UNIVERSITY OF SOUTHAMPTON
Faculty of Engineering, Science and Mathematics
School of Electronics and Computer Science

Mechanism Design for Eliciting Costly Observations in Next Generation Citizen Sensor Networks

by Athanasios Papakonstantinou

Supervisors: Prof. Nicholas R. Jennings and Dr. Alex Rogers
Examiners: Dr. Shaheen Fatima and Dr. Adam Prügel-Bennett

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

March 2010

Declaraction of Authorship

I, *Athanasios Papakonstantinou* declare that the thesis entitled *Mechanism Design for Eliciting Costly Observations in Next Generation Citizen Sensor Networks* and the work presented in it are my own. I confirm that:

- this work was done wholly or mainly while in candidature for a research degree at this University;
- where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
- where I have consulted the published work of others, this is always clearly attributed;
- where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
- I have acknowledged all main sources of help;
- where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
- parts of this work have been published as:

  A. Papakonstantinou, A. Rogers, E. H. Gerding, N. R. Jennings (2008) A Truthful Two-Stage Mechanism for Eliciting Probabilistic Estimates with Unknown Costs, *In Proceedings of the 18th European Conference on Artificial Intelligence (ECAI 2008)*, pages 448-452, Patras, Greece.

  A. Papakonstantinou, A. Rogers, E. H. Gerding, N. R. Jennings (2009) Mechanism Design for Eliciting Probabilistic Estimates from Multiple Suppliers with Unknown Costs and Limited Precision, *Proceedings of the Eleventh Workshop in Agent Mediated Electronic Commerce (AMEC 2009)*, pages 111-124, Budapest, Hungary.

Signed:

Date:

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF ENGINEERING, SCIENCE AND MATHEMATICS
SCHOOL OF ELECTRONICS AND COMPUTER SCIENCE

Doctor of Philosophy

by Athanasios Papakonstantinou

Citizen sensor networks are open information systems in which members of the public act as information providers. The information distributed in such networks ranges from observations of events (e.g. noise measurements or monitoring of environmental parameters) to probabilistic estimates (e.g. projected traffic reports or weather forecasts). However, due to rapid advances in technology such as high speed mobile internet and sophisticated portable devices (from smart-phones to hand-held game consoles), it is expected that citizen sensor networks will evolve. This evolution will be driven by an increase in the number of information providers, since, in the future, it will be much easier to gather and communicate information at a large scale, which in turn, will trigger a transition to more commercial applications. Given this projected evolution, one key difference between future citizen sensor networks and conventional present ones is the emergence of self-interested behaviour, which can manifest in two main ways. First, information providers may choose to commit insufficient resources when producing their observations, and second, they may opt to misreport them. Both aspects of this self-interested behaviour are ignored in current citizen sensor networks. However, as their applications are broadened and commercial applications expand, information providers are likely to demand some kind of payment (e.g. real or virtual currency) for the information they provide. Naturally, those interested in buying this information, will also require guarantees of its quality.

It is these issues that we deal with in this thesis through the introduction of a series of novel two-stage mechanisms, based on strictly proper scoring rules. We focus on strictly proper scoring rules, as they have been used in the past as a method of eliciting truthful reporting of predictions in various forecasting scenarios (most notably in weather forecasting). By using payments that are based on such scoring rules, our mechanisms effectively address the issue of selfish behaviour by motivating information providers in a citizen sensor network to, first, invest the resources required by the information buyer in the generation of their observations, and second, to report them truthfully.

To begin with, we introduce a mechanism that allows the centre (acting as an information buyer) to select a single agent that can provide a costly observation at a minimum cost. This is the first time a mechanism has been derived for a setting in which the centre has no knowledge of the actual costs involved in the generation of the agents' observations. Building on this, we then

make two further contributions to the state of the art, with the introduction of two extensions of this mechanism. First, we extend the mechanism so that it can be applied in a citizen sensor network where the information providers do not have the same resources available for the generation of their observations. These different capabilities are reflected in the quality of the provided observations. Hence, the centre must select multiple agents by eliciting their costs and the maximum precisions of their observations and then ask them to produce these observations. Second, we consider a setting where the information buyer cannot gain any knowledge of the actual outcome beyond what it receives through the agents' reports. Now, because the centre is not able to evaluate the providers' reported observations through external means, it has to rely solely on the reports it receives. It does this by fusing the reports together into one observation which then uses as a means to assess the reports of each of the providers.

For the initial mechanism and each of the two extensions, we prove their economic properties (i.e. incentive compatibility and individual rationality) and then present empirical results comparing a number of specific scoring rules, which includes the quadratic, spherical, logarithmic and a parametric family of scoring rules. These results show that although the logarithmic scoring rule minimises the mean and variance of an agent's payment, using it may result in unbounded payments if an agent provides an observation of poor quality. Conversely, the payments of the parametric family exhibit finite bounds and are similar to those of the logarithmic rule for specific values of the parameter. Thus, we show that the parametric scoring rule is the best candidate in our setting. We empirically evaluate both extended mechanisms in the same way, and for the first extension, we show that the mechanism describes a family of possible ways to perform the agent selection, and that there is one that dominates all others. Finally, we compare both extensions with the peer prediction mechanism introduced by Miller, Resnick and Zeckhauser (2007) and show that in all three mechanisms the total expected payment is the same, while for both our mechanisms the variance in the total payment is significantly lower.

Contents

List of Figures

List of Tables

Acknowledgements

I would like to sincerely thank my supervisors Nick Jennings and Alex Rogers for their help and guidance during my PhD. The time you have invested in me, is deeply appreciated. I am sure that your advice will prove to be invaluable as I try to find my place in the academic universe.

To Enrico Gerding, the co-author in the publications that were produced during the PhD, thank your for your contribution and your input. I am also grateful to EPSRC for their financial support through the MBC project without which my PhD would not have been possible.

To my parents, Costas and Valentini, thank you for believing in me, and making me believe more in myself. I hope that one day I will become the inspiration to my children that you are to me.

To Deppy, thank you for this amazing year. For your support, your patience, your bizarre optimism, and for putting up with me during the final stage of my PhD.

To Simon and 'Kuhn' for their friendship during the almost three years of living under the same roof and their commitment to our slightly dysfunctional, but still in perfect harmony, household.

Finally, to friends from the IAM lab and other labs around the University of Southampton, Archie, Rama, Zinovi, Raj, George, Nikitas, Bangelis, for three and a half great years.

To my parents

Chapter 1

Introduction

Advances in technology have made real-time information about the state of the world increasingly available through distributed online systems that are owned by different stakeholders and can be accessed by multiple users. An emerging application of these systems are *citizen sensor networks* (Goodchild, 2007; Sheth, 2009), a computing paradigm whereby large numbers of citizens share observations and information through mobile devices, in real-time, within open information systems[1]. Examples of such networks include social networking sites and the blogoshpere, where people post anything from their everyday routine to significant events they experienced such as fires[2], terrorist attacks[3] and aeroplane crashes[4].

Such networks typically consist of members of the public acting as information providers communicating observations which they acquire either directly through first hand experience, or indirectly through the control of portable devices such as mobile phones which have multiple uses (e.g. GPS navigator and on-line messaging through GPRS). For example, NoiseTube[5] is a project that monitors noise pollution across a city through the assignment of noise observations to members of the public. These contributors use their mobile phones to measure the noise levels in their surroundings and then communicate their measurement to a central agency which is responsible for the compilation of these observations into noise pollution maps (Figure 1.1). In both NoiseTube and the aforementioned applications of citizen sensor networks, the acts of the information providers are mainly altruistic since their contribution is voluntarily. In this context, the common characteristics of NoiseTube, the mapping of the Victoria fires and the reports during the terrorist attacks in Mumbai and the aeroplane landing in Hudson river, is that ordinary

[1] A recent Gartner group report predicted that by 2012 physical sensors will create 20 percent of non-video internet traffic see `http://www.gartner.com/it/page.jsp?id=876512`.

[2] In the Victoria fires in 2009, people could contribute by pointing fire locations in the Official Google Australia Blog: `http://google-au.blogspot.com/2009/02/mapping-victorian-fires.html`.

[3] GroundReport, an on-line organisation responsible for the coordination of 5000 amateur and professional reporters claimed that it had coverage of the Mumbai attacks before any of the mainstream media `http://www.groundreport.com/World/Terrorists-attack-Mumbai-target-posh-hotels/2874167`.

[4] *'There's a plane in the Hudson. I'm on the ferry going to pick up the people. Crazy.'* jkrums on Twitter `http://twitpic.com/135xa`.

[5] `http://noisetube.net/`

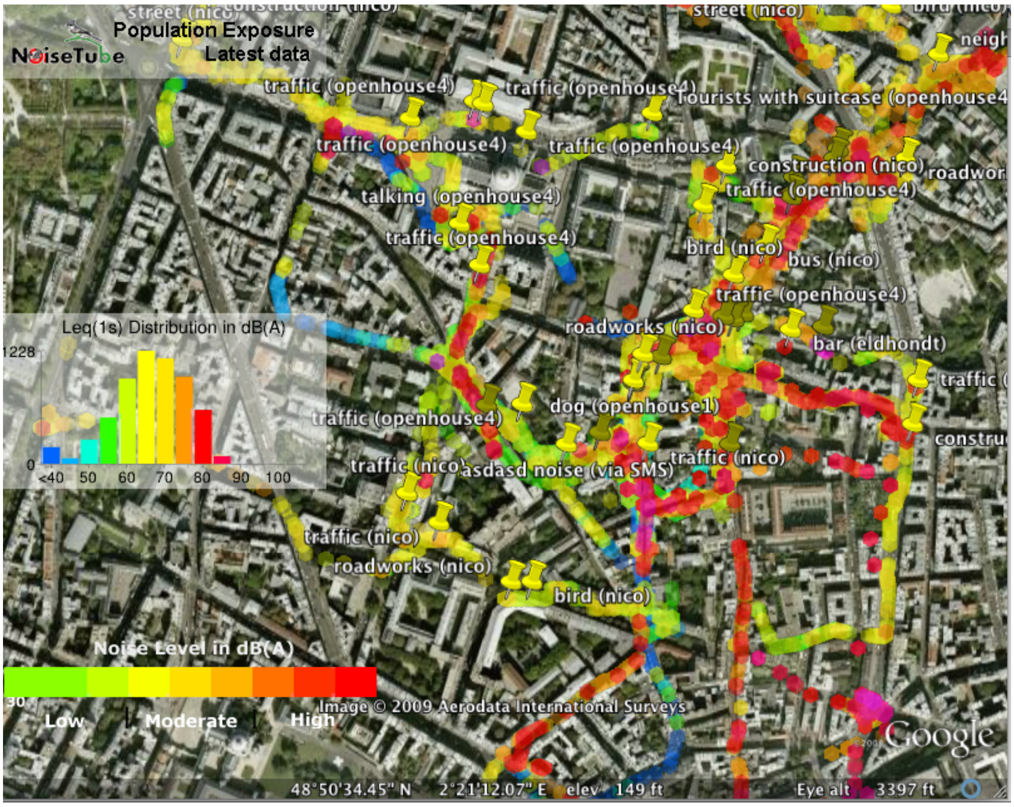

FIGURE 1.1: Noisemap of a neighbourhood of Paris generated by NoiseTube.

people offer information on extraordinary events, without expecting anything in return, even if their observations bring revenues to the media networks that use them (for example through advertising).

Such citizen sensor networks are easy to deploy and offer advantages over traditional means of gathering information, such as their ability to employ a large and diverse pool of information providers. Nevertheless, this expansion is likely to bring certain drawbacks. For example, it is expected that there will be a certain cost attached to the generation of the observation which reflects the effort a citizen puts into its generation. Furthermore, since costs are dependent on the specific resources available to each provider they represent private information. In terms of the noise pollution monitoring application, the cost of the observations is represented by the time the information providers invest in making their measurements. Contributors who invest a significant amount of time, produce measurements which are more representative of the situation around their location, hence more useful to those interested in acquiring it.

In this context, it is the existence of costs that signals a transition to more commercial applications of citizen sensor networks. Examples of commercial applications of citizen sensor networks are rare at present and focus primarily on traffic monitoring services such as traffic.com and Inrix, and services offering advice such as optimal routes to destinations (e.g. TrafficCast[6]), where the companies receive payments for the information they offer. The connection between

[6]`http://trafficcast.com/`

these applications and citizen sensor networks is not straightforward, since the roles for the citizens and those interested in acquiring their observations appear to have been inverted. In more detail, in this particular case members of the public do not act as information providers but instead as information buyers, since they are willing to pay for the estimates they will receive. The role of the information providers is undertaken by private companies, which face costs when generating their estimates of traffic routes depending on their computational resources, statistical methods and infrastructure. Indeed, they may well acquire information from citizens' mobile devices, and effectively act as aggregators of information. Therefore, in both existing altruistic applications such as NoiseTube and commercial citizen sensor networks such as TrafficCast, having knowledge of the involved costs is a challenge of significance importance since it is in every buyer's best interest to identify the provider that can provide an observation at a minimum cost.

Furthermore, costs have an impact on the quality of the observations, since they depend on the effort and resources invested in generating them. Indeed, contributors may choose to provide observations of low quality in order to minimise their costs. Therefore, it is important to develop processes that can evaluate the information provided to those interested in acquiring it (referred from this point as information buyers) and provide guarantees on its quality. This challenge becomes more significant when information providers have different motives, which may contradict each other and may exhibit selfish behaviour (i.e. by not investing sufficient resources in the generation of an observation, or by misreporting it). Selfish behaviour is a common problem in other open information systems, such as email protocol and the world wide web. For example, it has been estimated by the National Technology Readiness Survey that in 2004 the cost of email spam in the US economy alone, was more than \$21.9B (Aalberts et al., 2007). In this context, Sir Tim Berners-Lee, the inventor of the World Wide Web, considers the lack of processes that can address selfish behaviour within the original email protocol as the main cause behind spam's dominance; '*the people who designed the email system designed it for a world in which everyone was friendly with each other*' (Berners-Lee, 2008). Likewise, current citizen sensor networks (both altruistic and commercial applications) do not deal with the issue of selfish behaviour, and most of the relevant literature has yet to recognise its existence[7]. Therefore, when designing the protocols that will guide the interactions among information providers and buyers in next generation citizen networks, which will emerge as current systems grow and become more popular, it is crucial that selfish behaviour is taken into consideration.

In this thesis, we consider selfish behaviour that can manifest in two forms. First, the information providers may be inclined to allocate less than the sufficient resources, or even none, during the process of acquiring the information, especially when there is a certain cost attached to the generation of the information. More precisely, if a provider has the task of making and reporting an observation which has certain requirements, it may choose to allocate less than the required resources for the successful completion of the task, or even none by simply making up a report.

[7]In Section 2.1 we review current applications of citizen sensor networks and show that selfish behaviour is not addressed in any of the existing applications.

Within the noise pollution monitoring example, a sound measurement may have little value to the information buyer if it is measured from inside a vehicle, since the captured sound will be unrepresentative of the true state. In this case, the information providers effectively trick the buyers into believing that they have invested time in their observation, while in reality they have produced a low quality measurement which is not representative of reality.

The second sign of selfish behaviour for the information providers is opting to deliberately misreport their observations if they can benefit by doing so. For example, in the traffic feedback scenario, a company responsible for providing optimal routes to their subscribers may decide to manipulate them into following a route that has been suggested by other companies even if it knows that it may be congested. Indeed, it can be assumed that if multiple providers suggest a particular route to their subscribers that route will become congested, since most drivers will follow it. Therefore, if a company decides that it is more important to inflict losses on other companies by congesting their routes, it can do so by misreporting their traffic estimates to the subscribed drivers acting as information buyers in this commercial application of citizen sensor networks.

In addition to selfish behaviour, information providers will typically have different capabilities, and thus may only be able to provide information of a limited degree of precision. In terms of the sound monitoring example, the precision of the observations depends on their overall quality (i.e. clean sounds without any distortion are considered observations of high quality). In this context, it is expected that the information providers will have different capabilities as all of them are not able to invest the same amount of time in their observations, and therefore produce observations of various precisions. Likewise in a traffic feedback application, the companies acting as information providers have different capabilities which depend on the prediction methods they employ and the infrastructure (computational power and traffic data) available to them.

A final challenge commonly faced in citizen sensor networks is that information providers provide observations which are imprecise since they operate in environments where uncertainty is endemic. These environments include the streets of a city, which is the setting for both the noise monitoring and traffic feedback applications we discuss in this section. Due to the nature of the observed and reported information the information buyer cannot evaluate the received observations unless it chooses to invest heavily in infrastructure that will allow her to do so. For example, in the sound monitoring example, the agency responsible for the compilation of the noise pollution maps, in addition to the members of the public providing observations, may choose to refer to static stations with sophisticated microphones that could be used to evaluate those information providers operating in the surrounding area. In a similar vein, in the traffic feedback application, a driver may choose to subscribe to other services that monitor traffic through airborne means, and thus eliminate the effects of uncertainty, since they will be able to have a clear view of the traffic. However, for both cases, these extras come at a certain high cost. Indeed, one of the main reasons for using citizen sensor networks despite the drawbacks already outlined, is that they can be cheap, easy to deploy and do not rely on any special infrastructure besides equipment that can be available to most people in cities (i.e. mobile phones).

Now, besides uncertainty, those interested in acquiring information need to take into consideration another element of the environment in which the information providers operate: its dynamism. In dynamic environments the conditions change constantly as the environment around the providers evolves. This is the case for both the examples we consider, since a random event (i.e. a large vehicle passing near a provider's location while she measures, or a traffic accident) can affect an observation or an estimate at any time. Due to this dynamism, there will always be a delay between the event observed by the information provider, its evaluation and what actually happens. Therefore, the quality and accuracy of the reported observation cannot be verified through external means (even if they are available to the information buyer) and those interested in acquiring it must evaluate it based solely on the resources they find within the network such as the reports of the providers.

To address the above challenges of unknown costs, selfish behaviour, constrained resources and dynamic environments, we believe that citizen sensor networks should be viewed as multi-agent systems (MAS) (Lesser et al., 2003). We believe MAS offer a natural framework for describing such systems since citizens can be considered as self interested and autonomous entities interacting with each other and making decisions in possibly dynamic environments where uncertainty is endemic. Given this, in section 1.1 the aims of this research are explained in more detail. Section 1.2 then goes on to describe the research challenges, outline the techniques that currently exist for dealing with these challenges and explain how they are going to be extended so they can lead to the solution which is summarised in the end of the chapter. Finally, section 1.3 outlines the structure of the rest of the thesis.

1.1 Research Aims

In this research, our aim is to design the fundamental technologies that will govern the interactions among rational and self-interested information providers and buyers in next generation citizen sensor networks. The next generation citizen sensor networks will gradually replace current systems as the number of commercial applications and their scale increases. The main characteristic of the future networks will be the emergence of selfish behaviour among information providers and buyers, which will become apparent due to contradicting motives among the involved parts. In an environment where uncertainty is endemic and resources are constrained, this selfish behaviour becomes a dominating factor. In more detail, we consider networks, in which an information buyer (referred as centre) is interested in acquiring information (i.e. noise monitoring or traffic estimates) whose generation involves a certain cost. Now, as this cost is private information, known only to each individual provider, it is important to develop a process that would allow the information buyer to identify the agent that can provide it at a minimum cost.

Moreover, given that the agents may have contradicting motives, their selfish behaviour needs to be addressed by first incentivising them to allocate costly resources in generating an observation that is sufficiently precise and then by incentivising them to truthfully report it. It is particularly important to do so, in order to guarantee that the information buyer will receive accurate information at its desired degree of precision and the information suppliers will receive a fair reward that reflects their effort. Failing to do so, will result in the buyer receiving information of questionable quality which would not be of any use.

In the same context, the effect of agents' selfish behaviour is increased in an environment where they face constraints in the means they have at their disposal in order to produce the required observation. These constraints are reflected in the maximum precision of the observations that the information providers can provide, and therefore the information buyer may need to select multiple agents and fuse their individual results in order to obtain a more accurate observation at the precision it requires.

While the data fusion community has long studied the fusion of multiple noisy observations (see Zhao and Guibas (2004) for an overview), the move to multiple stakeholders introduces additional challenges. In more detail, fusing incorrect observations would degrade the network's overall performance since fusing deliberate erroneous observations would result in an incorrect final observation. This challenge becomes a dominating factor when the information buyer has no external means of evaluating the observations reported by the agents. In this case it has no access to knowledge of the state of the world and therefore has to evaluate an agent's report based only on the reports of the other agents. This makes citizen sensor networks more vulnerable to selfish behaviour, since evaluating an agent using an incorrect observation would widen the gap between the reported observations and the event that materialised, thus resulting in irrelevant and non-realistic results.

Therefore, against this background, a number of key requirements which form the objectives of this research become apparent. Specifically:

Requirement 1 Those interested in acquiring an observation need to elicit the costs of the various information providing agents, so they can identify those that can produce the estimate at a minimum cost.

Requirement 2 Agents should be expected to act in their own best interest. Thus they need to be incentivised to invest effort in the generation of the observation and report it truthfully to the information buyer.

Requirement 3 Information buyers need to be able to operate in environments whereby the information suppliers provide observations of limited precisions and hence combine multiple observations of possibly low quality.

Requirement 4 Information buyers need to be able to operate without knowledge of the state of the world and therefore develop processes to evaluate the agents based solely on their reports without access to external means.

In this context, it is clear that the agents' possible selfish behaviour is interconnected with all other requirements. Specifically, due to limited capabilities and lack of knowledge of the outcome, the consequences of selfish behaviour are magnified, when inaccurate and erroneous observations are fused with accurate ones. Therefore it is natural to assume a lack of trust on behalf of the information buyer towards the providers for two reasons. First, the information buyer cannot be sure about the amount of the resources invested by the providers for the generation of an observation, since there is always the possibility of a provider investing minimum resources and then after it is selected to provide that observation, to get paid for it disproportionally. Second, the information buyer cannot be sure whether that reported observation is indeed truthful, given that providers may misreport if they expect to benefit by doing so. Therefore, it is important to introduce *trust* between the information buyer and the providers through the elimination of selfish behaviour as a result of incentivising the information provider to not only truthfully report their observations, but also to invest effort and costly resources when producing them. In general, trust in an agent is an agent's belief that the other agents will fulfil their task as agreed (Dasgupta, 1988). This definition fits well in a MAS framework where we expect that some agents will attempt not to complete their task (i.e. an observation of an event) as a result of the two aspects of the selfish behaviour analysed earlier in this chapter. Therefore, trust-related solution concepts are briefly outlined in the next section and it is explained where existing approaches fail with respect to the above requirements. This, in turn, highlights the areas in which the contributions of this research will be made.

1.2 Research Contributions

We have identified four requirements that must be satisfied in order to design effective citizen sensor networks that can handle information providers' possible selfish behaviour and can face challenges such as unknown costs, constraints in their available resources and lack of knowledge of the outcome of the observed event. Now, since addressing selfish behaviour involves balancing this requirement against all others, the requirements cannot be addressed separately. Therefore, we focus on trust in MAS in order to address selfish behaviour, but we review existing approaches against all our requirements.

1.2.1 The Introduction of Trust in MAS

Against this background, as noted by Ramchurn et al. (2004), there are two main lines of research in trust within MAS. The first is the 'individual-level trust' whereby each agent maintains beliefs about the honesty or reliability of the other agents. Such trust models are often used when deciding between a number of alternative interaction partners. For such models to work, the agents must have access to a history of earlier interactions in order to make these decisions. This history could be their own direct experiences or information gathered by other sources. In either case, based on this information, an agent will decide whether it will or will not select a particular agent to produce the required estimate of forecast. Both these processes introduce a computational burden which results in additional complexity since now referring to the history of interactions is part of the decision process.

Furthermore, in some of these computational trust models, agents require multiple encounters with others in order to devise a strategy. However, this is inapplicable in the setting we consider, where an agent may not have the available resources to evaluate its encounters with other agents. This issue is magnified when citizen sensor networks operate in dynamic environments with constrained resources, since more resources are required to keep up with the constant change of the observed parameters so the agents can calculate their actions fast enough in order cope with their evolving surroundings.

Since several key requirements identified in section 1.1 are not met by this class of models, we move to 'system-level trust' as identified by Ramchurn et al. (2004). In this approach, participating agents are given incentives to be trustworthy by the rules of interactions governing the system. In more detail, in most cases, system-level trust is based on specifying the protocols regulating the interactions among agents, as opposed to the agent specific approach of individual-level trust. Given this, the use of mechanism design (MD) (Mas-Colell et al., 1995), which is the branch of game theory concerned with the design of a set of rules of interactions among agents in order to enforce desirable system-wide properties, is particularly appealing. Specifically, MD is useful in this case because, as opposed to individual-level trust models, the effects of an agent's actions on another can be modelled without concerns about each single agent.

In more detail, a common class of mechanisms are those based on auctions, such as VCG mechanisms (Vickrey, 1961; Clarke, 1971; Groves, 1973) and English and Dutch auctions (Krishna, 2002). Specially VCG mechanisms have a central role in mechanism design since they motivate truthful reporting on behalf of the agents through carefully selected payments and they have been widely used in resource or task allocation scenarios (see Section 2.4.2 for more details). However, they have also several negative aspects and cannot be directly applied in the cases we consider. In particular, they are not applicable in cases where an agent holds information which, if known to another, will affect its valuation (i.e. there is interdependent information). Now, this is often seen in citizen sensor networks, particularly in cases where agents have no precise knowledge about the state of the environment in which they operate. For example, in a noise pollution monitoring example, some of the participating members of the public are expected to be located within short distance from each other. The fact that multiple providers can generate an observation of the same event simultaneously without specific infrastructure is one of the main reasons that citizen sensor networks are preferred for these tasks over conventional sensor networks. However, since the observations are of the same event but from multiple angles, they will be dependent on each other and an information provider may be inclined to change its observation so it is similar to the observation of another agent if it expects to benefit by doing so.

This particular issue is being addressed to some extent by several mechanisms (Mezzetti, 2004, 2007; Ramchurn et al., 2009) that are closer to our aims. However, these mechanisms, either ignore the costs in the generation of an observation, or assume that agents will invest the maximum resources available to them for the completion of their task. In addition, all the above approaches assume that the information buyer will be able to determine whether an agent has successfully fulfilled its contract or not.

The latter assumption regarding agents' access to knowledge about the state of the world, is also found in mechanisms where payments are calculated with the use of *scoring rules* (Hendrickson and Buehler, 1971). Scoring rules incentivise an agent to truthfully reveal a probabilistic estimate of some future event. For example, in meteorology an agent that makes accurate sea-level pressure forecasts (Gneiting and Raftery, 2004) will get a higher score than an agent that did not. So far scoring rules have primarily been used in weather forecasting as a tool of eliciting truthful reports. In particular, a very popular scenario is that of rain prediction discussed in detail by Hendrickson and Buehler (1971) . In this case, a number of agents are asked to predict whether it is going to rain the next day and each agent is assigned a score depending on whether the event they predicted and reported materialised or not. Using this score as the basis of a payment in a mechanism, the agent is motivated to be truthful since the payment will be larger if his prediction is close to the truth. Such a scheme should therefore encourage rational self interested agents to be truthful in their reports if they want to maximise their profit.

While the above standard approaches based on scoring rules have several of the features we desire, it is assumed that all participating agents have access to information relevant to the state of the world. To overcome this, more recent work, for example (Miller, Resnick and Zeckhauser (2007); Jurca and Faltings (2005*b*) and Zohar and Rosenschein (2008)), estimate the true information from the existing reported observations. In so doing, however, they introduce less robust and generic solutions since they are based on a number of additional assumptions such as common priers among all agents, common knowledge of each agent's measured precision or knowledge of the distribution of the cost involved in the generation of the estimates (see Section 2.4.3 for more details).

1.2.2 Contributions

Against this background, we use the above general line of work and specifically Miller et al.'s (2007), as a point of departure in order to deal with self-interested agents making imprecise observations in a dynamic environment. In so doing, we contribute to the state of the art and address gradually all four requirements by developing three novel two-stage mechanisms.

1.2.2.1 Dealing with Unknown Costs

The first group of contributions in this thesis addresses requirements 1 and 2 by introducing the first mechanism that elicits both effort and truthful reporting of a single agent's observation, in a setting where the centre has no information about the agents' costs involved in the generation of that observation. In more detail:

- We describe a novel two-stage mechanism in which a centre uses a reverse second price auction in the first stage to elicit the true costs of the agents, and hence identify the agent that can provide an observation with a specified precision at the lowest cost. An appropriately scaled strictly proper scoring rule is then used in the second stage of the mechanism to incentivise this agent to generate and truthfully report its observation.
- We formally prove that this mechanism is incentive compatible in both costs and observations revealed, and that it is individually rational. That is, agents will truthfully report both costs and observations to the centre, and willingly participate within the mechanism.
- We are the first to empirically evaluate a parametric family of strictly proper scoring rules by comparing them to the standard quadratic, spherical and logarithmic scoring rules. We show that for certain values of the parameter, the resulting payment is similar to the logarithmic (optimal scoring rule) but also has finite lower bounds (as opposed to the logarithmic rule, which is potentially unbounded). Hence, the parametric scoring rule is the most appropriate choice for an information buyer who wants to minimise the payments it issues, while maintaining finite bounds.

Chapter 3 concentrates on this topic and the following work resulted in the following publication:

> A. Papakonstantinou, A. Rogers, E. H. Gerding, N. R. Jennings (2008) A Truthful Two-Stage Mechanism for Eliciting Probabilistic Estimates with Unknown Costs, *In Proceedings of the 18th European Conference on Artificial Intelligence (ECAI 2008)*, pages 448-452, Patras, Greece.

1.2.2.2 Dealing with Multiple Agents

In extending our initial mechanism we additionally address requirement 3 by presenting the first class of mechanisms that elicits private observations from multiple agents in a setting where the centre has to combine several observations of possibly low precision due to agents' restrictions in the quality of the observations they provide. In more detail:

- We present a novel two-stage mechanism in which a centre in the first stage asks the agents to report their costs and uses these costs to pre-select M agents. In the second stage the centre sequentially asks the M pre-selected agents to report their maximum precisions until it achieves its required precision (or it has run out of agents, in which case the centre will have to suffice with the closest precision to its required one). The agents selected through that process are asked to generate an observation of precision equal to their reported maximum precision, and then are paid for it using an appropriately scaled strictly proper scoring rule.

- We prove that our mechanism is incentive compatible in costs, maximum precisions and observations revealed, and that it is individually rational. Furthermore, we show that the agents maximise their expected utilities by generating observations of precision equal to their reported maximum precisions.

- We introduce a family of processes by which the centre may pre-select the M from N agents. Within this family, the centre divides the agents into groups of $n \leq N$ agents, asks them to reveal their costs, and then selects the m cheapest agents. The $(m+1)^{th}$ cost is then used within the subsequent scoring rule payment. We empirically evaluate this family for various values of the parameters n and m, and we calculate the total expected payment made by the centre, and the probability that it actually achieves its required precision, $P(\theta_0)$. We show both empirically and analytically that if the centre forms a single group of agents such that $n = N$ and $m = M$, it minimises its expected total payment.

This work is discussed in Chapter 4 and has produced the following publication:

> A. Papakonstantinou, A. Rogers, E. H. Gerding, N. R. Jennings (2009) Mechanism Design for Eliciting Probabilistic Estimates from Multiple Suppliers with Unknown Costs and Limited Precision, *Proceedings of the Eleventh Workshop in Agent Mediated Electronic Commerce (AMEC 2009)*, pages 111-124, Budapest, Hungary.

1.2.2.3 Dealing with the Centre's Lack of Access to Knowledge of the Outcome

Finally, in extending the above mechanism we address requirement 4 and introduce a novel mechanism in which the centre does not rely on knowledge of the realised outcome when calculating payments to the agents reporting their observations. Furthermore, we modify strictly proper scoring rules accordingly so they can motivate agents to truthfully report their observations, under the knowledge that they will be evaluated based on the other agents' reports. In more detail:

- We present a novel two-stage mechanism that is similar to that above in which a centre in the first stage pre-selects M agents from the available N agents by eliciting their costs. However, in the second stage, after it randomly selects the agents it needs to achieve its required precision, it uses a modified scoring rule to calculate the payment to each agent using the fused estimates of all the other selected agents.

- We prove that in our mechanism truthful reporting of costs is a dominant strategy, while truthful reporting of observations and maximum precisions is a Nash equilibrium. Finally we show that the mechanism is individually rational.

- We compare both this mechanism with the previous one and the peer prediction mechanism introduced by Miller, Resnick and Zeckhauser (2007) and show that the agents derive the same payment in all three mechanisms, hence the centre derives no additional penalty as a result of the lack of knowledge. However, we show that our approach yields significantly smaller variance in the total payments than the peer prediction in this case, and therefore it has more reliable and robust payments.

This work is discussed in Chapter 5 and has produced the following publication:

A. Papakonstantinou, A. Rogers, E. H. Gerding, N. R. Jennings (2009) Mechanism Design for the Truthful Elicitation of Costly Probabilistic Estimates in Citizen Sensor Networks, *Submitted in Artificial Intelligence, review pending*

1.3 Thesis Structure

The remainder of this thesis progresses as follows:

- In Chapter 2 relevant literature is discussed and analysed. We discuss citizen sensor networks related literature against our requirements set in this chapter and introduce and detail two applications of citizen sensor networks central to this thesis. We then review against the same requirements relevant research on trust by initially focusing on agent-level trust. After analysing various trust models and identifying their positive and negative aspects, we review the system-level trust models. We focus in particularly on Mechanism Design Trust and mainly on the class of second price auction based mechanisms due to its truth elicitation property. During that process, key definitions from mechanism design are introduced. An introduction to scoring rules and the review of existing mechanisms based on them follows and concludes the chapter.

- In Chapter 3, we describe and analyse the single agent two-stage mechanism with unknown costs. After a detailed description of our model, we introduce the mechanism and provide proofs of its economic properties. Finally, we present empirical results, where we compare the parametric scoring rule with the quadratic, spherical and logarithmic.

