanu, Neil Spring, and Bobby Bhattacharjee. Motivating participation in internet routing overlays. In *NetEcon '08: Proceedings of the 3rd international workshop on Economics of networked systems*, pages 91–96, New York, NY, USA, 2008. ACM.

[37] Steven L. Lima. Iterated prisoners dilemma: An approach to evolutionary stable cooperation. *American Naturalist*, 134(5):828–834, November 1989.

[38] Michael W. Macy and John Skvoretz. The evolution of trust and cooperation between strangers: A computational model. *American Sociological Review*, 63(5):638–660, 1998.

[39] Andreu Mas-Colell, Michael D Whinston, and Jerry R Green. *Microeconomic Theory.* Oxford University Press, 1995.

[40] D. G. McVitie and L. B. Wilson. The stable marriage problem. *Communications of the ACM*, 14(7):486–490, 1971.

[41] Nolan Miller, Paul Resnick, and Richard Zeckhauser. Eliciting informative feedback: The peer-prediction method. *Manage. Sci.*, 51(9):1359–1373, 2005.

[42] Ahuva Mu'alem. On decentralized incentive compatible mechanisms. In *EC '05: Proceedings of the 6th ACM conference on Electronic commerce*, pages 240–248, New York, NY, USA, 2005. ACM.

[43] Alan H. Murphy and Robert L. Winkler. Scoring rules in probability assessment and evaluation. *Acta Psychologica*, 34:273–286, 1970.

[44] Roger. B. Myerson and Mark. A. Satterthwaite. Efficient mechanisms for bilateral trade. *Journal of Economic Theory*, 29:265281, 1983.

[45] M. Nielsen, K. Krukow, and V. Sassone. A bayesian model for event-based trust. *Electronic Notes in Theoretical Computer Science*, 172:499–521, 2007.

[46] Noam Nisan and Amir Ronen. Algorithmic mechanism design (extended abstract). In *STOC '99: Proceedings of the thirty-first annual ACM symposium on Theory of computing*, pages 129–140, New York, NY, USA, 1999. ACM.

[47] Noam Nisan and Amir Ronen. Algorithmic mechanism design. *Games and Economic Behavior*, 35:166–196, 2001.

[48] Prize Committee of the Royal Swedish Academy of Sciences. Scientific background on the sveriges riksbank prize in economic sciences in memory of alfred nobel 2007: Mechanism design theory. Technical report, The Royal Swedish Academy of Sciences, 2007.

[49] T. N. Palmer. The economic value of ensemble forecasts as a tool for risk assessment: From days to decades. *Quarterly Journal of the Royal Meteorological Society*, 128:747–774, 2002.

[50] Athanasios Papakonstantinou, Alex Rogers, Enrico Gerding, and Nick Jennings. A truthful two-stage mechanism for eliciting probabilistic estimates with unknown costs. In *In Proceedings of the 18th European Conference on Artificial Intelligence (ECAI 2008), Patras, Greece*, May 2008.

[51] David Christopher Parkes. *Iterative combinatorial auctions: achieving economic and computational efficiency*. PhD thesis, University of Pennsylvania, 2001.

[52] J. Patel. *A Trust and Reputation Model for Agent-Based Virtual Organisations*. PhD thesis, University of Southampton, December 2006.

[53] A. Rapoport and A. Chammah. *Prisoners dilemma*. University of Michigan Press, 1965.

[54] Steven Reece, Alex Rogers, Stephen Roberts, and Nicholas R. Jennings. A multi-dimensional trust model for heterogeneous contract observations. In *Twenty-Second AAAI Conference on Artificial Intelligence*, pages 128–135, Menlo Park, California 94025, USA, 2007. AAAI Press.

[55] Steven Reece, Alex Rogers, Stephen Roberts, and Nicholas R. Jennings. Rumours and reputation: Evaluating multi-dimensional trust within a decentralised reputation system. In *The Sixth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-07)*, pages 1063–1070, New York, NY, USA, 2007. ACM.

[56] Paul Resnick and Richard Zeckhauser. Trust among strangers in internet transactions: Empirical analysis of ebays reputation system. In Michael R Baye, editor, *Advances in Applied Microeconomics*, volume 11. Amsterdam, Elsevier Science, 2002.

