

An overview of incentive contracting^{*}

Sarit Kraus^{*}

Department of Mathematics and Computer Science, Bar Ilan University, Ramat Gan, 52900 Israel

Received May 1994; revised January 1995

Abstract

Agents may contract some of their tasks to other agents even when they do not share a common goal. An agent may try to contract some of the tasks that it cannot perform by itself, or that may be performed more efficiently by other agents. One self-motivated agent may convince another self-motivated agent to help it with its task, by promises of rewards, even if the agents are not assumed to be benevolent. We propose techniques that provide efficient ways for agents to make incentive contracts in varied situations: when agents have full information about the environment and each other, or when agents do not know the exact state of the world. We consider situations of repeated encounters, cases of asymmetric information, situations where the agents lack information about each other, and cases where an agent subcontracts a task to a group of agents. Situations in which there is competition among possible contractor agents or possible manager agents are also considered. In all situations we assume that the contractor can choose a level of effort when carrying out the task and we would like the contractor to carry out the task efficiently without the need of close observation by the manager.

1. Introduction

Agents acting in non-collaborative environments may benefit from contracting some of their tasks to other agents. In this paper we present techniques for efficient contracting that can be used in different cases of multi-agent environments where the agents do not have a common goal and there is no globally consistent knowledge. We consider

^{*} This paper is based upon work supported by the National Science Foundation under Grants No. IRI-9423967 and the Israeli Ministry of Science and the Arts under grant 4210, and is an extension of [47]. I would like to thank Jonathan Wilkenfeld, Barbara Grosz and an anonymous referee for their comments, and Sean Engelson and Onn Shehory for useful discussions.

^{*} E-mail: sarit@bimacs.cs.biu.ac.il. Also affiliated with the Institute for Advanced Computer Studies, University of Maryland, College Park, MD, USA.

situations where a self-motivated agent that tries to carry out its own individual plan in order to fulfill its own tasks may contract some of its own tasks to another self-motivated agent(s). An agent may benefit from contracting some of its tasks that it cannot perform by itself, or when the task may be performed more efficiently by other agents.

The central question of this paper is how one agent can convince another agent to do something for it when the agents do not share a global task and the agents are not assumed to be benevolent. Furthermore, we consider situations where the contractor agent can choose different levels of effort when carrying out the task. The manager agent would like the contractor agent to carry out the task with the level of effort that the manager prefers without the need of close observation of the manager, enabling the manager to carry out other tasks simultaneously.

There are two main ways to convince another self-motivated agent to perform a task that is not among its own tasks: by threatening to interfere with the agent carrying out its own tasks, or by promising rewards [49]. This paper concentrates on subcontracting by rewards which may be accomplished in two forms: The first approach is a bartering system, where one agent may promise to help the other with future tasks in return for current help. However, as has long been recognized in economics, bartering is not an efficient basis for cooperation, particularly in a multi-agent environment. An agent wishing to subcontract a task to another agent may not have the ability to help it in the future, or one agent that can help in fulfilling another agent's task may not need help in carrying out its own tasks. The second approach is a monetary system which is developed for the provision of rewards, and which can later be utilized for other purposes.

In this paper we present a model of automated agents where incentive contracting is beneficial. We propose to use a monetary system in a multi-agent environment that allows for side payments and rewards between the agents, and where profits may be given to the owners of the automated agents. The agents will be built to maximize expected utilities that increase with the monetary values, as will be explained below. Assuming that each agent has its own personal goals, contracting would allow the agents to fulfill their goals more efficiently as opposed to working on their own.

The issue of incentive contracting has been investigated in economics and game theory during the last two decades (e.g., [2, 31, 40, 56, 88, 91]). These works in economics and game theory consider different types of contracts for different applications. Examples of these are contracts between: a firm and an employer or employers (e.g., [6, 7, 64, 78]); a government and taxpayers (e.g., [9]); a landlord and a tenant (e.g., [2]); an insurance company and a policy holder (e.g., [34, 58, 93, 102]); a buyer and a seller (e.g., [70, 77]); a government and firms (e.g., [72]); stockholders and managements (e.g., [2]); a professional and a client [98], etc. In these situations two parties usually exist. The first party (called "the agent" in the economics literature) must choose an action or a level of effort from a number of possibilities, thereby affecting the outcome of both parties. The second party (named "the principal") has the additional function of prescribing payoff rules. Before the first party (i.e., the agent) chooses the action, the principal determines a rule (i.e., a contract) that specifies the fee to be paid to the other party as a function of the principal's observations. Despite the similarity of the above applications, they differ in several aspects, such as the amount of information that is

available to the parties, the observations that are made by the principal, and the number of agents. Several concepts and techniques are applied to the principal-agent paradigm in the relevant economics and game theory literature.

We consider varied situations of automated agent environments; situations of certainty vs. uncertainty, full information vs. partial information, symmetric information vs. asymmetric information and bilateral situations vs. situations where there are more than two automated agents in the environment. For each of these situations we fit appropriate economics mechanisms and techniques, from the game theory or the economics literature, that can be used for contracting in environments of automated agents. We adjust these results to the automated agents environment and present all of them using uniform concepts that are appropriate to automated agents, i.e., translating the different concepts used in the various economics and game theory papers into a uniform framework. In all the situations that we consider, the agent that designs the contract is provided with techniques to maximize its personal expected utilities, given the constraints of the other agent(s). Throughout the paper, we use a robotics domain and an example of software agents to demonstrate the contracting techniques introduced above.

2. Related work in DAI

Research in DAI is divided into two basic classes: *cooperative distributed problem solving* and *multi-agent systems (MA)* [8, 28]. Research in cooperative distributed problem solving (e.g., [12, 18, 59, 61, 101]) considers how the work involved in solving a particular problem can be divided among a number of modules or “nodes”. The modules in a cooperative distributed problem solving system are centrally designed to improve the following properties of the system [8]:

- *Performance*: Concurrency may increase the speed of computation and reasoning, and may allow the system to solve large problems faster.
- *Reliability and stability*: The modules may provide redundancy, cross-checking and triangulation of the results. In case of failure of one of the modules, the other modules can fulfill its tasks.
- *Modularity*: Each module can be developed separately, making it easier to develop and extend the system.

The modules include the development of cooperating mechanisms designed to find a solution to a given problem.

Research in MA (e.g., [20, 29, 48, 104, 107, 110]) is concerned with coordinating intelligent behavior among a collection of autonomous (possibly heterogeneous) intelligent (possibly pre-existing) agents. In MA, there is no global control, and no globally shared goals or success criteria. There is, however, a possibility for *real competition* among the agents.

The MA and the cooperative distributed problem solving systems are the two poles of the DAI research. Our research falls closer to the MA systems pole. We consider the problem of a self-motivated agent (the manager) that tries to make another self-motivated agent (the contractor) fulfill one of its tasks. We assume that the contractor can choose between different levels of effort when trying to fulfill the task. The main

problem that we address is how the manager should motivate the contractor to choose a level of effort that the manager prefers.

The provision of incentives is in general not essential in cooperative distributed problem solving systems. It is assumed that it is in the agents' interests to help one another. This help can take the form of sharing tasks, results, or information [19]. In task sharing, an agent, which cannot fulfill a task on its own, will attempt to pass the task, in whole or in part, to other agents, usually on a contractual basis [101]. This approach assumes that an agent not otherwise occupied will readily take on the task and do it to the best of its abilities. Similarly, results and information are shared among agents in such environments with no expectation of reciprocation [12, 59, 61]. This benevolence is based on an assumption common to many approaches to coordination: that the system's goal is to solve the problem as best it can, thereby giving the agents shared, often implicit, global goals that they are all unselfishly committed to achieving.

One of the techniques that is used in cooperative distributed problem solving for task allocation is automated contracting. In this paper we concentrate on situations where the contractor needs to choose an effort level, and the main purpose of the contracting mechanism is to convince the contractor to agree to do the sub-task while choosing the effort level preferred by the manager. In contrast, in automated contracting the contractors do not need to choose effort levels when carrying out tasks, and thus there is no need for incentive contracts. However, work on automated contracting considers other problems essential to distributed problem solving as we discuss below.

A well-known framework for automated contracting is the contract net protocol [100, 101]. In the contract net protocol a contract is an explicit agreement between an agent that generates a task (the manager) and an agent that is willing to execute the task (the contractor). The manager is responsible for monitoring the execution of a task and processing the results of its execution, whereas the contractor is responsible for the actual execution of the task. The manager of a task announces the task's existence to other agents. Available agents (potential contractors) then evaluate the task announcements made by several managers and submit bids for the tasks they are suited to perform. As we explained above, since all the agents have a common goal and are designed to help one another, there is no need to motivate an agent to bid for tasks or to do its best in executing it if its bid is chosen. The main problems addressed by [99-101] are as follows:

- *Tasks decomposition*: how to break a large task into smaller ones.
- *Sub-tasks distribution*: how to match sub-tasks with problem solvers capable of handling them.
- *Synthesis of the overall solution*: how to synthesize the individual results of sub-tasks into a single overall solution.

In addition, they consider problems such as which information a possible contractor should send to a manager when it bids for a task and how the manager should evaluate bids. In this paper we consider incentive contracting in situations where there is only one task per contractor and therefore the problems of task decomposition and synthesis of the overall solution mentioned above do not arise. In some situations we consider the problem of distribution of a given task where there are several agents in the environment that compete for the job (see Section 5.7). However, while in the contract net the agents

that bid for a task voluntarily provide the manager with correct information about their capabilities and situation, in our framework the manager needs to construct a mechanism to make the possible contractors reveal their capabilities honestly.

The contract net protocol is a very general protocol for task distribution, and several refinements of the protocol were made in the last ten years. Malone et al. [67] developed a distributed scheduling protocol (DSP) based on the contract net protocol. The most important way in which DSP differs from the original contract net protocol is by its criteria for matching between tasks and agents (i.e., the problem of sub-tasks distribution). It includes two primary dimensions: (1) contractors select managers' tasks in the order of tasks' numerical *priorities*, and (2) managers select contractors on the basis of *estimated completion times* from among the contractors that satisfy the minimum requirements to perform the job. In addition to the problems addressed by Smith and Davis, Malone et al. considered problems such as how to estimate the processing time of a task and, if people supply their own estimation, how to encourage them to report honestly, and how to assign priorities to tasks in order to achieve various global scheduling objectives. Similar to the original contract net protocol, in Malone et al.'s model there is also no need to motivate the agents to bid or to make decisions in order to maximize the global expected utility of the system, and it is assumed that workstations voluntarily put their machines into a mode where the machines responds to requests for bids from the network. Also, the workstations don't need to choose effort levels; they either carry out a task or not. The DPS was tested using simulation of workstations on a network in a wide variety of situations (e.g., different processor speeds, system loads and message delay times). The results obtained in these simulations are as follows:

- (1) Substantial performance improvement results from sharing tasks among processors in systems with more than light loads.
- (2) In many cases these benefits are still present, even when message delay times are as much as 5 to 20 percent of the average task processing time.
- (3) In many cases, the additional benefits from pooling tasks among more than eight or ten machines are small.
- (4) Large errors in estimating task processing time cause little degradation in the scheduling performance.

A modified version of the contract net protocol for competitive agents in the transportation domain is presented in [94]. It provides a formalization of the bidding and the decision awarding processes, based on marginal cost calculations based on local agent criteria. More important, an agent will submit a bid for a set of delivery tasks¹ only if the maximum price mentioned in the tasks' announcement is greater than what the deliveries will cost that agent. A simple motivation technique is presented to convince agents to make bids; the actual price of a contract is half way between the price mentioned in the task announcement and the bid price. As in other automated contracting systems the contractor either honors its commitment to carry out a task or it does not. There are no

¹ Announcing one delivery at a time is not sufficient in general. This is due to the fact that the deliveries are dependent. For example, for two disjointed delivery sets T_1 and T_2 , the marginal costs that are saved by removing both T_1 and T_2 are usually larger than the sum of marginal cost that was saved by removing each of them alone.

different levels of actions in order to perform the task (e.g., the time it takes to carry out the task, the quality of the delivery), as in the incentive contracting framework that we consider. Therefore, there is no need for monitoring (beside checking whether the deliveries were done or not) or incentives to the contractors to choose an efficient level of action. On the other hand, in [94], Sandholm deals with the following challenging problems that are not considered part of our incentive contracting model:²

- how to choose which tasks to contract out,
- how to cluster tasks into sets to bargain over as atomic bargaining,
- how to bid when multiple bids and awards should be handled simultaneously,
- how to handle a large amount of messages consisting of bids and awards from other agents and how to prevent the agents from receiving announcements at a faster pace than they can process,
- how to decide to whom to award a set of tasks.

In [94] a set of experiments is described which demonstrates that the approach presented in that paper reduces the total transportation costs among autonomous dispatch centers.