- In Chapter 4, we extend our mechanism, and present a class of mechanisms where multiple agents can provide estimates of limited precisions. We empirically compare several approaches to perform the pre-selection, hence the class of mechanisms, and identify one that minimises the centre's expected total payment.

- In Chapter 5, we further extend our mechanism so the centre does not have to rely on knowledge of the actual outcome when calculating payments. We introduce a new set of strictly proper scoring rules, prove the economic properties of a mechanisms based on them, and empirically evaluate that mechanism.

- In Chapter 6, the conclusions and future directions of the research are presented.

Chapter 2

Literature Review

In this chapter we examine how trust is addressed initially in citizen sensor network and then in multi-agent systems in general. In particular, we first define citizen sensor networks and then proceed to discuss their applications and how they currently deal with the requirements that were identified in Section 1.1. In so doing, we identify the issues associated with trust as one of the key areas that requires further work. Given this, we then analyse the literature related to trust in MAS against the requirements we have set for this research. In this context, in Section 2.2 we review computational trust models and specify the main areas of research in that field. In Section 2.3 we adopt a more systematic approach to trust in MAS and review several reputation models that motivate participating agents to provide true feedback about the services they have experienced though a carefully selected set of global rules. We then consider the notion of trust within the mechanism design literature (section 2.4) and focus on auction and strictly proper scoring rule based truth elicitation mechanisms, as they provide the main point of departure for this Thesis. Finally, we conclude in section 2.5 and discuss which aspects of the existing research are going to be used in this work and how they need to be extended in order to fit the requirements identified in chapter 1.

2.1 Citizen Sensor Networks

As discussed in Chapter 1, advancements in science and technology have led to an unprecedented rise in communications and, as a consequence, to the emergence of citizen sensor networks, in which members of the public can share their observations over a vast network of personal computers, mobile phones, game consoles and home entertainment devices. In such networks, citizens typically act simultaneously both as information gatherers and processors as they collectively observe, report and analyse information regarding a vast range of activities from every-day trivial tasks (such as commuting and shopping) to world changing events that affect thousands of people (such as natural disasters and terrorist attacks), or instead communicate observations that can be used for commercial (such as weather forecasting services and traffic monitoring) or scientific purposes (such as astronomic and geographic observations or noise pollution monitoring). To this end, in Section 2.1.1 we discuss applications of citizen sensor networks related to observations of ongoing events, while in Section 2.1.2 we discuss applications related to scientific observations. Furthermore, in Section 2.1.3 we thoroughly describe two specific applications of citizen sensor networks, and finally in Section 2.1.4, we discuss how selfish behaviour is addressed in citizen sensor networks related literature.

2.1.1 Citizen Journalism

The increased participation of citizens in the process of gathering information has led to the appearance of a new, more direct, type of journalism, which describes events as they unfold and as they are experienced by ordinary people and not specialised journalists. Bowman and Willis (2003) define this type of journalism as *'participatory journalism'* and use it to describe the act of a citizen who takes part in the process of gathering, communicating and filtering news, in order to provide reliable and accurate news from multiple angles. Now, examples of participatory journalism include community powered news websites (Wikinews and Newsvine), social networking websites such as Facebook and Twitter and web portals belonging to mainstream media corporations such as BBC and CNN, where members of the public report their experiences.

In participatory or 'citizen' journalism, as opposed to traditional sensor networks, people are responsible for gathering and filtering the observed information. This is a significant improvement, since, even if human operators gather information through devices such as mobile phones and videocameras, we are still better at reacting to what we observe, due to our ability to identify what is interesting and put it into context. In addition to that, through citizen sensor networks it is possible for an event to be covered from multiple angles as many citizens may contribute to its coverage (e.g multiple reports in Flickr and Twitter of the Mumbai terrorist attacks). Subsequently, this satisfies the research requirement regarding the combination of multiple observations, since through that process the accuracy of the observation (in citizens journalism a report about an event) can be increased.

However, citizen journalism faces a number of drawbacks, which are closely connected with the reliance on humans as information suppliers and processors. Indeed, according to Sheth (2009), off-topic comments and disambiguity as a result of poor grammar and syntax sensitive to demographic characteristics such as nationality, race, religion and gender, are a common source of incomprehensible observations. Now, although these issues do not affect directly any of our research requirements, they do compromise the robustness of the citizen networks. Indeed, combining ambiguous observations with more accurate ones will decrease the overall performance of such a network, since the produced observation will not be an accurate representation of reality. Therefore, it is crucial to review literature that addresses this drawback by developing processes that will identify and disambiguate these observations.

2.1.2 Citizen Science

Sheth et al. (2008) discus citizen sensor networks in which the gathering of the observations is performed from machine sensors being controlled by humans and not directly from humans as in typical citizen journalism networks. In so doing, they primarily address the issue of disambiguation as the observations can now be represented numerically and hence are impervious to poor grammar and syntax. Secondarily, they highlight the importance of machine sensors in cases where data cannot be collected otherwise; due to hazardous environment or extreme conditions such as floods (Zhou and De Roure, 2006), glaciers (Martinez et al., 2004) or volcanoes (Werner-Allen et al., 2005). However, they identify the issue of lack of integration and communication among various sensor networks that often measure similar parameters in the same or neighbouring systems, that may occur in such systems. This, in turn, leads to the problem of isolation and neglect of important data which Sheth et al. (2008) define as the problem of '*too much data and not enough knowledge*'. They address this problem though the introduction of a '*semantic sensor web*' whereby sensor data is digitally enriched with semantic metadata in order to make sensor data from a particular network easier to interpretate and allow it to be combined with data from other sensor networks. These methods can be applied even in cases where sensor data from a particular network follows a distribution, but it is not consistent with the distributions of sensor data in other networks. As a consequence, these significant contributions result in citizen sensor networks in which data is consistent and therefore make it possible to combine multiple sources of information (Requirement 3). However, all these techniques create an additional overhead by introducing an extra layer of complexity through the manipulation of the existing data. This could make it impossible to deploy such citizen networks in dynamic environments where a sensor may not be able to both adapt in the constantly evolving conditions and process incoming data, thus contradicting requirement 4 of our research regarding dynamic environments of endemic uncertainty.

This issue is addressed in another application of citizen sensor networks. Goodchild (2007) refer to *'citizen science'* to describe groups (networks) of enthusiastic amateurs who enjoy taking part in scientific activities such as astronomy or noise pollution monitoring, in their free time. In such cases the issue of observations being of poor quality that bedevils citizen journalism applications appears to have been solved through the establishment and enforcement of unified standards (e.g. the use of standardised metrics to define distances among planets and probability distributions to measure physical parameters). Moreover, this is also ameliorated by the restricted nature of these communities which require if not a high level of skill, certainly a strong commitment before joining them. That is, it is more difficult for an individual to join an amateur science society, than it is to blog about a random experience.

In this context, Goodchild et al. (2007) focus on collecting and sharing geographic information through a platform, sponsored by the U.S. Federal Government: the Geospatial One-Stop

(GOS[8]), which provides access to geographic information services through directories that contain metadata harvested by an automated process from specific registered websites. Like all previous work, they also identify the need for interoperability among different sources of geographic information for each system they discuss. In a similar approach to this issue with the already reviewed literature, they propose the association of current data with extra information (metadata), or the conversion of all available data to more consistent forms that can be shared among various sources. However, as opposed to the literature reviewed so far, Goodchild et al. are the first to identify the costs attached to this process. Specifically, they acknowledge that the system designers may be inclined to limit their platform to a single source of information if the cost is passed to them, while if the costs are passed to the end users they may be discouraged from contributing their information. In both cases the result will be the same, the potential user base of such networks will be reduced to a minimum size. That does not contradict directly any of the research requirements we have set in Chapter 1, but it will reduce the effectiveness of a citizen sensor networks, as less sources of information may result to observations of low quality, given that it will not be possible to increase accuracy through the combination of multiple observation.

Against this background, we have given several applications of of citizen sensor networks both from citizen journalism and citizen science[9]. In this research, we will combine elements of both. That is, we intend to use an open large scale application where everybody is free to participate with no restrictions of citizen journalism networks, but, with observations of a standardised and uniformal format which can be modelled as probability distributions. In this context, in the following section we present and analyse two applications of citizen sensor networks. Through this analysis we intend to demonstrate which of the properties we seek in citizen sensor networks, as described in the Introduction, are met by the current state of the art applications of citizen sensor networks. Furthermore, these two applications will form the basis of the running examples which will be used throughout the rest of the thesis. Through these two examples we intend to establish how our contributions signal the move towards the next generation citizen sensor networks.

2.1.3 Applications of Citizen Sensor Networks

The first application considers a network responsible for the monitoring of noise pollution in urban areas. Under European Union regulations, each member state is responsible for the monitoring of sound levels in busy city areas such as airports, train and coach stations and motorways, in order to identify and document the sources of noise pollution and therefore improve the living

[8] `http://gos2.geodata.gov/wps/portal/gos`

[9] Further related literature to citizen sensor networks and in particularly to those related to geography, can be found in Maguire and Longley (2005), where there is a review of a number of web-portals that specialise on content and services related to geographic observations ('geoportals'), and several government initiatives that provide the infrastructure to promote and encourage geographical observations ('geospatials').

standards of European citizens by reducing noise in those busy areas. Their efforts are concentrated in the project NOISE[10] which is responsible for accumulating and managing data from the member states.

However, in the above project there are many restrictions. Most notably, the datasets are restricted to specific areas such as airports and coach stations which although suitable for some EU cities, are not suitable for others (i.e. Athens Airport is in a secluded area with no settlements around it, while Heraklion Airport is situated within a dense residential area). Therefore, the accumulated data is not accurate or representative of a city's true noise levels. Furthermore, data often is found to be outdated as the measurements are sporadic and not very frequent. Finally, budget constraints and other bureaucratic reasons occasionally halt the project.

These issues appear to be addressed by a research effort which introduced NoiseTube. Noisetube monitors noise throughout a city and not just in specific places. Noisetube gives an insight to citizen sensor networks of the future, as members of the public can easily download and install an application in their smartphones which uses the devices' microphones to record sounds and combine their GPS capabilities and internet access to apply a specific tag which ties the signal to a location through Google Maps. In this network none of the disambiguity or lack of integration issues, discussed above, are faced since the observations have a consistent format (sound decibel and a GPS tag). However it is based on the participation of a closed group of volunteers who are not concerned by the time involved in making a recording and communicating it. It is easy to see though, that if more people are to be involved, and thus utilise the true strength of citizen sensor networks, a payment as a way of covering the potential contributors' costs in time, could incentivise more to participate.

The second application involves monitoring of traffic conditions in cities. In particular Houston, Texas has deployed on top of traffic signs and under bridges, devices which collect real time information about traffic (denoted as receivers). The relevant traffic data is transmitted by transponders which belong to drivers subscribed in an automatic toll collection system using electronic tags. However, this system[11] has very limited functionality. It is restricted to specific motorways as the receivers have fixed position. Furthermore, the receivers are visible and therefore vulnerable to traffic accidents or vandalism. Finally, only a selected group of drivers (those who commute through the specific tolled road) are equipped with the transponders, so the system cannot create an accurate and realistic representation of the traffic conditions throughout the city. Therefore, it is not possible to produce estimates of traffic and suggestions of optimal routes.

The latter issue, has been addressed by the commercial application TrafficCast[12]. TrafficCast demonstrates how a commercial citizen sensor network could operate in the near future by collecting and combining various sources of information, from fixed sensors, to GPS smartphones and bluetooth transmitters. Although this application satisfies our requirement regarding the

[10]Noise Observation and Information Service for Europe `http://noise.eionet.europa.eu/`

[11]Houston TranStar AVI Traffic Monitoring System: `http://traffic.houstontranstar.org/aviinfo/avi-tech.html`

[12]`http://trafficcast.com/`

combination of multiple sources of observations, it does not satisfy our requirement regarding selfish behaviour. In more detail, to the best of our knowledge TrafficCast, and other similar services such as traffic.com and Inrix, are not in a position to provide robust and simple methods to guarantee the quality of the services they provide. Instead, they rely on statistical models and heuristics[13], which even if they are effective in particular cases, they cannot provide any guarantees of their overall effectiveness. Therefore, it is impossible for a subscribed driver, acting as the information buyer, to establish the difference between a misreporting information provider and a provider who produces a poor observation as a result of physical obstacles or limited resources. As a result of this, they are vulnerable to selfish behaviour on behalf of the providers, since they cannot identify whether an information provider has invested its maximum resources in the generation of that estimate or not. This is particular important in a commercial citizen sensor network, as failing to provide a service accessed through subscription will probably result in financial losses.

2.1.4 Discussion

In summing up, in this section we have reviewed existing literature regarding aspects of citizen sensor networks (citizen journalism, citizen science and geographical information databases) and have introduced and detailed two specific applications (noise and traffic monitoring). Now, although these approaches address the need for multiple sources of information and, in some cases, the existence of costs in the process of making this information accessible by multiple platforms, it does not consider the issue of selfish behaviour. In more detail, current applications of citizen sensor networks fail to cope with information providers who may choose to deliberately misreport their observations, or to not allocate sufficient resources in the generation of that information if they expect to gain benefits from it. As commercial applications of citizen sensor networks are still in embryonic stages, the implications of not addressing selfish behaviour are not visible and they appear not to be seriously affecting existing networks which are primarily based on the participants' enthusiasm and self imposed discipline. However, as commercial citizen sensor networks find increasingly more applications, and the number of the information providers increases, the emergence of selfish behaviour will severely degrade the efficiency and performance of such networks, since erroneous observations will be of minimal practical use anybody interested in acquiring them.

Given these consequences of selfish behaviour, it has become apparent that it is of great importance to address both aspects of selfish behaviour when designing the protocols that will guide the interactions among information buyers and providers in the next generation citizen sensor networks. Therefore, and after having identified the links between selfish behaviour and the absence of trust in Chapter 1, in the following section we proceed to explore the main approaches to trust in multi-agent

[13]In Section 2.2 in this chapter we provide several statistical approaches and identify their drawbacks with respect to our research requirements.

2.2 Individual-Level Trust Models

From this section we start our discussion regarding trust in MAS. Specifically, we review and discuss individual-level trust models. In these models, an agent has two alternatives when attempting to choose which agent it should trust from a pool of available agents. First, it has the option of directly interacting with the other agents in order to learn their behaviour after several encounters. Second, it can indirectly interact with other agents by referring their opinions, and based on third-party information, make a decision. In the first case the agent decides whether to trust an agent or not after it *learns* through its interactions, while in the second case an agent relies on external *reputation* sources. In addition to that, recently, elements of both approaches have been combined in order to introduce more theoretical foundations based on a probability theory framework. As a result, individual or agent-level trust models can be classified into learning and reputation models, based on an agent's intention to interact directly or indirectly with other agents and their combination, the probabilistic trust models.

2.2.1 Learning Models

In learning models, trust is considered to be a property that comes from direct interactions among agents. In this case, trust is achieved on the basis that if an agent systematically chooses to defect by not completing an agreed task (such as producing a probabilistic estimate), other agents learn to avoid it. Therefore, an agent tempted to defect in a single interaction (e.g. cheat and not pay the agreed sum in an on-line transaction), should reconsider in order to avoid the loss of profit that will come if the other agents stop referring to it.

In this context, Wu and Sun (2001) have developed a model in which trust emerges among agents using reinforcement learning if they are left to adapt in the environment. In such learning problems, agents choose their next state so they can maximise the feedback they receive after each action. In particular, Wu and Sun showed that the evolution of strategies helps good agents who want to cooperate to eventually isolate the malicious agents that do not have any intention to do so. However, this takes time and there might be a short term loss of profit when cooperating with third parties. Some of these issues were dealt by Sen and Sajja (2002) who showed that collaborating liars have more to gain after a few iterations in environments with large numbers of philanthropic agents. Furthermore, Mui et al. (2002) extend the previous work by designing a probabilistic model that identifies the number of interactions needed for achieving trust. Another approach is followed by Mukherjee et al. (2001), who allow their agents to reveal their strategies to their opponents in advance, and therefore possibly develop mutual trust in fewer interactions.

Now, although learning models satisfy some of the research requirements we set out in Chapter 1, since they deal with selfish behaviour, they fail to satisfy other requirements. In particular, the reviewed learning models promote trust among agents, and hence satisfy the requirement regarding selfish behaviour, by identifying which agents will be completing an agreed task and which agents will choose to defect from that task. However, in all the above models, agents need to evaluate a large number of available actions and interactions in order to adapt their strategies, something that is not possible for citizen networks since they operate in constantly changing environments. In such environments, due to their dynamism, agents do not necessarily have the opportunity of using learning in order to discover a trustworthy source. Indeed, in a dynamic setting an agent does not have the option of refining its strategy through multiple encounters, since the initial conditions may change after each interaction. Besides that, in cases where agents do not have the opportunity to verify if an encounter was successful or not (like those we consider), they must take into consideration, and therefore evaluate, all possible outcomes. Doing so, will add an overhead and will make such models less able to adapt to real time events.

2.2.2 Reputation Models

In order to address problems like those mentioned above, we consider reputation systems as another way of achieving trustworthy behaviour when there are few encounters among agents or there is a lack of prior information about an agent's incentives. In such cases, an agent does not refer to its own direct history of interactions with a particular agent, typically because it is impossible due to lack of previous interactions or perhaps because it is too computationally intense due to a very large number of agents. Although this is particularly prevalent in on-line markets and services, where someone must choose a product without relying on any prior information or experience, it can be also seen in citizen sensor networks where often it is impossible to identify the most accurate information provider based on previous reported observations. In such cases, reputation models attempt to provide a solution to this problem by providing third-party information, to which agents may refer during their decision making process.

Now, one specific line of research in reputation models is based on the idea of social networks. Such models draw an analogy between the arrangement of agents in networks and that of humans in societies. In these networks, each agent is able to validate the result of each interaction and act as a reputation source by sharing these results with their neighbours. In this vein, Singh et al. (2001) introduce the use of 'referrals' to devise a distributed reputation model. In an agent-based referral network, agents can either give referrals to each other or follow referrals given to them by others. Nevertheless, as this approach requires knowledge about the result of the interaction after it is finished, it contradicts the requirement, regarding the uncertainty in the environment, which we have explicitly discussed in section 1.1. In particular, requiring precise information about the state of the world after an interaction is complete is often not feasible in citizen sensor networks, as the information buyer may not be able to have access to that information in order to evaluate the providers. For example, in a traffic or a noise monitoring scenario, the information buyer has no other means of acquiring the information it requires, and that is the reason it has to rely solely on the citizens' sensors.

Another line of research, within reputation trust, is based on aggregating ratings about an agent's or a service's performance. However, this generates problems which have to do with misreporting during the process of gathering that information. Specifically, in such models agents may provide incorrect ratings or not provide any ratings at all as a way of inflicting losses or adding pressure (i.e. extract forcefully good feedback) to other agents. This is a common practise on eBay for example, where ratings are values of +1 or -1 summed up to give an overall profile. To this end, eBay's reputation system can be exploited when people do not report their bad experiences, since the absence of ratings cannot be regarded as a good or bad overall rating and thus it cannot be used in the decision process. Yu and Singh (2002) deal with this by considering the lack of belief as a state of uncertainty where all beliefs have an equal probability of being true. All possible states of an agent's reputation can be combined and mapped into two states (trustworthy or not). However, their approach relies on a large number of agents in order to cope with those who discredit others so that they can gain benefit by appearing more reliable, and does not consider the aspect of selfish behaviour related to lack of effort.

2.2.3 Probabilistic Trust Models

A common characteristic of the so far reviewed learning and reputation based trust models is that they do not have sound foundations in statistics or probability theory. Although this is not necessarily a negative aspect, it can be claimed that these models are restricted to the specific cases they consider. In a more general approach, Josang and Ismail (2002) propose a probabilistic centralised computational trust model in which the users of an on-line system rate the performance of the other members of an on-line community. They propose the use of beta probability density functions to combine users' feedback which again takes two values (i.e. positive and negative) and provide reputation ratings. However, this model does not take into consideration misleading feedback from users. This algorithm is extended by eliminating the unfair ratings by identifying the statistical pattern of the users' ratings, based on their previous feedback (Whitby et al., 2004). In so doing, Whitby et al. introduce elements from learning trust models, but still do not cope with cases where there are no previous interactions, or a user is systematically unfair. Therefore, it is clear that these models do not address effectively the issue of selfish behaviour.

This trend of combining elements from both learning and reputation models is continued by Yuan and Sung (2004) who propose a model that is appropriate for dynamic environments such as virtual and real markets and can take into account both public reputation available to all agents and private interaction history common in the learning trust models mentioned above. Thus allowing potential trading partners to be evaluated before the interaction. Then, and based on real experiences, the initial opinion can be modified or not. However, they contradict the requirements we set in Chapter 1 regarding costs and selfish behaviour, as they assume that there are no costs involved in the generation of the participating agents' ratings of the traded service or object and consequently assume that all agents will invest the maximum resources in the generation of their observations.

This appears to be a common problem in all the individual-trust models we have considered so far. Even in more recent work, like the TRAVOS model introduced by Teacy et al. (2006). In more detail, in TRAVOS, agents base their decisions on their previous experiences with the other agents, as well as on third-party experience with the trusters having a binary approach about a trustee's decision. That is, a truster makes the assumption that a trustee can only cooperate or defect, and therefore complete the assigned task or not. The same binary approach is followed by Wang and Singh (2007) who try to estimate a reputation rating when there is a lack of previous positive or negative feedbacks. This model also takes into consideration existing feedback when it is available in order to produce more accurate ratings when it is possible. Still it is restricted to the binary space and therefore may not be applicable in citizen sensor networks where often the observations have the form of continuous distributions. This particular shortcoming is addressed by Teacy (2006) and their extended TRAVOS-C model. In this, trustees have continuous action spaces, for example, the delivery time of a service or a metric about the quality of a product. Furthermore, there is an improvement in the assessment of the accuracy of the reputation sources by combining different sources and ignoring deliberately misleading reports. However, trusters can deliberately 'over-trust' a trustee if they have a large number of observations available and therefore the requirement of selfish behaviour is not fully satisfied, since the opinion of an agent is dependant on the quantity of the observations and not necessarily on their quality.

More recently, Jones et al. (2009) proposed a framework to address this issue by combining multiple trust models and not just one as commonly seen in the relevant literature. Furthermore, they propose the use of 'trust communities' whereby groups of agents can combine their observations and provide collective information of increased quality. Although this is directly relevant to our requirement regarding the fusion of multiple observations of possibly low quality, this work does not satisfy other requirements and specifically the one regarding the application of citizen sensor networks in dynamic environments, since a protocol that would have to use multiple trust models would not be able to monitor the constantly changing parameters in a dynamic system such as the streets of a city.

In total, these probabilistic trust models can be viewed as a significant improvement, given that they don not rely on the evaluation of a large number of actions in order to identify if an agent should be trusted or not, and hence are more robust than the learning models. Nevertheless, like all the reputation models we reviewed in this section, in relying on precise information or knowledge about the state of the world after the interaction is finished, they contradict our key requirement of outcome uncertainty.

In this context, there are advanced probabilistic trust models such as TRAVOS-C (Teacy, 2006) and IHRTM (Rettinger et al., 2008) that do not rely on this assumption as they use statistical methods to approximate the state of the world after the interaction is finished. However, although they acknowledge the existence of costs in the communication of the observation, they do not consider the costs involved in the generation of that observation. By doing so, they assume that agents will be willing to invest their maximum resources in giving their feedback regarding their experiences, which contradicts one element of selfish behaviour. Furthermore, they assume that agents would always be able to report their observations, which is something that does not always hold in citizen sensor network. Indeed, in such networks, there may be cases where this is impossible due to physical obstacles or random events, or even, given their rational and self-interested nature, by choice in order to save resources such as bandwidth or power.

2.3 System-Level Trust Models

From the analysis of all aspects of individual-level trust models, it is clear that they do not cover some of our key requirements identified in the introduction (section 1.1). Specifically, although they address some aspects of selfish behaviour, they either rely on a large number of agents or on precise knowledge about the outcome of the interactions among agents. Even the probabilistic trust models, which represent an improvement since they combine elements from both learning and reputation models and address some of the above requirements, allow agents to exhibit selfish behaviour and are limited to minimising the effects of this behaviour. Given these limitations, we need to move to a more systematic approach in trust that focuses on the design of protocols that guide the interactions among agents. Therefore, in this section we analyse and critique reputation models that introduce trust through a set of rules that endeavour to provider favourable global properties such as the investment of sufficient resources in the generation of an observation, or the truthful reporting of that observation. Although it may seem that by doing so we depart somewhat from the standard definition of trust in MAS presented in introduction (section 1.1), we do not, since truthful reporting of an agent's observation and trust are closely connected. In particular, if an agent's task is to produce and report an observation of an event, it successfully fulfils it, when it invests the necessary resources into that process and truthfully reports its observations.

In this context, Zacharia and Maes (2000) apply their models (SPORAS and HISTOS) in on-line communities like forums, mailing lists and chatrooms which can be regarded as citizens sensor networks (although their work precedes this term). Specifically, in SPORAS they propose a reputation mechanism in which new agents start with a minimum reputation value which is increased based on the ratings they receive. Also, the reputation of an agent already in the system can never be lower than the reputation of a new agent which just entered the system. Therefore, since they assume that two agents interact only once, it will always be better to trust the agent with the higher reputation which is always the agent that was in the system before the other agent. Zacharia and Maes claim that this biased treatment is a necessary trade-off for achieving anonymity in an on-line system. HISTOS is an extension where a more personalised reputation is used. In more detail, in HISTOS the reputation of each agent depends on how an agent rates other agents. HISTOS captures the common belief that in on-line communities people tend to value more the opinions of people they already know. The drawback of this system is that it does not penalise an agent that may lie and give good reports in an attempt to build good reputation since positive rating is more valued.

Another system-level reputation model which motivates agents to be truthful in their reports, is proposed by Jurca and Faltings (2005*a*). In their model, they motivate agents to be truthful through payments they receive when sharing feedback. This model is applied in a setting where service (i.e. P2P, grid computing, broadband) providers repeatedly offer their service to prospective clients. Providers can see whether the task was delivered or not, but those who constantly fail to deliver the task are not excluded from the market. However, like SPORAS and HISTOS, this model does not motivate agents to participate and act as reputation sources, which directly contributes one aspect of selfish behaviour, as agents do not invest any resources into the generation of their feedback.

To this end, in order to motivate a client agent to supply information after an interaction with a supplier, Jurca and Faltings (2006) extend their existing mechanism and guarantee that a truthful reporting agent will not lose utility and that those who report incorrect reputation will gradually lose utility. They achieve this by proposing a payment to the agents whose opinion about a provider matches with the opinion another agent gives about the same provider based on a Chi-square Independence test that examines the correlation between the two parties' opinions. So far, this model is the first model we consider within system-level reputation trust that does not rely on knowledge about the result of an interaction since it evaluates agents, based on other agents' subjective opinions. Furthermore, it guarantees that these subjective opinions are honest. However, they assume that all agents that share the same supplier, also have identical perception of the quality of this supplier's service. This assumption is not realistic in many citizen sensor networks where information providers make imprecise observations and hence can have different beliefs while observing the same event. Furthermore, they do not consider costs in the generation of the observation, and therefore contradict our requirement of selfish behaviour by subsequently assuming that given that there are no costs involved in the generation of an observation, agents will invest all their resources.

As it can be seen, the last two models we consider signal a transition to a more fundamental approach which is no longer restricted to a set of rules that can be applied only in specific scenarios. That is, through the application of techniques from the field of *mechanism design* (Mas-Colell et al., 1995), truthful reporting and voluntary participation of the agents in the reputation system can be guaranteed simply through an appropriately designed payment scheme. Thus, this provides the main point of departure for our work so we will expand upon it in more detail in the next section.

2.4 Mechanism Design based Trust

In this context, mechanism design is particularly appealing since it can be used as a tool to guarantee not only truthful reporting of an agent's observation, but also the fact that an agent will make the observation and invest sufficient resources in doing so. Although there are several incentive compatible mechanisms (such as the already discussed reputation mechanism proposed by Jurca and Faltings (2006)), we use auction-based[14] mechanisms as a starting point of our review. We do so because they have diverse applications from for task allocation problems that occur in MAS (Dash, Rogers and Jennings, 2004; Czumaj and Ronen, 2004), to information distribution in sensor networks (Klein et al., 2008) and allocation of multiple items in markets through sequential or simultaneous auctions (Fatima et al., 2005; Fatima, 2006) and also induce truthful reporting and voluntary participation. However, before we proceed into discussing and analysing relevant approaches in mechanism design, we need to define the fundamental terms in this area.

To this end, in Section 2.4.1 we provide the basic definitions that will be used in the rest of this chapter, while in Section 2.4.2 we review mechanisms based on incentive compatible auctions. After identifying which of the research requirements are met and which are not we introduce scoring rule based mechanisms (Section 2.4.3). Finally, since scoring rules are central in our work, we provide relevant background and a simple application that demonstrates how they can be used in mechanism design in order to induce effort on behalf of the agents in generating their observations and their truthful reporting (Sections 2.4.3.1 and 2.4.3.2 respectively).

2.4.1 Basic Definitions

Consider a set of n individual agents I which must make a joint choice from a set X of possible outcomes. Prior to the choice, each agent $i \in I$ privately observes its preferences over the set of possible outcomes X. This is formally modelled by supposing that agent i observes a parameter or a signal θ_i that determines its preference. Although θ_i represents private information available only to i, we assume, as it is standard, that it is drawn from a commonly known joint distribution. From now on we refer to θ_i as an agent's private type and to the set of possible types for agent i, as Θ_i, with $\Theta = \Theta_1 \times \cdots \times \Theta_n$ being the set of all the possible types of all agents.

In this context, a **social choice function** (SCF), is a function $f : \Theta \rightarrow X$, that for each possible set of agents' types, determines a joint choice $f(\theta_1, ..., \theta_n) \in X$.

Definition 2.1. Strategy
A strategy s_i is a complete contingent plan for an agent i. That is, a plan that specifies how agent i will act for every possible type in set Θ_i. Formally, a strategy for agent i is a mapping from Θ_i to M_i ($s_i : \Theta_i \rightarrow M_i$), where M_i, is the *message set* for agent i, and includes any messages that this agent communicates.

Definition 2.2. Strategy Profile
If $s_i(\theta_i)$ is one available strategy for agent i based on its type, the vector of agents' strategies is a *strategy profile* $s(\theta) \in M$, with $M = M_1 \times \cdots \times M_n$ being the set of joint messages.

[14]According to McAffe and McMillan (1987), an auction is a market institution with an explicit set of rules for determining resource allocation and prices on the basis of bids from the market participants.

Now, the function $g: M \rightarrow X$ that determines how agents' messages get turned into a social choice is called an **outcome function**, and the function $t: M \rightarrow \mathbb{R}^n$ that calculates monetary transfers issued to agents based on their messages, is called a **transfer action**.

Definition 2.3. Utility Functions
Now, $\forall x \in X$ an agent's utility function is defined as the function $u_i(x, \theta_i)$, with $u_i: X \times \Theta_i \rightarrow \mathbb{R}$. Each agent's utility function is publicly known and it represents a preference relation ($\succsim_i (\theta_i)$) over pairs of outcomes given the type θ_i. We also assume that each agent intends to maximise it utility in expectation.

Based on the above definitions of transfer and utility functions, we can define a specific type of utility function, which we will be using in the rest of this research, that is the *quasilinear utility function*. We focus on the quasilinear utility, given that it has been widely used in Auction Theory, since it ensures that agents can transfer utility through monetary side payments.

Definition 2.4. Quasilinear Utility
Given agent i's type θ_i, its quasilinear utility is denoted as follows:

$$u_i(x, \theta_i) = v(y, \theta_i) + t_i \tag{2.1}$$

where outcome x defines a choice $y \in Y$ from a discrete choice set and t_i is agent's payment ($\sum t_i \leq 0$). The function $v_i(.)$ is known as *valuation* function, and it represents agent $i's$ value for each choice $y \in Y$. In example 2.1, where we analyse the second price auction, y represents allocation and t_i are the payments made to the auctioneer.

Definition 2.5. Mechanism
A mechanism $\Gamma = \langle I, X, M, \Theta, g(.), t(.)) \rangle$ is a collection of agents, outcomes, agents' messages and types, together with an outcome and a transfer function.

In this context, the mechanism Γ is said to *implement* a SCF $f: \Theta \rightarrow X$ if there is an equilibrium strategy profile[15] $s^*(\theta)$, such that $g(s^*(\theta)) = f(\theta)$, with any $\theta = (\theta_1, \cdots, \theta_n) \in \Theta$.

Now, a special class of mechanisms are the *direct revelation mechanisms*, in which the only strategy available to the participating agents' is to report their types ($\widehat{\theta_i} = s(\theta_i)$, where $\widehat{\theta_i}$ is agent i's reported type). From a mechanism designer's perspective, these mechanisms are a valuable tool, as they can simplify the task of identifying the social choice functions which are implementable because it restricts the number of the possible mechanisms we have to consider (Mas-Colell et al., 1995).

Definition 2.6. Direct Revelation Mechanisms
A *direct revelation mechanism* is a mechanism in which $M_i = \Theta_i$ for all i and $g(\theta) = f(\theta)$ for all $\theta \in \Theta$.

[15] Later in this section we will introduce notions of equilibrium profiles such as dominant and ex post equilibrium profiles.

Although it is not explicit in the definition of a mechanism (Definition 2.5), many mechanisms not only consider transfers (payments or penalties) to the participating agents, but also the allocation of a good or a task (i.e. an observation of a physical parameter such as temperature or wind velocity). In order to show how allocations can be part of a mechanism, we provide an example of a second price auction mechanism[16].

Example 2.1 ***Second Price Auction Mechanism***
Suppose that there is a single unit of a good that is to be allocated to one of n agents by applying a second price auction. In a second price auction, the good is allocated to the agent with the highest bid[17], $b_i \in M_i$ with b_i being $s_i(\theta_i)$ for notational convenience. The winning agent then, is asked to pay the value of the second highest bid.