[57] Alex Rogers, Rajdeep K. Dash, N. R. Jennings, Steve Reece, and Stephen Roberts. Computational mechanism design for information fusion within sensor networks. In *Ninth International Conference on Information Fusion (Fusion 2006)*, 2006.

[58] Jordi Sabater and Carles Sierra. Reputation and social network analysis in multi-agent systems. In *AAMAS '02: Proceedings of the first international joint conference on Autonomous agents and multiagent systems*, pages 475–482, New York, NY, USA, 2002. ACM.

[59] V. Sassone, M. Nielsen, and K. Krukow. Towards a formal framework for computational trust. *Formal Methods for Components and Objects*, LNCS 4709:175–184, 2007.

[60] Leonard J. Savage. Elicitation of personal probabilities and expectations. *Journal of the American Statistical Association*, 66(336):783–801, December 1971.

[61] Rann Smorodinsky and Moshe Tennenholtz. Sequential information elicitation in multi-agent systems. In *AUAI '04: Proceedings of the 20th conference on Uncertainty in artificial intelligence*, pages 528–535, Arlington, Virginia, United States, 2004. AUAI Press.

[62] Hock Kim Tan and Luc Moreau. Trust relationships in a mobile agent system. In Gian Pietro Picco, editor, *Fifth IEEE International Conference on Mobile Agents*, volume 2240, pages 15—30. Springer-Verlag, December 2001.

[63] W. T. L. Teacy, G. Chalkiadakis, A. Rogers, and N. R. Jennings. Sequential decision making with untrustworthy service providers. In *Proceedings of the 7th International Conference on Autonomous Agents and Multiagent Systems*, May 2008.

[64] Laszlo Toka and Pietro Michiardi. A dynamic exchange game. In *PODC '08: Proceedings of the twenty-seventh ACM symposium on Principles of distributed computing*, pages 428–428, New York, NY, USA, 2008. ACM.

[65] Y. Zhu, Z. Toth, R. Wobus, D. Richardson, and K. Mylne. The economic value of ensemble-based weather forecasts. *Bulletin of the American Meteorological Society*, 83:73–83, 2002.

[66] Aviv Zohar and Jeffrey S. Rosenschein. Robust mechanisms for information elicitation. In *AAMAS '06: Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems*, pages 1202–1204, New York, NY, USA, 2006. ACM.

Appendices

Appendix A

Formal definitions

A.1 Mechanism design

Let us considering a system composed of Q agents, each having a particular type (signal or valuation) $t_i \in T_i$ and a strategy s_i that dictates what action $s_i(t_i) = A_i$ it will execute based on its type. Taking a recommender system as an example, an agent's type could be the quality it associates to a particular product while its strategy would indicate how precisely it should communicate it (a defective agent communicating a value that would maximise its profit, which does not have to be its real type).

In addition to these features, agents are equiped with utility functions $u_i : T_i \times o \rightarrow \mathbb{R}$, where $o \in \mathcal{O}$ represents the outcome generated by the mechanism, from the set $\mathcal{O}$ of possible outcomes. In a recommender system, the outcome could be a score the rater receives for the quality he communicated.

A mechanism $\mathcal{M}$ is a pair (A, g) composed of a set of actions $A = A_1 \times \cdots \times A_N$ made available to the agents and a mapping $g : A \rightarrow \mathcal{O}$, between the actions and outcomes. It is said that a mechanism implements a social choice function $f : T_1 \times \cdots \times T_N \rightarrow o$, if there exists a set of strategies $(s_1^*, \ldots, s_N^*)$ that is an equilibrium solution for the game induced by $\mathcal{M}$ and also allows it to generate the same outcome as$f$, $g(s_1^*(t_1), \ldots, s_N^*(t_N)) = f(T)$.

For a more extensive introduction into the MD literature, MasColell *et al.* [39] is a valuable reference, while the report compiled in 2007 by the Prize Committee of the Royal Swedish Academy of Sciences [48] details the concepts and historical development of the field.

A.2 Scoring rules

The following mathematical representations are adapted from Friedman [18]. Let us first consider a random variable $\mathcal{X}$, that has a distrubution F and range $\Omega \in R$ that describes the possible outcomes of $\mathcal{X}$. The distribution F is defined by a density function f that belongs to a set D

representative density functions defined on Ω.