In [82] a language for specification of complex relations among agents in cooperative distributed problem solving is described. By using this language, a designer of a system can define hierarchical relationships among the agents and specify to one agent the other agents' authority on it. The "authority" parameter indicates how much emphasis the agent should give to requests that arrive from different agents. Since the agents are not self-motivated, their willingness to help another agent will depend upon the designer's instructions. Pattison et al. suggest *focused addressing* as an additional mechanism of contracting to the one presented in the contract net protocol. This would mean that in addition to broadcasting requests for bids, an agent has the option of asking for help from another agent directly if it knows that the other agent can help it in its task and if it knows the other agent's address. In this paper, we also allow both of these addressing methods.

Subcontracting in cooperative distributed problem solving also appears in the paradigm of *planning for multiple agents*, where a single intelligent agent (usually called the master) constructs a plan to be carried out by a group of agents (the slaves), and then hands out the pieces of the plan to the relevant individuals [13, 60, 90]. Werner [109] presents a formal logical model for a master-slave relationship by one-way communication. Also in the master-slave model there is no need to choose a level of effort and there is no need for incentive contracting. That is, the main problem for a master is finding the best plan and synchronizing the agent's actions, rather than convincing other agents to carry out the plan appropriately without its observation. The simple master/slaves model was extended by Ephrati and Rosenschein [21] to allow the "slaves" more freedom in carrying out the plans. However, the slaves' main goal is still to satisfy their master's wishes.

In the last 35 years, mathematical economists have developed market mechanism models describing how resources in an economy may be optimally shared in informationally and computationally decentralized ways (e.g., [3, 4, 39, 45, 66]). Researchers

² Some of these problems were considered by [67, 99, 101] but they are revisited in [94] while taking into consideration the specific domain.

in distributed systems and distributed artificial intelligence (e.g., [41, 53, 106, 107]) applied these models to resource allocation and task distribution problems in computerized environments, where one of their main goals was to improve the overall performance of the system. For example, Wellman [107] uses market price mechanisms for coordination and task distribution in distributed planning systems. The agents are divided into consumers and producers and use an iterative method to adjust prices and reach an equilibrium. This method is applicable under the “perfect competition” assumption, which is appropriate when there are numerous agents, each of which is small in relationship to the entire economy. We consider incentive contracting when there is usually a small number of agents in the environment. We also deal with situations where agents are uncertain about the world, and the contractors (the producers in Wellman’s terminology) may not carry out the tasks as promised.

In our incentive contract model and in the automated contracting frameworks [99] there is a hierarchical relationship among the agents. In most of the multi-agent systems (MA) where agents are self-motivated, there is no hierarchy among the agents that communicate and cooperate. For example, Sycara [104, 105] presents a model of negotiation that combines case-based reasoning and optimization of the multi-attribute utilities. This model is used in labor management negotiations where two agents need to reach an acceptable agreement. In [50–52], a strategic negotiation model is presented for situations where a set of self-motivated autonomous agents have common goals that they want to satisfy as soon as possible. Each agent, while wanting to minimize its costs, prefers to do as little as possible and therefore tries to reach an agreement over the division of labor. This model is also applicable when the agents need to share a resource. Zlotkin and Rosenschein [111] present a theoretical negotiation model for two rational agents which have symmetric capabilities and identical costs for their actions.

Contracting in multi-agent systems was previously studied in [32]. A formal definition of the mental state of an agent (or a group of agents) that would like to contract out one of its tasks was presented. Contracting depends mainly on an agent which believes that by taking some action (and thus bringing about a certain state of affairs), it can get another agent to perform an action. However, a detailed algorithm for finding the “motivating” action and the appropriate contractor is not presented in [32]. Also, the issue of choosing the appropriate effort level by the contractor is not explicitly considered. The main contribution of the present paper is the presentation of techniques for drafting beneficial contracts in situations where the contractor agents need to choose an effort level when carrying a task.

3. A framework for incentive contracting

In the environments to be discussed below, there are two types of agents. We will refer to the agent(s) that subcontracts one of its tasks to another agent or agents as the *manager(s)*, and we will refer to the agent(s) that may agree to carry out the tasks as the *contractor(s)*. In order to convince the contractor to do the task and motivate it to do it well, the manager needs to provide the contractor with a beneficial contract. The contractor’s success in carrying out the task depends on the time and work intensity

which it will put into fulfilling the task. We will refer to the contractor's time and work intensity as its *effort* level. We propose constructing a monetary system in the multi-agent environment, which will provide a way for allocating rewards and evaluating outcomes.

The following are the conditions that a *contracting multi-agent* (CMA) framework should satisfy (for any specific distributed multi-agent domain), in order for it to be accepted by all the designers of agents (for that specific domain):

- *Simplicity*: The contract should be simple and there should be an algorithm to compute it. That is, the agents will be able to compute the details of the contract. For example, if finding the awards for the contractor requires solving a set of inequalities, then the agents need to have a procedure to state these inequalities and a procedure to solve them.
- *Pareto-optimality*: There should be no other contracted arrangement that is preferred by both sides over the one they have reached. This means that there will be no other contract where the utilities of *both* agents are greater than their utilities in the contract agreed upon.
- *Stability*: The results should be in equilibrium³ and the contracts should be reached and executed without delay.

3.1. Agents' utility functions

A designer of an automated agent, in any environment, needs to provide the agent with a decision mechanism based on some given set of preferences. Structures of symbolic goals provide the agents with a good framework for planning, when the world is perfectly controlled by the agent and the effects of all the operators are known completely and with certainty to the agent [17, 33]. Symbolic goals are easily communicated, they guide the search for alternative plans and the projection process, and they also solve the horizon problem (see [33] for detailed discussion). However, symbolic goals do not give any information about the relative merits of different desirable alternatives. In addition, when the agent is uncertain of the past, present, or future environment and is uncertain of the result of its actions, then the structures of symbolic goals are not satisfying. In such situations numeric utility functions and decision theory offer a normative model for choice under uncertainty by providing support in evaluating multiple objectives and value tradeoffs [46, 108].⁴ We therefore propose that each designer of autonomous agents develop a numerical utility function that it would like its agent to maximize.

In situations where there is uncertainty and the agents need to make decisions under risk, the designers need to decide on their agents' attitude toward risk. There are three types of behaviors toward risk. An agent is *risk averse* if it always prefers to receive an outcome equal to the expected value of an uncertain situation over entering an uncertain situation. An agent is *risk prone* if it always prefers to enter an uncertain situation over

³ A pair of strategies (σ, τ) is a Nash equilibrium if, given τ , no strategy of Agent 1 results in an outcome that Agent 1 prefers to the outcome generated by (σ, τ) and similarly for Agent 2 given σ . We discuss the notion of Nash equilibrium and other equilibria concepts in Section 3.2 below.

⁴ The problem of integrating goals and utility is considered in [33].

receiving an outcome equal to its expected outcome for entering an uncertain situation. An agent is *risk neutral* if it is indifferent between the two options. Decision theory offers a formalization for capturing risk attitudes. If an agent's utility function is concave, it is risk averse. If the function is convex, it is risk prone. A linear utility function yields risk neutral behavior [23,46].

We propose that a utility function of an automated agent in our contracting multi-agent (CMA) environment depends on the agent's monetary gain and effort. Developing a quantitative evaluation of effort and world states and assigning numerical values to these is a difficult problem. However, in situations where the agents (or their owners) are paid according to the outcomes of their activities and there is a direct relationship between effort and expense, it is easier to develop such evaluations and numerical utility functions. Examples of such domains include the transportation domain of [94] where agents may be paid according to the value of the deliveries they make and their expenses may depend on the number of miles they travel, their speed, weather, etc. In a software agents domain, where users query an information center (see Example 4.2 below), the value of the references and documents provided by the information center as a response to a query may depend on their monetary value to the user. The information center's efforts may be measured by the time and resources spent on searching for an answer.

Our framework does not restrict the designer of an agent to any specific utility function since we assume that the personality of the designer (e.g., his/her attitude toward risk) will affect his/her choice of the agent's utility function. However, we do provide the designer with ways to evaluate how the choice of a utility function may affect the possible outcomes of his/her agent's interactions with other agents, how the type of a utility function may affect the contract that will be reached, and the complexity of finding a contract.

3.2. *Equilibrium concepts in multi-agent environments*

The manager's strategy in our CMA environment specifies which contract to offer to the contractor and the contractor's strategy specifies how it should respond to a given offer. Our desire is to obtain strategies which are in equilibrium, since if the agents use these strategies, the interaction among the agents may become more stable. As we consider different situations, we use different concepts of equilibrium to gain stability.⁵

In simple situations, with complete information, we use the *Nash equilibrium* concept. If there are n agents in the environment, a set of strategies (s_1, s_2, \dots, s_n) is in *Nash equilibrium* if no agent can benefit from deviating from its strategy (i.e., choose another strategy), given that the other agents do not deviate. For example, suppose (s_m, s_c) are a pair of strategies for a manager and a contractor respectively. If (s_m, s_c) are in Nash equilibrium, then if s_m specifies a contract that the manager should offer the contractor, the contractor will not have a better response than to act according to s_c . On the other

⁵ We assume that if an agent is indifferent between two options, but the other agents prefer one of these options, then the agent will choose the option preferred by the other agents.

hand, given the possible responses of the contractor according to s_c , the manager's best strategy is to offer the contract indicated in s_m .⁶

When there is incomplete information, e.g., agents do not know their opponents' exact types, the notion of *Bayesian–Nash equilibrium* is useful. This equilibrium includes a set of beliefs (one for each agent) and a set of strategies. A strategy combination and a set of beliefs form a Bayesian–Nash equilibrium if the strategies are in Nash equilibrium given the set of beliefs, and the agents update their beliefs, according to Bayes' rule [37].

When there are several stages of interaction among the agents, the Nash equilibrium strategies may involve threats that in certain senses are not credible. In order to rule out such equilibria we use the concept of *perfect equilibrium* [97]. It can be said that a set of strategies is in *perfect equilibrium* if the agents' strategies induce an equilibrium at any stage of the interaction.

There are two approaches for finding equilibria for the type of situations we consider in this paper. The first is the straight game theory approach: a search for Nash strategies or for perfect equilibrium strategies. In this approach the researcher makes a guess that some strategy combination is an equilibrium and then checks to see that it is so. The second is the economist's standard approach: set up a maximization problem, and solve using calculus. The drawback of the game theory approach is that it is not mechanical and the number of possible guesses is very large (and possibly infinite) and therefore it is difficult to develop a computer program that will find the Nash equilibrium strategies.⁷ The maximization approach, on the other hand, is much easier to implement. However, the problem with the maximization approach in our context is that the players must solve their optimization problems together: the contractor's strategy affects the manager's maximization problem and vice versa.

In this paper we will use, whenever possible, the maximization approach, with some care. This means that the maximization problem of the designer of the contract (usually the manager) will include, as a constraint, its opponent's (usually the contractor) maximization problem. The maximization problem of the contract's designer agent can be solved automatically by the agent. That is, the contracts which we provide maximize the expected utility of the designer of the contract (usually the manager). However, when designing the contract, the agent must take into consideration the possible responses of its opponent, which is also trying to maximize its own expected utility.

3.3. Notation

We use the following notations in the rest of the paper. A summary of this notation is given in Fig. 1.

⁶ As we see in Section 7.1, there are situations where there is more than one equilibrium. In specific cases, an agent's strategy may belong to two equilibria. If it is the first to take an action, it needs to take into consideration the possible behavior of its opponent in all equilibria.

⁷ In our previous work on negotiation under time constraints, we have identified perfect equilibrium strategies and proposed to develop a library of meta strategies to be used when appropriate [50–52].

Notation	Meaning	Comments
<i>Effort</i>	Set of efforts of the contractor	$e, e_1, \dots, e_i \in \text{Effort}$
<i>Outcome</i>	Set of possible monetary outcomes for carrying out a task	$q, q_1, \dots, q_j \in \text{Outcome}; q(e) \in \text{Outcome}$ when q is a function of $e \in \text{Effort}$
<i>Rewards</i>	Set of possible monetary rewards to the contractor	$r, r_1, \dots, r_i \in \text{Rewards}$ $r(q)$ when r is a function of $q \in \text{Outcome}$
U^c	The contractor's utility function	
U^m	The manager's utility function	
$\hat{u} \in \mathbb{R}$	Contractor's utility from outside options	(Reservation price)
$e^* \in \text{Effort}$	Efficient effort level for the manager	Given contractor constraints
$q^* \in \text{Outcome}$	Efficient outcome for the manager	Given contractor constraints

Fig. 1. Notation used in the paper.