A particular element of the set of possible outcomes X in this particular case is represented by a vector $x = (y_1, ..., y_n)$, where $y_i = 1$ if agent i receives the good and $y_i = 0$ if it does not. In more detail, if $Y_i = \{0, 1\}$ is the set of possible results of the allocation for agent i and $M_{-i} = M_1 \times ... \times M_{i-1} \times M_{i+1} \times ... \times M_n$ is the set of joint messages of all agents besides ith, the function $y_i : M_i \times M_{-i} \rightarrow Y_i$ that allocates the goods to each agent based on its bid and the other agents' bids, is the **allocation function** and is denoted as follows:

$$y_i(b_i, b_{-i}) = \begin{cases} 1 & \text{if } b_i > \max_{j \neq i} b_j \text{ or } b_i = \max_{j \neq i} b_j \text{ and } i < j \\ 0 & \text{otherwise} \end{cases} \tag{2.2}$$

Moreover, in a second price auction agent i's transfer function $t_i : M_i \times M_{-i} \rightarrow \mathbb{R}$ depends on its bid and the other agents' bids and is denoted as follows:

$$t_i(b_i, b_{-i}) = \begin{cases} -\max_{j \neq i} b_j & \text{if } b_i > \max_{j \neq i} b_j \text{ or } b_i = \max_{j \neq i} b_j \text{ and } i < j \\ 0 & \text{otherwise} \end{cases} \tag{2.3}$$

In line with the general definition of the mechanism (Definition 2.5) and given that an element of X can be defined as $x = (y_1, ..., y_n)$, the outcome function of the mechanism $\Gamma = \langle I, X, M, \Theta, g(.), t(.) \rangle$ is the following:

$$g(\mathbf{b}) = (y_1(\mathbf{b}), ..., y_n(\mathbf{b}))$$

where $\mathbf{b} = (b_1, ..., b_n)$ is a vector of all agents' bids.

Now, in many cases (including the second price auction mechanism) the participation of the agents in the mechanism is voluntary, and therefore the mechanism must incentivise agents to participate in it. This property is known as '*individual rationality*' and for a mechanism to be

[16] Due to the central role of second price auction mechanisms in this research, we analyse and critique several variations of second price auction based mechanisms in section 2.4.2.

[17] Note that according to the literature (Krishna, 2002), we can safely ignore ties (when two or more bidders have the same highest bids), as it can get resolved by allocating the good to the agent with the lower index

individually rational, it must have an individually rational equilibrium strategy profile $s^*(\theta)$. In more detail, there are three individually rational strategy profiles based on which stage of the mechanism the agent chooses to participate in.

Definition 2.7. Individual Rational (IR) Strategy Profiles
As detailed by Mas-Colell et al. (1995), these are: ex post, interim and ex ante. Given that agents' utilities are quasilinear, IR strategy profiles are defined as follows:

- In **ex post** IR strategy profile $s^*(\theta)$, agents choose whether to stay or leave after announcing their types and learning an outcome from the set X of feasible outcomes. If $\bar{u}(\theta_i)$ is the utility for opting out[18], then in an ex post IR mechanism: $u_i(g(s_i^*(\theta_i), s_{-i}^*(\theta_{-i})), \theta_i) \geq \bar{u}(\theta_i)$ for all $\theta_i \in \Theta_i$ and $\theta_{-i} \in \Theta_{-i}$ with $\Theta_{-i} = \Theta_1 \times ... \times \Theta_{i-1} \times \Theta_{i+1} \times ... \times \Theta_n$.
- In **interim** IR strategy profile $s^*(\theta)$, the decision for each agent is made after learning its type, but before the outcome is calculated: $U(\theta_i|g) = E[u_i(g(s_i^*(\theta_i), s_{-i}^*(\theta_{-i})), \theta_i)|\theta_i] \geq \bar{u}(\theta_i)$ for all θ_i and $\theta_{-i} \in \Theta_{-i}$.
- In **ex ante** IR strategy profile $s^*(\theta)$, the agent makes its decision before knowing its type, therefore it must know the types' prior distribution: $U(\theta_i) = E[u_i(g(s_i^*(\theta_i), s_{-i}^*(\theta_{-i})), \theta_i)] \geq E(\bar{u}(\theta_i))$ for all θ_i and $\theta_{-i} \in \Theta_{-i}$.

An ex post IR strategy profile satisfies both the interim and ex ante conditions without the opposite being possible. Furthermore, interim IR is a stronger profile than ex ante, as in interim IR, agents agree to participate in the mechanism after learning their types, while in the ex ante, agents don't have access to this information when they consider participating in the mechanism and rely on knowledge about their types' prior distribution.

While individual rationality is a desirable property for a mechanism, it is not the only one. Often, a mechanism designer may require agents to be incentivised to truthfully report their types. This property is known as *'incentive compatibility'* and in an incentive compatible mechanism, each agent is motivated to reveal the truth about its type by having its utility function maximised. Furthermore, a social choice function is incentive compatible if truth telling by each agent is an equilibrium for mechanism Γ. More formally:

Definition 2.8. Incentive Compatible Mechanism
A mechanism is incentive compatible if it truthfully implements a social choice function. For direct revelation mechanisms, the social choice function $f(.)$ is truthfully implementable, if the mechanism has an equilibrium $s^*(\theta) \in M$ in which $s^*(\theta_i) = \theta_i$ for all $\theta_i \in \Theta_i$.

As our work progresses, we will encounter two types of strategy equilibrium profiles: *dominant strategy* and *ex-post Nash*. Therefore, we define them:

Definition 2.9. Dominant Strategy Equilibrium
In a *dominant strategy equilibrium*, an agent has the same utility maximising strategy to every

[18] The utility of an agent opting out from the mechanism does not depend on the outcome.

collection of strategies of other agents. Formally, the strategy profile $s^*(\theta) = (s_1^*(\theta_1), ..., s_n^*(\theta_n))$ is a *dominant strategy* equilibrium of the mechanism $\Gamma = \langle I, X, M, \Theta, g(.), t(.) \rangle$ if for every agent $i \in I$ and type $\theta_i \in \Theta_i$ and $\theta_{-i} \in \Theta_{-i}$:

$$u_i(g(s_i^*(\theta_i), s_{-i}(\theta_{-i})), \theta_i) \geq u_i(g(s_i'(\theta_i), s_{-i}(\theta_{-i})), \theta_i), \tag{2.4}$$

for all $s_i'(\theta_i) \in M_i$ and for all $s_{-i}(\theta_{-i}) \in M_{-i}$, with $M_{-i} = M_1 \times ... \times M_{i-1} \times M_{i+1} \times ... \times M_n$

Definition 2.10. Dominant Strategy Incentive Compatibility
A direct revelation mechanism Γ is *dominant strategy incentive compatible* (or *strategy proof*) if truth revelation is a dominant strategy equilibrium.

Now, in an ex post Nash equilibrium, each agent's chosen strategy is a best response to the strategies used by the other agents. In other words, no one can profitably deviate, given the actions of the other agents. In this particular case, it is required by the agents to select their strategy after they have knowledge about the strategies of all other agents, and therefore this equilibrium is weaker than the *dominant strategy* (Definition 2.9) which does not depend on other agents' strategies. A formal definition of *ex post Nash equilibrium* is the following:

Definition 2.11. Ex post Nash Equilibrium
The strategy profile $s^*(\theta) = (s_1^*(\theta_1), ..., s_n^*(\theta_n)) \in M$ is an *ex post Nash* equilibrium of the mechanism $\Gamma = \langle I, X, M, \Theta, g(.), t(.) \rangle$ if for every agent $i \in I$ and type $\theta_i \in \Theta_i$ and $\theta_{-i} \in \Theta_{-i}$:

$$u_i(g(s_i^*(\theta_i), s_{-i}^*(\theta_{-i})), \theta_i) \geq u_i(g(s_i'(\theta_i), s_{-i}^*(\theta_{-i})), \theta_i), \tag{2.5}$$

for all $s_i'(\theta_i) \in M_i$.

Definition 2.12. Ex post Nash Incentive Compatibility
A direct revelation mechanism Γ is ex post Nash incentive compatible if truth revelation is an ex post Nash equilibrium.

Finally, incentive compatibility in conjunction with direct mechanisms (definition 2.6) lead to a very important result in mechanism design: *'the revelation principle'*. According to the revelation principle, for any mechanism there exists a direct and incentive compatible mechanism that implements the same social choice functions in equilibrium.

2.4.2 Incentive Compatible Auction Mechanisms

Having given the basic definitions in mechanism design, now we start reviewing auction based mechanisms in order to identify whether they satisfy our research requirements or not. We focus on second price sealed auctions (Vickrey, 1961) (already described in Example 2.1) which have been widely used to incentivise truth-telling in dominant strategies, since bidders submit their sealed bids and the auctioneer allocates the resource to the highest bidder, who pays the second highest price. Now, there are several types of auctions considered in the literature (such as English, Dutch, first price, second price), which introduce trust in mechanism design from various perspectives. For example, some auctions motivate all agents to participate (i.e. second price auctions are individually rational), while other auctions incentivise the auctioneer (i.e. English and Dutch auctions) or the participating agents (second price auctions) to bid honestly based on their true valuation of a traded item. However, in this section we focus on second price auctions, since they are both incentive compatible and individually rational and therefore are a valuable solution to the relevant aspect of selfish behaviour in citizen sensor networks.

Nevertheless, besides the favourable properties of incentive compatibility and individual rationality for mechanisms that belong in this family, (i.e. the class of Vickrey-Clarke-Groves/(VCG) mechanisms (Vickrey, 1961; Clarke, 1971; Groves, 1973)) there are several shortcomings that are relevant to the context of this research.

First, second price auctions cannot prevent the auctioneer from lying. Specifically, since the bids are sealed and private, the auctioneer could ask for a higher price than the second price. The winner will pay the requested price, because he has no access to the bids of the other players. There are ways to address this, for example Hsu and Soo (2002) solve this without endangering the privacy of bids, just by randomly assigning the task of the auctioneer to one of the bidders and with the use of a public blackboard for agents to publish their bids, so the first and second highest winners can verify the results. Now, this solution introduces additional complexity and compromises the robustness of a citizen sensor network as the information buyer would have to constantly update this external source and the providers will have to refer to it. Although this does not directly contradict any of our research's requirements, it is a negative aspect of second price mechanisms which could be addressed by assigning the role of the centre randomly to a different agent. However, this is not viable in our setting, since the role of the centre is fixed to the agent interested in acquiring the information. In another approach, Suzuki and Yokoo (2004) propose an encryption method that will guarantee secure bidding without the participating agents knowing each others bids. Still, this method introduces multiple overlays of calculations which involve a new representation of the bids, an algorithm to identify the maximum bids, encryption and decryption processes distributed among different servers. Due to the complexity that is introduced as a direct result of these additional processes, this is not a suitable solution for a citizen sensor network which often has to rely on fast observations made by members of the public with portable devices and not dedicated computers with large processing power. Therefore, we leave this as an open question which will attempt to address in our future work.

The second, and perhaps most significant, drawback of VCG mechanisms, is the necessary assumption that agents can only hold valuations which do not depend on the valuations of other agents. However, this is often contradicted in the cases we consider. Specifically, in citizen sensor networks agents make noisy and uncertain observations, which lead to the possession of partial information about the value of the observed parameter. In such cases an agent may hold information, which if known to another agent might affect its valuation. For example, in a traffic monitoring scenario, the quality of a company's estimate of congestion (and hence its value), may be improved if it considers information from another company producing estimates about the traffic in the same area, but using different prediction methods. This type of information is called interdependent information (Krishna, 2002).

In this context, Jehiel and Moldovanu (2001) have shown that it is not only VCG mechanisms that are affected by this negative result, but that there can be no standard one-stage mechanism which is both efficient and incentive compatible for the procurement of estimates from multiple sources in an interdependent valuation setting. This is addressed to a certain degree by Mezzetti (2004) who shows that efficiency can be achieved in a two-stage mechanism, where in the first stage the final outcome is determined (i.e. the allocation of an observation), while in the second stage agents observe their payments and receive the final transfers from a centre, which acts as the information buyer in relevance to a citizen sensor network scenario. However, this work contradicts our requirements since it takes no consideration of the costs involved in the generation of an observation.

A solution to this problem comes from the sensor network related literature (Porter et al., 2008). In particular, they develop a mechanism which uses VCG payments in an interdependent information setting, where each agent reports its probability of failing to complete its assigned task, given that it will invest full effort into the completion of the task. Still, they restrict their mechanism by considering agents that can only report their own probability. In so doing, they do not consider the case where agents may observe the same probability from different perspective and combine them into one more accurate observation, thus contradicting our requirement regarding the combination of multiple sources of information. In addition to that, although agents' actions have costs, these costs do not depend on the success of their attempts and in all cases that they consider, they assume that the agents will attempt to complete the task without taking into consideration its cost. However, this assumption does not hold in our setting, where the agents are expected to invest less costly resources in a task, if they can increase their benefit by doing so (e.g. by being rewarded for a more precise measurement than the one they actually provide).

This particular shortcoming is addressed by Dash, Ramchurn and Jennings (2004), who proposed a mechanism in which a centre allocates suppliers to consumers based on the consumers' preferences. In this mechanism, consumers' preferences depend on their perception about the suppliers' actions. In more detail, consumers' decisions are influenced by others' opinions about the probability suppliers have of successfully fulfilling their assigned task. Consumers also provide all the information they possess about potential suppliers and, based on the effects of that information in the overall utility of the system, the reward or penalty for the agents is calculated by the centre. However, suppliers could possibly over report their costs or under report their valuations in order to increase their utilities since payments do not depend on whether a supplier succeeded or failed in completing its allocated task, and thus again this contradicts our requirement of selfish behaviour. Furthermore, it is assumed that buyers and sellers will have access to the full information about the state of the world after the task is allocated and completed, and therefore this contradicts our requirements regarding dynamism and uncertainty.

The issue of under-reporting valuation is addressed by Klein et al. (2008) who consider the allocation of shared bandwidth in a data network where participants are self interested and prefer to receive rather than share data. Although their mechanism copes with agents which under report their quality of information in order to avoid sharing it, it cannot handle reports of poor quality as agents can degrade their measurements without any implication. In more detail, agents may still report truthfully their information, but there is no means of identifying whether an agent has invested any resources in making this observation. This directly contradicts one of our requirements regarding selfish behaviour. Furthermore, the payments issued to the participants are based on the realised outcome, rather than the expected one, hence contradicting the requirement regarding uncertainty.

The latter requirement is fulfilled by a recent extension of the work of Dash, Ramchurn and Jennings (2004) by Ramchurn et al. (2009). In their work, they address both issues of under-reporting and uncertainty, by designing a payment scheme that ensures that over reporting costs or under reporting valuations is not a viable strategy. Now, regarding uncertainty about the execution of the task, they introduce a reputation system that contains information concerning the suppliers' performance that can be referred by the consumers. However, like all the previous work reviewed in this section, they assume full effort on behalf of the agents. In more detail, they assume that an agent responsible for completing an assigned task, will invest the maximum resources available to it for the completion of the task, without considering the involved costs. Hence, by not dealing with cases where agents are expected to invest less costly resources in a task, if they can increase their benefit by doing so, they don't address the relevant aspect of selfish behaviour as it was identified in Section 1.

A final negative aspect of VCG mechanisms is that the centre operates with a surplus, meaning that funds must be injected constantly in order to maintain wealth for the agents. In a citizen sensor network this will decrease the network's robustness, as it will depend on an external source for the support of the agents that may run out of funds. In terms of economics, this is called lack of budget balance. In more detail, there are several different types of budget balance, specifically in strict budget balance the total payments of agents should be equal to zero, while a weak budget balance only requires a non-negative total payment in order to achieve surplus and not deficiency. Furthermore, an ex ante budget balance mechanism is balanced on average, while an ex post budget balance is balanced all times, for all instances (Dash et al., 2003). Although budget balance is not in the requirements identified earlier, it is a desirable property, specifically in cases where citizen sensor networks cannot rely on external sources and must operate without human intervention (e.g. in a citizen science application where humans control sensors deployed in areas not accessible to them such as volcanoes and flooded rivers). Therefore, we review several methods of introducing budget balance, which will prove useful later in the analysis of our proposed model.

In particular, Cavallo (2006) has modified the standard VCG mechanism in order to achieve budget balance asymptotically by applying a redistribution mechanism after the centre makes its choice. This redistribution mechanism can be used for the allocation of single resources or bundles of products, and for more than ten agents it can achieve redistribution of more than 70% of the VCG surplus. However, the whole implementation is dependent on the existence of a centre to calculate the payments submitted back to the agents, which contradicts our research's requirements.

In this context, Petcu et al. (2006) expand the existing M-DPOP model which faithfully implements a distributed VCG outcome in order to achieve budget balance. They propose two distributed mechanisms which achieve different degrees of budget balance: R-M-DPOP reaches weak budget balance by redistributing the VCG tax, while BB-M-DPOP reaches exact budget balance at the expense of optimality. These distributed models implement a VCG payment which motivates agents to report truthfully and achieve a form of budget balance. However, they assume that agents have precise and perfect valuations, thus contradicting our requirement of operating in the presence of uncertainty.

As a conclusion, in reviewing several auction based mechanisms, we identified their flaws and showed which of our requirements are not satisfied. Therefore, we move from the realms of auction-based mechanism design. Specifically, we consider an alternative approach, still within the MD research domain, where agents are motivated to truthfully report their observations and invest effort into producing them through the use of strictly proper scoring rule to determine payments.

2.4.3 Using Scoring Rules as Reputation Mechanism

A number of researchers have proposed the use of *strictly proper scoring rules* to address several of the challenges (Matheson and Winkler, 1976; Friedman, 1983) that were not met by auction based mechanisms and, in particular, those regarding the interdependent information and the aspect of selfish behaviour related to the resources the agents commit in the generation of their observations.

In more detail, payments based on these rules reward accurate estimates or forecasts by depending on the difference between an event's predicted and actual outcome (observed at some later stage). Such mechanisms have been shown to incentivise agents to truthfully report their estimates in order to maximise their expected payment (see Section 2.4.3.1 for background on scoring rules). More recently, strictly proper scoring rules have been used in computer science to induce agents to report truthfully a probabilistic estimate and to commit costly resources into generating it at any required precision (Miller, Resnick and Zeckhauser, 2007; Miller, Pratt, Zeckhauser and Johnson, 2007). In more detail, Miller et al. apply a scoring system to assess agents' ratings against each other. In their mechanism, agents rate the quality of a service or a product and send their ratings to a central processing facility. The centre calculates an agent's payment (score) based on how close its reports is to the expected rating that the centre has calculated from the received reports of the other agents. As in most of the mechanism design trust models reviewed in this section, even if self-interest agents have the choice of revealing a different signal than their privately owned one, the use of strictly proper scoring rules motivates them to honestly report their private signal by maximising their expected utility, provided that all other agents report honestly too. The latter makes truthful reporting a *Nash equilibrium* solution.

Likewise, Jurca and Faltings (2005*b*) adjust their reputation mechanism (reviewed in subsection 2.3) by using the continuous scoring rule proposed by Miller et al. for calculating the payments to those agents that give feedback and therefore contribute to the reputation mechanism. Specifically, they divide agents into pairs and make them rate each other (i.e. agent i rates agent $i+1$). However both mechanisms still have more than one Nash equilibria in which agents collude and do not give a true feedback. In (Jurca and Faltings, 2007) they attempt to eliminate the unwanted Nash equilibria by proposing a method for enforcing the selection of a truthful strategy by introducing a small number of agents whose reports are always trustworthy. Although this is a significant contribution to the literature relevant to mechanism design reputation systems, it cannot be considered in a citizen sensor network as it could compromise the robustness of the network since it is not clear how these trustworthy agents would be selected or what would happen to the overall performance of the system if these agents stop operating.

Furthermore, strictly proper scoring rules have also been applied in a number of other similar contexts which, although different to ours, have similar objectives. For example, Zohar and Rosenschein (2008) focus on the *principal-agent problem* (Grossman and Hart, 1983; Rogerson, 1985) which is a more broad term for the mechanism design problem we consider where there is an asymmetry of information between two sides: a contractor who is interested in acquiring information it cannot evaluate itself and a contractee who is assigned to provide that information. In this context, they propose two mechanisms in which agents attempt to elicit information about each others' beliefs. In the first mechanism, they assume that each agent has access to a privately owned variable which takes a finite number of values. The principal agent wants to buy information from the rest and therefore motivates them to report truthfully by using a strictly proper scoring rule for payment. Since the buyer has no other means of assessing the information he receives, he evaluates the sellers based on his privately owned variable. In the second mechanism, they weaken the assumption that the designer of the mechanism has precise knowledge of the distribution of the probability of the event in question. Still, they rely on the fact that agents have a common probability distribution which expresses as 'close' notions of the governing probability distributions. This strong assumption has a severe impact since it restricts the applications of their method, as it makes it unrealistic and infeasible in real life scenarios where these assumptions rarely hold.

Now, although all the strictly proper scoring rule based mechanisms discussed in this section are effective in the specific cases that they consider, they assume that the costs of the traded information are known. Although our approach is similar, in our setting the costs represent private information known only to each individual agent (as it is explicitly stated in the introduction, where acquiring knowledge of the costs is considered a research requirement). Therefore, and due to the dependence of agent costs on the specific resources available only to it, we cannot apply any of these approaches directly.

Despite this drawback, the above mentioned mechanisms pose a significant contribution to relevant research. Specifically, they induce truthful reporting and effort on behalf of an agent even when there is no access to knowledge about the realised outcome, by using the other agents' opinions when calculating a payment to a particular agent. Furthermore, in the last two mechanisms, distributions can be used as agents' types, something that seems to be missing from most of the mechanism design models already reviewed. Indeed, in the majority of the interdependent information mechanisms that we have discussed so far, agents report discrete probability distributions (i.e. the probability of success or failure) and therefore cannot be applicable if agents' reports represent continuous probability distributions, which is the case in many citizen science applications responsible for the monitoring of environmental parameters. This is particularly important in our research, since the lack of interoperability and consistency of observations (common problem in citizen sensor networks, as mentioned in Section 2.1), can be addressed by modelling the traded information as continuous distributions. In so doing, we remove the disambiguity in the citizens' reported observations by applying a well defined standard of measurement such as the Gaussian distribution.

To this end, in the following section we provide the relevant background and key literature for the better understanding of scoring rules which have a central role in this work. In more details, in Section 2.4.3.1 we define various types of strictly proper scoring rules, while in Section 2.4.3.2 we describe a simple application of a strictly proper scoring rule based mechanism where agents' costs are known to the centre. Finally, in Section 2.4.3.3 we describe other types of scoring rules (not necessarily strictly proper) and how their applications satisfy or not our research requirements.

2.4.3.1 Background on Scoring Rules

One of the major purposes of statistical analysis is to make estimations and forecasts about the future. Although this is not directly relevant to our research, we are particularly interested in a task of equal importance that is closely connected to the task of making an observation of a current event or a forecast about a future event. These observations or forecasts (from this point referred as *probabilistic estimates*) must be evaluated and assessed regarding their quality, since poor estimates might lead to misleading actions. Many examples can be drawn from meteorology where wrong predictions often have devastating effects. To this end, scoring rules (Hendrickson and Buehler, 1971) are a system which assign a numerical score to a forecaster depending on how accurate her prediction is. Indeed, the closer the forecast is to the actual observation, the higher the assigned score. Thus, forecasters that seek to maximise their return are motivated to make careful assessments of high quality.

TABLE 2.1: The most popular binary scoring rules.

Scoring Rule	$G(\hat{p})$	$S(\hat{p},1)$	$S(\hat{p},0)$
Quadratic	$-2\hat{p}(1-\hat{p})$	$-2\hat{p}^2+4\hat{p}-2$	$-2\hat{p}^2$
Logarithmic	$\hat{p}\ln\hat{p}+(1-\hat{p})\ln(1-\hat{p})$	$\ln\hat{p}$	$\ln(1-\hat{p})$
Spherical	$\sqrt{\hat{p}}\sqrt{(1-\hat{p})}$	$-\frac{\sqrt{1-\hat{p}}}{2\sqrt{\hat{p}}}$	$-\frac{2\sqrt{\hat{p}}}{\sqrt{1-\hat{p}}}$

A rigorous definition of scoring rules can be drawn from Gneiting and Raftery (2004). In particular, if the forecaster quotes the predictive distribution P and the event x materialises, her reward is a function $S(P,x)$ which takes values from $-\infty$ to ∞. Suppose that the forecaster's best judgement is the distribution Q and that the expected value of $S(P,\cdot)$, $E_{x\sim Q}(S(P,x))$, is denoted as $S(P,Q)$. Then the forecaster has an incentive to report the truth (i.e. $P=Q$) if $S(Q,Q)\geq S(P,Q)$. A scoring rule with this property is defined as *strictly proper*. It can be seen that this truth eliciting property is equivalent to the incentive compatibility property in mechanism design, because in both cases participants maximise their payoff if they report truthfully.

Having defined strictly proper scoring rules, the simple rain prediction example presented in the introduction (Section 1.2) can be extended and formalised. In this example, various data aggregating on-line services quote a probability $\hat{p}\in[0,1]$ of rain and they are assigned a score depending on whether it rained or not. Specifically, the assigned score is $S(\cdot,1):[0,1]\rightarrow\mathbb{R}$ and $S(\cdot,0):[0,1]\rightarrow\mathbb{R}$, with $S(\hat{p},1)$ being the reward if a service succeeds in its forecast, and $S(\hat{p},0)$ the penalty if it fails.

However, we must choose the appropriate $S(\cdot,1)$ and $S(\cdot,0)$ in order to incentivise agents to quote their true probability. Hence, we must assign them a strictly proper scoring rule. Savage (1977) proved that strictly proper scoring rules can be generated if they satisfy the following requirements:

$$S(\hat{p},1)=G(\hat{p})+(1-\hat{p})G'(\hat{p})$$

and

$$S(\hat{p},0)=G(\hat{p})-\hat{p}G'(\hat{p})$$

where $\hat{p}$ is the agents' reported probability, $G:[0,1]\rightarrow\mathbb{R}$ is a bounded strictly convex function, preferably differentiable in (0,1) and G' is its derivative. If G is not differentiable, we must use the subgradient of G in [0,1].

These scoring rules use a binary approach depending on whether the event materialised or not, and therefore are called binary scoring rules. Different binary scoring rules can be created from the function $G:[0,1]\rightarrow\mathbb{R}$ with the most popular being the Brier or, as it is most commonly known, the quadratic. Following the above process, the quadratic, and two other popular scoring rules, the logarithmic and the spherical, listed in Table 2.1, can be derived.

A more general form for the binary scoring rules can be derived in order to apply them in cases where an event has more than two possible outcomes. Thus, in Table 2.2 we define the

TABLE 2.2: Discrete and continuous scoring rules.

Scoring Rule	Discrete	Continuous
Quadratic	$2r(x_i) - \sum r(x_i)^2$	$2r(x) - \int_{-\infty}^{\infty} r^2(x)dx$
Logarithmic	$\ln r(x_i)$	$\ln r(x)$
Spherical	$r(x_i)/(\sum r(x_i)^2)^{1/2}$	$r(x)/(\int_{-\infty}^{\infty} r^2(x)dx)^{1/2}$

discrete scoring rules and their continuous analogues which are useful when forecasters report distributions ($r(x_i)$) rather than probabilities. The continuous analogues (also found in Table 2.2) can be derived easily, if $r(x_i)$ which is the mass function of the discrete random variable x is replaced by density function $r(x)$ of the continuous random variable x.

Furthermore, the quadratic, spherical and logarithmic are not the only continuous strictly proper rules. Indeed other strictly proper scoring rules include the power rule family, which have been rigorously described more recently by Selten (1998). The discrete parametric scoring rule is defined as:

$$kr^{(k-1)}(x_i) - (k-1)\sum r^k(x_i) \tag{2.6}$$

while, its continuous analogue is:

$$kr^{(k-1)}(x) - (k-1)\int_{-\infty}^{\infty} r^k(x)dx \tag{2.7}$$

where k is a real number with $k > 1$

In addition to these scoring rules, Matheson and Winkler (1976) proposed a family of continuous scoring rules that can be derived from the binary case without being their analogues. In this particular case, where agents reveal their cumulative probability distributions, $R(x)$, instead of their density functions, the quadratic scoring rule is denoted as following:

$$S(R(x^*),u)) = \begin{cases} S_1(R(x^*)) = -2(1-R(x^*))^2 & \text{if } x^* \in (-\infty, u] \\ S_2(R(x^*)) = -2R^2(x^*) & \text{if } x^* \in (u, \infty) \end{cases} \tag{2.8}$$

where, x^* is the realised value of the outcome.

It is clear that the score $S(R(x^*),u)$ does not only depend on the realised value of outcome, x^*, but also on u. Although, u is just a random real number which divides the real line in two parts so x^* falls in one of the two intervals: $(-\infty,u]$ or (u,∞). In order to eliminate this dependency, we integrate equation 2.8 over $\mathbb{R}$. Therefore, according to the original scoring rule implementation and after x^* is revealed, the payment each agent receives is the following:

$$S^*(x^*) = -\int_{-\infty}^{x^*} 2R^2(u)du - \int_{x^*}^{\infty} 2(1-R(u))^2 du \tag{2.9}$$

Now, in this work we focus on the quadratic, spherical, logarithmic and parametric continuous strictly proper scoring rules, for two reasons. First, in the citizen sensor networks we consider, the distributed information is represented by continuous distributions (hence the use of binary strictly proper scoring rules is not applicable). Second, for the four already mentioned strictly proper scoring rules, as opposed to those introduced by Matheson and Winkler (1976), their expected counterparts can be analytically derived and expressed in close forms. Therefore, by using them, we will not restrict our analysis in solely empirical methods. In so doing, we will be able to provide theoretical results in order to satisfy the four research requirements set in Chapter 1, for a wide class of mechanisms and not just a specific instance.

2.4.3.2 An Application of Strictly Proper Scoring Rules in Mechanism Design

The purpose of this detailed example of strictly proper scoring rules, is to demonstrate their main properties through an application where costs involved in the generation of an observation are known to the centre. For this specific application, we assume that agents' noisy private observation, x, are modeled as continuous distributions and in particular as Gaussian random variables such that $x \sim N(x_0, 1/\theta)$, where θ is the observation's precision and x_0 is the true state of the parameter being estimated. After replacing the general probability density functions with the Gaussian distributions, we derive new expressions for each of the four scoring rules (expressed by $S(x_0; x, \theta)$ in Table 2.3). From these new expressions, and due to strictly proper scoring rule based payments being incentive compatible, we can simply integrate over the expected outcome in order to calculate the score an agent expects to derive, $\overline{S}(\theta)$, also shown in Table 2.3 given that it has generated an estimate of precision θ and has truthfully reported it to the centre (as it is incentivised to do).

Although incentive compatibility is one very desirable property, it is certainly not the only one. Particularly, in citizen sensor networks, some information providers may commit less than the required resources into the generation of the observation if they expect to increase their utility functions by doing so. To combat this, Miller, Resnick and Zeckhauser (2007) elicit effort through the use of appropriate scaling parameters, noting that any affine transformation of a strictly proper scoring rule does not affect its incentive compatibility property. Indeed, given knowledge of an agent's costs, they show that it is possible to induce an agent to make and truthfully report an estimate with a specified precision, θ_0. In this case, an agent's expected payment, $\overline{P}(\theta)$, is given by:

$$\overline{P}(\theta) = \alpha\overline{S}(\theta) + \beta \tag{2.10}$$

where α and β are the scaling parameters, and the expected utility of the agent is given by:

$$\overline{U}(\theta) = \alpha\overline{S}(\theta) + \beta - c(\theta) \tag{2.11}$$

The centre can now choose the value of α such that the agent's utility (its payment minus its costs) is maximised when it produces and truthfully reports an estimate of the required precision,

TABLE 2.3: Comparison of Quadratic, Spherical, Logarithmic and Parametric Scoring Rules

Scoring Rule:	Quadratic	Spherical
$S(x_0;x,\theta)$	$2\mathcal{N}(x_0;x,1/\theta)-\frac{1}{2}\sqrt{\frac{\theta}{\pi}}$	$\left(\frac{4\pi}{\theta}\right)^{\frac{1}{4}}\mathcal{N}(x_0;x,1/\theta)$
$\overline{S}(\theta)$	$\frac{1}{2}\sqrt{\frac{\theta}{\pi}}$	$\left(\frac{\theta}{4\pi}\right)^{\frac{1}{4}}$
$\overline{S}'(\theta)$	$\frac{1}{4\sqrt{\pi\theta}}$	$\frac{1}{4}\left(\frac{1}{4\pi\theta^3}\right)^{\frac{1}{4}}$
α	$4c'(\theta_0)\sqrt{\pi\theta_0}$	$4c'(\theta_0)\left(4\pi\theta_0^3\right)^{\frac{1}{4}}$
β	$c(\theta_0)-2\theta_0c'(\theta_0)$	$c(\theta_0)-4\theta_0c'(\theta_0)$

Scoring Rule:	Logarithmic	Parametric
$S(x_0;x,\theta)$	$\log(\mathcal{N}(x_0;x,1/\theta))$	$k\mathcal{N}(x_0;x,\theta)^{(k-1)}-\frac{k-1}{\sqrt{k}}\left(\frac{2\pi}{\theta}\right)^{\frac{1-k}{2}}$
$\overline{S}(\theta)$	$\frac{1}{2}\log\left(\frac{\theta}{2\pi}\right)-\frac{1}{2}$	$\frac{1}{\sqrt{k}}\left(\frac{2\pi}{\theta}\right)^{\frac{1-k}{2}}$
$\overline{S}'(\theta)$	$\frac{1}{2\theta}$	$\frac{k-1}{2\theta\sqrt{k}}\left(\frac{2\pi}{\theta}\right)^{\frac{1-k}{2}}$
α	$2c'(\theta_0)\theta_0$	$\frac{2c'(\theta_0)\theta_0\sqrt{k}}{k-1}\left(\frac{\theta_0}{2\pi}\right)^{\frac{1-k}{2}}$
β	$c(\theta_0)-2c'(\theta_0)\theta_0\left(\frac{1}{2}\log\left(\frac{\theta_0}{2\pi}\right)-\frac{1}{2}\right)$	$c(\theta_0)-\frac{2\theta_0}{k-1}c'(\theta_0)$

θ_0. To do so, it solves $\left.\frac{\mathrm{d}\overline{U}}{\mathrm{d}\theta}\right|_{\theta_0}=0$ to give:

$$\alpha=\frac{c'(\theta_0)}{\overline{S}'(\theta_0)} \tag{2.12}$$

In Table 2.3 we present this result, and the derivative of the expected score that is required to calculate it, for each of the four strictly proper scoring rules presented earlier in the cells $\overline{S}'(\theta)$ and α respectively.