A metric (distance function) assigns to any pair of densities f,g in D a real number $d(f,g)$ that represents the distance between the two densities. Consequently, a scoring rule $S(g,x)$ is a real valued function, where g is any density in D and x is a realisation of $\mathcal{X}$. S is strictly proper if, taking f as the true distribution, the f-expected value is maximised on D at $g = f$. In other words, for a proper scoring rule:

$$\begin{aligned} E_f S(g) &= \int_\Omega S(g,x) f(x)\,\mathrm{d}x \quad \text{and} \\ \max E_f S(g) &= E_f S(f) \end{aligned} \tag{A.1}$$

Metrics come in a wide variety and they can be accustomed to a problem's profile. Some of the most commonly used metrics are:

1. the L_1 metric, defined by $d_1(f,g) = \int_\Omega |f(x) - g(x)|\,\mathrm{d}x$ and representing the area between the graphs of the two densities.
2. the L_2 metric, defined by $d_2(f,g) = \left(\int_\Omega |f(x) - g(x)|^2\right)^{\frac{1}{2}}$ and being analogous to the standard deviation.
3. the renormalised L_2 metric, defined as $d^*(f,g) = d_2(\rho g, \rho f)$, where $\rho f = f/\|f\|$ represents the normalisation of f.

The mentioned scoring rules rely on this metrics, in the sense that the quadratic scoring rule is strictly proper on the L_2 metric, the spherical scoring rule is strictly proper on the renormalised L_2 metric while the logarithmic scoring rule is strictly proper on all these metrics.

When discussing about forecasts provided in the form of Gaussian distributions, an agent needs to report two values: the mean μ and the variance σ^2 of the distribution. For easier use, we can name the mean x_0 and instead of using variance, use the equivalent precision θ, where $\theta = 1/\sigma^2$. Thus, the three main scoring rules will have the formulas ([50]):

$$\text{Quadratic} : S(x_0, N(x_0; x, 1/\theta)) = 2N(x_0; x, 1/\theta) - \int_{-\infty}^{\infty} N^2(x, 1/\theta)\,\mathrm{d}x \tag{A.2}$$

$$\text{Spherical} : S(x_0, N(x_0; x, 1/\theta)) = N(x_0; x, 1/\theta)/(\int_{-\infty}^{\infty} N^2(x, 1/\theta)\,\mathrm{d}x)^2 \tag{A.3}$$

$$\text{Logarithmic} : S(x_0, N(x_0; x, 1/\theta)) = \log(N(x_0; x, 1/\theta)) \tag{A.4}$$

Finally, it should be noted that the logarithmic scoring rule has very attractive properties but due to its construction, it can generate $-\infty$ scores, while the other two scoring rules are more stable but not that intuitive. For exemplification, figure A.1 shows the scores an agent would obtain in case the true value of the modelled parameter would be 0 and its measurment and reported precision would be $\theta = 1$.

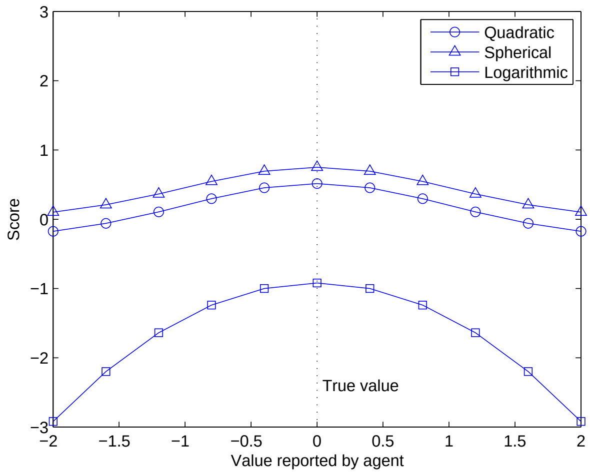

Figure A.1: Scores obtained when reporting forecasts of precision $\theta = 1$

Appendix B

Mechanism analysis

B.1 Mutual defection

In the more general case of α values being different from agent to agent, the precision of a fused estimate θ_f and the expected utility $\overline{U}$ will be:

$$\theta_f = \frac{\alpha_2(\alpha_1+1)^2}{\alpha_1\alpha_2+1}\theta_{CC} \tag{B.1}$$

$$\overline{U}(\alpha_1,\alpha_2,\theta_{CC}) = 2(\alpha_1+1)\sqrt{\frac{\alpha_2\theta_{CC}}{2\pi(2\alpha_1\alpha_2+\alpha_2+1)}} - \frac{1}{2}\sqrt{\frac{(\alpha_1+1)\theta_{CC}}{\pi}} - c\,\alpha_1\theta_{CC} \tag{B.2}$$

where α_1 is the defection coeficient of the agent computing the estimate while α_2 describes its counterpart.

B.2 Initial analysis on the use of the logarithmic scoring rule

The following analysis will use the Logarithmic scoring rule, presented in equation B.3.

$$S(x_0; x, 1/\theta) = \log N(x_0; x, 1/\theta) \tag{B.3}$$

$$\tag{B.4}$$

where θ_f is the precision of the fused estimate and θ is the precision employed by an agent in generating its personal measurement.

Considering a system composed of two agents, there can be five interaction scenarios: a)no interaction, b)reciprocal cooperation, c)defecting or d)being defected upon and e)mutual defection. The reminder of this paper will present a game theoretical analysis of these situations.

B.2.1 No interaction (N)

In this case, agents don't fuse estimates, meaning that $\theta_f = \theta$ and the utility $U(\theta)$ is:

$$U_N(\theta) = S(\theta) - c\theta \tag{B.5}$$

The optimal precision θ^*, is obtained by maximizing $\overline{U}(\theta)$:

$$\begin{aligned} \overline{U'}(\theta) &= \overline{S'}(\theta) - c = 0 \\ \overline{S'}(\theta) &= \frac{1}{2\theta} = c \\ \theta^*_N &= \frac{1}{2c} \end{aligned} \tag{B.6}$$

In this scenario, the maximum expected utility, $\overline{U}(\theta^*_N)$ is:

$$\begin{aligned} \overline{U}(\theta^*_N) &= \overline{S}(\theta^*_N) - c\theta^*_N \\ &= \frac{1}{2}\log\left(\frac{\theta^*_N}{2\pi}\right) - c\theta^*_N \\ &= \frac{1}{2}\log\left(\frac{1}{4\pi c}\right) - 1 \end{aligned} \tag{B.7}$$

Figure B.1 presents the relation between an agent's utility and its measurment precision, both from a analytical and experimental perspective.

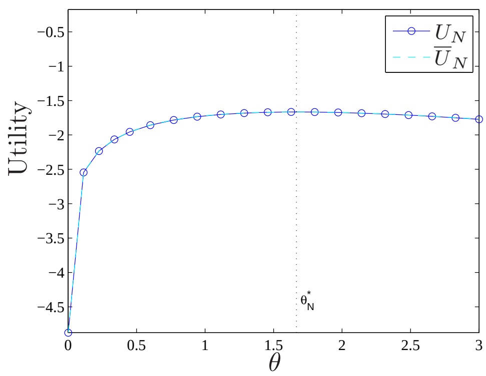

Figure B.1: Utility in case of no cooperation (N, c=0.3)

B.2.2 Cooperation - Cooperation (CC)

In this situation, the overall precision is $\theta_f = 2\theta$ and the utility:

$$U_{CC}(\theta) = S(2\theta) - c\theta \tag{B.8}$$

The optimal precision θ^*_{CC}, obtained by maximizing $\overline{U}(\theta)$ will be equal to θ^*_N, meaning that $\theta^*_{CC} = \frac{1}{2c}$. Thus, cooperation doesn't incentivise the use of higher precisions, but it provides higher utilities.

The maximised expected utility, $\overline{U}(\theta^*_{CC})$ is:

$$\begin{aligned}\overline{U}(\theta^*_{CC}) &= \overline{S}(2\theta^*_{CC}) - c\theta^*_{CC} \\ &= \frac{1}{2}\log\left(\frac{\theta^*_{CC}}{\pi}\right) - \frac{1}{2} - c\theta^*_{CC} \\ &= \frac{1}{2}\log\left(\frac{1}{2\pi c}\right) - 1 \qquad \text{(B.9)}\end{aligned}$$

Figure B.2 presents the expected and actual utility values obtained by agents with respect to their measurment precisions, in the case of reciprocal cooperation.