- *Effort level*: Given a task, there are several effort levels that the contractor may adopt when trying to perform that task. We denote the set of these efforts by *Effort*. We use $e, e_i \in \text{Effort}$ to denote specific effort levels. In all cases, the contractor will decide how much effort to expend, but its decision may be influenced by the contract offered by the manager.
- *Outcome*: While the contractor's expected utility depends on its effort level in performing a task, the expected utility of the contracting agent depends heavily on the outcome of the performed task. The set of possible outcomes is denoted by *Outcome*. We assume that in the CMA environment, the outcome depends on the effort level expended by the contractor and that it can be expressed using a monetary system. We denote the monetary value of performing a task by $q \in \text{Outcome}$. Given an effort level $e \in \text{Effort}$, $q(e)$ denotes the monetary outcome of performing a task, as a function of e . This function increases with the effort involved. That is, the more effort put in by the contractor, the better the outcome.
- *Rewards*: In order to convince the contractor to carry out a task, the manager offers to pay the contractor a reward using the CMA monetary system. We denote the set of possible rewards by *Rewards* and its elements by r . The reward $r \in \text{Rewards}$ may be a function of the outcome from carrying out the task (i.e., $q \in \text{Outcome}$).
- *Utility functions*: We denote the contractor's utility function by $U^c : \text{Effort} \times \text{Rewards} \rightarrow \mathbb{R}$. We assume that in the CMA environment the contractor prefers to do as little as possible and gain the highest rewards; therefore, U^c is a decreasing function in effort and an increasing function in rewards. We denote the manager's utility function by $U^m : \text{Outcome} \times \text{Rewards} \rightarrow \mathbb{R}$. The manager prefers to give lower rewards and obtain larger outcomes. Thus, U^m is an increasing function with the outcome and a decreasing function with the reward being paid to the contractor.
- *Outside options*: If the contractor does not accept the contract from the manager and does not carry out the task, then it can either perform another task (its own or others') or remain idle. Its expected utility in such a situation is its reservation price and we refer to it as \hat{u} .

In the rest of the paper, in order to simplify the presentation of formulas, when the scope of a variable is clear from the context and the above notations, we will omit the precise

definition of the variable. For example, when using r_i we will not always mention that r_i belongs to *Rewards*.

In our system, we assume that the manager rewards the contractor *after* the task is carried out. In such situations there should be a technique for enforcing these rewards. In the case of multiple encounters, reputational considerations may yield appropriate behavior. Some external intervention may be required to enforce commitments in a single encounter, e.g., the responsibility of the manager's owner for its contracts toward the contractor's owner. Our last definitions are concerned with the value of the contracts to the manager. The *first best* contract will provide the manager with a profit that is equal to a profit it could get when there is complete information and the manager can observe the contractor(s)' actions. The *second best* contract is Pareto-optimal given information asymmetry and constraints on writing contracts, e.g., the manager does not observe the contractor(s)' actions.

4. Full information

At first we assume that all the relevant information about the environment and the situation is known to both agents. In the simplest case the manager can observe the contractor's effort and actions and force it to perform at the effort level preferred by the manager by paying only when the required effort is made. The amount of effort required from the contractor will be the one that maximizes the manager's outcome, taking into account the task fulfillment and the rewards that need to be made to the contractor. However, in most situations it is either not possible or too costly for the manager to observe the contractor's actions and its level of effort. In some cases, the manager may be either trying to carry out another task at the same time, or it cannot reach the site of the action. We consider two cases in such situations:

- In Section 4.1 we consider the case where there is no uncertainty with respect to the result of the contractor's actions.
- In Section 4.2 there is uncertainty concerning the outcome of an action taken by the contractor.

4.1. Contracts under certainty

Suppose both agents have full information about the world and about each other, but the manager does not observe the contractor's actions. Under these circumstances, there is no uncertainty concerning the results of the contractor's actions, i.e., the outcome is a function of the contractor's effort. If this function is known to both agents, then the manager can offer the contractor a *forcing contract* [16, 34, 88], which means that the manager will pay the contractor only if it provides the outcome required by the manager. If the contractor accepts the contract, then it will perform the task with the effort level that the manager finds to be most profitable to itself, even without the manager's observation. Note, the outcome won't necessarily be a result of the highest effort on the part of the contractor, but rather a result of the effort which provides the manager with the desired outcome.

We assume that either the manager or the contractor is one of several similar perfect competitors. In the background other managers are competing to subcontract some tasks to the contractor, so that the manager's equilibrium profit equals zero, or many possible contractor agents compete for the manager's task, so that the agent's equilibrium utility equals its reservation price—the minimum that induces it to agree to perform the task.⁸ Suppose the contractor is one of many agents that compete for the manager's task. The manager should pick an effort level, $e^* \in \textit{Effort}$, that will generate the efficient output level, $q^* \in \textit{Outcome}$. As we explained above, since there are several possible agents available for contracting in equilibrium, the contract must at least provide the contractor with the utility \hat{u} . The manager needs to choose a reward function where $U^c(e^*, r(q^*)) = \hat{u}$ and $U^c(e, r(q)) < \hat{u}$ for $e \neq e^*$. \hat{u} is the minimal reward that will make the contractor accept the contract. Since the manager would like to pay the contractor as little as possible, but wants the contractor to accept the offer, then if the outcome reveals that the contractor provided the required effort level, the manager will pay the contractor \hat{u} . If the contractor accepts the contract, but does not choose the appropriate effort level, its reward will be even less than \hat{u} . We demonstrate this case in the following example.

Example 4.1 (*Contracting under certainty*). Two robotics companies, CompM and CompC⁹ are responsible for cleaning and garbage collection in adjacent cities (e.g., Tel-Aviv and Ramat-Gan). Each of the companies has several autonomous mobile robots that carry out the cleaning tasks in these cities.¹⁰

Most of the garbage collected by these companies is used for recycling, and therefore the companies are paid mainly according to the amount of garbage they collect and its value for recycling. The amount of garbage collected by a robot depends on the effort level with which it carries out the task, and the distribution of garbage in the area it tries to clean.

Suppose one of CompM's robots has to collect garbage far from the other robots of CompM, but close to several of CompC's robots. The CompM's robot would like to subcontract some of its garbage collection tasks and therefore approaches one of

⁸ Note that if this assumption is not made, there may be several equilibria. In such situations the designers of the agents may agree upon regulations that will make all agents in the environment focus on one of them. For example, they may agree that the manager will serve as a *focal arbitrator*. A focal arbitrator is an agent who can determine a focal equilibrium in the environment. In such a case, the equilibrium will be similar to the case where many possible contractor agents compete for the manager's task. One way of making the manager a focal arbitrator is by imposing regulations in which the contractor cannot negotiate the details of a contract; it can either accept the contract offered to it by the manager, or reject it.

⁹ The robots of company CompM will play the role of the managers and the robots of CompC will play the role of the contractors.

¹⁰ Most of the autonomous robots up today operate indoors (e.g., Flakey's of SRI, Polly's of MIT, Schimmer of Stanford [11,44,80]). Mobile robots that operate in rougher terrain are usually less autonomous (e.g., DANTE II that was developed by NASA and CMU and explored the crater on Mt. Spurr volcano in Alaska) or act in well-defined environments (e.g., CALMAN—a computerized articulated lawn mower with automatic navigation that was developed at Lulea University of Technology in Sweden). It seems that on-going research on perception, mapping, and navigation in a changing environment will contribute to the construction of "cleaning" automated agents, but it is likely to be a few years before such robots are operational.

CompC's robots. The CompC robot can collect garbage in three levels of effort (e): Low, Medium and High respectively denoted by 1, 2 and 3. CompM's robot cannot observe the effort of the CompC's robot since it wants to carry out another task simultaneously. The value of garbage collection is $q(e) = \sqrt{100e}$. The utility function of CompM's robot, if a contract is reached, is $U^m(q, r) = q - r$ and the utility function of CompC's robot in the case that it accepts the contract is $U^c(e, r) = 17 - 10/r - 2e$, where r is the reward to CompC's robot. If CompC's robot rejects the contract, it will busy itself with maintenance tasks and its utility will be 10. It is easy to calculate that the best effort level from CompM's robot's point of view is 2, in which there will be an outcome of $\sqrt{200}$. The contract that CompM's robot offers to the CompC's robot is $3\frac{1}{3}$ if the outcome is $\sqrt{200}$ and 0 otherwise. This contract will be accepted by CompC's robot and its effort level will be Medium.

Another issue of concern is how the manager will choose which agent to approach. In a situation of complete information (we consider the incomplete information case in Section 5) it should compute the expected utility for itself from each contract with each agent and choose the one with the maximal expected utility.

Our model is also appropriate in the case where there are several managers with the same utility functions, but only one possible contractor. In such cases, there should be information about the utilities of the managers in the event that they do not sign a contract, i.e., the managers' reservation price. The outcome to the manager in this case should be equal to its reservation price. In this case, the contractor¹¹ will offer a contract, trying to maximize its expected utility under the constraint that the manager will gain its reservation price.

4.2. *Contracts under uncertainty*

We continue to assume in this case that the agents have full information about each other, and that the manager does not observe the contractor's behavior. However, in most subcontracting situations, there is uncertainty concerning the possible outcome of an action. If the contractor chooses some effort level, then there are several possibilities for an outcome. For example, suppose a cleaning automated agent subcontracts its garbage collection task and suppose that there is uncertainty about the distribution of the garbage at the site. If the contractor chooses a high effort level and the garbage is distributed all over the area, the outcome may be similar to the case where the contractor chooses a low level of effort and the garbage is all in one place. However, if the contractor chooses a high effort level when the garbage is located in one area, the outcome may be higher and, thus, better to the manager. In such situations the outcome of performing a task does not reveal the exact effort level of the contractor, and consequently, choosing a stable and maximal contract is much more difficult.

Assuming that the world may be in one of several states, neither the manager nor the contractor knows the exact state of the world when agreeing on the contract. There is the possibility that the contractor may gain more information about the world during or after

¹¹ Here the contractor is the focal arbitrator.

completing the task, but only *after* signing the contract and choosing the effort level. The manager, however, is not capable of gaining more information about the world.

Following [34], we also assume that there is a set of possible outcomes to the contractor carrying out the task $Outcome = \{q_1, \dots, q_n\}$ such that $q_1 < q_2 < \dots < q_n$ depends upon the state of the world and upon the effort level of the contractor. Furthermore, we assume that, given a level of effort, there is a probability distribution attached to the outcomes that is known to both agents.¹² Formally, we assume that there is a probability function $\varphi : Effort \times Outcome \rightarrow \mathbb{R}$, such that for any $e \in Effort$, $\sum_{i=1}^n \varphi(e, q_i) = 1$ and for all $q_i \in Outcome$, $\varphi(e, q_i) > 0$.¹³ This characterizes the situations where the manager is not able to use the outcome to determine the contractor's effort level unambiguously.

The manager's problem is to find a contract that will maximize the manager's expected utility, knowing that the contractor may reject the contract or, even if it accepts the contract, the effort level will be chosen later [88]. The manager's reward to the contractor can be based only on the outcome. Let us assume that in the contract that will be offered by the manager, for any q_i , $i = 1, \dots, n$, the manager will pay the contractor the reward r_i . The maximization problem can be constructed as follows (see also [88]).¹⁴

$$\text{Maximize}_{r_1, \dots, r_n} \sum_{i=1}^n \varphi(\hat{e}, q_i) U^m(q_i, r_i) \quad (1)$$

with the constraints:

$$\text{(IR)} \quad \sum_{i=1}^n \varphi(\hat{e}, q_i) U^c(\hat{e}, r_i) \geq \hat{u}, \quad (2)$$

$$\text{(IC)} \quad \hat{e} = \text{argmax}_{e \in Effort} \sum_{i=1}^n \varphi(e, q_i) U^c(e, r_i). \quad (3)$$

Equation (1) states that the manager tries to choose the reward for the contractor, so as to maximize its expected utility subject to two constraints. First, the rewards for the contractor must be large enough to motivate the contractor to prefer the contract rather than to reject it. Constraint (2) is called the *individual rationality* (IR) constraint. This constraint requires that the expected utility of the contractor will be at least as much as its reservation price (\hat{u}). The second constraint (3), which is called the *participation*

¹² A practical question is how the agents find the probability distribution. It may be that they have preliminary information about the world. In the worst case, they may assume an equal distribution. The model can be easily extended to the case that each agent has different beliefs about the state of the world, i.e., has its own probability function, which is known to its opponent [81].

¹³ The formal model in which the outcome is a function of the state of the world and the contractor's effort level, and in which the probabilistic function gives the probability of the state of the world which is independent of the contractor's effort level, is a special case of the model described here [34, 81, 91].

¹⁴ As we mentioned above, we omitted the definitions of the variables in some of the formulas. In the formulas below, as well as in the rest of the paper, $r_i \in \mathbb{R}$ and $q_i \in Outcome$.

constraint (IC), provides the contractor with the motivation it needs to choose the effort level that the manager prefers, given the contract it is offered. This means that given the agreed rewards, \hat{e} will provide the contractor with the highest outcome.

In order to be able to use the above framework in the CMA environment, the agents should be able to solve the above maximization problem. The algorithms that should be used depend primarily on the utility functions of the agents, as we will describe in the next two sections.