Having defined the α parameter of the affine transformation that elicits effort and honest reporting, we calculate parameter β which motivates agents to participate in the mechanism by ensuring that their expected utility is always positive. In more detail, we now note that in order for a self-interested agent to incur the cost of producing a forecast, it must expect to derive positive utility from doing so. Thus, the centre can use the constant β to ensure that it makes the minimum payment to the agent, hence ensuring that the mechanism is individually rational. When costs are known, the centre can do so by making the agents indifferent between producing the forecast or not, by setting $\overline{U}(\theta_0)=0$, thus giving:

$$\beta=c(\theta_0)-\frac{c'(\theta_0)}{\overline{S}'(\theta_0)}\overline{S}(\theta_0) \tag{2.13}$$

Again, cells β in Table 2.3 show this result for each of the four scoring rules.

Finally, it should be noted that the expected scores of the quadratic, spherical and logarithmic scoring rules are concave depending on the precision θ (expressed as $\overline{S}(\theta)$ in Table 2.3). While this property of the expected scores forms the basis of the calculations and proofs that we present in the following chapters, it does not hold for any strictly proper scoring rule. For example, in the parametric scoring rule for $k > 3$ the second derivative of the expected score (denoted by Equation 2.14) becomes positive and therefore the expected score convex, which in turn, as we will show in the following chapter, results in a payment that fails to incentivise an agent to produce an estimate at the required precision, θ_0.

$$\overline{S}''(\theta) = \frac{(1-k)(3-k)}{4\theta^2\sqrt{k}}\left(\frac{2\pi}{\theta}\right)^{\frac{1-k}{2}} \tag{2.14}$$

Therefore, and in order to guarantee the concavity of the expected score we will restrict the parameter k to the space $(1,3)$.

2.4.3.3 Other Types of Scoring Rules

In the scoring rules literature, strictly proper scoring rules are emphasised, since their truth elicitation property is very important. However, Savage (1977) suggests that scoring rules should have a stronger *monotonicity* property which will motivate agents to reveal a distribution very close to their true one, if dishonest reporting cannot be avoided. Friedman (1983) formalises this property and refers to such scoring rules as *effective scoring rules*. In these rules, the expected score is higher when the reported distribution is 'closer' to the distribution of the actual outcome and therefore it can be easily proved that every effective scoring rule is strictly proper. Specifically, Friedman proves that within an Euclidean space, the quadratic and spherical scoring rules are effective, while the logarithmic is not. Note that this is not the only difference between the three most popular strictly proper scoring rules. In particular, the logarithmic scoring rule has no lower bound which can result in infinite payments. This may be problematic in a citizen sensor network, as it may be impossible to constantly inject funds to citizens that have to face a relatively large payment.

Despite this negative aspect, a variation of the logarithmic scoring rule has been used in prediction systems in electronic markets. In more detail, *market scoring rules* (Hanson, 2003, 2007) have found applications in systems where forecasters are allowed to change their initial reported forecast and then get paid according to their new report after taking into account the benefits they gained from that change. This makes market scoring rules particularly appealing in dynamic systems with rich interactions among the participating agents (Guo and Penock, 2009; Chen et al., 2009). Although we are interested in a similar setting where there is a lack of knowledge about the state of the world, we adopted a less complex approach, in which agents communicate their estimates to a centre and then receive their payments. This setting is more amenable to strictly proper scoring rules than market scoring rules, since the centre can calculate the agents' payments in a single round simply by comparing it with the fused reported estimates, instead of implementing a repetitive process which would allow each agent to modify its reported observation.

2.5 Summary

This chapter has reviewed literature relevant to selfish behaviour in citizen sensor networks and has identified the lack of literature related to trust in such networks. This realisation motivated our focus on trust applications in MAS. Therefore, the two principal lines of research in trust were analysed: individual and system-trust. After discussing the main agent-based trust models, and specifically, learning, reputation and combined probabilistic models, it was apparent that they failed to satisfy one or more of the requirements of our research. For that reason, we focused on the design of rules that regulate the interactions among agents, rather than on trying to identify all the possible actions of each agent. In taking a more systematic approach in reputation trust, we were able to address the issues of selfish behaviour, more effectively than the reputation systems in agent-level trust. However, we found that they were efficient only in the specific scenarios they considered as they lacked theoretical foundations that would make them suitable for open systems such as citizen sensor networks. Therefore we discussed mechanism design based solutions and argued that although many of the attributes we desire in our solution were met, there is still a number of issues that need to be resolved. Within the MD line of research, we argued that scoring rules can be used to solve many of the problems identified in this section. Given this, the next chapter presents our implementation of a mechanism based on scoring rules, which elicits both effort and honest reporting of private signals, in a setting where costs are unknown to the centre.

Chapter 3

A Mechanism for Dealing with Unknown Costs

In the previous chapter we discussed several applications of scoring rules in mechanism design. Now, while these approaches are effective in the specific cases that they consider, they all rely on the fact that the cost of the agent providing the observation or estimate is known by the centre. However, in many citizen sensor networks, this is not the case, as these costs represent private information known only to each individual information provider and not to the information buyer. For example, in a noise monitoring network, the costs involved in the generation of a sound observation represent the amount of time the citizen has invested in making the observation. Subsequently, if a contributor has spent a significant amount of time in making her observation, it will be more accurate and more realistic. Now, different people, have different levels of commitment in a project like that, and it is not possible for the buyer to know in advance how much time the providers invest. In a traffic information service, multiple companies providing advice regarding traffic conditions such as estimates of optimal routes of possibly low congestion to subscribed drivers act as information providers. They face costs depending on the their technological capacity and the statistical models used both for the the planning of the optimal routes and the estimation of traffic in locations where there are minimum reports. A subscriber of such a service who acts as the information buyer is not expected to know these costs.

To this end, we introduce a mechanism in which the centre (effectively the information buyer) does not rely on any knowledge of costs involved in the generation of the agents' (information providers) observations or probabilistic estimates, and hence we address research requirement 1 identified in Section 1.1 regarding unknown costs. In so doing, we contribute to the state of the art by developing a novel two-stage mechanism in which the centre in the first stage elicits the agents' unknown costs and identifies the agent that can provide an observation at the minimum cost, while in the second stage, through appropriately scaled strictly proper scoring rules, it incentivises that agent to invest sufficient resources into generating an observation or a probabilistic estimate at a specific precision and report it truthfully (hence address Requirement 2 regarding selfish behaviour). In the setting we consider, we assume that there are multiple agents, all of which are capable of producing an estimate of at least the required precision (we relax this assumption in Chapter 4, where we consider agents that have limitations in the maximum precision of the estimate they produce), and that the centre has access to knowledge about the state of the world after it receives the selected agent's observation (we relax this assumption in Chapter 5).

The rest of this chapter is organised as follows: In Section 3.1 we describe our model, while in Section 3.2 we proceed to describe our novel two stage mechanism. In Section 3.3 we prove its economic properties (i.e. incentive compatibility and individual rationality). Furthermore, in Section 3.4 we empirically evaluate our mechanism by comparing the quadratic, spherical and logarithmic scoring rules with a parametric family of strictly proper scoring rules in a setting where costs depend linearly on precision. We show that while the logarithmic rule results in the centre making the lowest expected payment to the agent, this payment is unbounded. The other rules are bounded, but result in higher expected payments. Hence, we find that for certain values of the parameter, the parametric scoring rule is preferred, since the resulting payment is similar to the logarithmic both in its expected value and its variance, which has been identified as the optimal rule, but with finite lower bound. Finally, in section 3.5 we summarise and particularly discuss how to relax the assumption that a single agent can provide an estimate at any precision the centre requires.

3.1 Eliciting Information from a Single Agent

We now describe the information elicitation problem, outlined in Chapter 1. Specifically, we consider a citizen sensor network in which an information buyer wants to acquire an observation or a probabilistic estimate of an event at a certain quality, denoted by a minimum precision θ_0 (henceforth referred to as the required precision). In the context of the applications of citizen sensor networks introduced in Section 2.1.3, in a noise monitoring scenario, the organisation responsible for compiling the noise maps of European cities acts as information buyer and is interested in acquiring noise observations of a certain standard (which reflects the precision of the observation) from members of the public equipped with mobile phones with microphones, who in turn act as information providers. In a traffic report service, the role of information buyer is undertaken by the subscribed drivers that receive information regarding traffic conditions and estimates of optimal routes, by multiple competing private companies acting as information providers.

Against this background, we assume that an information buyer derives no additional benefit if the observation is of precision greater than θ_0 and that there are $N \geq 2$ rational, risk neutral providers that can provide the buyer with a noisy observation, x, of precision θ. We model the providers' private observations as Gaussian random variables such that $x \sim N(x_0, 1/\theta)$, where

x_0 is the true state of the parameter being observed. Furthermore, as mentioned above, this true state is unknown to both the buyer and the providers at the time that the observation is requested, but becomes available to the buyer at some time in the future. We note that this assumption is unrealistic, as the purpose of the citizen sensor networks we consider, is to provide observations of events that could not be provided with conventional stationary sensors and we hence intend to relax it in Chapter 5. However, there are examples where it is possible to evaluate the observations, depending on whether the information buyers are willing to invest a substantial amount of resources for this task. For example, in a noise monitoring scenario, the information buyer may be able to deploy its own equipment in few specific locations to verify the observations of some providers. Having assessed some of the providers it can use these evaluations to identify whether they truthfully report their observations or not, and given that information providers are not static, after some time we can reasonably assume that most providers will have had their observations evaluated at some point. In a traffic reports service the subscribed drivers may have the ability of verifying the conditions on the streets by subscribing into a service that gives live feed on the situation on the streets through airborne means. That is why, this mechanism cannot be simply considered as a first step towards a complete mechanism presented at a later stage. Instead, it is a stand alone mechanism that can address specific issues in specific settings[19]

As already mentioned, in the setting we consider, the providers incur a cost, $c(\theta)$, in producing their estimate, which we assume to be a function of the estimate's precision. In addition to that, while the buyer has no information regarding the providers' cost functions, we assume that all cost functions are convex (i.e. $c_i''(\theta) \geq 0$), and we note that this is a realistic assumption in all cases where there are diminishing returns as the precision increases. For example, in the noise monitoring scenario, a contributing member of the public may dedicate more of her time and provide an observation that will be an accurate representation of the noise pollution in that area. Likewise, in a traffic reports service, the cost of the service increases as more advanced models are used for the generation of estimates of traffic when there is scarce real time data, given that after a certain level of precision, advancements are minimal and costs in processing power due to the increased complexity are higher.

Finally, we do not assume that all providers use the same cost function, but we do demand that the costs of different providers and their derivatives do not cross (i.e. the ordering of the agents' costs and their derivatives is maintained over all precisions). Although these assumptions are not directly relevant to a specific citizen sensor network application, they are necessary in order to prove incentive compatibility and individual rationality of the following mechanism and hence address the research requirement regarding selfish behaviour (Requirement 2). Indeed, in Section 3.3 we show that these assumptions are necessary, since there can be no mechanism based on scaled strictly proper scoring rules, that can be incentive compatible with respect to

[19]In Conclusions and Future Work, we provide a summary of Mechanisms 1, 2 and 3 described in Chapters 3, 4 and 5 respectively and explicitly mention that when the information buyer has access to such knowledge, it is better to use Mechanisms 1 and 2 since truthful reporting is a stronger solution concept (dominant strategy, as opposed to the Nash equilibrium in Mechanism 3 in Chapter 5).

costs revealed, and incentivise the selected agent to generate an estimate at the optimal precision, without them being satisfied (Corollaries 3.1 and 3.2).

Given this model[20], the challenge is to design a mechanism that enables an information buyer to identify the provider that can provide the observation at the lowest cost, and to provide a payment to this provider such that it is incentivised to generate the observation with a precision at least equal to the required one and to report it truthfully.

3.2 The Mechanism

Given that Jehiel and Moldovanu (2001) have shown that there is no one-stage mechanism that can incentivise agents to truthfully report their valuations in an interdependent information setting, we address the above mentioned challenges by designing a two-stage mechanism (Mechanism 1). In the first stage of the mechanism, the centre incentivises agents to truthfully reveal their costs and identify the agent with the lowest one, while in the second stage, the centre uses an appropriately scaled strictly proper scoring rule in order to incentivise that agent to generate an estimate at the required precision, and to truthfully report it.

In more detail, in the first step (Step 1.1) of the first stage of Mechanism 1 the centre announces that it requires an estimate at a specific precision and asks all agents to report their cost functions. In practise the centre only requires the cost function and its derivative at that required precision, and not the entire functions. However, for notational convenience we request the agents to reveal their entire cost function. Now, in the second step (Step 1.2), the centre uses a reverse second price auction to elicit the agents' costs, and hence identify the agent that can provide an estimate with a specified precision at the lowest cost. The centre then allocates the estimate to that agent. In the the second stage of the mechanism, in the first step (Step 2.1) the centre announces a payment based on a scaled strictly proper scoring rule and determines the scaling parameters α and β:

$$\alpha = \frac{c'_j(\theta_0)}{\overline{S}'(\theta_0)} \quad \text{and} \quad \beta = c_j(\theta_0) - \frac{c'_j(\theta_0)}{\overline{S}'(\theta_0)}\overline{S}(\theta_0) \tag{3.1}$$

where $\overline{S}'(\theta_0)$ is the derivative of expected score, $\overline{S}(\theta)$, at the required precision, with $\overline{S}(\theta)$ being a strictly concave[21] function, and $c_j(\theta_0)$ being the second lowest cost determined in step 1.2. To this end, in step 2.2 the selected agent from step 1.2 generates the required estimate and finally at step 2.3 it receives its payment from the centre.

[20]Note that in this research we restrict our model to continuous parameter distributions and continuous cost functions, however, in our future work, we intend to extend the model so it can deal with discrete distributions and costs.

[21]We note that the quadratic, spherical, logarithmic and parametric scoring rules satisfy this property (see row 2 of Table 2.3).

Mechanism 1 The mechanism for dealing with unknown costs:

1. First Stage

1.1 The centre announces that it needs an estimate of required precision θ_0, and asks all agents $i \in \{1,\ldots,N\}$, where $N \geq 2$, to report their cost functions $\widehat{c}_i(\theta)$.

1.2 The centre assigns the task of generating the estimate to the agent who reported the lowest cost at the required precision, i.e., agent i such that $\widehat{c}_i(\theta_0) = \min_{k\in\{1,\ldots,N\}} \widehat{c}_k(\theta_0)$.

2. Second Stage

2.1 The centre announces a scoring rule $\alpha S(x_0;x,\theta)+\beta$, where:

   (i) $S(x_0;x,\theta)$ is a strictly proper scoring rule,

   (ii) $\overline{S}(\theta)$ is strictly concave as a function of precision θ, and

   (iii) α and β are determined using equations 3.1, based on the second-lowest reported cost functions (i.e. $\widehat{c}_j(\theta)$ such that $\widehat{c}_j(\theta_0) = \text{argmin}_{k\neq i} \widehat{c}_k(\theta_0)$).

2.2 The agent selected in the first stage produces an estimate x with precision θ and reports $\widehat{x}$ and $\widehat{\theta}$ to the centre.

2.3 Once the actual outcome has been observed, the centre then gives the following payment to the agent:

$$P(x_0;\widehat{x},\widehat{\theta}) = \alpha S(x_0;\widehat{x},\widehat{\theta})+\beta \tag{3.2}$$

Now regarding this mechanism, at first glance it may appear that it is similar to a reverse second-price or Vickrey auction, where the agents' rewards are equal to the second-lowest reported costs. This is certainly true. However, the mechanism here is more complex as the two stages of our mechanism are interconnected and the selected agent's reward in the second stage is determined by the scaled scoring rule, with the scaling parameters depending on the second lowest cost identified in the first stage.

Having detailed the mechanism, in the next section we identify and prove its economic properties.

3.3 Economic Properties of the Mechanism

Specifically, in this section we show that:

1. The mechanism outlined above is incentive compatible in the first stage regarding the costs. In particular, truthful revelation of the agents' cost functions is a weakly dominant strategy.
2. The mechanism is incentive compatible regarding the selected agent's reported estimate and precision in the second stage.
3. There can be no incentive compatible mechanism regarding the agents' cost functions revealed when the cost functions overlap.
4. The mechanism is individually rational.
5. The centre motivates the selected agent to make an estimate with a precision which is at least as high as θ_0, the precision required by the centre. We refer to the actual precision produced as the 'optimal precision' (from the perspective of the agent) θ^*, since for this precision the expected payment is maximised.
6. For some concave cost functions, or for those cost functions whose derivatives overlap, the selected agent will not be able to generate an estimate at the optimal precision. For these cases, the centre will not be able to achieve the precision it requires.

We now prove these properties. To do so, we first derive two lemmas which are then used in the proofs that follow. The first lemma shows that if the true costs of the agent making the prediction are higher than the costs used for the scaling of the scoring function, then the agent's utility will always be negative. More formally:

Lemma 3.1. If $c_t(\theta)$ and $c_s(\theta)$ are convex functions with $c_t(\theta) > c_s(\theta)$, $c'_t(\theta) > c'_s(\theta)$ and $c_t(0) = c_s(0) = 0$, where $c_t(\theta)$ is the agent's true cost function, $c_s(\theta)$ is the cost function used to scale the scoring function and $c'_t(\theta)$ and $c'_s(\theta)$ their respective derivatives, then $\overline{U}(\theta) < 0$ for any θ.

Proof. Concavity of the expected score $\overline{S}(\theta)$ implies:

$$\overline{S}'(\theta_0)(\theta - \theta_0) \geq \overline{S}(\theta) - \overline{S}(\theta_0) \tag{3.3}$$

Similarly, convexity of the cost function $c_s(\theta)$ gives:

$$c_s'(\theta_0)(\theta - \theta_0) \leq c_s(\theta) - c_s(\theta_0) \tag{3.4}$$

Given that by definition $\overline{S}(\theta)$ and $c_s(\theta)$ are strictly increasing (as stated in the model description in Section 3.1), dividing with $\overline{S}'(\theta_0)$ and $c_s'(\theta_0)$ maintains the sign in inequalities 3.3 and 3.4. Therefore, from:

$$(\theta - \theta_0) \geq \frac{\overline{S}(\theta) - \overline{S}(\theta_0)}{\overline{S}'(\theta_0)}$$

and

$$(\theta - \theta_0) \leq \frac{c_s(\theta) - c_s(\theta_0)}{c_s'(\theta_0)}$$

it follows that:

$$\frac{\overline{S}(\theta) - \overline{S}(\theta_0)}{\overline{S}'(\theta_0)} \leq \frac{c_s(\theta) - c_s(\theta_0)}{c_s'(\theta_0)}$$

or

$$\frac{c_s'(\theta_0)}{\overline{S}'(\theta_0)}(\overline{S}(\theta) - \overline{S}(\theta_0)) + c_s(\theta_0) - c_s(\theta) \leq 0 \tag{3.5}$$

Now, the expected utility, is given by: $\overline{U}(\theta) = \alpha\overline{S}(\theta) + \beta - c(\theta)$ (Equation 2.11), with the scaling parameters α and β already defined using Equations 2.12 and 2.13 as $\alpha = \dfrac{c'(\theta_0)}{\overline{S}'(\theta_0)}$ and $\beta = c(\theta_0) - \dfrac{c'(\theta_0)}{\overline{S}'(\theta_0)}\overline{S}(\theta_0)$, where θ_0 is the centre's required precision and $\overline{S}(\theta)$ is the expected score.

Therefore, an agent's expected utility is given by:

$$\overline{U}(\theta) = \frac{c_s'(\theta_0)}{\overline{S}'(\theta_0)}(\overline{S}(\theta) - \overline{S}(\theta_0)) + (c_s(\theta_0) - c_t(\theta)) \tag{3.6}$$

Therefore, since $c_t(\theta) > c_s(\theta)$, *for any* θ the following holds:

$$\frac{c_s'(\theta_0)}{\overline{S}'(\theta_0)}(\overline{S}(\theta) - \overline{S}(\theta_0)) + c_s(\theta_0) - c_s(\theta) \leq 0 \Rightarrow$$

$$\frac{c_s'(\theta_0)}{\overline{S}'(\theta_0)}(\overline{S}(\theta) - \overline{S}(\theta_0)) + c_s(\theta_0) - c_t(\theta) < 0 \Rightarrow$$

or

$$\overline{U}(\theta) = \frac{c_s'(\theta_0)}{\overline{S}'(\theta_0)}(\overline{S}(\theta) - \overline{S}(\theta_0)) + c_s(\theta_0) - c_t(\theta) < 0$$

□

The next lemma shows that if the true costs of the agent performing the estimate are less than the costs which are used to scale the scoring rule, the optimal precision θ^* will be greater than θ_0.

Lemma 3.2. If $c_t(\theta)$ and $c_s(\theta)$ are convex functions with $c_t(\theta) < c_s(\theta)$, $c_t'(\theta) < c_s'(\theta)$ and $c_t(0) = c_s(0) = 0$, where $c_t(\theta)$ is the agent's true cost function, $c_s(\theta)$ is the cost function used to scale the scoring function and $c_t'(\theta)$ and $c_s'(\theta)$ their respective derivatives, then $\theta^* > \theta_0$.

Proof. The agent's optimal precision, θ^*, which maximises its expected utility is formally denoted by $\theta^* = \text{argmax}_\theta \overline{U}(\theta)$, with $\overline{U}'(\theta^*) = 0$.

Now, the agent's expected utility is already defined by Equation 3.6 as:

$$\overline{U}(\theta) = \frac{c_s'(\theta_0)}{\overline{S}'(\theta_0)}(\overline{S}(\theta) - \overline{S}(\theta_0)) + (c_s(\theta_0) - c_t(\theta))$$

Given that the optimal precision, θ^*, maximises the expected score, we have $\overline{U}'(\theta^*) = 0$, and hence, after replacing θ with θ^* and calculating the derivative of the expected utility (Equation 3.6):

$$\frac{c_s'(\theta_0)}{\overline{S}'(\theta_0)}\overline{S}'(\theta^*) - c_t'(\theta^*) = 0 \Leftrightarrow \frac{\overline{S}'(\theta^*)}{\overline{S}'(\theta_0)} = \frac{c_t'(\theta^*)}{c_s'(\theta_0)} \tag{3.7}$$

Let $f(\theta) = \overline{S}'(\theta)/\overline{S}'(\theta_0)$ and $g(\theta) = c_t'(\theta)/c_s'(\theta_0)$. In the model description (Section 3.1) it is explicitly stated that $\overline{S}(\theta)$ is (strictly) concave, strictly increasing and twice differentiable. Therefore, $f'(\theta) \leq 0$ for all θ_0.

In Section 3.1 we also listed the assumptions regarding the agents' cost functions. In more detail, we assumed that the cost functions are convex, strictly increasing and that they and their derivatives maintain the same order, without overlapping for all θ. Now, since $c_t''(\theta) \geq 0$ (due to costs' convexity) and $c_s'(\theta) \geq 0$ (since cost functions are strictly increasing), $g'(\theta) \geq 0$ for all θ.

Furthermore, according to the initial assumptions regarding costs: $c_t'(\theta) < c_s'(\theta)$ for all θ, and hence $g(\theta_0) < 1$. Since $g(\theta)$ is strictly increasing and $f(\theta_0) = 1$, the functions $f(\theta)$ and $g(\theta)$ must cross at $\theta' > \theta_0$ (i.e. $f(\theta') = g(\theta')$). Thus $\theta^* > \theta_0$

At this point it should be noted that although the lemma's conditions regarding the derivatives of the cost functions seem irrelevant to a citizen sensor network setting, they are critical to the proofs of the mechanism's economic properties[22]. Indeed, there can be two cost functions $c_t(\theta)$ and $c_s(\theta)$ such that $c_t(\theta) < c_s(\theta)$ and $c_t'(\theta) > c_s'(\theta)$ for which $\theta^* < \theta_0$.

Specifically, for $c_t(\theta) = \frac{3\theta^2}{2}$ and $c_s(\theta) = e^{\theta} - 1$, it is easy to show graphically that $c_t(\theta) < c_s(\theta)$ for every $\theta > 0$, while there are some values of θ such that $c_t'(\theta) > c_s'(\theta)$ (i.e. for $\theta = 1$ $c_t'(1) = 3$ and $c_s'(1) = e - 1$).

Now, $\overline{S}''(\theta) \leq 0$, since S is convex, and therefore S' is a decreasing function. That is, for $\theta^* > \theta_0 \Rightarrow \overline{S}'(\theta^*) < \overline{S}'(\theta_0)$, hence Equation 3.7 will take the following form for $\theta_0 = 1$:

$$\frac{\overline{S}'(\theta^*)}{\overline{S}'(\theta_0)} = \frac{3\theta^*}{e} < 1 \Leftrightarrow 3\theta* < e \Rightarrow \theta^* < 1$$

which does not hold given that initially we have assumed that $\theta^* > \theta_0$ with $\theta_0 = 1$ $\square$

Having presented these two key lemmas, we now proceed to prove the four economic properties of our mechanism.

Theorem 3.1. Truthful revelation of agents' cost functions in the first stage of the mechanism is a weakly dominant strategy.

Proof. We prove this by contradiction. Let $c_t(\theta)$ and $\widehat{c}(\theta)$ denote an agent's true and reported cost functions respectively. Furthermore, let $c_s(\theta)$ denote the cost function used to scale the scoring function *if the agent wins* (i.e. if $\widehat{c}(\theta_0) < c_s(\theta_0)$).

First, suppose that the agent misreports, but this does not affect whether it wins or not. In this case, since the costs are based on the second-lowest costs, this does not affect the scoring rule if the agent wins. Moreover, if the agent loses the payoff is always zero. Therefore, there is no incentive to misreport.

Second, suppose that the agent's misreporting affects whether that agent is pre-selected or not. There are now two cases:

1. The agent wins by misreporting, but would have lost when truthful.

2. The agent loses by misreporting, but would have won when truthful.

In this context:

- Case (1) can be formally denoted as $c_t(\theta_0) > c_s(\theta_0)$ and $\widehat{c}(\theta_0) < c_s(\theta_0)$. Now, since the true cost $c_t(\theta_0) > c_s(\theta_0)$, it follows directly from Lemma 3.1 that the expected utility $\overline{U}(\theta)$ is strictly negative, irrespective of θ. Therefore, the agent could do strictly better by reporting truthfully in which case the expected utility is zero.

[22]We intend to show that all our assumptions regarding the cost functions (i.e. non crossing condition and concavity) are necessary for the proofs of the economic properties

- Case (2) can be formally denoted as $c_t(\theta_0) < c_s(\theta_0)$ and $\widehat{c}(\theta_0) > c_s(\theta_0)$. In this case the agent would have won by being truthful, but now receives a utility of zero. To show that this type of misreporting is suboptimal, we need to show that, when $c_t(\theta_0) < c_s(\theta_0)$, an agent benefits from being selected and generating the (optimal) estimate (i.e. $U(\theta^*) > 0$ when $c_t(\theta_0) < c_s(\theta_0)$). Now, since θ^* is optimal by definition, then $U(\theta^*) \geq U(\theta_0)$. From the expected utility in equation 3.6, we have $\overline{U}(\theta_0) = c_s(\theta_0) - c_t(\theta_0) > 0$ when $c_t(\theta_0) < c_s(\theta_0)$, and hence $U(\theta^*) > 0$ at true costs reporting.

□

Corollary 3.1. Incentive compatibility with respect to agents' reported costs and precisions does not hold if the agents' cost functions cross at θ'.

Proof. The proof regarding the agents' reported costs comes directly from the above theorem, as we need to show only one example where an agent is incentivised to misreport its cost function.

Against this background, following the notation of the above theorem, let $c_t(\theta)$ and $\widehat{c}(\theta)$ denote an agent's true and reported cost functions respectively, while $c_s(\theta)$ denotes the cost function used to scale the scoring function and θ' is the point where two cost functions (suppose $c_s(\theta)$ and $c_t(\theta)$) intersect. In this context, we intend to show that an agent can do better by misreporting and losing, rather than by reporting truthfully and winning. In more detail, since $c_s(\theta)$ and $c_t(\theta)$ overlap at θ', $c_s(\theta) < c_t(\theta)$, for every $\theta > \theta'$. Therefore, according to Lemma 3.1, the expected utility will be strictly negative. If the agent misreports its cost function so it is not selected, its utility will be zero. Therefore, the mechanism is no longer incentive compatible with respect to reported cost functions.

Now, regarding the agent's reported precision, if at θ' $c_t(\theta)$ and $c_s(\theta)$ intersect then $c_t(\theta') > c_s(\theta')$ and therefore $\overline{U}(\theta') < 0$. Given that this agent is the cheapest one, at least for $\theta \leq \theta'$ it is in its best interest to report a precision lower that θ' even if it makes an estimate with precision greater than θ', in order to maintain to positive utility. That is, $\widehat{\theta} < \theta'$, while $\theta > \theta'$. Therefore, the mechanism is no longer incentive compatible with respect to reported precision.

□

Theorem 3.2. The mechanism is incentive compatible regarding the agent's reported estimate and precision in the second stage.

Proof. Given the scaling of the scoring rules described in step 1 in the second stage of Mechanism 1, the expected utility of the agent, if it produces an estimate of precision θ, which it reports with precision $\widehat{\theta}$, is denoted by $\overline{U}(\theta, \widehat{\theta})$, and is given by:

$$\overline{U}(\theta, \widehat{\theta}) = \frac{c_s'(\widehat{\theta})}{\overline{S}'(\widehat{\theta})} \left(\overline{S}(\theta, \widehat{\theta}) - \overline{S}(\widehat{\theta}) \right) + c_s(\widehat{\theta}) - c_t(\theta) \tag{3.8}$$

where $\overline{S}(\theta,\widehat{\theta})$ is the agent's expected score for producing an estimate of precision θ and reporting its precision as $\widehat{\theta}$. Furthermore, $\overline{S}(\widehat{\theta})$ is the agent's expected score for producing and truthfully reporting an estimate of precision $\widehat{\theta}$, $c_t(.)$ is the true cost function of the agent, and $c_s(.)$ is the cost function used to produce the scoring rule (i.e. the second lowest revealed cost).

Now, the proof of mechanism's incentive compatibility regarding the agent's reported estimate follows directly from the definition of the strictly proper scoring rule. Indeed, from the general form of the utility function:

$$\overline{U}(\theta,\widehat{\theta}) = \alpha(\widehat{\theta})\overline{S}(\theta,\widehat{\theta}) + \beta(\widehat{\theta}) - c_t(\theta) \tag{3.9}$$

it is clear that since the expected score, $\overline{S}(\theta,\widehat{\theta})$, is maximised when the agent reports truthfully its estimate, the same will happen to the expected utility as it is an affine transformation of the expected score.

However, proving that the mechanism is incentive compatible with respect to agent's reported precision is not as straightforward since the scaling parameters α and β depend on the reported precision. Therefore, we must explicitly show that the expected utility will be maximised at $\widehat{\theta} = \theta$.

In this context, taking the first derivative of the expected utility (Equation 3.8) with respect to $\widehat{\theta}$ gives:

$$\frac{\partial \overline{U}(\theta,\widehat{\theta})}{\partial \widehat{\theta}} = \left[\frac{d}{d\widehat{\theta}}\left(\frac{c_s'(\widehat{\theta})}{\overline{S}'(\widehat{\theta})}\right)\right]\left(\overline{S}(\theta,\widehat{\theta}) - \overline{S}(\widehat{\theta})\right) + \frac{c_s'(\widehat{\theta})}{\overline{S}'(\widehat{\theta})}\left(\overline{S}'(\theta,\widehat{\theta}) - \overline{S}'(\theta)\right) + c_s'(\widehat{\theta})$$

or

$$\frac{\partial \overline{U}(\theta,\widehat{\theta})}{\partial \widehat{\theta}} = \left[\frac{d}{d\widehat{\theta}}\left(\frac{c_s'(\widehat{\theta})}{\overline{S}'(\widehat{\theta})}\right)\right]\left(\overline{S}(\theta,\widehat{\theta}) - \overline{S}(\widehat{\theta})\right) + \frac{c_s'(\widehat{\theta})}{\overline{S}'(\widehat{\theta})}\overline{S}'(\theta,\widehat{\theta}) \tag{3.10}$$

Now, since S is a strictly proper scoring rule, then $\overline{S}(\theta,\widehat{\theta}) = \overline{S}(\widehat{\theta})$ and $\overline{S}'(\theta,\widehat{\theta}) = 0$[23] when $\widehat{\theta} = \theta$. Hence:

$$\left.\frac{\partial \overline{U}(\theta,\widehat{\theta})}{\partial \widehat{\theta}}\right|_{\widehat{\theta}=\theta} = 0 \tag{3.11}$$

and thus, the utility of the agent is maximised when its reported precision is equal to the precision of the estimate that it subsequently produces ($\widehat{\theta} = \theta$).