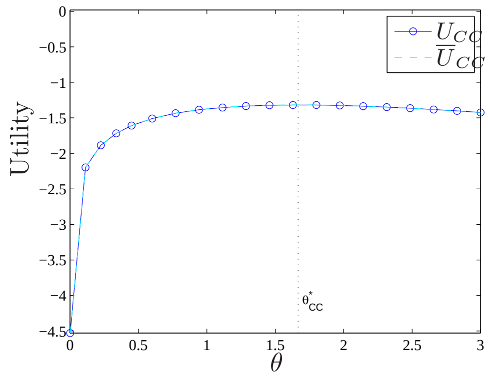

Figure B.2: Utility when cooperating (CC, c=0.3)

B.2.3 Defection - Cooperation (DC)

In case of defection, an agent is considered to be using a precision $\theta_{DC} = \alpha\theta_{CC}$, where $\alpha \in (0, 1]$ is a coeficient defining the degree of defection. The fused estimate will be then of precision $\theta_f = (\alpha + 1)\theta_{CC}$ and the utility:

$$U(\theta_{DC}) = S((\alpha + 1)\theta_{CC}) - c\alpha\theta_{CC} \qquad \text{(B.10)}$$

Maximising the utility, we find the expression for the optimal precision θ^*_{DC} and the corresponding expected utility value:

$$\theta^*_{DC} = 0 \qquad \text{(B.11)}$$

$$\overline{U}(\theta^*_{DC}) = \frac{1}{2}\log\left(\frac{1}{2\pi c}\right) \qquad \text{(B.12)}$$

$$\qquad \text{(B.13)}$$

Figure B.3 presents expected and actual utility values obtained by defecting agents, when they are interacting with cooperating ones. A new parameter, α, has been introduced in order to control defection ($\theta_{DC} = \alpha\theta_{CC}, \alpha \in (0, 1]$).

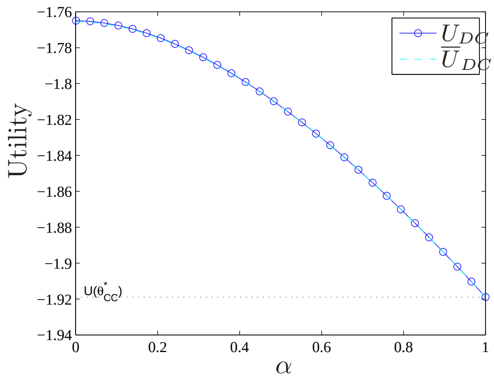

Figure B.3: Utility when defecting agains a cooperator (DC, c=1)

B.2.4 Cooperation - Defection(CD)

The precision of the cooperator's fused estimate is dependent on the degree of defection α of its interaction partner, dependency that is shown in figure B.5 from both an analytical and experimental point of view. Derivation of the fused estimate precision θ_f can be obtained in the following way:

$$
\begin{aligned}
x_1 &= N(x_0, \theta) & \sigma_f^2 &= \frac{1}{4}(\frac{1}{\theta} + \frac{1}{\alpha\theta}) \\
x_2 &= N(x_0, \alpha\theta) & & \\
x_f &= \frac{x_1 + x_2}{2} & \theta_f &= \frac{4\alpha}{\alpha+1}\theta
\end{aligned}
\tag{B.14}
$$

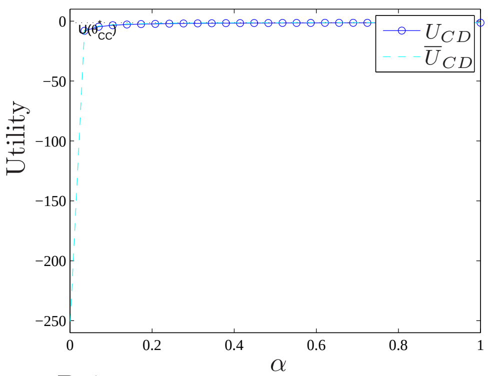

Figure B.4: Utility when being defected upon (CD, c=0.03)