4.2.1. Risk neutral agents

If the manager and the contractor are risk neutral, then solving the maximization problem can be done using any linear programming technique (e.g., simplex, see for example [83, 103]). Furthermore, in most situations, the solution will be very simple: the manager will receive a fixed amount from the outcome, and the rest will go to the contractor. That is, $r_i = q_i - C$ for $1 \leq i \leq n$, where the constant C is determined by constraint (IR)(2) [98].

Example 4.2 (*Risk neutral software agents under uncertainty*). Suppose there is an information center that has several large databases (e.g., the Earth Science Data and Information System (ESDIS) of the National Aeronautics and Space Administration (NASA)). The information center receives queries from users (possibly automated agents) and answers the queries by providing references and documents that are relevant to the query. Given a query, both the information center and the user are uncertain about the number of documents in the information center's databases that are relevant to the query. However, they both know that if the information center uses more resources (e.g., CPU time) searching its databases, then its probability of finding more documents will increase. The amount of resources that the information center uses in answering a query will be referred to as its effort level. In particular, based on previous experience, the user and the information center have some probabilistic estimation of the number of documents that will be found given a specific effort level of the information center.

In order to simplify the problem we assume that there are only two effort levels possible for the information center, Low ($e = 1$) and High ($e = 2$). Suppose the user asked a query such that the user and the information center estimate that there are either 30 or 100 related documents.¹⁵ In addition, both the information center and the agent estimate that if the information center chooses the Low effort level, then the probability that it will find 30 documents is $\frac{2}{3}$ and the probability that it will find 100 documents is $\frac{1}{3}$. On the other hand, if it searches with the High effort level, then the probability that it will find 30 documents is $\frac{1}{12}$, and the probability that it will find 100 documents is $\frac{11}{12}$.

If the user gets 30 documents it is worth 50, while locating 100 documents is worth 75 to the user. The user's¹⁶ utility function is $U^m(30, r) = 50 - r$ and $U^m(100, r) = 75 - r$.

¹⁵ In real situations we expect that the set of possible numbers of documents will be much larger (but finite and discrete) and also that the number of possible effort levels will be much larger. However, this small example demonstrates the technique.

¹⁶ Note that the user plays the role of the manager and the information center plays the role of the contractor.

The information center's utility function is $U^c(r, e) = r - 10e$; if it doesn't respond to the user's query, it works on maintaining its databases, and its expected utility will be 5, i.e., $\hat{u} = 5$.

In solving the maximization problem above, we reach the conclusion that the user should offer the information center the reward $2\frac{1}{12}$ if it provides only 30 documents and $27\frac{1}{12}$ if it provides the user with 100 documents. The information center should choose the High level of effort and the user will always gain a profit of $47\frac{11}{12}$.

Similar situations may occur between the cleaning automated agents.

Example 4.3 (Risk neutral robots under uncertainty). Suppose the utility function of the CompC's robot from Example 4.1 is $U^c(r, e) = r - e$, and suppose that it can choose between two effort levels: Low ($e = 1$) and High ($e = 2$), and suppose that its reservation price is $\hat{u} = 1$. There are then two possible monetary outcomes to the garbage collection scenario: $q_1 = 8$ and $q_2 = 10$. The utility function of CompM's robot remains as it was in the previous example, i.e., $U^m(q, r) = q - r$.

If CompC's robot chooses the Lower level of effort then the outcome will be q_1 with probability $\frac{3}{4}$ and q_2 with probability $\frac{1}{4}$. If it takes the High level effort the probability of q_1 is $\frac{1}{8}$ and of q_2 it is $\frac{7}{8}$. In such situations, CompM's robot is able to ensure itself a profit of $6\frac{3}{4}$. That is, $r_1 = 1\frac{1}{4}$ and $r_2 = 3\frac{1}{4}$. The robot of CompC will choose the High level effort.

4.2.2. The contractor is risk averse

When the agents are not neutral toward risk, then solving the manager's maximization problem becomes much more difficult. However, if the agents' utility functions are carefully chosen, then an algorithm does exist.

Suppose the contractor is risk averse and the manager is risk neutral (the methods are also applicable when both are risk averse). Grossman and Hart [31] present a three-step procedure in order to find appropriate contracts in such situations. The first step of the procedure is to find for each possible effort level, the set of reward contracts that will induce the contractor to choose that particular effort level. The second step of the procedure is then to find the contract which supports that effort level at the lowest cost to the manager. The third step of the procedure is to choose the effort level that maximizes profits, keeping in mind the need to support that effort with a costly reward contract. Formally, step one and two are as follows: Suppose the manager wants the contractor to choose the effort level $e' \in \text{Effort}$, it will need then to solve the following:

$$C(e') = \text{Minimize}_{r_1, \dots, r_n} \sum_{i=1}^n \varphi(e', q_i) r_i \tag{4}$$

with the constraints:

$$(\text{IR}) \quad \sum_{i=1}^n \varphi(e', q_i) U^c(e', r_i) \geq \hat{u}, \tag{5}$$

$$(IC) \quad \sum_{i=1}^n \wp(e', q_i) U^c(e', r_i) \geq \sum_{i=1}^n \wp(e, q_i) U^c(e, r_i) \quad \text{for all } e \in \text{Effort}. \quad (6)$$

The first constraint (5) requires that the expected utility for the contractor will be at least as good as its outside options (its reservation price). The second constraint (6) requires that given the contract, the contractor will prefer to take the effort level e' . The minimization problem states that the manager is looking for a contract where it can pay as little as possible to induce the contractor to choose e' . For this minimization problem there is an algorithm if U^c satisfies several properties, including the property that the preferences of the contractor over entering uncertain situations are independent of its actions [31, 83, 89].¹⁷ That is, the contractor's preferences over reward lotteries are independent of its actions and effort level.

After finding a set of possible values, r_1, \dots, r_n for every $e \in \text{Effort}$ (where the set may be empty since there could be effort levels which the manager cannot make the contractor choose), and after finding the minimum expected reward $C(e)$, for any effort level, the manager is ready to move to the third step, which is easy to compute. The manager will then choose the effort level that will provide it with the maximal outcome:

$$\text{Maximize}_{e \in \text{Effort}} U^m \left(\sum_{i=1}^n \wp(e, q_i) q_i, C(e) \right). \quad (7)$$

The contractor's computational task is easier. After being offered a contract, the contractor only needs to check the validity of the inequalities that appear as constraints in the manager's maximization problem. That is, when the contractor needs to check the validity of the individual rationality constraint (IR) in order to decide whether to accept the contract or not. When the contractor needs to decide which effort level to provide, it should consider its expected utility from its effort level, similar to the maximization problem described in the *participation* constraints (IC). In both cases, since all variables are known, based on the suggested contract, these checks are very easy.

Example 4.4 (*Risk averse contractor under uncertainty*). Suppose the situation is exactly as in Example 4.3 but the designer of the robot determines that the contractor will be risk averse and its utility function is as in Example 4.1: $U^c(r, e) = 17 - 10/r - 2e$ and $\hat{u} = 1$.

The maximization problem that the manager should solve is:

$$\text{Maximize}_{r_1, \dots, r_n} \sum_{i=1}^n \wp(\hat{e}, q_i) (q_i - r_i) \quad (8)$$

with the constraints:

¹⁷ In [89] the problem of finding a contract when the manager can choose an effort level from a real interval is considered. Rogerson identifies the sufficient condition in which the constraints (IC) can be replaced with the requirement that the effort level be a stationary point for the contractor. In such situations a solution can be calculated using the Kuhn-Tucker Theorem.

$$(IR) \quad \sum_{i=1}^2 \wp(\hat{e}, q_i) \left(17 - \frac{10}{r_i} - 2e \right) \geq 1, \quad (9)$$

$$(IC) \quad \hat{e} = \operatorname{argmax}_{e \in \{1,2\}} \sum_{i=1}^2 \wp(e, q_i) \left(17 - \frac{10}{r_i} - 2e \right). \quad (10)$$

Grossman and Hart's three-step procedure [31] requires that the manager first determine the minimal reward needed to make the contractor choose $e_1 = 1$ and what the minimal reward is that will make it choose $e_2 = 2$:

$$C(e_1) = \operatorname{Minimize}_{r_1, r_2} \frac{3}{4}r_1 + \frac{1}{4}r_2 \quad (11)$$

with the constraints:

$$(IR) \quad \frac{3}{4} \left(17 - \frac{10}{r_1} - 2 \right) + \frac{1}{4} \left(17 - \frac{10}{r_2} - 2 \right) \geq 1, \quad (12)$$

$$(IC) \quad \frac{3}{4} \left(17 - \frac{10}{r_1} - 2 \right) + \frac{1}{4} \left(17 - \frac{10}{r_2} - 2 \right) \\ \geq \frac{1}{8} \left(17 - \frac{10}{r_1} - 4 \right) + \frac{7}{8} \left(17 - \frac{10}{r_2} - 4 \right). \quad (13)$$

The results of solving this minimization problem using Lagrangian multipliers is that the minimal reward to make the contractor choose $e_1 = 1$ is $r_1 = r_2 = \frac{5}{7}$. A similar minimization problem can be stated and solved for $e_2 = 2$. In this case the minimal reward to make the contractor choose effort level $e_2 = 2$ is $r'_1 = 1$ and $r'_2 = 1\frac{8}{17}$. Finally, the manager should check which effort level it prefers, given the above rewards, i.e., it should compare $\wp(e_1, q_1)(q_1 - r_1) + \wp(e_1, q_2)(q_2 - r_2)$ and $\wp(e_2, q_1)(q_1 - r'_1) + \wp(e_2, q_2)(q_2 - r'_2)$. The conclusion is that the manager can obtain the largest expected utilities by offering $r'_1 = 1$ and $r'_2 = 1\frac{8}{17}$. The contractor will then compute its expected utility from choosing effort level e_1 (i.e., $\frac{3}{4}(17 - 10/r'_1 - 2) + \frac{1}{4}(17 - 10/r'_2 - 2)$) and from choosing effort level e_2 (i.e., $\frac{1}{8}(17 - 10/r'_1 - 4) + \frac{7}{8}(17 - 10/r'_2 - 4)$), and it will then realize that its expected utility from both effort levels is the same. The contractor will then verify that its expected utility from the offered contract is greater than \hat{u} (i.e., $\frac{1}{8}(17 - 10/r'_1 - 4) + \frac{7}{8}(17 - 10/r'_2 - 4) \geq 1$), and will then accept the contract and choose effort level e_2 since its expected utility from both effort levels are the same and e_2 is preferred by the manager.¹⁸

4.2.3. Obtaining imperfect information about the contractor's behavior

Even in situations where the manager cannot observe the actions of the contractor, it may be able to gain some information about its behavior. For example, it can gain

¹⁸ In the rest of the paper we will not specify the contractor's computation procedures, since in most of the situations, given a contract, the contractor needs only to check the validity of the inequalities that appear as constraints in the manager's maximization problem, similar to the check done in this example. Since all variables are known, based on the suggested contract, this check is straightforward.

information by setting up a camera in the garbage collection site. This information may be imperfect, and the process of getting this information is called an imperfect (noisy) monitoring process. In particular, if the contractor takes effort level e , then the result of such a monitoring mechanism may be $e + \delta$ where δ is a random variable drawn from $[\alpha_0, \alpha_1]$ for some finite α_0, α_1 . These results will enable the manager to obtain some estimation of the contractor's effort level. The main question is, however, whether using such monitoring is beneficial.

We continue to assume that the assumptions described in the beginning of Section 4.2 hold. That is, the agents have full information about each other, the manager does not observe the contractor's behavior, there is uncertainty concerning the state of the world and neither agent knows the state of the world, but both agents observe the outcome of the contractor carrying out the task. Under the above conditions, it has been shown that if the contractor is risk neutral, there are no gains (to either agent) from the use of any monitoring mechanism [35]. This claim holds when the manager is either risk neutral or risk averse.¹⁹ However, according to the above conditions, if the contractor is risk averse, there are potential gains to monitoring. This is the case, particularly, if a contract of the following form is an optimal monitoring contract: If the contractor's action is judged acceptable on the basis of the monitored outcome, the contractor will then be paid according to a prespecified schedule. Otherwise, it will receive less preferred, fixed rewards [35]. To demonstrate this idea we use a modification of an example that appears in [35].

Example 4.5. Suppose the utility function of CompC's robots from the previous examples is $U^c(e, r) = r^{0.25} - \frac{4}{5}e^{1.25}$, its reservation price is $\hat{u} = 0$ and the utility function of CompM's robot is, as in previous examples, $U^m(q, r) = q - r$. Suppose the world's situation is θ which is uniformly distributed on $[0, 1]$ and the outcome function is $q(e, \theta) = e + \theta$. The monitoring technology then includes only monitors, which are uniformly distributed on $[e - \varepsilon, e + \varepsilon]$ for some $\varepsilon > 0$. That is, if the contractor chooses effort level e , the monitor will provide an equal probability number α , between $e - \varepsilon$ and $e + \varepsilon$.