At this point, for completeness, we confirm that the second derivative is negative at $\widehat{\theta} = \theta$. To this end, after deriving Equation 3.10, the second derivative of a selected agent's expected utility, is given by

$$\frac{\partial^2 \overline{U}(\theta,\widehat{\theta})}{\partial (\widehat{\theta})^2} = \left[\frac{d^2}{d\widehat{\theta}^2}\left(\frac{c_s'(\widehat{\theta})}{\overline{S}'(\widehat{\theta})}\right)\right]\left(\overline{S}(\theta,\widehat{\theta}) - \overline{S}(\widehat{\theta})\right) + \left[\frac{d}{d\widehat{\theta}}\left(\frac{c_s'(\widehat{\theta})}{\overline{S}'(\widehat{\theta})}\right)\right]\left(\overline{S}'(\theta,\widehat{\theta}) - \overline{S}'(\widehat{\theta})\right) +$$

[23] Note that for notational convenience we denote $\frac{\partial \overline{S}(\theta,\widehat{\theta})}{\partial (\widehat{\theta})}$ and $\frac{\partial^2 \overline{S}(\theta,\widehat{\theta})}{\partial (\widehat{\theta})^2}$ as $\overline{S}'(\theta,\widehat{\theta})$ and $\overline{S}''(\theta,\widehat{\theta})$ respectively

$$+\left[\frac{d}{d\widehat{\theta}}\left(\frac{c_s'(\widehat{\theta})}{\overline{S}'(\widehat{\theta})}\right)\right]\overline{S}'(\theta,\widehat{\theta})+\frac{c_s'(\widehat{\theta})}{\overline{S}'(\widehat{\theta})}\overline{S}''(\theta,\widehat{\theta}) \tag{3.12}$$

Now, S is a strictly proper scoring rule, and therefore $\overline{S}(\theta,\widehat{\theta})=\overline{S}(\widehat{\theta})$ and $\overline{S}'(\theta,\widehat{\theta})=0$ when $\widehat{\theta}=\theta$. Hence, Equation 3.12 takes a simpler form for $\widehat{\theta}=\theta$:

$$\left.\frac{\partial^2\overline{U}(\theta,\widehat{\theta})}{\partial(\widehat{\theta})^2}\right|_{\widehat{\theta}=\theta}=\frac{c_s'(\widehat{\theta})}{\overline{S}'(\widehat{\theta})}\overline{S}''(\widehat{\theta})-c_s''(\widehat{\theta})+\frac{c_s'(\widehat{\theta})}{\overline{S}'(\widehat{\theta})}\overline{S}''(\theta,\widehat{\theta}) \tag{3.13}$$

Now, at $\widehat{\theta}=\theta$ the first term of equation 3.13 is negative because S is strictly proper, and this implies that $\overline{S}''(\theta,\widehat{\theta})$ is negative at $\widehat{\theta}=\theta$ since $\overline{S}(\theta,\widehat{\theta})$ is maximised at $\widehat{\theta}=\theta$. Furthermore, $c_s''(\widehat{\theta})$ is positive, assuming convexity of the cost function, and $\overline{S}''(\widehat{\theta})$ is negative assuming concavity of the scoring rule. Hence, the second derivative is negative at $\widehat{\theta}=\theta$.

To this end, we have shown that an agent is incentivised to report a precision equal to the actual precision of its estimate, and hence, the mechanism is incentive compatible with respect to agent's reported precision. □

Theorem 3.3. The two-stage mechanism is individually rational.

Proof. Having shown in Theorem 3.1 that the true reporting of cost functions in the first stage is a weakly dominant strategy, we only have to examine whether the selected agent is incentivised to participate into the the second stage of the mechanism and report its estimate and its precision to the centre.

In more detail, since agents that do not win in the first stage receive zero utility (Theorem 3.1), we only consider the case of the selected agent. For that agent, its true cost function is less or equal to the cost function used for the scaling of the expected score ($c_t(\theta)\leq c_s(\theta)$).

Given that the selected agent's expected utility is: $\overline{U}(\theta)=\frac{c_s'(\theta_0)}{\overline{S}'(\theta_0)}(\overline{S}(\theta)-\overline{S}(\theta_0))+c_s(\theta_0)-c_t(\theta)$ (equation 3.6), it follows that $\overline{U}(\theta_0)=c_s(\theta_0)-c_t(\theta_0)\geq 0$.

Furthermore, in Lemma 3.2, we have shown that the agent may produce an estimate $\theta^*>\theta_0$. However, since θ^* is optimal by definition, then $\overline{U}(\theta^*)\geq\overline{U}(\theta_0)$, and thus $\overline{U}(\theta^*)\geq 0$. □

Theorem 3.4. For the agent selected in the first stage of the mechanism, it is optimal to produce an estimate with a precision equal or higher than the precision required by the centre, i.e., $\theta^*\geq\theta_0$.

Proof. This proof follows directly from Lemma 3.2 where we show that there is an optimal precision, θ^*, such that $\theta^*\geq\theta_0$ if $c_t(\theta)$ and $c_s(\theta)$ are convex functions with $c_t(\theta)<c_s(\theta)$, $c_t'(\theta)<c_s'(\theta)$ and $c_t(0)=c_s(0)=0$.

In more detail, given that the agents reveal their true cost functions, we have $c_t(\theta)\leq c_s(\theta)$ and from the model description we assume that $c_t(0)=c_s(0)=0$ and that the order of the cost functions is maintained for their derivatives $c_t(\theta) < c_s(\theta) \Rightarrow c'_t(\theta) < c'_s(\theta)$. Therefore, and after covering all Lemma 3.2's conditions, it follows that $\theta^* \geq \theta_0$. □

Corollary 3.2. There can be no optimal precision, for certain concave cost functions.

Proof. We intend to show that our assumption regarding the convexity of the cost functions, is not only a realistic one, driven by the nature of the information providers' costs in citizen sensor networks, but also critical for the achievement of one of the mechanism's goals. In more detail, by showing that there are certain concave functions, for which the optimal precision, θ^*, cannot be defined, we show, that concave cost functions cannot be considered in general, thus justify our approach of employing convex cost functions.

In this context, we need to show that there are concave cost functions so that the optimal precision, θ^*, with $\theta^* = \text{argmax}_\theta \overline{U}(\theta)$ does not exist. We will show this, by identifying at least one type of concave cost functions so that $\overline{U}''(\theta) > 0$, for any value of θ. As a result of this, $\overline{U}(\theta)$ will not be maximised at θ^*, or at any θ.

Against this background, $\overline{U}''(\theta)$ is derived by Equation 3.6 as:

$$\overline{U}''(\theta) = \frac{c'_s(\theta_0)}{\overline{S}'(\theta_0)}\overline{S}''(\theta) - c''_t(\theta) = \delta\overline{S}''(\theta) - c''_t(\theta) \tag{3.14}$$

Now, given that both cost functions and expected score are concave functions, $c''_t(\theta) < 0$ and $\overline{S}''(\theta) < 0, \overline{U}''(\theta) > 0$ for any concave functions, that satisfy the following assumption:

$$c''_t(\theta) < \delta\overline{S}''(\theta) \tag{3.15}$$

□

Note that these proofs indicate that the two stages of the mechanism are inextricably linked and cannot be considered in isolation of one another. Indeed, apparently small changes to the second stage of the mechanism can destroy the incentive compatibility property of the first stage. For example, it is important to note that our mechanism is more precisely known as *interim* individually rational (Mas-Colell et al., 1995), since the utility is positive in expectation. In any specific instance, the payment could actually be negative if the prediction turns out to be far from the actual outcome. An alternative choice for the second stage of the mechanism would be to set β such that the payments are always positive, thus making the mechanism ex-post individually rational. However, this would then violate the incentive-compatibility property since the agents could then receive positive pay-offs by misreporting their cost functions. Likewise, it might be tempting to imagine that the centre could use the revealed costs of the agents in order to request a lower precision, confident in the knowledge that the selected agent will actually produce an estimate of the required precision. However, by effectively using the lowest revealed cost within the payment rule in this way, the incentive-compatibility property of the mechanism would again be destroyed.

3.4 Empirical Evaluation

Having proved the economic properties of the mechanism in the general case for any convex cost function, we now present empirical results for a specific scenario in which costs are linear functions, given by $c_i(\theta) = c_i\theta$, where the value of c_i is drawn from a uniform distribution $c_i \sim \mathcal{U}(1,2)$ and $\theta_0 = 1$. This is the simplest scenario that satisfies our assumptions regarding the cost functions and therefore, allows us to focus on the comparison of the four scoring rules' performance in order to identify the differences between them and select the scoring rule that minimises the centre's expected payment and its variance. To this end, for a range from 2 to 20 agents participating in the first stage, we simulate the mechanism 10^6 times and, for each iteration, record the payment made to the agent that provided the observation and the precision of this observation. Due to the number of iterations that we perform, the standard errors in the mean values plotted are much smaller than the symbol size shown in the plot, and thus for clarity, we omit them.

The payment the agent expects to derive, P, and its actual precision, θ^*, for every value of $N \in [2,20]$ are shown in Figure 3.1. Initially we restrict our analysis to the quadratic, spherical and logarithmic scoring rules for a varying numbers of agents, N. In so doing, we are able to observe the expected payment and required precision as the number of agents increases and then chose a specific value of N in order to focus on the parameter k and how it affects the mechanism's payments. Now, first note that as expected, as the number of agents increases, the mean payment decreases toward the lower limit of the uniform distribution from which the costs were drawn. Furthermore, note that there is a fixed ordering over the entire range, with the payment resulting from the quadratic scoring rule being the highest, and that of the logarithmic scoring rule being the lowest.

In this figure, we also show the mean of the lowest and second lowest costs evaluated at the required precision θ_0 (denoted by $\bar{c}_1\theta_0$ and $\bar{c}_2\theta_0$ respectively). The first cost represents the minimum payment that could have been made if the costs of the agents were known to the centre. While, the second represents the payment that would have been made, had the agent produced an observation of the required precision θ_0 rather than its own optimal precision θ^*. The gap between $\bar{c}_1\theta_0$ and $\bar{c}_2\theta_0$ is the extra amount that must be paid as a result of the costs being unknown and is the same regardless the scoring rule used. On the other hand, the gap between $\bar{c}_2\theta_0$ and the mean payment of any particular scoring rule, depends on the choice of the scoring rule as it represents the loss that the centre has to cover, as a result of the agent producing an estimate at its optimal precision, θ^*. The goal in selecting scoring rules is clearly to minimise this gap, and it can be seen that the logarithmic scoring rule is closest in succeeding in this. The reason for this can be seen in Figure 3.2 where the precision of the observations that

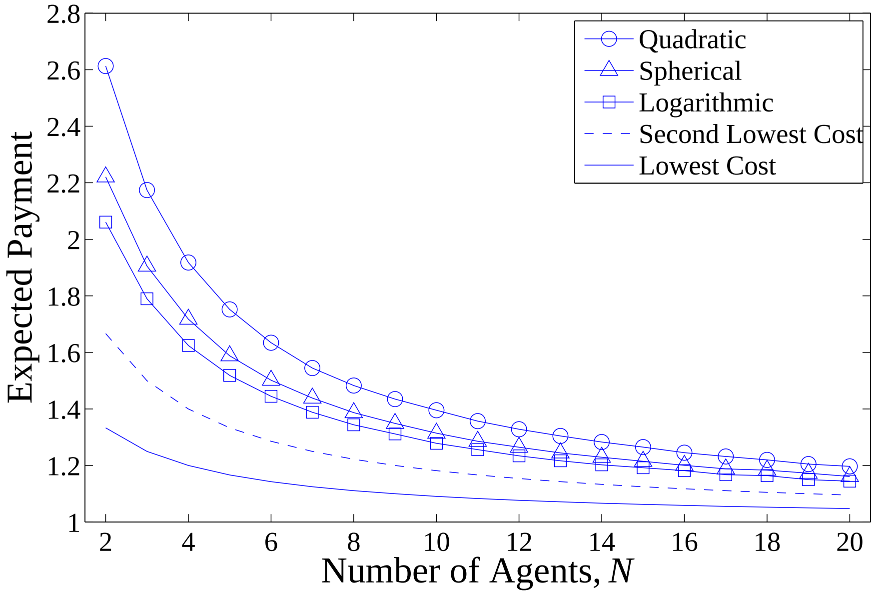

FIGURE 3.1: Selected agent's expected payment.

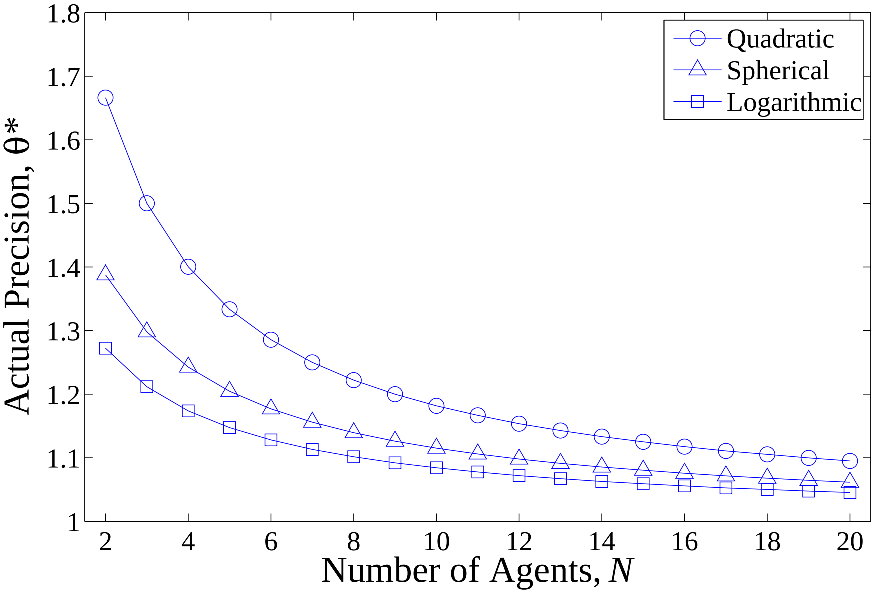

FIGURE 3.2: Selected agent's optimal precision.

were actually made are shown. Note that in this figure the logarithmic scoring rule is shown to induce agents to produce observations closer to the required precision than both the spherical and the quadratic scoring rules, and that is why the gap between the expected payment and $\overline{c}_2\theta_0$ is minimised when the logarithmic scoring rule is used.

In order to emphasise the above result, we derived the analytical expressions of the expected payment, $\overline{P}(\theta_0)$, and the optimal precision, θ^*, for a single run of the mechanism in this specific setting where cost functions are represented by linear functions (Table 3.1). In more detail, initially we calculated the optimal precision, which is the value of θ that maximises the expected utility, $\overline{U}(\theta) = \alpha\overline{S}(\theta) + \beta - c_t(\theta)$. The expected utility was calculated based on Equations of α

TABLE 3.1: Analytical calculation of Quadratic, Spherical, Logarithmic and Parametric Scoring Rules with linear cost functions for an instance of the mechanism

SR:	Quadratic	Spherical	Logarithmic	Parametric
$\overline{P}(\theta_0)$	$c_2\theta_0(2\frac{c_2}{c_1}-1)$	$c_2\theta_0(4\left(\frac{c_2}{c_1}\right)^{\frac{1}{3}}-3)$	$c_2\theta_0(1+\log\left(\frac{c_2}{c_1}\right))$	$\frac{c_2\theta_0}{k-1}\left[2\left(\frac{c_2}{c_1}\right)^{\frac{k-1}{2}}+k-3\right]$
θ^*	$\left(\frac{c_2}{c_1}\right)^2\theta_0$	$\left(\frac{c_2}{c_1}\right)^{\frac{4}{3}}\theta_0$	$\left(\frac{c_2}{c_1}\right)\theta_0$	$\left(\frac{c_2}{c_1}\right)^{\frac{2}{3-k}}\theta_0$

Costs are given by linear functions, $c(\theta) = c\theta$, and c_1 and c_2 are the lowest and second lowest costs.

and β and $\overline{S}(\theta)$ from Table 2.3, after replacing $c_t(\theta)$ and $c_s(\theta)$ by the linear functions $c_1\theta$ and $c_2\theta$, with c_1 and c_2 being the lowest and second lowest cost function. For precision equal to the optimal precision ($\theta = \theta^*$), we then calculated the expected payment $\overline{P}(\theta_0) = \alpha\overline{S}(\theta_0) + \beta$ as a function of the costs c_1, c_2 and the centre's required precision, θ_0.

Against this background, in Table 3.1, we show that when the payment is based on the logarithmic scoring rule, the agent's expected payment is minimised in comparison to the other two scoring rules, while its optimal precision is closer to the required one.

In Table 3.1 we also introduce the parametric scoring rule (already described in Section 2.4.3.1) in our setting. It can be seen that for $k = 2$, $k = 1.5$, and $k \to 1$ the expected payment and optimal precision for the parametric scoring rule, is equal to those of the quadratic, spherical and logarithmic scoring rules. Despite this similarity and the fact that the logarithmic scoring rule is the best choice as it minimises the payment the centre expects to issue to the selected agent and its optimal precision (as indicated by Figures 3.1 and 3.2 and Table 3.1), the parametric scoring rule has a significant advantage over the logarithmic. That is, the existence of finite lower bound for $k \in (1,3)$, since it is a polynomial expression of a distribution density function.

These results are analysed in Table 3.2, where we have analytically calculated the lower and upper bounds of the scoring rules (S^-, S^+) and the scaled payments (P^-, P^+), based on the principle that the lower bound is derived from the scoring rule $S(x_0; x, \theta)$, as the probability of the event is 0 ($\mathcal{N}(x_0; x, 1/\theta) = 0$), while for the upper bound, that probability is 1, ($\mathcal{N}(x_0; x, 1/\theta) = 1$). In this table it can be seen that the logarithmic scoring rule, as opposed to the parametric scoring rule, does not have a finite lower bound (denoted as S^- in row 1 in Table 3.2), which means that if the agent's estimate is far from the actual outcome (hence its p.d.f equal to 0), then a payment (denoted as P^- in row 3 in Table 3.2), based on the logarithmic scoring rule will go to $-\infty$ since the agent's probability density function goes to 0.

However, as $k \to 1$ the payment based on the *scaled* parametric scoring rule does not have a finite lower bound since the scaling parameter α goes to $-\infty$ (see rows 3 and 5 in Table 3.2). Therefore, if the centre requires finite payments, the parametric scoring rule is an equally inappropriate choice to the logarithmic scoring rule. Nevertheless, for the parametric family the effect of an agent's imprecise estimate on its payment can be minimised, if the parameter's value is chosen appropriately.

TABLE 3.2: The upper and lower bounds for the parametric ($k \rightarrow 1$) and logarithmic scoring rules and their resulting payments

Scoring rule:	k-power for $k \rightarrow 1$	Logarithmic
S^-:	0	$-\infty$
S^+:	1	0
P^-:	$-\infty$	$-\infty$
P^+:	$c_s(\theta_0) - 2c_s'(\theta_0)\theta_0$	$c_s'(\theta_0)\theta_0 \left[1 - \log\left(\frac{\theta_0}{2\pi}\right)\right] + c_s(\theta_0)$
α:	∞	$2c_s'(\theta_0)\theta_0$
β:	$c_s(\theta_0) - \infty$	$c_s(\theta_0) - 2c_s'(\theta_0)\theta_0 \left[\frac{1}{2}\log\left(\frac{\theta_0}{2\pi}\right) - \frac{1}{2}\right]$

$c_s(\theta_0)$ is the cost function of the agent with the second lowest cost.

In this context, and in order to determine that value of the parameter, we empirically evaluate the parametric scoring rule for a parameter space $(1, 2.5]$[24] and compare it with the quadratic, spherical and logarithmic scoring rule as the parameter k increases. We choose $N = 10$, since the results of the simulations of the three other scoring rules have already indicated that for high values of N the expected payments converge to the lower limit of the cost's distribution and low values of N are less likely to occur in the applications that we are targeting. The results of this analysis are demonstrated in Figures 3.3 and 3.4, where we show the expected payment made by the centre and its variance. We first note that the expected payment that results from the quadratic, spherical and logarithmic is equal to the expected payment based on the parametric scoring rule, for $k \rightarrow 1$, $k = 1.5$ and $k = 2$ respectively. This is also seen in the analytical results after calculating the expected payment and the optimal precision for any values c_1 and c_2, given linear cost functions ($c(\theta) = c\theta$) with c_1 the lowest cost, and c_2 the second lowest cost (as seen in Table 3.1).

In more detail, as $k \rightarrow 1$ the optimal precision and the resulting payment is asymptotically equal to the optimal precision and payment of the logarithmic scoring rule, while for $k = 1.5$ the optimal precision and expected payment are equal to the optimal precision and expected payment of the spherical scoring rule. Furthermore, the same results apply for $k = 2$ and the quadratic scoring rule, as expected, given that for $k = 2$ the power scoring rule takes the quadratic scoring rule's form.

Now regarding Figure 3.4, we show that the logarithmic scoring rule results in the payment with the lowest variance, while a payment based on the quadratic scoring rule has the highest variance and the spherical is between them. Following the previous results, as $k \rightarrow 1$ and for $k = 2$ the variance of a payment based on the parametric scoring rule is equal to the variance of a payment based on the logarithmic and quadratic scoring rules respectively. However, as opposed

[24]We do not simulate the mechanism for $k > 2.5$ since as k increases the expected payment increases exponentially and the results are skewed towards infinity.

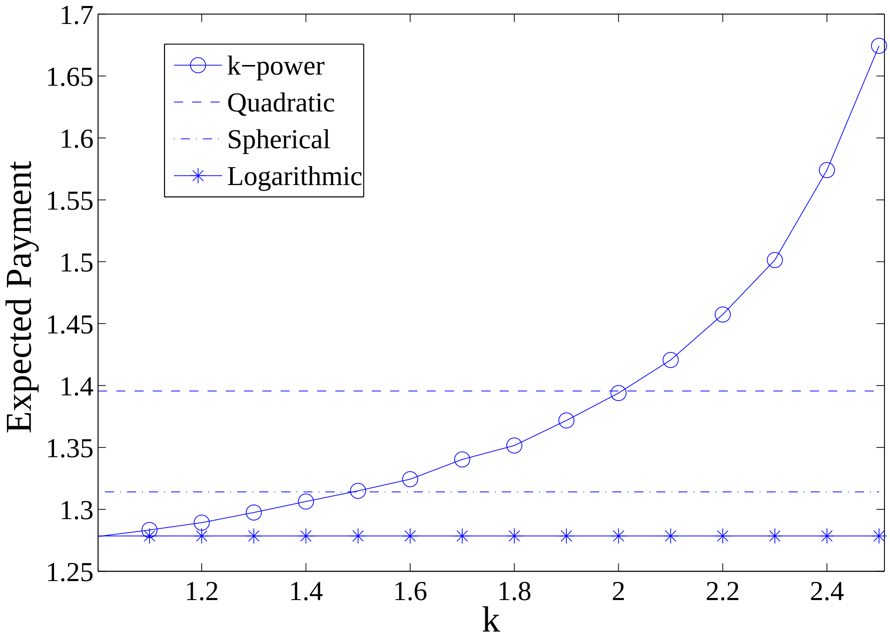

FIGURE 3.3: The mean of the centre's payment.

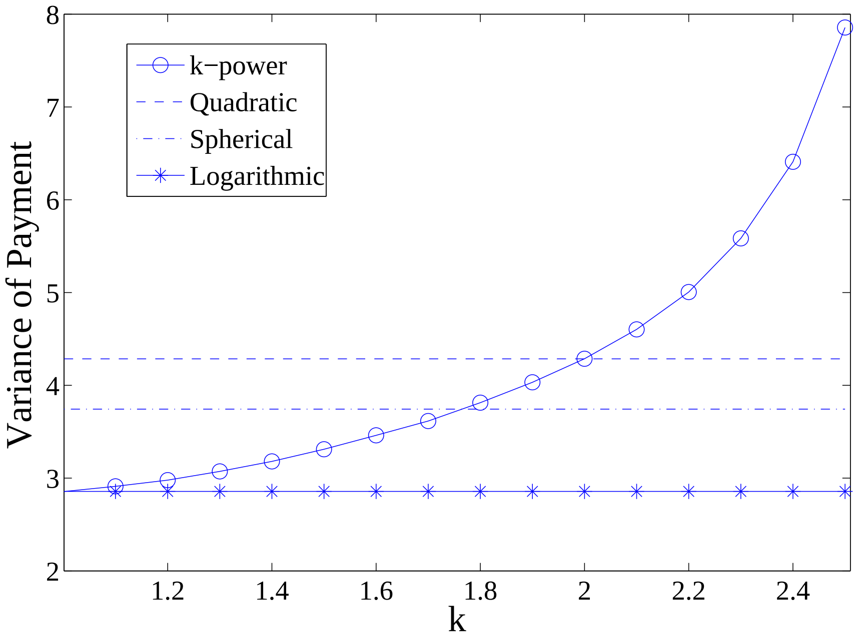

FIGURE 3.4: The variance of the centre's payment.

to the expected payments and optimal precisions, the results are not the same for $k = 1.5$ and the spherical rule, with the variance of a payment based on the spherical rule being greater.

To sum up, so far the logarithmic and parametric scoring rules appear to share both their advantages and disadvantages, and the large negative payments to the agents producing imprecise observations cannot be avoided for these two scoring rules. However, the value of the parameter can be chosen in order to minimise the lower bound of the centre's payments which reflects the maximum fine imposed on an agent providing an extremely inaccurate estimate. In Figure 3.5

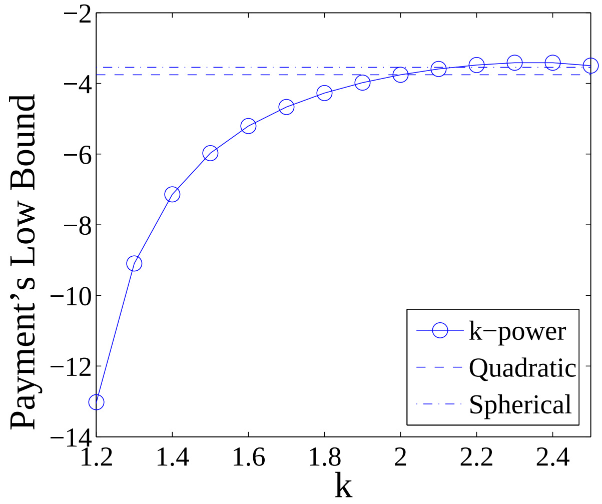

FIGURE 3.5: The fine the agent has to pay to the centre on the worst case scenario.

we plot the lower bound of the payments based on the quadratic, spherical and parametric scoring rules (we omit the logarithmic scoring rule as it goes to $-\infty$), and based on this result, we select k to be equal to 1.2 in our future experiments. This value results in an expected payment and variance that is close to the logarithmic scoring rule, whilst not penalising the agent excessively in the worst case. The choice of parameter value here is is somewhat arbitrary, however, in practise, it will depend on the details of the particular application domain.

3.5 Summary

This chapter introduced a two-stage mechanism based on strictly proper scoring rules that motivates self-interested rational agents to invest sufficient resources in making a costly observation of a specified precision and report it truthfully to a centre, in a setting in which a centre is faced with multiple agents but has no knowledge about the costs involved in the generation of the observations. This mechanism can be applied in citizen sensor networks where the information buyer has to elicit the costs the providers have to face when generating their observations and hence addresses research requirement 1. Furthermore, by proving that the mechanism was incentive compatible in costs, estimates and precisions revealed and that the selected agent will generate an estimate of at least the required precision, we showed that it can successfully address the issue of selfish behaviour in citizen sensor networks and satisfied research requirement 2.

It should be noted that although this mechanism fully satisfies the requirements regarding the providers' unknown costs and selfish behaviour which occurs in a citizen sensor network, it only partially satisfies our requirements regarding operating in an environment where the centre has to combine several possibly inaccurate observations as we discussed in Section 1.1. Therefore, in the next chapter we intend to extend this mechanism so it does not rely in the assumption that a single agent is capable of producing an observation at the centre's required precision. In so doing, we will be able to address the more realistic problem where information providers face restrictions in the precisions of their observations or estimates due to their limited capabilities. Indeed, in the noise monitoring scenario, it is not expected for the members of the public to have the same level of commitment to the project, therefore some will invest less time than others, resulting in observations of lower accuracy. Restrictions for the information providers may also appear in a traffic update service application, where companies offering such a service, may have a limited source of available data due to budget restrictions and prediction models which are not state of the art. In such cases, the produced estimates of the traffic conditions in a particular location, will have low quality and it is very likely to be inaccurate. Hence in the next chapter we propose a mechanism where, the centre (acting as information buyer) will have to combine multiple observations or probabilistic estimates of possibly low quality in order to achieve a higher degree of accuracy.

Chapter 4

A Mechanism for Dealing with Multiple Agents that have a Limited Degree of Precision

In the previous chapter we considered the case where a single agent is able to generate an estimate of at least the precision required by the centre. Now, as already mentioned in Chapter 1, this may not always be the case in many citizen sensor networks where information providers face restrictions in the resources available to them for the generation of their observations and therefore have a limited ability to produce estimates of arbitrary quality. In such cases the information buyer may have to procure observations from multiple providers and fuse them together in order to achieve a sufficiently high precision. For example, for the noise monitoring scenario, where members of the public measure the noises around them with their mobile phones, it is not expected for all them to invest the same time in producing the observations, hence there will be some observations of low quality, which are not necessarily representative of the sound pollution levels in a particular area. Likewise, in a traffic update service, where companies provide optimal travel routes and estimates of the congestion on the streets based on data they collect from various sources (such as stationary sensors in traffic lights or mobile sensors in taxis), they may have limited sources of incoming data or poor prediction methods. Therefore, they are not able to provide estimates of the traffic conditions that are of satisfactory precision.

To this end, we revise the mechanism in the previous chapter by relaxing the assumption that a single agent is capable of estimates of any precision the centre may require and therefore address research Requirement 3 set in Chapter 1. In so doing we contribute to the state of the art, by proposing a parametrised iterative mechanism (Mechanism 2), which although it may appear to be similar to the previous mechanism (i.e. two stages, first stage to elicit costs, second to calculate payments), it is significantly different, specifically regarding the process used to elicit those costs and calculate the payments.

We first describe the extension of the information elicitation problem outlined in Section 3.1 so it takes into consideration citizen sensor networks which consist of agents facing limitations in the precisions of the observations they can provide (Section 4.1), while in Section 4.2 we provide a two-stage mechanism that deals with this requirement effectively. In more detail, in the first stage the centre pre-selects M from N agents through a series of $m+1$ reverse price auctions with $m \leq M$, while in the second stage it elicits the pre-selected agents' observations, after sequentially approaching them in a random order. Having detailed the mechanism, we formally prove that it is incentive compatible regarding the reported costs, maximum precisions and observations, that pre-selected agents generate an observation at that maximum precision and that it is an individually rational mechanism (Section 4.3). Furthermore, in Section 4.4 we empirically compare several approaches by which the centre may pre-select M from N agents and show both empirically and analytically that if the centre forms a single group of agents in the first stage it minimises its expected payment. Finally, in Section 4.5 we summarise and discuss how to relax, and therefore address the final requirement (Requirement 4), regarding the centre's access to knowledge about the observed event after it receives the agents' observations.

4.1 Eliciting Information from Multiple Sources

We now extend the model described in Chapter 3, as we consider a citizen sensor network where the information providers face restrictions in the precision of the observations they can provide. As a consequence of these restrictions, there is now a limit in the maximum precisions of the providers' observations, denoted by θ_i^c. Thus, information providers can produce observations of any precision up to and including this maximum value (i.e. $0 \leq \theta_i \leq \theta_i^c$). We consider this requirement very realistic as in many citizen sensor networks, information providers cannot produce observations at any precision a buyer may require. In a noise monitoring scenario this limit can be dictated by the time available to the member of the public making the sound observation, while in the traffic information service, the limit in the precision of the estimates of traffic congestion is defined by the sources available to the companies and the statistical models they use to estimate the traffic.

In both these examples, restrictions enforce limitations on the quality and accuracy of the observations generated by the information providers, whether they collect this information themselves or rely on external sources. In order to deal with these shortcomings, those interested in acquiring this information have to combine multiple observations. That is in the noise monitoring scenario several members of the public on nearby locations, while in the traffic scenario drivers can subscribe to multiple traffic information services. In all these cases, combining several observations of possibly low quality, would result in an observation of satisfactory quality, which would assist the information buyer in achieving its desired degree of accuracy. To this end, in order to fuse k conditionally independent and unbiased probabilistic estimates, $\{\widehat{x}_1, \ldots, \widehat{x}_k\}$ of possibly different precisions $\{\widehat{\theta}_1, \ldots, \widehat{\theta}_k\}$, into a single estimate with mean $\bar{x}$ and precision $\bar{\theta}$, the information buyer uses the standard result (see DeGroot and Schervish (2002)):

$$\bar{x} = \frac{\sum_{i=1}^{k} \widehat{x}_i \widehat{\theta}_i}{\bar{\theta}} \quad \text{and} \quad \bar{\theta} = \sum_{i=1}^{k} \widehat{\theta}_i \tag{4.1}$$

By fusing the providers' reported estimates, the information buyer manages to acquire an estimate of higher precision than the precision of any of the individual agents. Indeed, it can be seen that $\bar{\theta} \geq \theta_i$ for any agent i. However, note that for this fusion to be appropriate, agents must be incentivised to truthfully report both their estimates and their precisions. Otherwise, misreported estimates which are not an accurate representation of the state of the world, will be fused with truthful estimates and thus the efficiency of the network will be degraded as the estimated from the reports state of the world will not match the actual one.

To this end, given this model, the challenge is to design a mechanism in which the centre will be able to initially identify those agents that can provide their estimates at the lowest cost (Requirement 1), then motivate these agents to truthfully report their maximum precisions and finally generate and truthfully report their estimates with precisions equal to their reported maximum precisions (Requirement 2) in an environment where agents have limitations in the precisions of the observations they can provide (Requirement 3). By addressing these challenges, we will be solving research requirements 1, 2 and 3 introduced in Section 1.1.