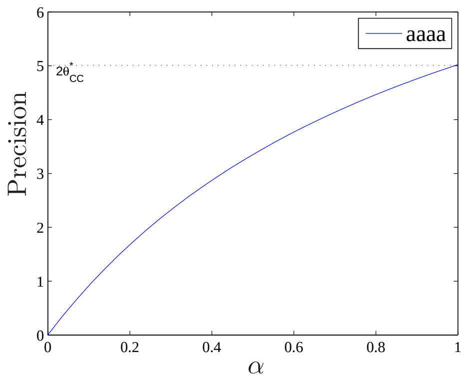

Figure B.5: Fused estimate's precision($\theta_{CC} = 1$)

The expected score has a different formula, due to the difference between the reported precision and the one resulted from defection. Equation B.15 details its value.

$$
\overline{S}(x_0; x, \theta_b) = \frac{1}{2}\log\left(\frac{\theta_b}{2\pi}\right) - \frac{\theta_b}{2\theta_f} \tag{B.15}
$$

The expected utility, for the CD scenario, will be:

$$\overline{U}_{CD}(\alpha, c) = \frac{1}{2}\log\left(\frac{\theta_b}{2\pi}\right) - \frac{\theta_b}{2\theta_f} - c\theta \tag{B.16}$$

Maximising the utility, through the use of θ^*_{CC}, we obtain:

$$\overline{U}^*_{CD}(\alpha, c) = \frac{1}{2}\log\left(\frac{1}{2\pi c}\right) - \frac{3\alpha + 1}{4\alpha} \tag{B.17}$$

B.2.5 Defection - Defection (DD)

In the case of reciprocal defection, both agents use precisions in the form of $\alpha\theta_{CC}$. Considering the simple case where agents use an identical α, the precision of a fused estimate will not be $\theta_f = (\alpha + 1)\theta_{CC}$ as believed by each of them, but:

$$\theta_f = \frac{\alpha(\alpha + 1)^2}{\alpha^2 + 1}\theta_{CC} \tag{B.18}$$

Respecting equation B.15, the expected score in this situation will be given by equation B.19:

$$\overline{S}(x_0; x, \theta) = \frac{1}{2}\log\left(\frac{(\alpha + 1)\theta}{2\pi}\right) - \frac{\alpha^2 + 1}{2\alpha(\alpha + 1)} \tag{B.19}$$

The expected utility will be then:

$$\overline{U}_{DD}(\alpha, c) = \frac{1}{2}\log\left(\frac{(\alpha + 1)\theta}{2\pi}\right) - \frac{\alpha^2 + 1}{\alpha(\alpha + 1)} - c\alpha\theta \tag{B.20}$$

Figure B.6 presents analytical and experimental results for the utility obtained in case of mutual defection case, while figure B.7 depicts the relation between the defection coeficient α and the .

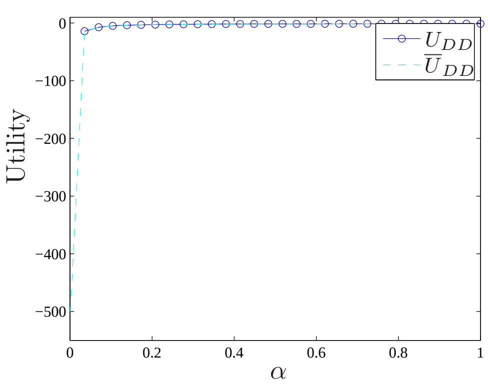

Figure B.6: Utility of mutual defection (DD, c=0.3)

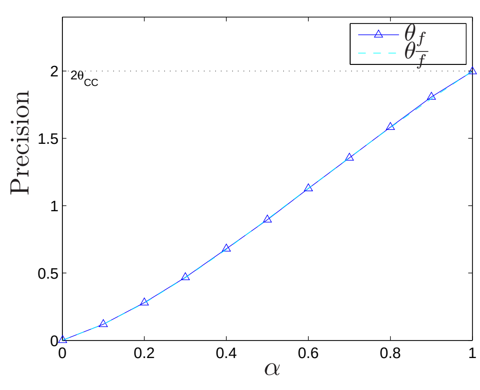

Figure B.7: Fused estimate's precision($\theta_{CC} = 1$)