The contract that will be offered by the CompM robot is a function of the outcome and the monitored information α :

$$r(q, \alpha) = \begin{cases} \frac{5}{4}\varepsilon, & \text{if } \alpha \geq 2e + 2^{-6}e^{-3} - \varepsilon, \\ 0, & \text{otherwise.} \end{cases}$$

The effort level chosen by the CompC's robot depends on ε . If $\varepsilon < 2^{-1.25}$, then it will choose $2e + 2^{-6}e^{-3}$. In such situations the CompC's robot will always get the reward $\frac{5}{4}\varepsilon$ and its expected utility is 0. The expected utility of CompM's robot is $\frac{1}{2} + 2^{-5} + \varepsilon^{-3} + \frac{3}{4}\varepsilon$. If $\varepsilon \geq 2^{-1.25}$, then CompC's robot will not choose the required level of effort, but rather will take a lower level effort, $5 * 2^{-6}\varepsilon^{-3}$. It may be that the monitoring value α will be lower than $2e + 2^{-6}e^{-3} - \varepsilon$ and CompC's robot won't get any reward. The probability

¹⁹ The manager's utility function should be monotone increasing with $q - r$, concave and continuously differentiable. The proof to the claim appears in [35, Proposition 3].

In addition, in each of the n contracts offered by the manager the contractor's utility should be higher than its reservation price. The manager should find a set of such self-selection contracts that will maximize its expected utility, based on its probabilistic beliefs. Formally:

$$\text{Maximize}_{(q_1, r_1), \dots, (q_n, r_n)} \sum_{i=1}^n \phi_i U^m(q_i, r_i) \quad (15)$$

subject to:

$$(SS) \quad \text{Eq. (14),}$$

$$(IR) \quad U^c(e_i, r_i) \geq \hat{u}, \quad \text{where } f(\theta_i, e_i) = q_i. \quad (16)$$

We demonstrate this maximization problem in the next example.

Example 5.1 (*Contracting under asymmetric information (software agents)*). Similar to Example 4.2, the user asks the information center a query. However, the user is uncertain as to whether the databases of the information center were updated recently or not. That is, the user believes that the databases can be either in state $\theta_1 = 1$ or in state $\theta_2 = 2$. The information center, of course, knows the state of its databases. The number of documents that will be found by the information center depends on the state of its databases and the effort level it will choose to search with. The outcome function is $f(e, \theta) = e\theta$, the user's utility function is $U^m(q, r) = q - r$ and the information center's utility function is $U^c(e, r) = r - e^2$. Hence, with $f(e, \theta) = e\theta$, the information center's utility function is a function of the output, reward and the state of the databases is $U^c(q, r, \theta) = r - (q/\theta)^2$. We also assume that the information center's (i.e., the contractor's) reservation price is $\hat{u} = 1$, and the user (manager) believes with probability 0.25 that the state of the databases is θ_1 (i.e., $\phi_1 = 0.25$), and it believes with probability 0.75 that the state of the world is θ_2 .

In such a situation the user should solve the following maximization problem:

$$\text{Maximize}_{(r_i, q_i), i=1,2} 0.25(q_1 - r_1) + 0.75(q_2 - r_2) \quad (17)$$

subject to:

$$r_1 - q_1^2 \geq r_2 - q_2^2,$$

$$r_2 - (q_2/2)^2 \geq r_1 - (q_1/2)^2,$$

$$r_1 - q_1^2 \geq 1,$$

$$r_2 - (q_2/2)^2 \geq 1,$$

$$0 \leq r_i \leq q_i, \quad i = 1, 2.$$

If the output function f is twice differentiable in e , with $f_e > 0$ and $f_{ee} < 0$ for all θ ,²¹ then there is an interesting result concerning the manager's preference over the

²¹ f_e denotes the first derivative of f by e and f_{ee} is the second derivative.

information available to the contractor. If the contractor has full information about the state of the world *before* signing the contract, then the manager's expected utility is *lower* than in the case where it and the contractor have symmetric beliefs (either perfect or imperfect) about the state of the world *before* signing the contract [6, 15]. This conclusion is a result of the fact that when they share the same (perfect or imperfect) state of information, the contractor can be held to its reservation level of expected utility.

5.2. Asymmetric information after reaching an agreement

In some situations, the contractor is able to collect more information *before* it performs the agreed upon task but only after signing the contract. For example, when CompC's robot reaches the garbage collection site, it may find out what the exact state of the world is and know for sure what the outcome will be if it takes a specific level of effort.

If agreements are enforced, i.e., if the contractor cannot opt out of the agreement after it is signed, then the only difference between the previous case and the current one is, that constraints (IR) (16) should be about the *expected* utility of the contractor, rather than its eventual utilities, since at the time of the contract, the exact utility is not known to the contractor. If the agents have similar probabilistic beliefs about the state of the world when signing the contract (i.e., ϕ_i), then the constraint is as follows:

$$(IR) \quad \sum_{i=1}^n \phi_i U^c(e_i, r_i) \geq \hat{u}, \quad \text{where } f(\theta_i, e_i) = q_i. \quad (18)$$

We demonstrate this in the following example.

Example 5.2. (*Risk neutral agents under asymmetric information (cleaning automated agents)*). Suppose the situation is exactly as in Example 4.3, and CompC's robot can find out more information after the robots have reached a contract, but before choosing its level of effort. As in Example 4.3 the contractor can choose between two effort levels Low ($e = 1$) and High ($e = 2$) and its reservation price is $\hat{u} = 1$. There are then two possible monetary outcomes to the garbage collection: $q_1 = 8$ and $q_2 = 10$. The agents' utility functions are the same as in Example 4.3. The world can be in one of eight possible states $\theta_1, \dots, \theta_8$ with equal probability. The outcome function is defined as follows: For $1 \leq i \leq 6$, $f(1, \theta_i) = q_1$, for $7 \leq i \leq 8$ $f(1, \theta_i) = q_2$, $f(2, \theta_1) = q_1$ and for $2 \leq i \leq 8$, $f(2, \theta_i) = q_2$. Note that this yields the same probabilistic outcome as in Example 4.3.

There are two possibilities for constructing the contracts, depending on which effort level the contractor will choose if the state of the world is either $\theta_2, \dots, \theta_6$. It is clear that if the state is θ_1, θ_7 or θ_8 the contractor will choose the Low effort level. If the manager would like the contractor to choose High effort level in states $\theta_2, \dots, \theta_6$, then the manager should solve the following minimization problem (we list only the binding constraints):

$$\text{Minimize}_{r_1, r_2} \quad \frac{1}{8}r_1 + \frac{7}{8}r_2 \quad (19)$$

subject to:

of this happening is $1 - 2^{-5}\epsilon^{-4}$, and CompC's robot's expected utility is still 0, while the expected utility of CompM's robot in this case is $\frac{1}{2} + 5 * 2^{-7}\epsilon^{-3}$. In both cases, CompC's expected utility is more than $\frac{1}{2}$, which is what it can expect if it does not use a monitoring mechanism.

From the above results, it follows that when $\epsilon > 2^{-1.25}$, the rewards to CompC's robot increase with ϵ , its effort level decreases with ϵ , and the expected utility of CompM's robot decreases with ϵ . These results fit the belief that as monitoring becomes less precise (i.e., ϵ increases), the manager's expected utility decreases.

5. Asymmetric and incomplete information

There are some situations in which the contractor may have more information than the manager. First, the contractor may have obtained more information concerning the environment, e.g., the information center from Example 4.2 may know the exact state of its databases, while the user in that example may only have some probabilistic beliefs about the databases based on previous experience. Second, in other situations the manager may not know the utility function of the contractor. The contractor then may be one of several types that reflect the contractor's ability to carry out its task, its efficiency or the cost of its effort. However, we assume that given the contractor's type, its utility function would be known to its party. For example, suppose the cleaning company CompC builds robots of two types. The specifications of the robots are known to CompC's robots and to CompM's robots; however, CompM's robots do not know the specific types of CompC's robots they will encounter.

In both cases, the manager could simply ask the contractor for the additional information, i.e., its type or the state of the world, however the contractor will not tell the truth unless the manager provides it with a monetary incentive to do so. This will often cause inefficiency from the manager's point of view. The search for an equilibrium in such situations may often be extremely difficult, but there is a useful technique that, in using it, the manager can reduce the number of contracts it needs to consider, as we explain below. The manager should search for an optimal mechanism [14] as follows: the manager offers the contractor a menu of contracts indexed by the agent's type (or the state of the world). The contractor can then decide whether to accept the menu of contracts or not. If it accepts the offer it sends a message to the manager reporting its type. The manager is then committed to the contract indexed by this type. The rewards of the contractor in each of these contracts are the functions of the outcomes.²⁰

The big advantage of this mechanism is the *revelation principle*: For every contract that leads to lying, there is a contract with the same outcome for the contractor (given its type or the state of the world) but without inducement for the contractor to lie. Therefore, without loss of generality, it is enough for the manager to consider only

²⁰ Given the chosen contract, the contractor chooses an effort level which maximizes its own expected utility. In each of the menu's contracts, the contractor's expected utility should be at least as high as its expected utility if it does not sign the contract.

contracts where it is in the contractor's interest to honestly report its type [76]. There are two main limitations in using the revelation principle. First, there is a need for communication since the contractor needs to send a message to the manager specifying its type. Second, this mechanism requires strong precommitment capability on the part of the manager. After the contractor reveals its type honestly, it is often in the manager's advantage (sometimes the contractor's as well) to re-negotiate the contract, and offer a different one. We discuss these issues in Sections 5.4 and 5.6. We will consider several situations of asymmetric information.

- In Section 5.1 we consider the case where the state of the world is known to the contractor, but not to the manager.
- In Section 5.2 neither agent knows the state of the world *before* signing the contract. The contractor finds out that information *after* signing the contract, but *before* choosing its effort level.
- In Section 5.3 the contractor's information is initially better than that of the manager, but it knows the exact state of the world only after a contract is signed (but *before* choosing the effort level).
- In Section 5.4 the contractor cannot predict the outcome, based on its private information, either before or after signing the contract.
- In Section 5.5 both agents have some private information, e.g., they have some private information about their types.

5.1. Asymmetric information about the state of the world

Suppose the world can be in one of several states, $\theta_1, \dots, \theta_n$. If the contractor chooses a level of effort e and the state of the world is θ , then the outcome will be $f(e, \theta)$ [36]. As in previous cases the contractor's utility function ($U^c(e, r)$) increases with the reward it gets from the manager (r), and decreasing with its effort (e). The manager's utility function ($U^m(q, r)$) increases with the outcome, and decreases with its reward to the contractor. We assume that the contractor knows the state of the world θ , but the manager has no definite knowledge about the state of the world, having only a probabilistic belief. We denote its belief that the world is in state θ_i , $i = 1, \dots, n$ by ϕ_i and assume that $\sum_{i=1}^n \phi_i = 1$.

As we described above, in the first step of the agents' interaction, the manager will offer the contractor n pairs (one for each state) for an outcome and a payoff (q_i, r_i) . The contractor will then report its private information, i.e., the state of the world, to the manager. According to this message, the corresponding contract is implemented. In the third step the contractor chooses its effort level, and is paid according to the chosen contract and the outcome. As was mentioned above, based on the revelation principle, we will restrict our attention to *direct mechanisms* under which the contractor reports the situation of the world honestly, motivated by the contract. That is, if the state of the world is θ_i , then (q_i, r_i) is the best contract among the ones offered by the manager. This constraint is called "self-selection". Formally,

$$\begin{aligned}
 (\text{SS}) \quad & \forall i \in \{1, \dots, n\} \quad U^c(e_i, r_i) \geq U^c(e_j, r_j) \\
 & \text{where } 1 \leq j \leq n, \quad f(\theta_i, e_i) = q_i, \quad f(\theta_i, e_j) = q_j.
 \end{aligned}
 \tag{14}$$

$$(IR) \quad \frac{1}{8}(r_1 - 1) + \frac{5}{8}(r_2 - 2) + \frac{2}{8}(r_2 - 1) \geq 1, \quad (20)$$

$$(IC) \quad r_2 - r_1 \geq 1. \quad (21)$$

By solving this problem we can conclude that the manager can always keep $7\frac{1}{8}$ of the outcome and pay the contractor $r_1 = \frac{7}{8}$ and $r_2 = 2\frac{7}{8}$. Similarly, we can formalize the problem where the contractor chooses effort level *Low* in states $\theta_2 - \theta_6$. The rewards should be $r'_1 = r'_2 = 2$ and the expected utility for the manager is $6\frac{1}{2}$. In order for the manager to maximize its expected utility, the first option is better since it yields the manager an expected outcome of $7\frac{1}{8}$. This is higher than in Example 4.3, where its expected outcome is $6\frac{3}{4}$.

We would like to consider the option of monitoring in such situations. It was proved in [35] that if the contractor is risk neutral, and if it is able to get information about the exact state of the world *after* signing the agreement, then monitoring is not valuable. If the contractor is risk averse, monitoring may be beneficial as we will explain in Section 5.6. The manager can design a contract that will make the contractor choose the Pareto-efficient effort level for the real state of the world.