4.2 The Mechanism

In Mechanism 2 we extend the mechanism discussed in the previous chapter by relaxing the assumption that the centre can select a single agent that can provide the estimate at the required precisions. The centre can now elicit estimates from multiple agents which have limited precisions in the estimates they can provide. In order to address this issue, the centre in the first stage, iteratively pre-selects M of the N available agents based on their reported costs, while in the second stage agents belonging in a subset of the pre-selected agents identified in the first stage are incentivised to generate an estimate at their reported maximum precision and truthfully report it to the centre.

In more detail, in the first step (Step 1.1) of the first stage the centre divides all the available agents, N, into groups of $n \leq N$ agents and asks the first group of n agents to reveal their cost functions. Then in Step 1.2 it implements a reverse $m+1$ auction in order to select the m cheapest agents, with $m < n$. This procedure is repeated until all N agents have been asked to report their cost functions and by the end of the first stage M agents are pre-selected.

In the second stage, the centre initialises the iterative process by setting its required precision equal to θ_r (Step 2.1) and in the next step asks a random agent from the M pre-selected agents to reveal its private maximum precision (Step 2.2). In Step 2.3, the centre checks whether by asking that agent to report its maximum precision it achieves the precision it requires. If not, it sets the required precision equal to the remaining needed precision ($\theta_r = \theta_r - \widehat{\theta}_j^c$) and repeats Steps 2.2 and 2.3 until it satisfies its required precision. In Step 2.4 those agents selected through the iteration of Steps 2.2 and 2.3 are asked to produce an estimate at their reported maximum precision and the centre announces the scaled strictly proper scoring rule. The scaling parameters α and β are given by:

$$\alpha_j = \frac{c_s'(\widehat{\theta}_j^c)}{\overline{S}'(\widehat{\theta}_j^c)} \quad \text{and} \quad \beta_j = c_s(\widehat{\theta}_j^c) - \frac{c_s'(\widehat{\theta}_j^c)}{\overline{S}'(\widehat{\theta}_j^c)}\overline{S}(\widehat{\theta}_j^c) \tag{4.2}$$

where $\widehat{\theta}_j^c$ is the reported maximum precision, and c_s is the $(m+1)^{th}$ cost which was associated with that agent at an iteration of Step 1.2. Then it proceeds to the final step (Step 2.5) where the agents that were selected during the iterations of Steps 2.2 and 2.3 to produce an estimate at their reported maximum precision, produce these estimates at any precisions and report them. Although in this step we could restrict agents to report their estimates with precision $\widehat{\theta}_j^c$, we shall show in Section 4.3, under this mechanism the agents are automatically incentivised to report $\widehat{\theta}_j = \widehat{\theta}_j^c$ anyway. Finally, those agents receive their payments after the centre observes the actual outcome, while the rest of the M agents are discarded at Step 2.5.

We now proceed to prove that this mechanism leads the agents to truthfully reveal their costs in the first stage (so that those which can produce the estimate at the lowest cost can be identified), and that the M pre-selected agents are incentivised to truthfully report their maximum precisions to the centre and subsequently make and truthfully report estimates of these precisions in the second stage. These properties are not obvious, and as in the single agent mechanism (Mechanism 1 in Chapter 3), they depend rather subtly on the details of the mechanism. For example, we note that if after asking all M agents for their maximum precisions, the centre does not achieve its required precision, the mechanism must proceed to the payment phase (Step 2.5). That is, the centre must commit to paying all pre-selected agents for their estimates at their reported maximum precisions, even if it does not acquire its required precision. Failure to observe this policy would lead agents to over-report their maximum precision, in order that some payment was received, and thus, the mechanism would no longer be incentive compatible in terms of maximum precisions.

Furthermore, note that in Step 1.2, the centre chooses the m agents with the lowest reported costs, and discards the remaining $n-m$ agents. If these agents were not discarded, but were placed back into the pool of available agents, then the mechanism would no longer be incentive compatible in terms of costs; agents would have an incentive to over-report their costs, such that when they are eventually pre-selected, their payment rule will be calculated using a higher cost. Finally, in Step 2.2, the centre must randomly ask the pre-selected agents to report their maximum precisions using an ordering which is independent of their reported costs. Failing to do so will undermine incentive compatibility in terms of costs of the first stage of the mechanism, thereby illustrating how the two stages interact.

Mechanism 2 The mechanism for dealing with multiple agents that can provide estimates of a limited precision:

1. First Stage

1.1 The centre selects $n \geq 2$ agents from the available N and asks them to report their cost functions $\widehat{c}_i(\theta)$ with $i \in \{1, \ldots, n\}$.

1.2 The centre selects the m, $(1 \leq m < n)$, agents with the lowest costs, associates the $(m+1)^{th}$ cost with these agents and discards the remaining $n - m$ agents.

1.3 The centre repeats the above two steps until it has asked all N agents to report their cost functions. Note that when N is not exactly divisible by n and we have a single remainder, it is discarded. Otherwise in the final round the centre modifies n and m such that $n = N \bmod n$ and $m = \min(m, n-1)$.

1.4 We denote the total number of the agents pre-selected in this stage as M and note that its value depends on N, n and m.

2. Second Stage

2.1 The centre sets its required precision θ_0 equal to θ_r.

2.2 The centre randomly selects one of the pre-selected agents and asks it to report its maximum precision $\widehat{\theta}_k^c$, with $k \in \{1, ..., M\}$.

2.3 The centre sets $\theta_r = \theta_r - \min(\theta_r, \widehat{\theta}_k^c)$ and if $\theta_r > 0$ it repeats Step 2.2 of the second stage.

2.4 The centre asks agent j, belonging in the agents selected through the iteration of Steps 2.2 and 2.3, to produce an estimate at its reported maximum precision, $\widehat{\theta}_j^c$, and presents it with a scaled strictly proper scoring rule with the scaling parameters α_j and β_j being determined using equation 4.2.

2.5 Each of these agents produces an estimate x_j with precision θ_j and reports $\widehat{x}_j$ and $\widehat{\theta}_j$ to the centre, which after observing the actual outcome, x_0, issues the following payments:

$$P_j(x_0; \widehat{x}_j, \widehat{\theta}_j) = \alpha_j S_j(x_0; \widehat{x}_j, \widehat{\theta}_j) + \beta_j \tag{4.3}$$

with α_j and β_j being already determined in Step 2.4.

4.3 Economic Properties of the Mechanism

Having detailed the mechanism, we now identify and prove its economic properties. Specifically, we show that:

1. The mechanism is incentive compatible with respect to the pre-selected agents' reported estimates and maximum precisions.
2. The actual precision of the generated estimate is equal to the reported maximum precision.
3. The mechanism is incentive compatible with respect to the agents' reported costs.
4. The mechanism is individually rational.

Theorem 4.1. The mechanism is incentive compatible with respect to the pre-selected agents' reported estimates and maximum precisions which is also the precision of the generated estimate.

Proof. Regarding the pre-selected agents' reported estimates, the proof of incentive compatibility follows directly from the definition of the strictly proper scoring rules. As per in Theorem 3.2, a pre-selected agent's expected utility (Equation 4.4) is an affine transformation of the expected score which is maximised when a pre-selected agent reports truthfully its private estimate.

Now, regarding the pre-selected agents reported maximum precision, although the proof may look similar to that of Theorem 3.2, in this case we intend to show that a pre-selected agent will not only truthfully report its maximum precision, but that this will be the actual precision of its generated estimate.

In this context, when the agent reports its estimate, it must do so with the precision that it claimed was its maximum. Thus, $\hat{\theta} = \widehat{\theta}^c$. Now, given the scaling of the scoring rules described in Step 2.4 the mechanism, the expected utility of the agent, if it reports its maximum precision as $\widehat{\theta}^c$, and subsequently produces an estimate of precision θ, which it reports with precision $\widehat{\theta}^c$, is denoted by $\overline{U}(\theta, \widehat{\theta}^c)$, and is given by:

$$\overline{U}(\theta, \widehat{\theta}^c) = \frac{c_s'(\widehat{\theta}^c)}{\overline{S}'(\widehat{\theta}^c)} \left(\overline{S}(\theta, \widehat{\theta}^c) - \overline{S}(\widehat{\theta}^c) \right) + c_s(\widehat{\theta}^c) - c_t(\theta) \tag{4.4}$$

where $\overline{S}(\theta, \widehat{\theta}^c)$ is the agent's expected score for producing an estimate of precision θ and reporting its precision as $\widehat{\theta}^c$. Furthermore, $\overline{S}(\widehat{\theta}^c)$ is the agent's expected score for producing and truthfully reporting an estimate of precision $\widehat{\theta}^c$, $c_t(.)$ is the true cost function of the agent, and $c_s(.)$ is the cost function used to produce the scoring rule (i.e. the $(m+1)^{th}$ lowest revealed cost in the group from which the agent was pre-selected). Taking the first derivative of this expression with respect to $\widehat{\theta}^c$ gives:

$$\frac{d\overline{U}(\theta,\widehat{\theta}^c)}{d\widehat{\theta}^c} = \frac{d}{d\widehat{\theta}^c}\left(\frac{c_s'(\widehat{\theta}^c)}{\overline{S}'(\widehat{\theta}^c)}\right)\left(\overline{S}(\theta,\widehat{\theta}^c)-\overline{S}(\widehat{\theta}^c)\right)+\frac{c_s'(\widehat{\theta}^c)}{\overline{S}'(\widehat{\theta}^c)}\overline{S}'(\theta,\widehat{\theta}^c) \tag{4.5}$$

Now, since S is a strictly proper scoring rule, then $\overline{S}(\theta,\widehat{\theta}^c) = \overline{S}(\widehat{\theta}^c)$ and $\overline{S}'(\theta,\widehat{\theta}^c) = 0$ when $\theta = \widehat{\theta}^c$. Hence:

$$\left.\frac{d\overline{U}(\theta,\widehat{\theta}^c)}{d\widehat{\theta}^c}\right|_{\widehat{\theta}^c=\theta} = 0 \tag{4.6}$$

and thus, the utility of the agent is maximised when it reveals as its maximum precision, the precision of the estimate that it subsequently produces.

For completeness, we confirm that the second derivative is negative at $\widehat{\theta}^c = \theta$. To this end, the second derivative is given by:

$$\frac{d^2\overline{U}(\theta,\widehat{\theta}^c)}{d(\widehat{\theta}^c)^2}(\widehat{\theta}^c=\theta) = \frac{c_s'(\widehat{\theta}^c)}{\overline{S}'(\widehat{\theta}^c)}\overline{S}''(\theta,\widehat{\theta}^c) - c_s''(\widehat{\theta}^c) + \frac{c_s'(\widehat{\theta}^c)}{\overline{S}'(\widehat{\theta}^c)}\overline{S}''(\widehat{\theta}^c) \tag{4.7}$$

Now, the first term of equation 4.7 is negative because S is strictly proper, and this implies that $\overline{S}''(\theta,\widehat{\theta}^c)$ is negative at $\widehat{\theta}^c = \theta$. Furthermore, $c_s''(\widehat{\theta}^c)$ is positive, assuming convexity of the cost function, and $\overline{S}''(\widehat{\theta}^c)$ is negative assuming concavity of the scoring rule. Hence, the second derivative is negative at $\widehat{\theta}^c = \theta$.

We now show that it will actually produce an estimate of precision equal to its reported maximum precision. To this end, we note that when $\widehat{\theta}^c = \theta$, the expected utility of the agent is given by:

$$\overline{U}(\theta) = c_s(\theta) - c_t(\theta) \tag{4.8}$$

Since $c_s(.)$ and $c_t(.)$ and their derivatives do not cross or overlap and maintain their order (as required by the initial model Section 3.1), then $\overline{U}(\theta)$ is a strictly increasing function. Thus the agent will maximise its expected utility by producing an estimate at its maximum precision, and thus, $\theta = \theta^c$, and hence, $\widehat{\theta}^c = \hat{\theta} = \theta^c$, as required. □

Theorem 4.2. The mechanism is incentive compatible with respect to the agents' reported costs.

Proof. We prove this by contradiction and consider two cases depending on whether or not an agent is pre-selected in the first stage of the mechanism as a result of its misreporting. Let $c_t(.)$ and $\widehat{c}(.)$ denote an agents' true and reported cost functions respectively. Furthermore, let $c_s(.)$ denote the cost function used to scale the scoring rule if that agent is among the m agents with the lowest reported costs in its group of n agents in the first stage of the mechanism (i.e. $c_s(.)$ is the $(m+1)^{th}$ cost of that group).

First, suppose that the agent's misreporting does not affect whether it is pre-selected or not. In this case, had the agent been pre-selected, its payment would have been based on the $(m+1)^{th}$ cost of its group and therefore independent of its own report. Conversely, had the agent not been pre-selected, it would have received zero utility, since the remaining $n-m$ agents, of a group of initially n agents, that are not pre-selected are discarded. Hence, there is no incentive to misreport.

Second, suppose that the agent's misreporting affects whether that agent is pre-selected or not. There are now two cases:

1. The agent is pre-selected by misreporting, but would have not been if it was truthful.
2. The agent is not pre-selected by misreporting, but would have been if it was truthful.

In this context:

- Case (1) can be formally denoted as $c_t(\widehat{\theta}^c) > c_s(\widehat{\theta}^c)$ and $\widehat{c}(\widehat{\theta}^c) < c_s(\widehat{\theta}^c)$. Now, since the true cost $c_t(\widehat{\theta}^c) > c_s(\widehat{\theta}^c)$, it follows directly from Theorem 4.1 that the expected utility $\overline{U}(\theta) = c_s(\theta) - c_t(\theta)$ is strictly negative, irrespective of θ. Therefore, the agent could do strictly better by reporting truthfully in which case the expected utility is zero.
- Case (2) can be formally denoted as $c_t(\widehat{\theta}^c) < c_s(\widehat{\theta}^c)$ and $\widehat{c}(\widehat{\theta}^c) > c_s(\widehat{\theta}^c)$. In this case the agent would have been pre-selected if it was truthful, but now receives a utility of zero since it hasn't been pre-selected due to its misreporting. To show that this type of misreporting is suboptimal, we need to show that, when $c_t(\widehat{\theta}^c) < c_s(\widehat{\theta}^c)$, an agent benefits from being pre-selected, since it may then be asked to generate an estimate at its reported maximum precision, $\widehat{\theta}^c$. It follows directly from Theorem 4.1 that $\overline{U}(\widehat{\theta}^c) = c_s(\widehat{\theta}^c) - c_t(\widehat{\theta}^c) > 0$ when $c_t(\widehat{\theta}^c) < c_s(\widehat{\theta}^c)$, and therefore there is no incentive for an agent that would have been pre-selected to misreport its cost function.

Hence, we have shown that a pre-selected agent is incentivised to report its truthful cost function, as required. □

Theorem 4.3. The mechanism is interim individually rational.

Proof. Due to Theorem 4.2, we can assume that all agents, and consequently those pre-selected, will report their true cost functions, and therefore $c_t(\theta) \leq c_s(\theta)$. In Theorem 4.1 we show that the expected utility $\overline{U}(\theta) = c_s(\theta) - c_t(\theta)$ is strictly non-negative, irrespective of θ. Therefore, the expected utility of a pre-selected agent that generates an estimate of precision equal to its reported maximum precision $\widehat{\theta}^c$, is strictly non-negative (i.e. $\overline{U}(\widehat{\theta}^c) \geq 0$), and hence the mechanism is *interim* individually rational. Note that the mechanism is *interim* individually rational, since the utility is non-negative in expectation but there may be instances in which the payment could be negative if the prediction turns out to be far from the actual outcome. Setting β in the second stage, such that the payments are always positive, would make the mechanism ex-post individually rational. However, this would then violate the incentive-compatibility property since the agents could receive positive pay-offs by misreporting their cost functions in the first stage. □

In this section we have proved the economic properties for the class of mechanisms that is defined by the different approaches in performing the pre-selection (i.e. the values of n and m) in the mechanism's first stage. The mechanism's economic properties are not affected by the process followed in the stage of pre-selection, as long as some global rules are followed, such as agents that are not pre-selected are discarded and do not end up in the initial pool of agents. However, the same does not happen for the total payment issued by the centre, or its probability of achieving it required precision, which depend on how the groups of n and m agents are formed in Steps 1.1, 1.2 and 1.3. Now, the issued total payment and the probability define the efficiency of a mechanism of that class. Therefore, in the following section, through the empirical evaluation of Mechanism 2, we determine which are the optimal values of the parameters n and m in order to identify those that minimise the payment and maximise the probability of achieving the required precision.

4.4 Empirical Evaluation

Against this background, in this section we present empirical results for a specific scenario in order to explore the effect that the parameters n and m have on the centre's total payments, and on the probability of achieving its required precision. In more detail, as before, the cost functions are represented by linear functions, given by $c_i(\theta) = c_i\theta$, where c_i are independently drawn from a uniform distribution $c_i \sim \mathcal{U}(1,2)$. The maximum precisions of the selected agents, θ_i^c, are independently drawn from another uniform distribution $\theta_i^c \sim \mathcal{U}(0,1)$ and finally the centre's required precision, θ_0, is equal to 1.7 in order to generate representative results whereby the probability of achieving the required precision, $P(\theta_0)$, covers a broad range of values in $[0,1]$. Finally, we restrict our analysis to the use of the parametric scoring rule for $k = 1.2$ as we have shown in Section 3.4, that among the common rules and for various values of the parameter k, this rule is a good choice for a centre intending to issue low payments with low variance, that still remain bounded.

Given this, and for $N = 7$, we explore all possible combinations of n and m given the constraints that $2 \leq n < N$ and $1 \leq m < n$. We choose $N = 7$ since for this number of agents the difference between the scoring rule based payment and the lower bound denoted by the average $(m+1)^{th}$ cost is still significant, as it can be seen in Figure 3.1, for the single agent case (Mechanism 1), where it shows that as N increases the payments converge to that lower bound. In this context, for each combination, we simulate the mechanisms 10^6 times and for each iteration we record whether the centre was successful in acquiring an estimate at its required precision, and the sum of all the payments it issued to those agents that were asked to produce an estimate. In Figure

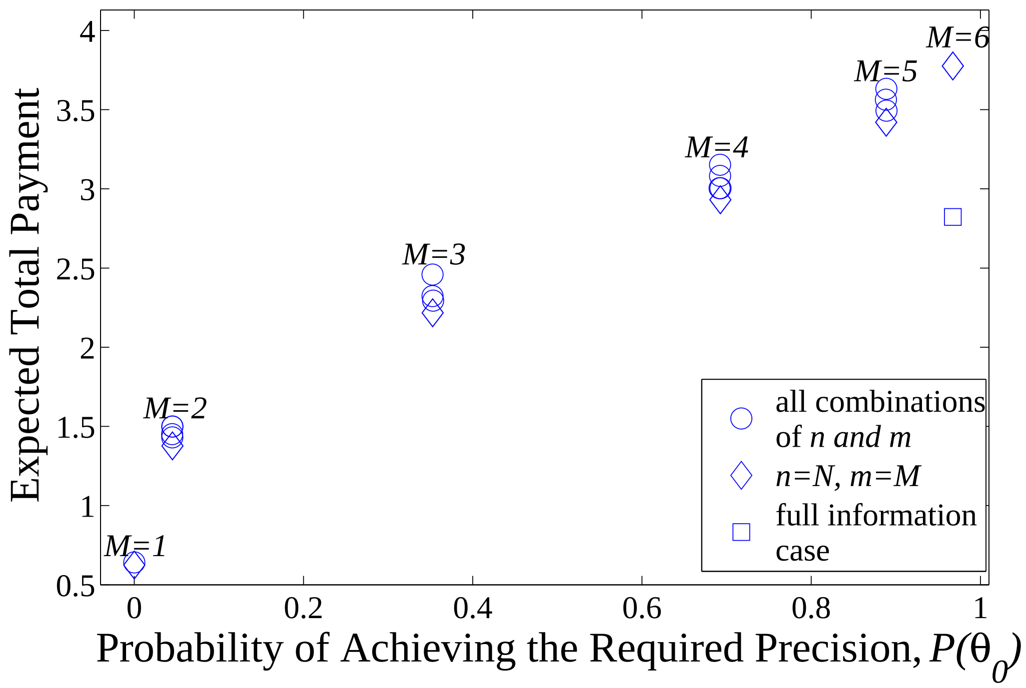

FIGURE 4.1: Centre's probability of achieving the required precision and the mean total payment it has to issue.

4.1 we plot, for each possible combination of n and m, the probability of acquiring the required precision and the total payment made by the centre. We note that again the standard error of the mean values is insignificant and for clarity we omit it.

The circles in Figure 4.1 depict the results for all possible combinations of n and m, except where $n = N$ and $m = M$, which is indicated by a diamond (the reason for this will become clear shortly). Furthermore, the square indicates the case where the centre has full information of the agents' costs. This case results in significantly lower total payments to the agents since the centre is able to select those agents with the lowest costs to generate the estimates, and it needs only pay them sufficient to ensure that their expected utility is infinitesimally greater than zero. In more detail, the centre minimises its expected payment by not having to implement the $m+1$ reverse auctions in the first stage, since it directly chooses the agents with the lower costs. Indeed, implementing $m+1$ reverse auctions and choosing the $(m+1)^{th}$ cost to determine the scaling parameters of the scoring rule payments, leads to higher payments since the $(m+1)^{th}$ cost is always larger than the costs of the m agents. Hence, when costs are known, the payment is less, given that it depends on the scaling parameters α and β which in turn are linear expressions of the cost functions. Therefore, this case represents a lower bound for the payments of the mechanism, since the centre can minimise the payment that it expects to issue to the selected agents given that it only has to cover their cost which is known in advance.

Now, regarding Figure 4.1, we first note that many possible combinations of n and m give rise to the same value of $P(\theta_0)$, and thus the family of possible pre-selection methods fall into 6 distinct columns. This is because this probability depends only on the number of agents that are pre-selected (denoted by M) and many of these combinations result in the same number of agents being pre-selected (e.g. if $N = 7$, both $n = 4$, $m = 2$ and $n = 5$, $m = 3$ result in $M = 4$). In the following theorem, we show that this probability increases, as the number of the pre-selected agents increases:

Theorem 4.4. The probability of achieving the required precision by the centre increases as M increases.

Proof. Although in the empirical evaluation of the mechanism we rely on uniform distributions for the reported precisions, in this proof we do not follow these restrictions, and show that the probability of achieving the required precision increases as the number of agents increases for any positive distribution of the reported precisions.

In this context, we consider M independent random variables to represent the M agents' maximum precisions with support $[a,b]$ such that $a \geq 0$ and $b > 0$. If Θ_M is a random variable such that $\Theta_M = \theta_1^c + ... + \theta_M^c$, then the probability, p_M, of achieving the required precision, θ_0, using M agents is denoted as following:

$$p_M = \int_{\theta_0}^{\infty} f_{\Theta_M}(\Theta_M) d\Theta_M \tag{4.9}$$

where $f_{\Theta_M}(\Theta_M)$ is the distribution of Θ_M.

Likewise, the probability of achieving the required precision, if we consider $M+1$ independent random variables is the following:

$$p_{M+1} = \int_{\theta_0}^{\infty} f_{\Theta_{M+1}}(\Theta_{M+1}) d\Theta_{M+1} \tag{4.10}$$

where $f_{\Theta_{M+1}}(\Theta_{M+1})$ is the distribution of Θ_{M+1}.

However, $f_{\Theta_{M+1}}(\Theta_{M+1})$ takes the following form:

$$f_{\Theta_{M+1}}(\Theta_{M+1}) = \int_0^{\infty}\int_0^{\infty} \delta(\Theta_{M+1} - \Theta_M - \theta_{M+1}) f_{\theta_{M+1}}(\theta_{M+1}) d\theta_{M+1} f_{\Theta_M}(\Theta_M) d\Theta_M$$

Now, given that $\Theta_M \geq 0$:

$$f_{\Theta_{M+1}}(\Theta_{M+1}) = \int_0^{\infty} f_{\Theta_M}(\Theta_{M+1} - \theta_{M+1}) f_{\theta_{M+1}}(\theta_{M+1}) d\theta_{M+1}$$

We apply this result in Equation 4.10, and hence p_{M+1} takes the following form:

$$p_{M+1} = \int_0^{\infty} d\theta_{M+1} \int_{\theta_0}^{\infty} f_{\Theta_M}(\Theta_{M+1} - \theta_{M+1}) f_{\theta_{M+1}}(\theta_{M+1}) d\Theta_M$$

$$= \int_0^{\infty} f_{\theta_{M+1}}(\theta_{M+1}) d\theta_{M+1} \int_{\theta_0}^{\infty} f_{\Theta_M}(x) dx$$

However, since $\Theta_{M+1} \geq 0$:

$$\int_{\theta_0-\theta_{M+1}}^{\infty} f_{\Theta_M}(x)dx \geq \int_{\theta_0}^{\infty} f_{\Theta_M}(x)dx = p_M$$

Hence,

$$p_{M+1} \geq \int_0^{\infty} f_{\theta_{M+1}}(\theta_{M+1})d\theta_{M+1} \cdot p_M \geq p_M$$

□

Second, note that for each possible value of M, the case where $n = N$ and $m = M$ dominates all other combinations of n and m (i.e. it results in the lowest mean total payment). This case corresponds to a single selection stage in which M agents are pre-selected directly from the original N in a single step. This result is demonstrated empirically in Figure 4.1, and we present an analytical proof for the specific case that we consider here:

Theorem 4.5. In a setting with linear cost functions, where agents' costs and maximum precisions are independently drawn from uniform distributions, for a given probability of achieving θ_0, the centre minimises its expected total payment when $n = N$ and $m = M$.

Proof. Given the mechanism and setting described above, we first note that when the costs of the agents are represented by linear functions, then $c_i(\theta) = c_i\theta$, and hence, $c_i'(\theta) = c_i$. Using this result within the scaling parameters of the payment rule described in Step 2.4, gives the result:

$$\alpha_j = \frac{c_s}{\overline{S}'(\widehat{\theta}_j^c)} \quad \text{and} \quad \beta_j = c_s\widehat{\theta}_j^c - \frac{c_s}{\overline{S}'(\widehat{\theta}_j^c)}\overline{S}(\widehat{\theta}_j^c) \tag{4.11}$$

Thus, both α and β are proportional to c_s, and hence the payment to any agent is also proportional to the cost used in the calculation of the scaling parameters. Secondly, we note that due to the random selection of agents within the second stage of the mechanism, the precision of the estimate generated by any agent is independent of the cost used to generate its payment rule. Hence, the expected total payment to the agents is proportional to the mean cost used to generate their payment rules.

Now, the costs used to generate the payment rule of any agent is the $(m+1)^{th}$ lowest reported cost when m agents are pre-selected from n. Thus, in order to show that setting $n = N$ and $m = M$ minimises the expected total payment of the centre, we must simply show that the expected value of the $(M+1)^{th}$ cost when pre-selecting M agents from N, is lower than any other combination. To do so, we note that if the costs of the agents are i.i.d. from the standard uniform distribution[25], and the agents report truthfully (as they are incentivised to do), then the density function that describes the $(m+1)^{th}$ cost, denoted by $C_{m+1:n}$, is given by:

$$c_{m+1:n}(u) = \frac{n!}{m!(n-m-1)!}u^m(1-u)^{n-m-1}, \quad 0 \leq u \leq 1 \tag{4.12}$$

[25]For notational simplicity we shall assume that the costs are drawn from $\mathcal{U}(0,1)$, but we note that the proof is valid for a uniform distribution of any support.

and Arnold et al. (2008) show that $c_{m+1:n}(u) \sim B(m+1, n-m)$ and therefore the mean of this distribution is simply $\frac{m+1}{n+1}$. Thus, we now prove that the $(M+1)^{th}$ cost when pre-selecting all M agents directly from N is less than the expected cost that results from first pre-selecting m agents from n and then pre-selecting the remaining $M-m$ agents from $N-n$. Therefore, and given that $c_{M+1:N}(u) \sim B(M+1, N-M)$ and $c_{M-m+1:N-n}(u) \sim B(M-m+1, N-n-M+m)$, we must prove the inequality:

$$\left(\frac{M+1}{N+1}\right) \leq \frac{m}{M}\left(\frac{m+1}{n+1}\right) + \frac{M-m}{M}\left(\frac{M-m+1}{N-n+1}\right) \tag{4.13}$$

subject to the constraints that $M < N$, $m < n$ and $N-n > M-m$, and we note that if it holds in this case, then it holds for all possible combinations of n and m.

A first step towards the proof of equation 4.13, is performing the following substitutions:

$$a = m, \;\; b = M-m, \;\; c = n, \;\; d = N-n$$

and now equation 4.13 takes the following form:

$$\frac{(a+b)(a+b+1)}{c+d+1} \leq \frac{a(a+1)}{c+1} + \frac{b(b+1)}{d+1} \tag{4.14}$$

with $a,b,c,d \geq 0$, $a < c$, and $b < d$.

Now, by multiplying all fractions in equation 4.14 to obtain a common denominator,$(c+1)(d+1)(c+d+1)$, and noting that this denominator is positive, translates equation 4.14 into the following condition:

$$\begin{aligned}(a+b)(a+b+1)(c+1)(d+1)& \\ -a(a+1)(c+d+1)(d+1)& \\ -b(b+1)(c+d+1)(c+1) &\leq 0\end{aligned} \tag{4.15}$$

We can rearrange this expression into the form:

$$F_1(a,b,c,d) + F_2(a,b,c,d) + F_3(a,b,c,d) \leq 0 \tag{4.16}$$

where:

$$F_1(a,b,c,d) = -(d \cdot a - b \cdot c)^2 \tag{4.17}$$

$$F_2(a,b,c,d) = -b(c-a)^2 - b^2(c-a) - a(d-b)^2 - a^2(d-b) \tag{4.18}$$

$$F_3(a,b,c,d) = a(b-d) + b(a-c) \tag{4.19}$$

Now, it is easy to verify that F_1, F_2, and F_3 are all negative given the initial constraints that $a,b,c,d \geq 0$, $a < c$ and $b < d$. Hence, it follows that equation 4.16 is negative. □

Finally, we note that given that it is always preferable to set $n = N$, the choice of the value of m is determined by the trade-off between the total payment made by the centre and the probability of it acquiring an estimate of its required precision. If the distributions of cost and maximum precisions are known, this can be evaluated prior to running the mechanism through simulation. However, if these distributions are unknown, setting $m = N - 1$ ensures that the probability of acquiring the required precision is maximised (but doing so will also incur the greatest expected payment).

4.5 Summary

In this chapter, we extended our original two-stage mechanism introduced in Section 3 by relaxing the assumption that a single agent can provide an estimate of infinite precision and introduced limitations on the agents' maximum precisions. As a result of this, we have introduced an extended mechanism which, in addition to Requirements 1 and 2 regarding elicitation of unknown costs and addressing selfish behaviour, already satisfied by the previous mechanism, successfully addresses research Requirement 3 regarding the need for multiple observations. In so doing, this mechanism extends the state of the art by being the first mechanism that incentivises agents to invest the maximum resources available to them for the generation of their estimates, despite the involved costs, and in turn, report them truthfully to a centre. In addition to that, the agents face restrictions in the precisions of the estimates they can provide, and the centre does not have any knowledge of the costs of the estimates.

Against this background, this mechanism can be applied in citizen sensor networks where information providers might not be able to produce observations of a sufficient precision to individually meet the information buyer's needs, hence leaving the buyer no option but to combine multiple such estimates and fuse them into a more accurate one. However, although this mechanism solves research Requirements 1, 2 and 3 as introduced in Section 1.1, it does not deal with the research requirement regarding the information buyer's knowledge of the outcome of the observed event (Requirement 4), since it relies on the assumption that the centre will observe the state of the world after receiving the agents' estimates. Now, since this is not always the case, in the following chapter we relax this assumption and hence satisfy the final research requirement. In so doing, we will be able to apply our mechanism in more realistic and efficient citizen sensor networks where the information buyer will not have spare resources available in order to evaluate the providers' observations, hence it will not have access to any information about the state of the world after it receives the observations from the providers. For example, for both the noise monitoring and traffic information scenarios, it is rather unrealistic to assume that the true state of the world is available to the centre at some time in the future in order to calculate the agent payments, since for both cases it is very unlikely for the information buyers to have access to the additional infrastructure needed to evaluate the providers.

Therefore it is crucial, to design a mechanism which in not relying on external sources for the evaluation of the actual outcome it will save the information buyer from an additional cost, and by evaluating the observations as soon as they are reported, will result in fair payments which reflect the providers' efforts which do not depend on a delayed observed state of the world . In order to address these issues effectively we will design a new mechanism based on a modified strictly proper scoring rule which will evaluate each agent against all other agents' fused observations immediately upon receiving their reports and not at some time the future after receiving the actual outcome of the observed event. Thus this mechanism will lead towards robust, effective and fair citizen sensor networks, which, first, will not have to rely on external sources of the evaluation of the providers' observations, and, second, will be able to adapt in dynamic environments where conditions evolve constantly.

Chapter 5

A Mechanism for Dealing with the Centre's Lack of Access to Knowledge of the Outcome

In each of the previous two chapters we have considered mechanisms whereby the centre has access to the actual outcome of the estimated event which it can use to evaluate the agents' reported estimates when calculating their payments. In Mechanism 1, where we introduced this assumption, we clarified that it is not realistic in general, but it can hold for specific cases we described. Specifically, when the information buyer has enough resources available to invest for the infrastructure needed to verify the providers' reported observations. For example, in the noise monitoring scenario, these resources may be sophisticated microphones being located in numerous fixed locations needed to assess the providers' observations. Similarly, in terms of the traffic feedback service these extra resources may be services which provide live views of the streets.

However, as already mentioned in Chapter 1, these cases are rare, since citizen sensor networks are deployed and operate in dynamic environments (such as the busy streets of a city, in both the examples we consider). In these dynamic environments, the state of the world constantly changes and evolves. In our examples, these changes could be a result of a random and unpredicted event that is likely to occur in a city such as a heavy vehicle passing next to a member of the public generating a noise observation, or a traffic accident that makes the estimate of an optimal route obsolete. Now, unless the information buyers have a substantial amount of available resources to invest in evaluating the observations, the difference between the reported and the observed outcomes will be substantial as there will be a time delay between the reported observation and the realised outcome. Therefore, it is very important to develop a mechanism that will not rely on knowledge of the actual outcomes of the observed or estimated events.

To this end, in this chapter we revise the previous mechanism by removing the assumption that the centre (which represents the information buyer) will have knowledge of the estimated event, after receiving the estimates from the agents. In so doing we address the final research requirement (Requirement 4) set in Chapter 1, regarding operating in dynamic environments of uncertainty where an information buyer cannot or is not willing to evaluate the received reports. In satisfying this requirement, we also contribute to the state of the art by proposing a novel two stage mechanism in which the centre uses the fused reported estimates to calculate the payments to the pre-selected agents in the second stage, rather than the final revealed outcome. In order to ensure incentive compatibility for this mechanism too, we modify the standard strictly proper scoring rules so they can cope effectively with the centre's lack of knowledge of the estimated event. Without this change, the standard rules would no longer motivate the agents to truthfully report the precisions of their estimates.

In more detail, we first update our initial model, introduced in Section 3.1, so it can now relate to citizen sensor networks where information buyers have no means of evaluating the providers' observations due a lack of knowledge about the outcome of the observed event. In this context, in Section 5.1 we describe the updated information elicitation problem, show how standard strictly proper scoring rules cannot address this problem effectively, and hence introduce the modified strictly proper scoring rules that do not rely on knowledge of the outcome. In Section 5.2 we present a new two stage mechanism based on these modified scoring rules. In Section 5.3 we show that our mechanism is incentive compatible in maximum precisions and estimates revealed. In more detail, we show that using the fused reported estimates instead of the outcome when calculating the scaling parameters and the payments to the pre-selected agents in the second stage of the mechanism, results in truthful reporting of these agents' maximum precisions and estimates being a Nash equilibrium. Furthermore, in Section 5.4 we empirically evaluate our mechanism and compare it with the one we introduced in the previous chapter, in which the centre has access to the actual outcome and with a modified version of the 'peer prediction mechanism' proposed by Miller, Resnick and Zeckhauser (Section 2.4.3) so that the centre has no knowledge of the costs, in order to examine the effects of the fusion in the payments of the pre-selected agents. We show both analytically and empirically that for all the mechanisms we simulate, the agents expect to derive the same payment, which means that the centre incurs no additional cost as a result of its lack of knowledge of the outcome. However, we identify a significant difference between the fusion and the peer prediction methods, by showing that in our mechanism the variance of the total payment issued to the selected agents by the centre is significantly lower than the total payment's variance in the 'prediction mechanism' for $M > 2$ (although both are greater than the case where the outcome is known). Finally, in Section 5.5 we summarise and conclude.

5.1 Evaluating Information without Knowledge of the Outcome

This model is the same as the model discussed in Section 4.1, where we introduced limits in the precisions of the agents' generated estimates. Due to these restrictions, the centre had to fuse multiple estimates of possible low precision in order to acquire an estimate of higher precision.

However, here we additionally remove the assumption regarding the centre's access to knowledge of the outcome of the estimated event. As a consequence, the centre now has to rely solely on the observations it receives from the agents. In relaxing this assumption, we satisfy the final requirement of this research and can now consider citizen sensor networks where information buyers either cannot or do not want to necessarily evaluate the accuracy of the received information. As already mentioned above, this is very likely to happen in real world applications of citizen sensor networks deployed in dynamic environments. In such environments conditions change constantly and information buyers need to have access to costly equipment in order to evaluate the reported observations within minimum elapsed time before they become obsolete and not relevant to the evaluation.

Against this background, an information buyer relying solely on the reported observations of the providers, has two options when using these reports in order to shape its perception of reality: *peer prediction* and *fusion*. In more detail, in the peer prediction mechanism (Section 2.4.3) the centre's perception of reality is represented by a single agent's reported estimate. That is, in this mechanism the centre scores an agent directly against each one of the rest of the individual agents and then calculates its payment by averaging the scaled scores. However, in the mechanism we introduce in this section, the centres uses the fused reported estimates and precisions of all the other agents and excludes from the fusion process the agent that is currently receiving the payment. This is necessary, since fusing its estimate with the rest would incentivise that agent to report a very high precision in order to influence the final result[26].

In more detail, when there are $K \geq 2$ available agents, the centre calculates the payment to agent i, after fusing the K_{-i} agents' conditionally independent and unbiased probabilistic estimates, $\{\widehat{x}_1, .., \widehat{x}_{i-1}, \widehat{x}_{i+1}, .., \widehat{x}_k\}$ of possibly different precisions $\{\widehat{\theta}_1, .., \widehat{\theta}_{i-1}, \widehat{\theta}_{i+1}, .., \widehat{\theta}_k\}$, into one estimate with mean $\overline{x}$ and precision $\overline{\theta}$ by using the standard result (see DeGroot and Schervish (2002)):

$$\overline{x}_{-i} = \frac{\sum_{j \in K_{-i}} \widehat{x}_j \widehat{\theta}_j}{\overline{\theta}_{-i}} \quad \text{and} \quad \overline{\theta}_{-i} = \sum_{j \in K_{-i}} \widehat{\theta}_j \tag{5.1}$$

where $K_{-i} = \{1, .., i-1, i+1, .., k\}$ is the set that contains all k agents besides agent i, which is the agent that is receiving the payment from the centre.

In this case, it is in any selected agent's best interest to consider its belief about the fused observations of all the other agents when reporting its precision. Now, this means that agent i's expected score $\overline{S}(\overline{x}; x_i, \theta_i)$ is maximised not at $\widehat{\theta}_i = \theta_i$ but at $\widehat{\theta}_i = \theta_i + \overline{\theta}_{-i}$. Indeed, if $N(\overline{x}_{-i}; x_i, 1/\theta_i + 1/\overline{\theta}_{-i})$ and $\mathcal{N}(\overline{x}_i; \widehat{x}_i, 1/\widehat{\theta}_i)$ are Gaussian distributions with mean and variance $(x_i, 1/\theta_i + 1/\overline{\theta}_{-i})$ and $(\widehat{x}_i, 1/\widehat{\theta}_i)$, which represent agent i's true and reported estimate's distributions, agent i's expected score, which is given by:

$$\overline{S}(\overline{x}_{-i}; \widehat{x}_i, \widehat{\theta}_i) = \int_{-\infty}^{\infty} \mathcal{N}(\overline{x}_{-i}; x_i, 1/\theta_i + 1/\overline{\theta}_{-i}) S(\overline{x}_{-i} | \mathcal{N}(\overline{x}_i; \widehat{x}_i, 1/\widehat{\theta}_i)) d\overline{x}_{-i} \tag{5.2}$$

[26]Given that the fused parameter depends on the precision of the agents' estimates, if an agent claims that its observation is very precise and reports an increased precision, the value of the fused parameter will be very close to its reported observation.

will be maximised at $\widehat{\theta}_i = \theta_i + \overline{\theta}_{-i}$, since for that value of the reported precision, $\widehat{\theta}_i$, the two distributions become identical.

Subsequently, an agent wanting to maximise its expected score (equation 5.2), will have to report $\theta_i + \overline{\theta}_{-i}$ instead of θ_i, which is impossible since it does not have access to other agents' precisions ($\overline{\theta}_{-i}$). However, given that the centre, when calculating the payments, has access to both θ_i and $\overline{\theta}_{-i}$, it can modify the strictly proper scoring rule so that the agent is only required to report θ_i but the payment is calculated using $\theta_i + \overline{\theta}_{-i}$. Thus, we introduce the *modified strictly proper scoring rule*, $S(\overline{x}_{-i}; \widehat{x}_i, \widehat{\theta}_i + \overline{\theta}_{-i})$ and in the following theorem we show that it incentivises agent *i* to report its true parameters (estimate and its precision).

Theorem 5.1. Truthful revelation of an agent's parameters is a Nash Equilibrium.

Proof. Given that an agent *i*'s estimate is represented by the Gaussian distribution $\mathcal{N}(x_0; x, 1/\theta)$, under the modified strictly proper scoring rules, the score it expects to derive is the following:

$$\overline{S}(\overline{x}_{-i}; \widehat{x}_i, \widehat{\theta}_i + \overline{\theta}_{-i}) = \int_{-\infty}^{\infty} \mathcal{N}(\overline{x}_{-i}; x_i, 1/\theta_i + 1/\overline{\theta}_{-i}) S(\overline{x}_{-i} | \mathcal{N}(\overline{x}_i; \widehat{x}_i, 1/\widehat{\theta}_i + 1/\overline{\theta}_{-i})) d\overline{x}_{-i} \quad (5.3)$$

where $\mathcal{N}(\overline{x}_{-i}; x_i, 1/\theta_i + 1/\overline{\theta}_{-i})$ and $\mathcal{N}(\overline{x}_i; \widehat{x}_i, 1/\widehat{\theta}_i + 1/\overline{\theta}_{-i})$ are Gaussian distributions with mean and variance $(x_i, 1/\theta_i + 1/\overline{\theta}_{-i})$ and $(\widehat{x}_i, 1/\widehat{\theta}_i + 1/\overline{\theta}_{-i})$ respectively, which will be denoted as Q and R respectively. Now equation 5.3 takes the following form:

$$\overline{S} = \int_{-\infty}^{\infty} Q(\overline{x}_{-i}) S(\overline{x}_{-i} | R(\overline{x}_{-i})) d\overline{x}_{-i} \quad (5.4)$$

At this point it should be noted that S is a strictly scoring function, as defined by Hendrickson and Buehler (1971) and Savage (1977), with the strictly proper scoring rule, $S(\overline{x}_i; \widehat{x}_i, \widehat{\theta}_i)$ and the modified strictly proper scoring rule, $S(\overline{x}_{-i}; \widehat{x}_i, \widehat{\theta}_i + \overline{\theta}_{-i})$ being two different expressions. The key property of the strictly scoring function $S(Q,R)$ is that its expected value is maximised when $Q = R$.

Now, given the definition of Q and R, for $\widehat{x}_i = x_i$ and $\widehat{\theta}_i = \theta_i \Rightarrow Q = R$. Therefore, for $\widehat{x}_i = x_i$ and $\widehat{\theta}_i = \theta_i$, a payment based on the modified strictly scoring rule, $S(\overline{x}_{-i}; \widehat{x}_i, \widehat{\theta}_i + \overline{\theta}_{-i})$, is incentive compatible, since an agent will maximise its expected payment if it reports truthfully its parameters, assuming that all other agents also report their true parameters. The latter makes truthful reporting a Nash equilibrium since an agent will maximise its utility, thus making this strategy the optimal, if all other agents report truthfully their parameters too. □

5.2 The Mechanism

Having defined the *modified strictly proper scoring rules*, in this section, we extend the two-stage mechanism introduced in Section 4.2 by relaxing the assumption that the centre will have access to the actual outcome when calculating the payments to the pre-selected agents (Mechanism 3). To this end, in the first stage the centre pre-selects M out of N agents and identifies their cost functions, while in the second stage it calculates their payments.

In more detail, the first stage is a variation of the first stage of Mechanism 2. However, in this mechanism, instead of dividing the available agents into groups of n agents and then initiating multiple reverse $(m+1)^{th}$ price auctions to pre-select M agents, the centre pre-selects them directly from the N available through a single reverse $(M+1)^{th}$ auction. This specific approach to pre-selection is based on the empirical and theoretical results presented in Section 4.4. Therefore, we modify the first stage of this mechanism accordingly, so that the centre minimises the total payment it has to issue to the pre-selected agents, by performing that pre-selection instantly, and not through multiple, iterations.

In the beginning of the second stage, the centre initialises an iterative process by setting the required precision equal to θ_r (Step 2.1). In Step 2.2, the centre randomly asks each of the pre-selected agents to report their maximum precisions, until it achieves its required precision (Step 2.3). Each of the selected agents through the iteration of the previous steps, is asked to produce an estimate at its reported maximum precision and the centre determines the scaling parameters (Step 2.4). That is, for agent j, from that group, parameters α_j and β_j, now, are determined using $\overline{\theta}_{-j}$, the fused reported precisions of every agent besides agent j. The cost used in the scaling parameters now is fixed to c_{M+1}, since for this mechanism, the centre implements a single reverse $(M+1)^{th}$ price auction in the first stage. In the final step of this mechanism, each of the agents selected in Steps 2.2 and 2.3 produces an estimate at any precision (and not necessarily their reported maximum precision) and report its parameters to the centre. The centre calculates their payments by using the scaled modified scoring rule. However, in this mechanism we deny the centre of any such knowledge, therefore when agent j receives its payment, the centre uses the fused reported estimates of all the other agents, instead of the actual outcome x_0 that has been used for both the previous mechanisms (Mechanisms 1 and 2).

Mechanism 3 The mechanism for dealing with the case where the centre does not have access to the actual outcome:

1. First Stage

1.1 The centre asks $N \geq 2$ agents to report their cost functions $\widehat{c}_i(\theta)$ with $i \in \{1, \ldots, N\}$.

1.2 The centre selects the M, $(1 \leq M < N)$, agents with the lowest costs, associates the $(M+1)^{th}$ cost with these agents and discards the remaining $N-M$ agents.

2. Second Stage

2.1 The centre sets its required precision θ_0 equal to θ_r.

2.2 The centre randomly selects one of the pre-selected agents and asks it to report its maximum precision $\widehat{\theta}_k^c$, with $k \in \{1, ..., M\}$.

2.3 The centre sets $\theta_r = \theta_r - \min(\theta_r, \widehat{\theta}_k^c)$ and if $\theta_r > 0$ it repeats Step 2.2.

2.4 The centre asks agent j, belonging in the agents selected through the iteration of Steps 2.2 and 2.3, to produce an estimate at its reported maximum precision, $\widehat{\theta}_j^c$, and presents it with a scaled strictly proper scoring rule. Scaling parameters α_j and β_j are now based on the expected value of the *modified scoring rule*, $\overline{S}(\widehat{\theta}_j^c, \overline{\theta}_{-j})$, and its derivative. The parameters are given by:

$$\alpha_j = \frac{c'_{M+1}(\widehat{\theta}_j^c)}{\overline{S}'(\widehat{\theta}_j^c, \overline{\theta}_{-j})} \quad \text{and} \quad \beta_j = c_{M+1}(\widehat{\theta}_j^c) - \frac{c'_{M+1}(\widehat{\theta}_j^c)}{\overline{S}'(\widehat{\theta}_j^c, \overline{\theta}_{-j})} \overline{S}(\widehat{\theta}_j^c, \overline{\theta}_{-j}) \tag{5.5}$$

where, c_{M+1} is the $(M+1)^{th}$ cost identified the first stage and $\overline{\theta}_{-j}$ is the fused precisions of all the agents that are asked to produce an estimate except agent j and is defined in Equation 5.1.

2.5 Each of these agents produces an estimate x_j with precision θ_j and reports $\widehat{x}_j$ and $\widehat{\theta}_j$ to the centre, which in turn issues the following payment:

$$P_j(\overline{x}_{-j}; \widehat{x}_j, \widehat{\theta}_j + \overline{\theta}_{-j}) = \alpha_j S_j(\overline{x}_{-j}; \widehat{x}_j, \widehat{\theta}_j + \overline{\theta}_{-j}) + \beta_j \tag{5.6}$$

with α_j and β_j already determined in Step 2.4.

5.3 Economic Properties of the Mechanism

Having detailed the mechanism, we now prove the incentive compatibility of the mechanism with respect to the revealed maximum precisions and estimates. Note that truthful reporting of the agents' cost functions in the first stage is still a dominant strategy and that this mechanism is also individually rational, like all the previous mechanisms. We refrain from re-writing the proofs of both incentive compatibility with respect to reported costs and individual rationality, since they are identical to the proofs of Theorems 4.2 and 4.3 respectively as they do not depend on the fused reported precisions and estimates. However, the same does not hold for the pre-selected agents' reported maximum precisions and estimates, as we have shown in Theorem 5.1, truthful reporting of an agent's parameters is a Nash equilibrium. Putting this result in the context of Mechanism 3, where pre-selected agents report their maximum precisions and estimates leads to the following theorem:

Theorem 5.2. Truthful reporting of the maximum precisions and estimates is a Nash equilibrium, with the reported maximum precision being the actual precision of the estimate.

Proof. In the mechanism described above, when agent j reports its estimate, its reported precision, $\widehat{\theta}_j$, is equal to its reported maximum precision, $\widehat{\theta}^c_j$. Indeed, $\hat{\theta}_j > \widehat{\theta}^c_j$ is not possible given that the centre is already informed of the agent's maximum precision, and $\hat{\theta}_j < \widehat{\theta}^c_j$ would not be in the agent's best interest since under reporting its precision would lead to a smaller payment. Therefore, $\hat{\theta}_j = \widehat{\theta}^c_j$. Now, given the scaling of the scoring rules described in Step 2.4, the expected utility of the agent, if it reports its maximum precision as $\widehat{\theta}^c_j$, and subsequently produces an estimate of precision θ_j, which then reports with precision $\widehat{\theta}^c_j$, is denoted by $\overline{U}_j(\theta_j, \widehat{\theta}^c_j)$, and given by:

$$\overline{U}_j(\theta_j, \widehat{\theta}^c_j) = \frac{c'_s(\widehat{\theta}^c_j)}{\overline{S}'_{f,j}(\widehat{\theta}^c_j, \overline{\theta}_{-j})} \left(\overline{S}_{f,j}(\theta_j, \widehat{\theta}^c_j, \overline{\theta}_{-j}) - \overline{S}_{f,j}(\widehat{\theta}^c_j, \overline{\theta}_{-j})\right) + c_s(\widehat{\theta}^c_j) - c_t(\widehat{\theta}^c_j) \tag{5.7}$$

where $\overline{S}_{f,j}(\theta_j, \widehat{\theta}^c_j, \overline{\theta}_{-j})$ is agent j's expected score for producing an estimate of precision θ and reporting to the centre, $\widehat{\theta}^c$ and $\overline{S}_{f,j}(\widehat{\theta}^c_j, \overline{\theta}_{-j})$ is agent j's expected score for producing and truthfully reporting an estimate of precision $\widehat{\theta}^c_j$. Furthermore, $c_t(.)$ is the true cost function of the agent, and $c_s(.)$ is the $(M+1)^{th}$ lowest revealed cost which was determined in the first stage.

Note that $\overline{S}_{f,j}$ is the expected value of the modified scoring rule $\overline{S}(\overline{x}_{-j}; \widehat{x}_j, \widehat{\theta}_j + \overline{\theta}_{-j})$, already defined in Theorem 5.1:

$$\overline{S}(\overline{x}_{-j}; \widehat{x}_j, \widehat{\theta}_j + \overline{\theta}_{-j}) = \int_{-\infty}^{\infty} \mathcal{N}(\overline{x}_{-j}; x_j, 1/\theta_j + 1/\overline{\theta}_{-j}) S(\overline{x}_{-j} | \mathcal{N}(\overline{x}_j; \widehat{x}_j, 1/\widehat{\theta}_j + 1/\overline{\theta}_{-j})) d\overline{x}_{-j} \tag{5.8}$$

Now, taking the first derivative of expected utility (Equation 5.7) with respect to $\widehat{\theta}^c_j$ gives:

$$\frac{d\overline{U}_j(\theta_j,\widehat{\theta}_j^c)}{d\widehat{\theta}_j^c} = \frac{d}{d\widehat{\theta}_j^c}\left(\frac{c_s'(\widehat{\theta}_j^c)}{\overline{S}_{f,j}'(\widehat{\theta}_j^c,\overline{\theta}_{-j})}\right)\left(\overline{S}_{f,j}(\theta_j,\widehat{\theta}_j^c,\overline{\theta}_{-j}) - \overline{S}_{f,j}(\widehat{\theta}_j^c,\overline{\theta}_{-j})\right) + \\ + \frac{c_s'(\widehat{\theta}_j^c)}{\overline{S}_{f,j}'(\widehat{\theta}_j^c,\overline{\theta}_{-j})}\overline{S}_{f,j}'(\theta_j,\widehat{\theta}_j^c,\overline{\theta}_{-j}) \tag{5.9}$$

We have already shown that truthful revelation is a Nash equilibrium for the modified scoring rule, $S_{f,j}$ (Theorem 5.1). Hence, when $\theta_j = \widehat{\theta}_j^c$:

$$\overline{S}_{f,j}(\theta_j,\widehat{\theta}_j^c,\overline{\theta}_{-j}) = \overline{S}_{f,j}(\widehat{\theta}_j^c,\overline{\theta}_{-j}) \quad \text{and} \quad \overline{S}_{f,j}'(\theta_j,\widehat{\theta}_j^c,\overline{\theta}_{-j}) = 0$$

and subsequently:

$$\left.\frac{d\overline{U}_j(\theta_j,\widehat{\theta}_j^c)}{d\widehat{\theta}_j^c}\right|_{\widehat{\theta}_j^c=\theta_j} = 0 \tag{5.10}$$

Therefore, a preselected agent's expected utility is maximised when it reveals as its maximum precision, the precision of the estimate that it subsequently produces, given that all other agents do the same.

We now show that it will actually produce an estimate of precision equal to its reported maximum precision. To this end, we note that when $\widehat{\theta}_j^c = \theta_j$, the expected utility of the agent is given by:

$$\overline{U}_j(\theta) = c_s(\theta_j) - c_t(\theta_j) \tag{5.11}$$

Since $c_s(.)$ and $c_t(.)$ and their derivatives do not cross or overlap, and their order is maintained, $\overline{U}_j(\theta_j)$ is a strictly increasing function (given that $\overline{U}_j'(\theta_j) > 0$. Thus the agent will maximise its expected utility by producing an estimate at its maximum precision, and thus, $\theta_j = \theta_j^c$, and hence, $\widehat{\theta^c}_j = \hat{\theta}_j = \theta_j^c$, as required. □

We have shown now that in our modified scoring rule based mechanism, truthful reporting of maximum precisions and estimates is a Nash equilibrium. Although this result is significant as it proves a general property for this class of mechanisms (for any values of M and N and any convex cost functions), it does not give any information regarding the performance of the mechanism. That is, there is no information about the payment's mean and variance, or how they are affected as the parameters (M, N) of the mechanism change. This is critical information, as it forms the foundation which will allow us to compare our mechanism with other mechanisms in the literature, and see whether it outperforms them. In more detail, we can now compare our mechanism with the mechanism introduced in the previous chapter (Mechanism 2) in order to see the effects of centre's lack of knowledge of the outcome, and the 'peer prediction mechanism' after modifying it so it satisfies our requirements (i.e. unknown costs and limited maximum precisions) in order to compare different approaches in replacing the centre's knowledge of the observed event.

5.4 Empirical Evaluation

In this section we present empirical results for a specific scenario, in order to compare it with Miller et al.'s peer prediction mechanism so we can demonstrate the differences between fusion and peer prediction and identify which one is more appropriate for a centre that intends to minimise the mean and variance of the total payment issued to the agents. In this scenario, like the previous ones, the cost functions are represented by linear functions, given by $c_i(\theta) = c_i\theta$, where c_i are independently drawn from a uniform distribution $c_i \sim \mathcal{U}(1,2)$. Also, the maximum precisions of the selected agents, θ_i^c, are independently drawn from another uniform distribution $\theta_i^c \sim \mathcal{U}(0,1)$ and, as before, the centre's required precision, θ_0, is equal to 1.7. Finally, as before, we use the parametric family of scoring rules and set the parameter equal to 1.2. Within this scenario our intention is to explore the effects of fusion and peer prediction in the pre-selected agents' payments and therefore we will compare the payments' mean and variance, while using as benchmark the case where the centre has access to knowledge of the actual outcome (Mechanism 2 from Chapter 4).

For the purposes of this analysis, the peer prediction mechanism had to be slightly modified in order to eliminate the assumption that the centre has knowledge of the agents' costs. Hence, we transform the peer prediction mechanism to a *two-stage peer prediction mechanism*, in which the centre in the first stage asks all agents, N, to report their cost functions and then pre-selects M of them, while in the second it allocates the payments to the agents providing the estimates. The first stage is identical to the first stage of Mechanism 2 presented in Section 4.2, while in the second stage the centre calculates the payment to an agent not by using the fused reported estimates, but by scoring that agent against each one of the remaining $M-1$ agents and then by averaging over the $M-1$ respective payments.

To this end, for $N = 7$, we evaluate our mechanisms (both Mechanisms 2 and 3) and the two-stage peer prediction mechanism and define the lower bound of the mechanisms' payments which is represented by the case where the centre has access to the agents' cost functions (denoted as the 'full information' case) and thus can optimally allocate the estimate to the agent it needs in order to achieve the required precision. In more detail, we simulate the mechanisms 10^6 times and for each iteration we record whether the centre was successful in acquiring its required precision and the sum of the payments issued to the selected agents and we calculate the mean and variance of that total payment (Figures 5.1 and 5.2). In Figure 5.1 we calculate for $M \in \{1,2,3,4,5,6\}$ the probability the centre has of achieving the required precision and the total payment (we omit errorbars, since given the number of the iterations of the mechanism the standard error in the plotted values is smaller than the symbol size), while in Figure 5.2 we calculate the variance of the total payment the centre issues to the selected agents.

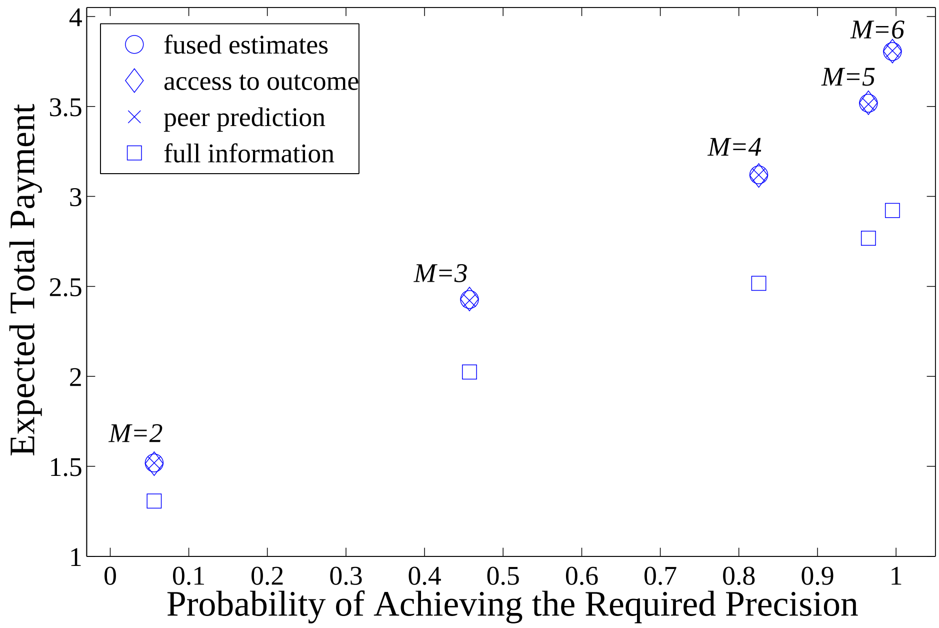

FIGURE 5.1: Centre's probability of achieving the required precision and the mean total payment.

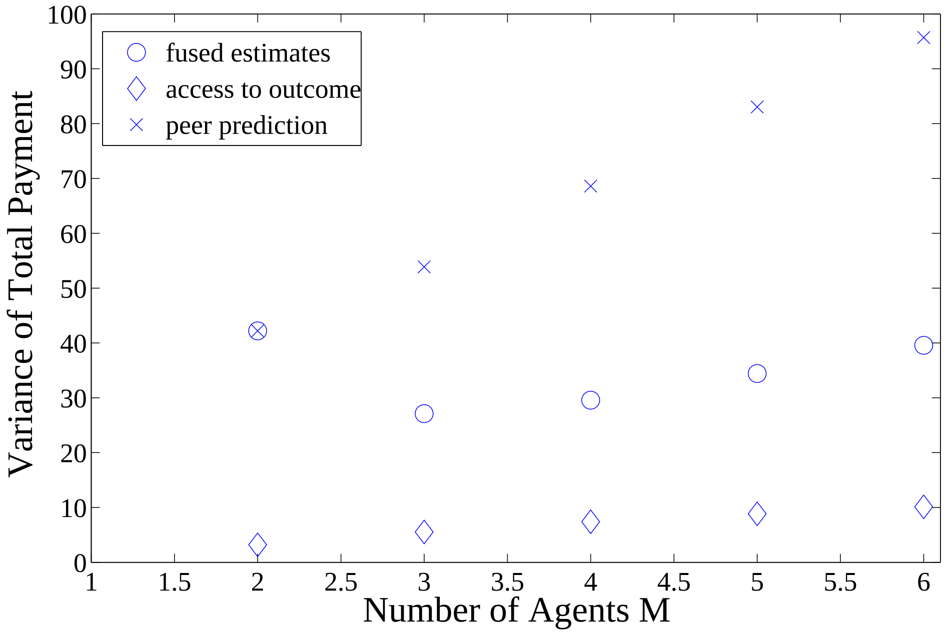

FIGURE 5.2: Centre's total payment's variance.

Considering Figures 5.1 and 5.2, each one of the three mechanisms is indicated with a different marker, and we also indicate the case where the centre has full information of the agents' costs which represents the lower bound for all mechanisms' payments (since the centre can optimally achieve its required precision at minimum cost). Specifically, we note that for every value of M, the sum of the expected payment for every mechanism is the same. This means that the centre expects to derive no additional penalties as a result of its lack of knowledge of the actual outcome. Effectively this result shows that the uncertainty that has been introduced in our setting due to the lack of knowledge of the actual outcome has no impact on the expected payments. The explanation of this suspicious result lies within the properties of the mechanism itself. More precisely, we have shown that our mechanism is incentive compatible and that the payment a pre-selected agent expects to derive is equal to the cost used for the scaling of the modified strictly proper scoring rules. We provide the intuition behind this result in the following theorem, where we calculate the total expected payment for the three evaluated mechanisms for the general case of any convex cost function.

Theorem 5.3. The sum of the expected agents' payments is independent of the knowledge of the actual outcome for the three mechanism that elicit probabilistic estimates from multiple suppliers with limited precisions and unknown costs.

Proof. In the mechanism in which the centre has access to the actual outcome, the payment agent j expects to derive, after the centre observes the actual outcome, is the following:

$$\overline{P}_j(\theta_j,\widehat{\theta}_j^c) = \frac{c_s'(\widehat{\theta}_j^c)}{\overline{S}_j'(\widehat{\theta}_j^c)}\left(\overline{S}_j(\theta_j,\widehat{\theta}_j^c) - \overline{S}_j(\widehat{\theta}_j^c)\right) + c_s(\widehat{\theta}_j^c) \tag{5.12}$$

In this context, given that agents produce estimates with precisions equal to their reported maximum precisions (Theorem 4.1), $\theta_j = \widehat{\theta}_j^c$. Thus, agent j's expected payment is:

$$P_j(\theta_j) = c_s(\theta_j) \tag{5.13}$$

where $c_s(\theta_j)$ is the scaling cost used for in the calculation of agent j's payment.

Now, in the mechanism in which the centre has no access to the actual outcome, the payment the selected agent j expects to derive is given by:

$$\overline{P}_j(\theta_j,\widehat{\theta}_j^c) = \frac{c_s'(\widehat{\theta}_j^c)}{\overline{S}_{f,j}'(\widehat{\theta}_j^c,\overline{\theta}_{-j})}\left(\overline{S}_{f,j}(\theta_j,\widehat{\theta}_j^c,\overline{\theta}_{-j}) - \overline{S}_{f,j}(\widehat{\theta}_j^c,\overline{\theta}_{-j})\right) + c_s(\widehat{\theta}_j^c) \tag{5.14}$$

Thus, the payment agent j expects to derive, given that $\theta_j = \widehat{\theta}_j^c$ is:

$$P_{f,j}(\theta_j) = c_s(\theta_j) \tag{5.15}$$

Finally, in the peer prediction mechanism the centre scores agent j against every other agent in pairs, and then calculates its payment after averaging over the $M-1$ payments that correspond to each one of the selected agents. Miller et al. also modified the existing scoring rules in order to maintain incentive compatibility. Hence, the payment agent j expects to derive in the peer prediction mechanism is the following:

$$\overline{P}_{p,j}(\theta_j,\widehat{\theta}_j^c) = \frac{1}{M-1}\sum_{i\in M_i}\frac{c_s'(\widehat{\theta}_j^c)}{\overline{S}_{p,j}'(\widehat{\theta}_j^c,\overline{\theta}_i)}\left(\overline{S}_{p,j}(\theta_j,\widehat{\theta}_j^c,\overline{\theta}_i) - \overline{S}_{p,j}(\widehat{\theta}_j^c,\overline{\theta}_i)\right) + c_s(\widehat{\theta}_j^c) \tag{5.16}$$

where $\overline{S}_{p,j}$ is the expected modified scoring rule $\overline{S}(\overline{x}_i;\widehat{x}_j,\widehat{\theta}_j+\overline{\theta}_i)$ which scores agent j only against the estimate of another (agent i). Miller, Resnick and Zeckhauser (2007) have shown that this scoring rule is incentive compatible, therefore $\overline{S}_{p,j}(\theta_j,\widehat{\theta}^c_j,\overline{\theta}_i)=\overline{S}_{p,j}(\widehat{\theta}^c_j,\overline{\theta}_i)$ when $\theta_j=\widehat{\theta}^c_j$. Hence, the payment agent j expects to derive is:

$$P_{p,j}(\theta_j)=\frac{1}{M-1}\sum_{i\in M_{-j}} c_s(\theta_j)=c_s(\theta_j) \tag{5.17}$$

Since the $\overline{P}_j(\theta_j)=\overline{P}_{f,j}(\theta_j)=\overline{P}_{p,j}(\theta_j)=c_s(\theta_j)$, all three payments are equal as required. □

Finally, regarding Figure 5.2, we show that the variance of the payments is lower when the centre has access to the actual outcome, therefore it becomes apparent that there is an effect in the payment due to the introduced uncertainty. Moreover, for $M=2$ the variance of the payments is the same in both mechanisms since our mechanism becomes identical with the peer prediction mechanism, given that when the centre calculates the payment to the selected agent, it has access to only one agent's estimate. However, for $M>2$ the variance of the payments the centre issues in the peer prediction mechanism is much greater than the variance of payments in our mechanism. This comes as a result of the peer prediction mechanism's increased sensitivity to the agents' reported maximum precisions, as opposed to our mechanism, in which an agent's reported maximum precision is fused with the others and therefore its impact is dampened. In more detail, after recording the payments after every iteration of the peer prediction mechanism, we realised that the increased variance of the agents' reported estimates, results in an increased variance in the total payments. However, in our mechanism the estimates are fused and therefore it is less likely for significant variations to appear at the end of each iteration.