If it is possible for the contractor to cancel the contract after obtaining the information about the state of the world, then this possibility should be taken into consideration when the agents agree on the contract [95]. When the contractor can opt out of an agreement, the question is what are its alternatives at that point. It may be that it can still get its original outside options, i.e., its reservation price \hat{u} . In other situations, however, it may have already lost the original outside option, and therefore gain less from a new option. Let us denote the contractor's new reservation price by \hat{u}^{new} . In such situations, the manager needs to add an additional constraint to its maximization problem. That is, in addition to constraints (14) and (18), the following constraint should be added:

$$\forall i, 1 \leq i \leq n \text{ such that } f(\theta_i, e_i) = q_i, \quad U^c(e_i, r_i) \geq \hat{u}^{\text{new}}. \quad (22)$$

This constraint verifies that even when the contractor finds out more information about the environment *before* it chooses its level of effort, it will benefit from choosing the level e_i and will consequently keep the agreement. Of course, these constraints reduce the manager's expected utility, and it will need to suggest to the contractor higher payments to make sure it won't opt out. We will demonstrate this in the case where the contractor is risk neutral as in Example 4.3.

Example 5.3. (*Risk neutral agents under asymmetric information with opting out (cleaning automated agents).*) Suppose the situation is exactly as in Example 5.2, but before choosing its level of effort, CompC's robot can opt out of the agreement and get its original reservation price (i.e., $\hat{u}^{\text{new}} = \hat{u} = 1$). Therefore, instead of constraint (20), the following should be stated:

$$r_1 - 1 \geq 1, \quad r_2 - 2 \geq 1. \quad (23)$$

The manager should then offer $r_1 = 2$ and $r_2 = 3$. The expected outcome for the manager will be 6.875 which is lower than in the case where the contractor cannot opt out.

5.3. Asymmetric and imperfect information before contracting

We consider the situation where the contractor's information is initially better than that of the manager, but that it knows the exact state of the world only after a contract is signed. For example, CompC's robot may initially have better information about the garbage distribution than CompM's robot. However, it does not have full information about the state of the world. Only after reaching the garbage collection site (*after* signing an agreement), does it find out about the real garbage distribution. Note that in the previous section (Section 5.2), both agents have the same preliminary beliefs about the state of the world, and the asymmetry in information arises only *after* reaching an agreement. On the other hand, in Section 5.1, the contractor already knows the state of the world *before* signing the contract. That is, the situation of this section is between that of Section 5.1 and Section 5.2.

As in previous situations, we assume that the outcome is a function of the contractor's effort level and the state of the world, i.e., $q = f(e, \theta)$. At no time can the manager observe either e or θ . Suppose that the possible states of the world are $\theta_1, \theta_2, \dots, \theta_n$, such that $\theta_i < \theta_{i+1}$ for $1 \leq i \leq n$. Furthermore, the manager does not know the exact probability distribution of θ , but rather knows that there are D possible probability distributions ϕ^d , and it believes with probability ϕ^d that the real distribution is ϕ^d . Before signing the contract, the contractor does not know the actual state of the world either, but it does know which probability distribution function is the correct one. We assume that the utility function of the contractor can be written as a function of q and r as follows: $U^c(q, r) = r - e(q, \theta)$ where $f(e(q, \theta), \theta) = q$. In such situations the optimal strategy for the manager [36] is to design at most D distinct contracts from which the contractor can make a binding choice by sending a message to the manager. Thus the maximization problem of the manager is as follows [96]:

$$\text{Maximize}_{\{(q_1^1, r_1^1), \dots, (q_n^1, r_n^1)\}, \dots, \{(q_1^D, r_1^D), \dots, (q_n^D, r_n^D)\}} \sum_{d=1}^D \phi^d \sum_{i=1}^n \phi^d(\theta_i) U^m(q_i, r_i) \quad (24)$$

subject to:

$$\text{(IR)} \quad \sum_{i=1}^n \phi^d(\theta_i) (r_i^d - e(q_i^d, \theta_i)) \geq \hat{u} \quad \forall d = 1, \dots, D, \quad (25)$$

$$\text{(SS)} \quad \sum_{i=1}^n \phi^d(\theta_i) (r_i^d - e(q_i^d, \theta_i)) \geq \sum_{i=1}^n \phi^d(\theta_i) (r_i^r - e(q_i^r, \theta_i)) \quad \forall r, d = 1, \dots, D, \quad (26)$$

$$\text{(IC)} \quad r_i^d - e(q_i^d, \theta_i) \geq r_j^d - e(q_j^d, \theta_i) \quad \forall i, j = 1, \dots, n \text{ for each } d = 1, \dots, D, \quad (27)$$

where $\phi^d(\theta_i)$ is the probability that the state of the world is θ_i according to distribution d ($\phi^d(\theta_i) > 0 \forall i, d$), q_i^d is the output produced by the contractor in state θ_i under contract $\{(q_i^d, r_i^d)\}$ and r_i^d is the reward to the contractor under that contract.

The first set of constraints (IR) (25) guarantees that any contract selected by the agent provides him with a level of expected utility that is at least as good as its reservation price. The second set of constraints (SS) (26) ensures that the contractor will report honestly about the actual distribution (i.e., will choose contract $\{(q_i^d, r_i^d)\}$ when φ^d is the actual distribution). The third set of constraints (IC) (27) guarantees that the agent will produce q_i^d in state θ_i if it chooses contract $\{(q_i^d, r_i^d)\}$. Note, that if $D = 1$ the maximization problem is as in Section 5.2.

5.4. Asymmetric information and uncertainty

There are some situations that are characterized by both private information *and* uncertainty. This means that the contractor cannot predict the outcome based on its private information, since the private information only provides a better estimation of what the outcome may be. One example of such a situation is as follows [10]. In the first stage of the interaction, the manager offers the contractor a menu of contracts based on a message it will send in addition to the observed outcome. The contractor may reject the offer or agree to it and sign a contract. In the second stage, the contractor may gain some private information ξ about the world, after signing a contract, but before sending a message or choosing an effort level. This information will help the contractor to improve its prediction as to what the outcome will be, given its level of effort. For example, when the robot of CompC reaches the area that it needs to clean, it determines the garbage distribution of this area (i.e., it collects information about the world's state). This information may not be complete, but it is not known to the robot of CompM at all. In the third stage, the contractor sends a message to the manager and chooses a level of effort. In the fourth stage the outcome is observed by both agents, and the contractor is paid according to the outcome and its earlier message. Note that in such situations, the contractor has committed itself not to leave the agreement once it has observed ξ .²² Also in this case [10], the agents can concentrate on the class of contracts that induce the contractor to send a truthful message to the manager. This is due to the fact that it has been shown [10] for any untruthful contracts, a truthful one can be found in which the expected utility of the agents is the same. The maximization problem of the manager is similar to the one in Section 5.2; the contractor's utility that appears in the constraints is replaced by its expected utility given ξ .

5.5. Both parties have private information

There are some situations where both the manager and the contractor have private information, e.g., both agents have private information about their own types. To be able to concentrate on the effect of the private information of the agents, we assume that the actions taken by the contractor are observable by the manager. However, we continue to assume that there is uncertainty about the outcome. That is, we assume that, given a level of effort, there is a probability distribution of φ which is attached to the possible

²² In most of the situations the manager is better off making such a commitment. However, in some situations, both agents can be made better off through re-negotiation [14, 26, 38, 55].

outcomes that is known to both agents (as in Section 4.2). Furthermore, we assume that the agents can agree on probabilistic actions, i.e., they will agree that the contractor will choose its level of effort, using an agreed upon probability distribution.

Suppose that each of the agents has some probabilistic beliefs about its opponent's private information, then, in order for an informed manager to do better than an uninformed one, it must actively participate in the contract selection and not only in the mechanism design. We describe here an interaction procedure that satisfies the following properties: The revelation principle holds, there exists a perfect Bayesian equilibrium which is Pareto-optimal for the different types of managers, and the manager generically does strictly better than when the contractor knows the manager's private information [68]. There are up to four possible stages in an interaction.

- (1) In the first stage of the interaction, the manager offers a mechanism to the contractor which specifies:
 - (a) a set of possible messages that each party can choose,
 - (b) for each pair of messages m_m, m_c that can be chosen simultaneously by the manager and the contractor respectively, a corresponding probabilistic function of the effort level will be chosen by the contractor (note that the probabilistic choice mechanism and the effort level are observable by the manager),
 - (c) pairs of outcomes and rewards.
- (2) In the next stage the contractor accepts or refuses the mechanism. If it refuses the mechanism, it receives its reservation price \hat{u} , and the interaction ends.
- (3) The agents can send each other the messages simultaneously.
- (4) The contractor performs the task at the appropriate effort level and is paid according to the outcome.

For example, suppose there are two types of managers (a and b) and two types of contractors (1 and 2). The set of possible messages can include the agents' types (i.e., the manager can send the messages " a " and " b " and the contractor can send the messages "1" and "2"). The manager should offer a menu of contracts that includes four possibilities, one for each combination of the agents' types. For example, $Cont^{a,1}: [a, 1: e^{a,1}, (q_1, r_1^{a,1}), \dots, (q_n, r_n^{a,1})]$ indicates that if the manager sends the message " a " in step (3) and the contractor sends the message "1", then the contractor will choose effort level $e^{a,1}$ (which can also be a probabilistic function of possible effort levels) and its reward will depend on the outcome. For example, if the outcome is q_n , its reward will be $r_n^{a,1}$. Similarly, $Cont^{b,1}: [b, 1: e^{b,1}, (q_1, r_1^{b,1}), \dots, (q_n, r_n^{b,1})]$ specifies a contract when the manager sends the message " b " and the contractor sends the message "1".

As in previous cases, the agents can limit themselves to honest reports. In situations where the exact type of the manager does not directly influence the contractor's utilities, [68, 77] show that the manager can profit from the contractor's incomplete information. The intuition behind these results is as follows. When the manager proposes a contract, it is subject to two types of constraints. The (IR) constraint requires that the expected utility of the contractor, when accepting the contract, will be higher than the contractor's reservation price. There are also constraints to ensure that when the contract is carried out, the contractor behaves in the appropriate way, given its private information

(IC). When the manager does not have private information, the constraints must hold individually for the manager's specific type. If the contractor has incomplete information about the manager, the constraints need to only be held in "expectation" over the suggested contracts which are functions of the manager's type. That is, the expected utility of the contractor that appears in the constraints is the sum of the expected utilities for each of the manager's types, multiplied by the probability that the manager is of this type. For example, suppose in the example described above, $EU^{c1}(Contr^{a,1})$ denotes the contractor's expected utility if its type is 1, and it accepts the contract $Contr^{a,1}$ above. Similarly, $EU^{c1}(Contr^{b,1})$ denotes the contractor's expected utility from the contract $Contr^{b,1}$. Then, if the contractor knows that the manager's type is a , the constraint (IR) with respect to the contractor of type 1 will be $EU^{c1}(Contr^{a,1}) \geq \hat{u}$, and if the contractor knows that the manager's type is b , the constraint (IR) will be $EU^{c1}(Contr^{b,1}) \geq \hat{u}$. However, if the contractor believes that with probability p_a the manager's type is a and with probability p_b its type is b , then the constraint (IR) is $p_a EU^{c1}(Contr^{a,1}) + p_b EU^{c1}(Contr^{b,1}) \geq \hat{u}$.

For this reason, if the contractor is not informed about the manager's type, the manager of a given type can increase its utility above its possible utility in situations where the contractor is fully informed, by violating some of its constraints, as long as they are offset by constraints of the other types of the managers. Actually, it was proved in [68] that in most of these situations, there exists a mechanism in which all types of managers do strictly better than in the instances where the contractor is fully informed. However, in order to take advantage of the contractor's incomplete information, the manager must refrain from revealing its type at the mechanism proposal stage (i.e., stage (1) above). Otherwise, the constraints must hold for the revealed type, rather than for just the expected types. Note, that since *all* types of managers do better in the case that the contractor is not informed, the manager can't benefit from pretending to be a different type. This means that if the selection of the mechanism by the manager depends in any way upon the manager's individual type, then the selection of the mechanism itself will convey information about its type to the contractor. Therefore, any manager, regardless of its type, should offer the same mechanism.²³

Cases in which the manager's private information influences the contractor's utilities are more complex [69]. In such situations it is no longer true that, without loss of generality, the manager can postpone revealing its type until the third stage of the interaction. The manager may wish to disclose information about itself in order to influence the contractor's actions; if so then the manager's proposal should balance between total disclosure and complete concealment. Furthermore, the manager's expected utility when it has private information which influences the contractor's utility, may be even lower than in a case where the manager does not have any private information at all. This is because the contractor's expected utility may be low, given some of the manager's types denoted by "bad" types. Therefore, when the contractor's probabilistic belief that its opponent's type is "bad" is high (even if the actual

²³ Maskin and Tirole [68] show that any equilibrium of the mechanism design presented here can be computed as a Walrasian equilibrium of a fictitious economy. In this economy, the traders are the different types of manager. For more technical and formal details see [68].

type is not “bad”), the contractor must be paid correspondingly high rewards to encourage it to accept the contract. Note that in the first case we considered, where the contractor is not directly influenced by the manager’s type, its original beliefs do not play an important role since the contractor cares only about how the manager’s type will affect its behavior in the implementation of the mechanism, but no more than that.