5.5 Summary

In this chapter, we provided a non-trivial extension to our previous mechanism (Chapter 4) by eliminating the assumption that the centre has access to the realised outcome of the estimated event when calculating the payments to the pre-selected agents. As a result of this, now the centre uses the fused reported estimates of the pre-selected agents when calculating their payments. Therefore, in addition to all other requirements addressed also in previous mechanisms (Chapters 3 and 4), we have addressed the final research requirement regarding the information buyer not being willing or able to evaluate the received observations in dynamic environments of uncertainty. In so doing, this mechanism contributes to the state of the art by being the first mechanism that can elicit agents' unknown costs and incentivise them to commit costly resources to the generation of these estimates which, in turn, they report truthfully. This is the first mechanism with these properties that can cope with agents facing restrictions in the precisions of the estimates they can provide, and with the centre not having external means to evaluate their reports.

Given this background, the mechanism introduced in this chapter can be applied in citizen sensor networks where the information buyers either do not want to invest the necessary extra resources needed in order verify the accuracy of the providers' reports, or do not have the means to do so. This contribution becomes more significant for citizen sensor networks operating in dynamic environments, where there is always a difference between the observed and the actual state of the world used for the evaluation of the observation, due to constantly changing conditions. For example, in a noise monitoring scenario, where members of the public measure the noise of their surroundings and transmit it to a central agency responsible for planning and designing noise level maps of a city, the sound levels constantly change as the environment constantly evolves. Therefore, it is impossible for the information buyer, to verify the providers' observations, since even if it had sophisticated equipment to do so, the observed state of the world would change immediately. Likewise, in a traffic monitoring and planning service, a driver subscribed to this service (acting as the information buyer) is not able to verify the traffic reports she receives, and even if at some point she is at the location of the reported traffic observation it is very likely that the conditions will have changed, specifically in a big city where unpredicted events such as accidents, works and road blockages occur often.

To sum up, with this mechanism we have satisfied all the requirements set in Section 1.1 and therefore in the following chapter we conclude and discuss future work.

Chapter 6

Conclusions and Future Work

In this chapter we present an overarching view on the contributions of this thesis towards the stated aim of designing robust and effective protocols to govern interactions among information providers and buyers in next generation citizen sensor networks. To begin, in Section 6.1, we initially summarise the research carried within each specific topic addressed in this thesis. In more detail, we analyse the challenges that occur in current citizen sensor networks and those that are likely to occur as citizen sensor networks evolve. In this context, we focus on the issues associated with selfish behaviour that are bound to appear in the future, as citizen sensor networks find more commercial applications and increase in size. After briefly summarising key parts of the related literature, we discuss our contributions to the state of the art and how they satisfied the research requirements we set in Section 1.1. Then in Section 6.2 we outline some general lines of future work that follow from this work which can be expanded in independent lines of research.

6.1 Conclusions

The goal of this thesis was to address the selfish behaviour that we expect to appear as citizen sensor networks find more commercial applications. We argued that selfish behaviour will emerge as the interests of those acting as information buyers, intending to purchase observations or probabilistic estimates of an event, contradict with those generating the required observations who in turn, act as information providers. In more detail, such selfish behaviour can manifest itself in two main forms: (i) the allocation of insufficient resources on behalf of the citizens in the generation of the observations; and (ii) the deliberate misreporting of the citizens' observations. Now, current applications of citizen sensor networks do not focus on selfish behaviour as they mostly consist of voluntarily projects driven by the enthusiasm of those participating or restricted to specialists with no intention of exhibiting selfish behaviour. However, such systems still have to face issues such as a lack of knowledge of the costs involved in the generation of the observations on behalf of the information buyers and the fact that providers generate observations of possibly low quality due to limited resources available. Finally, an additional challenge is the information buyers' lack of access to knowledge of the actual state of the world after receiving the observations due to their unwillingness or inability to invest in the extra resources that would assist them in evaluating the observations they receive. This issue is magnified, when citizen sensor networks are applied in dynamic environments, where even if the information buyers are willing to make such an investment, the reported observation cannot be verified based on the current state of the world, as it is very likely that the conditions will be different and the reported and observed states of the world will not match.

We argued that the effect of the above challenges, common in current citizen sensor networks, will be further emphasised when they are considered in conjunction with selfish behaviour. As citizen sensor networks gain in popularity, increase in size and find even more diverse applications, emerging selfish behaviour may well degrade the efficiency and overall performance of the citizen sensor networks unless it is addressed. Subsequently, we posited the need for an approach that takes into consideration all these challenges appearing or likely to appear in future citizen sensor networks. Hence, in Section 1.1 we identified four research challenges which should be addressed in order to achieve robust and efficient citizen sensor networks that are immune to the effects of selfish behaviour. Specifically, we consider: (i) how to elicit the unknown costs of the information providers; (ii) how to address both aspects of selfish behaviour given that the information providers are expected to act on their best interest; (iii) how to deal with providers having restrictions in their available resources and hence providing observations of limited precision; and (iv) how to deal with information buyers that cannot evaluate the providers' observations based on knowledge of the actual state of the world.

With this problem in mind, we reviewed two lines of research regarding applications of citizen sensor networks, *citizen journalism* and *citizen science* and brought up the lack of consideration for the emergence of selfish behaviour and its effects. Furthermore, we highlighted the link between the appearance of selfish behaviour in citizen sensor networks and the lack of trust among information buyers and providers. Therefore, we reviewed the relevant trust literature against our research's requirements which were explicitly specified in chapter 1 (Section 1.1). Specifically, after considering learning and reputation trust models, as well as models that combine elements of both, it became apparent that they are not suitable for dynamic environments due to their complexity, and that they do not address effectively all aspects of selfish behaviour. Furthermore, most of these models rely on assumptions of full knowledge of the costs involved in the observations and the actual outcome of the observed event.

Therefore, we turned to applications of mechanism design, in which a set of rules can motivate agents to be truthful, and hence, trust can be an intrinsic property of the system. However, many of the existing mechanisms, although they had carefully selected incentive compatible payments in order to motivate agents to be honest with their reports, contradicted several of our key requirements. In particular, some of the reviewed incentive compatible auction-based mechanisms assume that all agents hold independent information. In assuming that the information is perfect and not imprecise, they contradict our requirement regarding uncertainty. Furthermore, the few mechanisms that do address the issue of interdependency, assume that agents will attempt to complete the assigned task (i.e. an observation or a probabilistic estimate of an event) without taking into consideration its cost. Moreover, the assumption that agents will ignore costs, contradicts our assumption of selfish behaviour in citizen sensor networks. Therefore we had to consider other means of addressing both issues of selfish behaviour in such networks, but still within the general mechanism design context, as it can fully satisfy the requirements of selfish behaviour without the complexity of the approaches in trust in multi-agent systems already reviewed.

In more detail, we turned to the notion of scoring rules, which to date have primarily been used as a tool for evaluating probabilistic forecasting, and then we reviewed their applications in mechanism design. Specifically, the use of strictly proper scoring rules based payments in mechanisms, motivates agents to truthfully report their private observations or probabilistic estimates. Furthermore, an appropriate scaling of the payment can motivate participating agents to invest costly resources for the generation of that observation. Now, while the reviewed applications of scoring rules were significant contributions in the specific cases they considered, they assumed that the costs involved in the generation of the observations were common knowledge available to the information buyer.

We addressed this issue in Chapter 3, where we introduced a two-stage mechanism, based on strictly proper scoring rules. That mechanism (Mechanism 1) motivated a single agent (acting as the information provider) to make a costly probabilistic estimate at a required precision and report it truthfully to the centre (effectively the information buyer). In this line of work, we extended the state of the art by being the first to deny the centre any knowledge regarding the agents' costs and hence satisfied the relevant research requirement (Requirement 1). Moreover, by motivating the selected agent to invest sufficient resources in the generation of the estimate and report it truthfully, we addressed effectively both issues of selfish behaviour and hence satisfied the requirement regarding selfish behaviour (Requirement 2). However, we made two key assumptions. First, that the agent assigned with the task of providing the estimate will be able to do so at any required precision by the centre. Second, that the centre will be able to observe the state of the world upon receiving the estimate. Although, both these assumptions hold in some applications of citizen sensor networks (i.e. networks with information providers with enough resources to produce observations at any required precision, or information buyers having the infrastructure to evaluate the providers' reported observations), we consider them unrealistic because of the amount of the required resources. Therefore, in the chapters that followed, we relaxed both these assumptions.

The first assumption was addressed in Chapter 4, where we contributed to the state of the art by introducing a mechanism (Mechanism 2) that did not rely on a single agent to provide an estimate at any precision. This satisfied the research requirement regarding the combination of multiple observations of possibly low precision (Requirement 3). In more detail, in this mechanism, due to agents' limited capabilities, the centre had to acquire more than one estimate and fuse them into one more accurate. Therefore, the centre in the first stage had to pre-select a number of agents, and then in the second stage ask them to generate an estimate at any precision they could until it achieved its required precision. Mechanism 2 was a non-trivial extension of the previous mechanism, as it defined a class of mechanisms, depending on multiple options to perform the pre-selection in the first stage. Also, we described the concept of fusing estimates of possible low precision in order to acquire one more precise.

Finally, in Chapter 5, we relaxed the assumption regarding centre's knowledge of the realised outcome by denying the centre access to such knowledge when calculating the payments to the agents reporting their estimates. In so doing, we addressed the final research requirement (Requirement 4). Specifically, we contributed to the state of the art by initially modifying the strictly proper scoring rules so they could incentivise agents to truthfully report their estimates when they were expecting to be scored against the fused reports of all the other agents. Then, based on these modified scoring rules, we introduced a novel two-stage mechanism, where the centre for the first time elicited multiple agents' private cost functions and estimates, in a setting where, first, it had no access to knowledge regarding both the agents' costs and the outcome of the estimated event and, second, the agents faced restrictions in the precision of the estimates they provided.

For all three mechanisms we provided both theoretical and empirical results. Specifically, we proved that for all mechanisms truthful reporting of costs in the first stage is a weakly dominant strategy. However, the same does not apply for the reported estimates and their precision in the second stage, as for Mechanisms 1 and 2 truthful reporting is a weakly dominant strategy, while for Mechanism 3 we showed that it is a Nash equilibrium. Now, regarding the empirical analysis, we were the first to thoroughly compare the quadratic, spherical and logarithmic scoring rules with a parametric family of strictly proper scoring rules (also analytically), and identify a sub-optimal value for the parameter k. Furthermore, by evaluating Mechanism 2, we compared several approaches to perform the pre-selection, and identified one that minimises the centre's expected total payment. Indeed, we showed that forming one group of pre-selected agents in the first stage, and effectively performing the pre-selection in one iteration, instead of pre-selecting the same number of agents in multiple iterations, the centre minimises its payment in the second stage. Based on this result, we performed the pre-selection for Mechanism 3. With the empirical evaluation of Mechanism 3, our contribution to the state of the art was completed, by showing that the increase of uncertainty in our model, as introduced by the centre's lack of knowledge of the realised outcome, did not have an impact on the centre's expected total payment, but it was restricted only to the total payment's variance.

Speaking more general, the work presented in this thesis contributes towards the design of protocols that will guide the interactions among information buyers and providers in present and future citizen sensor networks, as they evolve and selfish behaviour becomes a more prevalent factor. We provided three mechanisms, each of which addresses specific issues that reflect on current and expected challenges in citizen sensor networks. In this context, the three mechanisms cannot be considered as just an extension of each other. Hence, it should be noted that although Mechanism 3 is a generalisation of Mechanisms 1 and 2, both of them are contributions in their own right and would be used in preference to the more general one (Mechanism 3) in specific settings. Indeed, Mechanism 1 is the optimal choice in a setting where all agents can provide estimates of precisions higher than the centre's requirement (with out necessarily being infinite), and therefore the centre does not have to acquire multiple estimates. In that case, the centre guarantees that it will receive an estimate of precision at least equal to its required one, while in Mechanism 2 this is not certain, as the sum of the pre-selected agents' maximum precisions may be less than the centre's requirement. Furthermore, in pre-selecting only one agent, instead of multiple agents for an observation of equal precision, it minimises its expected payment (as seen in Section 4.4). In addition to that, we have seen that there may be some applications of citizen sensor networks, where the information buyers have means of verifying the reported observations based on knowledge of the observed event. In these cases, Mechanism 2 is more favourable since truthful reporting is a stronger solution concept (i.e. dominant strategy) when compared to the final mechanism, and the payments are more robust since their variance is minimised.

The above described techniques are not restricted to applications related to citizen sensor networks. They can be applied in traditional sensor networks responsible for monitoring environmental parameters. For example, there are many applications of sensor networks operating in extreme environments, such as flooded areas (Zhou and De Roure, 2006), glaciers (Martinez et al., 2004) or volcanoes (Werner-Allen et al., 2005), where it is impossible to ascertain the 'ground truth' through external means. In this cases, sensors make imprecise and noisy measurements with the possibility of emergence of selfish behaviour which mainly manifests through the investment of minimum resources in the generation of the measurements. This behaviour is bound to appear since sensors face multiple restrictions in bandwidth and power as they have to operate unattended for a large period of time and it is crucial to maintain a balance between making and communicating a measurement so they can continue functioning.

In addition to that, the fact that in our thesis we model information as a traded commodity, which involves a cost in generating it and a payment in acquiring it, makes it possible to link our work to literature related to e-commerce applications. In more detail, the use of on-line markets for trading items and services of any type has created the need for the evaluation of these items or services. Now, we already documented this when reviewing reputation systems (Section 2.2.2), and although, in this thesis we do not focus on such applications (i.e. rating and ranking mechanisms in on-line sites such as ebay.com and amazon.com), we identified the flaws within existing solutions and effectively showed how they can be addressed through the mechanisms we introduced. Furthermore, as on-line markets expand, multiple websites that only rate and rank services without necessarily selling them appear (examples include music rating sites such as rateyourmusic.com and cinema pictures rating sites such as imdb.com). A key characteristic of such systems is the emergence of selfish behaviour which becomes a dominating factor due to the scale of such systems, the anonymity of those involved, their fierce competition and possible financial benefits through concealed advertisement. Therefore, techniques which evaluate those evaluating services are necessary.

To sum up, although in some cases we can apply directly the results of this thesis in the wide application domains outlined above, it cannot be done without extending it so it can address extra challenges which may also appear to some extent to the citizen sensor networks of the future. Therefore, in the following section, we outline these research challenges that will appear in extending this work to a wider set of applications and the corresponding directions of future work which they represent.

6.2 Future Work

In the future there are three main directions the research covered in this thesis could take. In particular, one direction would be to address the issue of an untrustworthy centre that uses a cost function lower than the second lowest for Mechanism 1, or the $(m+1)^{th}$ and $(M+1)^{th}$ lowest costs for Mechanisms 2 and 3 respectively, when calculating payments in the second stage. In terms of citizen sensor networks, this means, that selfish behaviour may not be restricted to information providers. In some cases, information buyers may be inclined to select an agent without taking its cost into consideration. Specially, when dealing with members of the public, this issue becomes more important as an information buyer may exhibit favouritism over particular providers. Additionally, an information buyer may choose to use a lower cost when calculating the payments, than the one defined by the reverse price auction in the first stage of the mechanism. A second direction would be to address the issue of collusion among information providers that will appear in Mechanism 3, where the providers are expected to be paid for their service based on other providers' observations. In this case, according to Jurca and Faltings (2007), truth telling is not the only Nash equilibrium, since the payments to the agents are calculated based on the other agents' reports. Finally, a third direction would be to modify the existing model and mechanisms so the exchanged information is no longer modeled by continuous distributions, but by discrete distributions (such as the recommender systems similar to the ranking web sites introduced in Section 6.1). This extension will make it possible to find wider applications for our mechanisms and address similar problems in other domains (i.e. ranking systems, environmental monitoring).

We will now discuss each one of these directions in more detail:

Information buyers' dishonesty: In the future we intend to address one of the key problems of second price auctions and consequently of the reverse second and $(m+1)^{th}$ price auctions we use in the first stage of the mechanisms in this thesis. As already seen in Section 2.4.2, second price auctions, as opposed to English and Dutch auctions, cannot prevent the centre (acting as the auctioneer) from lying. The same problem exists in our mechanisms, as the cost functions are private information known only to each of the auction's participants. Specifically in citizen sensor networks, the information buyer can use lower cost functions than the second and $(m+1)^{th}$ costs when calculating payments in the second stage of our mechanism, exactly like the auctioneer in the standard second price auctions, who can ask for a higher price than the second bid. In order to address this issue, in Section 2.4.2 we discussed two approaches: the public blackboard by Hsu and Soo (2002) and the heavy encryption protocols by Suzuki and Yokoo (2004). However, we dismissed them both as they were either not applicable in a citizen sensor network, or introduced additional complexity through multiple layers of calculations. Still, the combination of these techniques could be a starting point for future research, which will relax some of these heavy restrictions. For example, a public blackboard system with simpler encrypted methods which will guarantee the providers' anonymity.

Information providers' collusion: Jurca and Faltings (2007) have shown that when payments are calculated based on other agents' reported observations and not the observed actual outcome, agents are expected to collude. This is the case in Mechanism 3 where the centre relies on the agents' reported estimates as it has no means of observing the state of the world when issuing the payments. Now, by depriving communication among information providers we address this issue, since the providers will not be able to know each others' observations in advance and therefore it will be impossible for them to collude. However, as citizen sensor networks evolve, and the available technology advances, this may not be the case. Citizen journalism is particularly vulnerable to this issue, as members of the public may agree to misreport information deliberately in order to achieve their goals. For example, the residents of a small town may agree to fabricate news and use social networking websites to publicise them in order to attract attention to their town and benefit from the temporary attention. Now, in addressing this issue, it is important not to introduce additional complexity to the network by employing specialists in order to supervise the information providers and evaluate their reports, as opposed to what Jurca and Faltings suggest. Such an approach would contradict the main purpose of such a mechanism, which is to be applied in networks, where the information buyers do not have external means to evaluate the providers observations, and therefore rely solely on them. A solution for this issue, would be to try to use the possibility of collusion among providers for the buyer's benefit by introducing the notion of cooperation among providers. For example, the information providers could collectively produce an observation through sharing resources. However, in so doing we also introduce the possibility of defection and we have to change the model and mechanism drastically so it can cope with both these options. Initially, we have to modify the fusion process to include a parameter that would measure the percentage of cooperation and then calculate the providers' utility functions in order to show whether the new mechanism will remain incentive compatible.

Discrete Distributions: Finally, for our future work we would like to extend our model so it can also include cases where the cost functions and estimates are modelled by discrete distributions. In so doing, we will be able to apply the techniques introduced in this thesis in a wider set of problems and not restrict to citizen sensor networks. The information observed and reported by contributors in other information systems, besides citizen sensor networks, may be represented by discrete probability distributions (e.g. ratings and rankings of services or items in recommender systems (Adomavicius and Tuzhilin, 2005)). Therefore, it is important to develop the processes to guide us towards effectively addressing issues such as selfish behaviour, unknown costs, observations of low quality and lack of access to the ground truth. However, in this case, several of the assumptions we make regarding the order of the cost functions and their derivatives, will not hold. Specifically, none of the conditions of Lemmas 3.2 and 3.1 in Chapter 3 will hold, and consequently proving the economic properties of a new mechanism based on discrete costs will be a challenging task. The complexity of this task lies within the difficulties involved in linking the cost of producing an observation with its precision and the resources invested in its generation. So an initial step, in dealing with discrete costs, is to design a model that will link the cost to the effort of making an estimate or an observation. Specifically, this can be done through a scheme where a contributor makes multiple attempts to generate an observation (i.e. rate a service) prior to reporting it. Given that a certain fixed cost can be assigned in making a single observation, multiple efforts will show more resources being invested in the final task of reporting the overall observation.

Bibliography

Aalberts, R. J., Poon, P. and Thistle, P. (2007), 'Trespass, nuisance, and spam: 11th century common law meets the internet', *Communications of the ACM* **50**(12), 40–45.

Adomavicius, G. and Tuzhilin, A. (2005), 'Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions', *IEEE Transactions on Knowledge and Data Engineering* **17**(6), 734–749.

Arnold, B. C., Balakrishnan, N. and Nagaraja, H. N. (2008), *A First Course in Order Statistics*, SIAM.

Berners-Lee, T. (2008), 'Interview with silicon republic'.
URL: *www.siliconrepublic.com/news/news.nv?storyid=single10624*

Bowman, S. and Willis, C. (2003), We media: How audiences are shaping the future of news and information, Technical report, The Media Center at the American Press Institute.

Cavallo, R. (2006), Optimal decision-making with minimal waste: strategyproof redistribution of VCG payments, *in* 'Proceedings of the 5th International Joint Conference on Autonomous agents and multiagent systems (AAMAS 2006)', Hakodate, Japan, pp. 882–889.

Chen, Y., Dimitrov, S., Sami, R., Reeves, D. M., Pennock, D. M., Hanson, R. D., Fortnow, L. and Gonen, R. (2009), 'Gaming prediction markets: Equilibrium strategies with a market maker', *Algorithmica* .

Clarke, E. (1971), 'Multipart pricing of public goods', *Public Choice* **11**(1), 17–33.

Czumaj, A. and Ronen, A. (2004), On the expected payment of mechanisms for task allocation, *in* 'Proceedings of the 23rd annual ACM symposium on Principles of Distributed Computing (PODC 2004)', St. John's, Newfoundland, Canada, pp. 98–106.

Dasgupta, P. (1988), *Trust: Making and Breaking Cooperative Relations*, Blackwell Publishers, chapter Trust as a Commodity.

Dash, R. K., Jennings, N. R. and Parkes, D. C. (2003), 'Computational mechanism design:a call to arms', *IEEE Intelligent Systems* pp. 40–47.

Dash, R. K., Ramchurn, S. D. and Jennings, N. R. (2004), Trust-based mechanism design, *in* 'Proceedings of the 3rd International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2004)', Vol. 2, New York, USA, pp. 748–755.

Dash, R., Rogers, A. and Jennings, N. (2004), A mechanism for multiple goods and interdependent valuations, *in* 'Proceedings of 6th International Workshop on Agent-Mediated E-Commerce (AMEC 2004)', New York, USA, pp. 197–210.

DeGroot, M. H. and Schervish, M. J. (2002), *Probability and Statistics*, Addison Wesley.

Fatima, S. S. (2006), Sequential versus simultaneous auctions: A case study, *in* 'Proceedings of the 8th International Conference on Electronic Commerce (ICEC 2006)', New Brunswick, Canada, pp. 82–91.

Fatima, S. S., Wooldridge, M. and R., J. N. (2005), Sequential auctions for objects with common and private values, *in* 'Proceedings of the 4th International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2005)', Utrecht, The Netherlands, pp. 635–642.

Friedman, D. (1983), 'Effective scoring rules for probabilistic forecasts', *Management Science* **29**(4), 447–454.

Gneiting, T. and Raftery, A. E. (2004), Strictly proper scoring rules, prediction and estimation, Technical Report 463, University of Washington.

Goodchild, F. M., Fu, P. and Rich, P. (2007), 'Sharing geographic information: An assessment of the geospatial one-stop', *Annals of the Association of American Geographers* **97**(2), 250–266.

Goodchild, M. F. (2007), 'Citizens as sensors: the world of volunteered geography', *GeoJournal* **69**(4), 211–221.

Grossman, S. J. and Hart, O. D. (1983), 'An analysis of the principal-agent problem', *Econometrica* **51**(1), 7–45.

Groves, T. (1973), 'Incentives in teams', *Econometrica* **41**(4), 617–631.

Guo, M. and Penock, D. M. (2009), Combinatorial prediction markets for event hierarchies, *in* 'Proceedings of 8th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2009)', Budapest, Hungary, pp. 201–208.

Hanson, R. (2003), 'Combinatorial information market design', *Information Systems Frontiers* **5**(1), 107–119.

Hanson, R. (2007), 'Logarithmic market scoring rules for modular combinatorial information aggregation', *The Journal of Prediction Markets* **1**(1), 3–15.

Hendrickson, A. D. and Buehler, R. J. (1971), 'Proper scores for probability forecasters', *The Annals of Mathematical Statistics* **42**(6), 1916–1921.

Hsu, M. and Soo, V. (2002), A secure multi-agent vickrey auction scheme, *in* 'Proceedings of the 5th International Workshop on Deception, Fraud and Trust in Agent Societies', Bologna, Italy, pp. 86–91.

Jehiel, P. and Moldovanu, B. (2001), 'Efficient design with interdependent valuations', *Econometrica* **69**(5), 1237–1259.

Jones, K., Janicke, H. and Cau, A. (2009), A property based framework for trust and reputation in mobile computing, *in* 'Proceedings of the 23rd International Conference on Advanced Information Networking and Applications (AINA 2009)', pp. 1031–1036.

Josang, A. and Ismail, R. (2002), The beta reputation system, *in* 'Proceedings of the 15th Bled Conference on Electronic Commerce (BledECC 2002)', Bled, Slovenia.

Jurca, R. and Faltings, B. (2005*a*), Reputation-based pricing of p2p services, *in* 'Proceeding of the 2005 ACM SIGCOMM Workshop on Economics of Peer-to-peer Systems (P2PECON 2005)', Philadelphia, Pennsylvania, USA, pp. 144–149.

Jurca, R. and Faltings, B. (2005*b*), Reputation-based service level agreements for web services, *in* 'In Service Oriented Computing (ICSOC 2005)', Vol. 3826 of *Lecture Notes in Computer Science*, Springer Berlin / Heidelberg, pp. 396–409.

Jurca, R. and Faltings, B. (2006), Using chi-scores to reward honest feedback from repeated interactions, *in* 'Proceedings of the 5th international joint conference on Autonomous agents and multiagent systems (AAMAS 2006)', Hakodate, Japan, pp. 1233–1240.

Jurca, R. and Faltings, B. (2007), Collusion resistant, incentive compatible feedback payments, *in* 'Proceedings of the 8th ACM Conference on Electronic Commerce (ACM-EC 2007)', San Diego, California, USA, pp. 200–209.

Klein, M., Moreno, G. A., Parkes, D. C., Plakosh, D., Seuken, S. and Wallnau, K. (2008), Handling interdependent values in an auction mechanism for enhanced bandwidth allocation in tactical data networks, *in* 'Proceedings of the 3rd international Workshop on Economics of Networked Systems (NetEcon 2008)', Seattle, Washington, USA, pp. 73–78.

Krishna, V. (2002), *Auction Theory*, Academic Press.

Lesser, V., Ortiz, C. L. J. and Tambe, M., eds (2003), *Distributed Sensor Networks*, Kluwer Academic Publishers.

Maguire, J. D. and Longley, A. P. (2005), 'The emergence of geoportals their role in spatial data infrastructures', *Computers, Environment and Urban Systems* **29**(1), 3–14.

Martinez, K., Ong, R., Hart, J. K. and Stefanov, J. (2004), Glacsweb: A sensor web for glaciers., *in* 'Proceedings of the 1st European Workshop on Wireless Sensor Networks.', Springer.

Mas-Colell, A., Whinston, M. D. and Green, J. R. (1995), *Microeconomic Theory*, Oxford University Press.

Matheson, J. E. and Winkler, R. L. (1976), 'Scoring rules for continuous probability distributions', *Management Science* **22**(10), 1087–1096.

McAffe, R. P. and McMillan, J. (1987), 'Auctions and bidding', *Journal of Economic Literature* **XXV**, 699–738.

Mezzetti, C. (2004), 'Mechanism design with interdependent valuations: Efficiency', *Econometrica* **72**(5), 1617–1626.

Mezzetti, C. (2007), 'Mechanism design with interdependent valuations: Surplus extraction', *Economic Theory* **31**(3), 473–488.

Miller, N. H., Pratt, J. W., Zeckhauser, R. J. and Johnson, S. (2007), 'Mechanism design with multidimensional, continuous types and interdependent valuations', *Journal of Economic Theory* **136**, 476–496.

Miller, N., Resnick, P. and Zeckhauser, R. (2007), 'Eliciting honest feedback: The peer prediction method', *Management Science* **51**(9), 1359–1373.

Mui, L., Mohtashemi, L. and Halberstadt, A. (2002), A computational model of trust and reputation, *in* 'Proceedings of the 35th Hawaii International Conference on System Sciences (HICSS 2002)', Vol. 7, Hawai, USA, pp. 188–196.

Mukherjee, R., Banerjee, B. and Sen, S. (2001), Learning mutual trust, *in* 'Trust in Cyber Societies', Springer-Verlag, Berlin, pp. 145–158.

Petcu, A., Faltings, B. and Parkes, D. C. (2006), MDPOP: Faithful distributed implementation of efficient social choice problems, *in* 'Proceedings of the 5th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2006)', Hakodate, Japan, pp. 1397–1404.

Porter, R., Ronen, A., Shoham, Y. and Tennenholtz, M. (2008), 'Fault tolerant mechanism design', *Artificial Intelligence* **172**(15), 1783–1799.

Ramchurn, S. D., Huynh, D. and Jennings, N. R. (2004), 'Trust in multi-agent systems', *The Knowledge Engineering Review* **19**, 1–25.

Ramchurn, S. D., Mezzetti, C., Giovannucci, A., Rodriguez, J. A., Dash, R. K. and Jennings, N. R. (2009), 'Trust-based mechanisms for robust and efficient task allocation in the presence of execution uncertainty', *Journal of Artificial Intelligence Research* **35**, 119–159.

Rettinger, A., Nickles, M. and Tresp, V. (2008), A statistical relational model for trust learning, *in* 'Proceedings of the 7th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2008)', Estoril, Portugal, pp. 763–770.

Rogerson, W. P. (1985), 'The first-order approach to principal-agent problems', *Econometrica* **53**(6), 1357–1367.

Savage, L. J. (1977), 'Elicitation of personal probabilities and expectations', *Journal of the American Statistical Association* **66**(336), 783–801.

Selten, R. (1998), 'Axiomatic characterization of the quadratic scoring rule', *Experimental Economics* **1**(1), 43–61.

Sen, S. and Sajja, N. (2002), The evolution and stability of cooperative traits, *in* 'Proceedings of the 1st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2002)', Bologna, Italy, pp. 1114–1120.

Sheth, A. (2009), 'Citizen sensing, social signals, and enriching human experience', *IEEE Internet Computing* **13**(4), 87–92.

Sheth, A., Henson, C. and Sahoo, S. S. (2008), 'Semantic sensor web', *IEEE Internet Computing* **12**(4), 78–83.

Singh, M. P., Yu, B. and Venkatraman, M. (2001), 'Community-based service location', *Communications of the ACM* **44**(4), 49–54.

Suzuki, K. and Yokoo, M. (2004), Secure generalized vickrey auction using homomorphic encryption, *in* 'Financial Cryptography', Vol. 2742 of *Lecture Notes in Computer Science*, Springer Berlin / Heidelberg, pp. 239–249.

Teacy, W. T. L. (2006), Agent-Based Trust and Reputation in the Context of Inaccurate Information Sources, PhD thesis, University of Southampton.

Teacy, W. T. L., Patel, J., Jennings, N. R. and Luck, M. (2006), 'TRAVOS: Trust and reputation in the context of inaccurate information sources', *Autonomous Agents and Multi-Agent Systems* **12**(2), 183–198.

Vickrey, W. (1961), 'Counterspeculation, auctions and competitive sealed tenders', *The journal of finance* **16**(1), 8–37.

Wang, Y. and Singh, M. P. (2007), Formal trust model for multiagent systems, *in* 'Proceedings of the 20th International Joint Conference on Artificial Intelligence (IJCAI 2007)', Hyderabad, India, pp. 1551–1556.

Werner-Allen, G., Johnson, J., Ruiz, M., Lees, J. and Welsh, M. (2005), Monitoring volcanic eruptions with a wireless sensor network., *in* 'Proceedings of the 2nd European Workshop on Wireless Sensor Networks (EWSN 2005)', Instabul, Turkey, pp. 108–120.

Whitby, A., Jsang, A. and Indulska, J. (2004), Filtering out unfair ratings in bayesian reputation systems, *in* 'Proceedings of the 7th International Workshop on Trust in Agent Societies (TRUST 2004)', New York, USA.

Wu, D. and Sun, Y. (2001), The emergence of trust in multi-agent bidding: a computational approach, *in* 'Proceedings of the 34th Hawaii International Conference on System Sciences (HICSS 2001)', Vol. 1, pp. 1041–1048.

Yu, B. and Singh, M. P. (2002), An evidential model of reputation management, *in* 'Proceedings of the 1st International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2002)', Vol. 1, Bologna, Italy, pp. 295–300.

Yuan, S.-T. and Sung, H. (2004), 'A learning-enabled integrative trust model for e-markets', *Applied Artificial Intelligence* **18**, 69–95.

Zacharia, G. and Maes, P. (2000), 'Trust management through reputation mechanisms', *Applied Artificial Intelligence* (14), 881–907.

Zhao, F. and Guibas, L. J. (2004), *Wireless Sensor Networks: An Information Processing Approach*, Morgan Kaufmann.

Zhou, J. and De Roure, D. (2006), Designing energy-aware adaptive routing for wireless sensor networks, *in* 'In Proceedings of The 6th International Conference on ITS Telecommunications', Chengdu, China, pp. 680–685.

Zohar, A. and Rosenschein, J. S. (2008), 'Mechanisms for information elicitation', *Artificial Intelligence* **172**, 1917–1939.