5.6. Value of information and communication

There are two important questions related to asymmetric information situations [10, 74]:

- (1) Will the manager always be better off, the more the contractor knows about the world?
- (2) Is communication beneficial to the manager? Meaning, is it better for the manager to suggest a menu of contracts to the contractor and ask it to send a message informing the manager of the current state of the world, or will it be better off offering only a single contract, based only on the joint observed outcome?

The second question is essential when communication is costly to the manager. Intuitively, it seems that both communications and a knowledgeable contractor will allow for more efficient contracting. The contractor may use its knowledge to choose the correct actions, and with a menu of contracts the contractor may select the rewards tailored to the actual situation. Surprisingly, the answer to both questions is that it is not always the case that communications and knowledgeable contractors will improve the managers’ benefits, rather their effect depends on the exact details of the situation. There are even situations when less information by the manager is preferred to more [30].

As we explained in Section 5.1, when the contractor has full private information before signing the contract, the manager’s expected utility is lower than if they have symmetric beliefs. If the contractor acquires its information *after* signing the agreement, then its effect on the manager varies. The contractor may use its additional information in two ways: It may use its information to take a low effort level, thereby reducing the benefits for the managers, or it may use the information to improve the outcome (see two demonstrating examples in [10]). If the manager gains information *after* signing an agreement, then the information is only valuable if it is affected by the contractor’s level of effort (see Section 4.2.3), and can therefore be used to estimate the contractor’s effort level [30]. For example, information gained by setting up a camera in a garbage collection site provides the manager with an estimation of the contractor’s effort level and may therefore be useful to the manager.

The disadvantage of communications is that the “self-selection” constraint can sometimes be very restrictive so that the information received by the manager is not beneficial. This occurs particularly if the contractor has perfect private information about the world, i.e., given an action, it can anticipate the exact outcome, for any “appropriate” menu of contracts. The manager can then replicate its benefits, using a single contract. Furthermore, even if the contractor does not have perfect information, there are many situations in which there is no value for communication [14, 74]. These situations are such that the stochastic outcome is informative.

If the outcome is not informative, however,²⁴ then communication is valuable. It is valuable for two reasons; because it allows the manager to implement a more efficient level of effort without having to pay the contractor for making it choose correctly, and alternatively, menu contracts can be valuable even though the contractor's action choices are unchanged. In such situations, the value of communication results from the rewards given to the contractor. There are, of course, situations where the manager can use the information gathered in the menu contracts for other purposes (e.g., later contracts with other agents). In such a case, it may prefer the menu of contracts, even if it cannot benefit in the current interaction.

5.7. Several contractors compete for the job

There may be a situation where there are several agents in the environment, and the manager can choose one of them to do the job. The agents may each be of a different type (measuring, for example, efficiency and ability), or independently drawn from a set of possible types. If the manager does not know the types of the other agents, the following mechanism is appropriate: The manager announces a set of contracts indexed by agents' types and asks the potential contractors to report their types. On the basis of these reports, the manager chooses one agent [73].²⁵ The agent that is chosen, chooses a level of effort that is not observable by the manager. The rewards to the chosen contractor depend upon the contractor's reported type and the observed outcome. As in previous cases, the manager can use, without loss of generality, contracts in which the agents report their types honestly [76].²⁶

An important aspect in the design of the contracts is the marginal return to the manager by increasing the probability that a specific type (e.g., z_i) will be chosen. This marginal return consists of the outcome minus the rewards that the contractor receives, and minus the increase in the expected rewards to the other types of agents. The latter effect arises because, by increasing the probability that a report of z_i will be chosen, the manager makes it more attractive for higher types to pretend to be z_i . To prevent this, the manager must improve the rewards for all the types that are higher than z_i .

If the agents' types satisfy the appropriate conditions (see details in [73]) that are related to the above described aspect, and if the highest reported type is chosen, then the contract may be optimal for the manager. However, the manager's benefits will be lower than in the case where it can observe the contractor's effort level (i.e., it gets only the "second best" benefits).

²⁴ See [74] for exact conditions.

²⁵ There are situations where the agents' types are multi-dimensional. That is, the manager is uncertain about different aspects of the contractor that are independent; for example, its capabilities and its disk space. Techniques to formalize the maximization problem in such situations, and methods to solve it can be found in [54, 71].

²⁶ The measure of risk aversion will influence the agents' behavior when there are more than one possible contractor in the environment. A less risk averse agent will usually have the ability to win over more risk averse agents in service of any risk averse manager [92].

6. Repeated encounters

Suppose the manager wants to subcontract its tasks several (finite) times. Two types of contracts are possible in such situations: Long term contracts, where one contract is signed before the repeated encounters start, and short term contracts, i.e., in the beginning of each encounter a new contract is agreed upon by the agents.

6.1. Short term contracts

Repetition of the encounters between the manager and the contractors enables the agents to reach efficient short term contracts if the number of encounters is large enough²⁷ and if the contractor can be “punished” sufficiently [27, 43, 65, 85, 86].

Based on the average outcome, the manager could form an accurate estimate of the contractor’s effort over a certain amount of time. That is, if the manager wants the contractor to make a certain effort level of $\hat{e} \in \text{Effort}$ in all the encounters, it can compute the expected outcome over that certain amount of time if the contractor actually performs the task with that effort level. The manager can keep track of the cumulative sum of the actual outcomes and compare it with the expected outcome. If after several encounters the manager realizes that the cumulative outcome is below a given function of the expected outcome, it should impose a severe “punishment” on the contractor. If the function over the expected outcome is chosen carefully [85], then the probability of imposing a “punishment” when the contractor is in fact carrying out the desired effort level, can be made very low. Meanwhile, the probability of eventually imposing the “punishment” if the agent does not do \hat{e} is 1.0.

Suppose there is asymmetric information where we assume that in each of the encounters the situation is similar to that of Section 5.1, meaning that in each encounter t , the outcome q^t is a function of the contractor’s effort level e_t and the state of the world’s θ_t (which may change from one encounter to the other). The outcome at time t does not depend on the contractor’s actions in previous encounters, and the states of the world in the encounters are independently and identically distributed.²⁸ In each encounter, the manager offers a reward function of $r_t(q^t)$, and the contractor chooses its effort level based on the state of the world, i.e., $e_t(\theta_t)$. If there is a single encounter, then only second best contracts can be achieved and we denote the reward function and the effort

²⁷ In [85] the number of encounters should be larger than some thresholds, but finite and known to the agents. Fudenberg et al. [27] assume that there is a terminal date, T , such that after T the manager’s profit will no longer depend on the contractor’s actions, that there will be no additional information arriving, and that the manager won’t give the contractor any further rewards. However, the contractor may be inactive for some of the periods, and in particular, the contractor may opt out before date T . They don’t assume that T is large, but rather make other assumptions such as that there is common knowledge of technology and preferences, and equal access to banking. Also, Holmstrom and Milgrom [43] Malcomson and Spinnewyn [65] don’t assume that T is large, but make additional assumptions about the agents’ utility functions and about the environment. For example, Holmstrom and Milgrom assume that the contractor has access to unlimited saving and borrowing at the same interest rate as the manager.

²⁸ In [27] it is assumed that past actions and signals can affect current outcomes and signals, as long as these dependencies are publicly revealed.

level function by (r^*, e^*) . We denote the first best solution by (\hat{r}, \hat{e}) and the expected outcome in this case for the manager and the contractor by \hat{v} and \hat{x} , respectively.

The notion of an *epsilon equilibrium* [85] will be used, although it imposes weaker restrictions on the agents' strategies than the restrictions imposed by the Nash equilibrium. For any positive number epsilon, an *epsilon equilibrium* is a pair of strategies that allows the average of each agent's expected utility to be within epsilon from the expected utility of the best response to the other agent's strategy. One rationale for epsilon equilibrium is, that if the agents have sufficient inertia, they will not bother to realize possible small gains [24]. The main motivation for using the epsilon equilibrium concept is as follows: In every perfect equilibrium (defined in Section 3.2) of the (finite) T -period game, the outcome in every period is a Nash equilibrium of the one-period game. On the other hand, in infinite multiple stage games (i.e., T is infinite), in which each agent can observe the other agent's one-period strategies, there are perfect equilibria of the game which result in the use of "cooperative" pairs of strategies (in our situations, the first best strategies) in each one-period game, particularly in the use of Pareto-optimal pairs of strategies. In the same situation, it was shown that for any positive epsilon, if T is sufficiently large, then there are epsilon equilibria of the T -period game (i.e., T is finite) which results in cooperative behavior in all or most of the component one-period games. That is, for epsilon equilibria, infinite horizon repeated games may be well approximated by long finite horizon games. Unfortunately, the number of perfect equilibria in infinite horizon repeated games is very large, as is indicated by the "Folk Theorem".²⁹ However, the number of possible equilibria strategies can be limited by considering "trigger strategies" [84], and first best strategies can be sustained in epsilon equilibria of the multiple encounters situation by the "trigger strategies". The trigger strategy for the contractor, denoted by ρ , is very simple: It uses the effort level function \hat{e} until the first encounter where the manager does not use the reward function \hat{r} ; at that encounter and in each encounter thereafter the contractor will optimize against the reward function announced for each encounter.

The suitable trigger strategy for the manager is a little more complicated. In each encounter t , based on the history of outcomes through encounter $t - 1$, the manager must decide whether to make the reward \hat{r} or switch to the reward function r^* . If its switching rule is too lax, then the contractor may be able to accumulate a large enough extra expected utility by cheating before getting caught, thereby making cheating attractive. On the other hand, if the switching rule is too strict, then there will be a substantial probability that the manager will switch to r^* before the contractor ever starts cheating. We define $C_t = f(e_t(\theta_t), \theta_t)$, i.e., C_t is the outcome in encounter t if the contractor uses the effort level function e_t and the state of the world is θ_t . We define S_n to be the sum of outcomes in periods 1 to n , that is, $S_n = C_1 + \dots + C_n$. We let \hat{C}_t denote the outcome in period t if the contractor uses \hat{e} , and let \hat{S}_n be the corresponding cumulative sum of outcomes by the end of encounter n . The random variables \hat{C}_t are independent

²⁹ This theorem is called "Folk Theorem" because no one remembers who should get credit for it [5]. The theorem says that under certain conditions (see [1, 5, 25]) in any infinitely repeated n -person game, with finite action sets at each repetition, any combination of actions observed in any finite number of repetitions is a unique outcome of a sub-game perfect equilibrium.

and identically distributed since the θ_t are so. Their expected value is \hat{c} . We let b_n be a strictly increasing sequence of positive numbers ($n \geq 1$), and define the random variables \tilde{N} and N by:

$$\tilde{N} = \min\{n \geq 1 \mid S_n - n\hat{c} \leq -b_n\}, \quad N = \min\{\tilde{N}, T\}. \quad (28)$$

The following trigger strategy should be used by the manager: Pay the contractor \hat{r} in each period through N and thereafter use the reward function r^* . We shall denote this strategy by $\sigma((b_n))$. We define \mathbf{B} as the class of positive sequences (b_n) that satisfy ([85]):

- b_n are strictly increasing, and $\lim_{n \rightarrow \infty} b_n/n = 0$.
- There exists $\lambda > 1$ such that $b_n \geq \lambda b_n^0$, $n \geq 1$.

The main result of [85] on these strategies is as follows: For any $\varepsilon > 0$ there exists a sequence (b_n) in \mathbf{B} and T_ε , such that for all $T \geq T_\varepsilon$ the pair of strategies $(\sigma((b_n)), \rho)$ is an ε equilibrium, and yields the manager and contractor average expected utilities respectively of at least $(\hat{v} - \varepsilon)$ and $(\hat{x} - \varepsilon)$.³⁰

6.2. Long term contracts

In the previous section we assumed that the number of encounters between the manager and contractors may be very large. This enables the manager's strategy for offering a contract in a given time period t , to depend on the average outcome in the prior $t - 1$ encounters. If there is a limited number of encounters the contracts need to be more complicated since there is not enough information that has accumulated.

For example, suppose that the agent is evaluated according to its average performance, there is uncertainty about the state of the world (i.e., each single encounter is as in Section 4.2) and the number of encounters is small (e.g., two encounters). If the contractor is "lucky" in the first encounter, the outcome will be high, and in the second encounter it can take a low effort level without adversely affecting the sum of both encounters. The contractor, therefore, is motivated to adjust its effort over time as a function of its previous performance. As a result of this phenomenon, the optimal contracts in situations where the number of encounters is small, will not be a simple function of the average outcomes [57] in general. The problem of the contractor trying to adjust its effort over time as a function of its previous performance, may also arise when the number of encounters is very large. However, if the number of encounters is very large, such behavior will eventually be detected.

The problem of subcontracting when the number of repeated encounters is small is considered in [57]. It is assumed that the manager can commit itself before the first encounter to a long term contract that will be implemented during all their encounters. The outcome of each encounter depends on the contractor's effort level (which is unobservable to the manager), and the state of the world in that encounter, which is not known to either agent, as in Section 4.2. Suppose there are only two encounters [57],

³⁰ In [86] the situation of symmetric information with uncertainty is considered. That is, the situation of a single encounter is as in Section 4.2. It provides Pareto-optimal strategies only in the case that there are infinite encounters.

and before the first encounter the manager offers a binding contract. Then the reward in the first encounter will depend upon the outcome of that encounter, but the reward of the second encounter will depend upon the outcomes of the first and second encounters. If the contract is accepted by the contractor then it should choose the effort level of the first encounter. The outcome of the first encounter is observed by both agents, and the contractor is paid according to the contract. In the second encounter, the contractor chooses an effort level which is a function of the outcome of the first encounter. The outcome of the second encounter is also observed by both agents and the rewards are given. When the manager chooses the contract, it should solve a maximization problem similar to that of Section 4.2. However, the manager's expected utility that appears in the maximization expression (1) should be replaced by its expected utility in both encounters. Similarly, it should consider the appropriate constraints (i.e., IR and IC) on the effort levels chosen by the contractor in both encounters. Subject to these constraints, the manager is able to update the contractor's rewards over time in any fashion that it desires. It was shown in [57] that the rewards in the second encounter should be an increasing function of the outcome of the first encounter.

7. Subcontracting to a group

Suppose that the task the manager wants to contract out can be performed by a group of agents. Each of the contractors is independent in the sense that it tries to maximize its own utility. The manager offers a contract to each of the possible contractors. If one of them rejects the offer, then the manager cannot subcontract the task.³¹ Otherwise, the contractors can simultaneously choose effort levels. As in previous sections, the manager cannot observe the effort levels and the members of the group while they carry out the task.

7.1. Individual outcome is observed

In this section we assume that each contractor yields an observable outcome of q_i and that the overall outcome will be equal to the sum of the q_i . The advantage of using the multiple outputs to form the basis for a reward to each agent is that usually some information about the state of the world can be concluded from observing the whole array of q_i 's [79], i.e., in such a situation, the individual actions can be estimated by comparing the performances of the different agents.

7.1.1. One agent's effort does not influence the other agents' outcomes

7.1.1.1. *The contractors have symmetric information.* Suppose the outcome for an agent is a probabilistic function of its effort level e_i , that the state of the world is θ , and that the individual aspects are ε_i , i.e., $q_i = f(e_i, \theta, \varepsilon_i)$. For example, in the cleaning automated agents case, θ could reflect the garbage distribution in the whole site, while

³¹ We will also consider below the situation where, if an agent accepts the contract, it will be implemented regardless of the other agents' responses.

ε_i represents the garbage distribution in the exact location of contractor i . Each of the contractors observes θ before it chooses its effort level, but it does not observe ε_i before making its choice.³² We assume that the contractors are identical, i.e., have the same utility function $U^c(e, r) = v(r) - c(e)$ and the same abilities. We will assume that $f(e_i, \theta, \varepsilon_i) = e_i\theta + \varepsilon_i$ and that φ_{ε_i} are the distribution functions of ε_i .

In the first model, there is no exchange of messages between the agents. Since only the outcome is observed, this is the only thing the rewards depend upon. The main question to be asked is: Is it better to make a contract based on all the outcomes, or is it better for a contractor's reward to depend only on its own outcome?

When the contractors' outcomes are independent, then observing all the q_i provides no additional information about the contractor's effort. In this case, the rewards should depend only on the individual outcome. Sometimes it is possible to find enough statistics from q_1, \dots, q_n , denoted by $T(\{q_1, \dots, q_n\})$, about the state of the world. The rewards of a specific agent should then depend upon its individual outcome and on $T(\{q_1, \dots, q_n\})$ [79]. For example, if both θ and ε are normally distributed random variables, then the average value of $\{q_1, \dots, q_n\}$ provides sufficient statistical information for θ . When the number of contractors becomes very large, the estimation of θ converges to the true value. In such situations, the rewards should depend on q_i and on the estimation of θ .

Another option when designing a contract for a group of contractors is to pay the contractors according to their ordinal positions alone and not according to the actual size of their output, i.e., to encourage a contest among the agents. Suppose there are two contractors; using the contest approach, there is a winner's reward r_w and a loser's reward r_l . The winner's output q_w is not necessarily worth r_w , so that the winner is actually paid more than its contribution to the overall outcome. This is done in order to motivate the contractors to choose greater effort levels. A larger prize for the winner motivates greater effort by all agents and increases the manager's outcome [79].

If the first contractor chooses effort level e_1 , and the second chooses effort level e_2 , then the first one will "win" if $\theta e_1 + \varepsilon_1 > \theta e_2 + \varepsilon_2$. Each of the contractors tries to choose higher levels of effort in order to be paid r_w . However, even though they both choose higher effort levels, it does not increase their probability of winning (which is, if we speak of symmetric equilibrium, $\frac{1}{2}$). The expected utility of a contractor i is therefore,

$$\frac{1}{2}[v(r_w) + v(r_l)] - c(e_i). \quad (29)$$

The details of how to compute r_w and r_l in a given situation are described in [79]. An interesting result from this is that in some situations it is possible to make the contractors choose an effort level, using the above "contest" mechanism, which is even larger than when the manager can observe the agent's effort levels, i.e., better than the first best contract. A variation of this method is when the "winner" must win by an amount greater than a certain margin. That is, instead of ranking contractors solely on the basis of the relative position of their outcomes, the manager can rank one contractor above another if that agent's outcome is greater than its opponent's by a positive margin.

³² We consider the case where a contractor can alter its effort level after observing ε_i in the next section.

The introduction of “margins” can lower the probability that any “prize” will be paid while maintaining the same level of motivation for choosing high levels of effort.

There are several other methods for possible rewards for members of a group. For example, giving a reward only to the agent whose output is the highest or punishing the agent that came in last [79]. Rewards that are based on relative performance are generally more flexible, and reduce the risk taken by the contractors [78].

7.1.1.2. The contractors have private information. In this case we assume that each of the contractor’s outcomes is affected by different aspects of the state of the world, in which each agent can only observe its own private “aspect” of that world. There is a probabilistic correlation between these aspects, but agents cannot observe each other’s aspects and the manager cannot observe any of them. For example, if a robot of CompM subcontracts its garbage collection task to a CompC’s robot and a robot of a third company, then each of them can observe the garbage distribution in its own garbage collection site *before* signing the contract, and since they gather the garbage in adjacent sites, the garbage distribution at their sites are correlated. CompM’s robot, however, does not know either distribution.

Suppose there are only two agents, A and B , and two output functions $f^l(e^l, \theta^l)$, $l = A, B$ [63]. Then we also assume that θ^l can be θ_1^l or θ_2^l (i.e., the world can be in four different states with two possibilities for each variable). For $l = A, B$ let $\varphi(\theta_i^l)$ be the probability that $\theta^l = \theta_i^l$ for $i = 1, 2$. We denote this probability by p_i^l and assume that $p_i^l > 0$ and that $\varphi(\theta_1^l) + \varphi(\theta_2^l) = 1$. As in previous sections, the level of effort, e^l , is not observable. We do assume, however, that for each l , $f^l(e^l, \theta_1^l) < f^l(e^l, \theta_2^l)$ for all e^l , therefore, θ_2^l represents a “good” state and θ_1^l a “bad” state. The state variables are positively but imperfectly correlated. We denote by s_i^A the probability of $\theta^B = \theta_i^B$, given that $\theta^A = \theta_i^A$ and similarly s_i^B denotes the probability of $\theta^A = \theta_i^A$ given that $\theta^B = \theta_i^B$. We assume that $1 > s_1^A > s_2^A > 0$.

Agent l ($= A, B$) *privately* observes θ^l *before* signing a contract with the manager. The manager is risk neutral and the contractors are risk averse. Their utility functions are similar to that which appears in Section 7.1.1. Given the utility function of the contractor l , and the state of the world, one can compute the “disutility” of producing an outcome such as q^l . The contractor’s utility can, therefore, be expressed as a function of the rewards and the outcome (as we did, for example, in Section 5.3). We will assume that $U^{c,l}(q^l, r^l) = v^l(r^l) - d^l(q^l, \theta^l)$ and that the contractor’s reservation price is \hat{u}^l . A typical contract that can be offered by the manager to agent A in this case, is of the following form [63]:

You may choose to produce either q_1^A or q_2^A . Your reward, r^A will depend not only on your output, but also on what agent B will produce. If you choose to produce q_i^A ($i \in \{1, 2\}$), then

- if agent B produces q_1^B , you will be paid r_{i1}^A ,
- if agent B produces q_2^B , you will be paid r_{i2}^A ,
- if agent B does not sign the contract, you will be paid r_{i0}^A .

In [15] the maximization problem of the manager was stated. It restricted the contractor’s output choices to a Bayes–Nash equilibrium, given that they are guaranteed at

least their reservation price (conditional on their private information). This is done for $l = A, B$.

$$\text{Maximize}_{q_i^l, r_{ij}^l, i, j \in \{1,2\}} p_1^l [s_1^l(q_1^l - r_{11}^l) + (1 - s_1^l)(q_1^l - r_{12}^l)] + p_2^l [s_2^l(q_2^l - r_{21}^l) + (1 - s_2^l)(q_2^l - r_{22}^l)] \tag{30}$$

subject to:

$$\text{(IR)} \quad s_i^l v^l(r_{i1}^l) + (1 - s_i^l) v^l(r_{i2}^l) - d^l(q_i^l, \theta_i^l) \geq \hat{u}^l, \quad i = 1, 2, \tag{31}$$

$$\text{(IC)} \quad s_i^l v^l(r_{i1}^l) + (1 - s_i^l) v^l(r_{i2}^l) - d^l(q_i^l, \theta_i^l) \geq s_i^l v^l(r_{j1}^l) + (1 - s_i^l) v^l(r_{j2}^l) - d^l(q_j^l, \theta_j^l), \quad i, j = 1, 2; i \neq j. \tag{32}$$

The result of this maximization provides the manager with rewards that discourage a contractor from choosing the output q_1^l when it observes θ_2^l . The reward will satisfy $r_{11}^l > r_{12}^l$ and $r_{21}^l = r_{22}^l = r_2$. These contracts yield to the manager the highest possible expected outcome. If the manager offers each agent $l = A, B$ the choice of

- producing q_1^l and receiving a probabilistic reward of $\{r_{11}^l, r_{12}^l\}$, or
- producing q_2^l and receiving a sure reward of r_2^l ,

then the manager will get the maximum outcome if both agents respond as the manager desires, i.e., sign their respective contracts and produce the output q_i^l when they observe θ_i^l . In the case of a single agent, the constraints ensure that the contractor will choose the desired effort level. However, if there are two agents, there exists another pair of equilibrium strategies whose outcome, from the contractors' point of view, is better to both agents than the outcome in the equilibrium that the manager wants to implement. The outcome for the manager, if they choose that level of effort is, however, low [15]. In particular, there is an equilibrium for both contractors to always choose the outcome q_1^l (regardless of their observed state), and in all states they will both be strictly better off than in the equilibrium preferred by the manager (i.e., choose q_1^l if the state is θ_1^l and q_2^l if the state is θ_2). Of course, in this case, the manager will definitely be worse off. It was suggested in [15] to strengthen the incentive constraints of one contractor, so that its chosen strategy will provide a better outcome for the manager. But although this method does guarantee a unique equilibrium, it is also costly to the manager.

A costless method of making the contractors choose the "correct" strategies is suggested in [63]. This method, however, makes the contracts more complicated. The main idea is that the manager offers one of the contractors, e.g., A , a range of extra possible output options $q_1^A(\epsilon)$, indexed by ϵ , where $0 < \epsilon \leq 1 - s_1^A$. If agent A chooses one of these options $q_1^A(\epsilon)$, then it essentially produces q_1^A , except that $q_1^A(\epsilon)$ has some inconsequential modification " ϵ " which is costless for agent A to effect. The importance of ϵ is that it acts as a signal that agent A sends to the manager:

"Agent B is cheating; from my perspective, the probability that B is choosing q_1^B is at least $s_1^A + \epsilon$."

In light of such a signal from agent A , if agent B chooses q_1^B , the manager will pay it an amount \bar{r}_1^B , where $v^B(\bar{r}_1^B) = s_2^B v^B(r_{11}^B) + (1 - s_2^B) v^B(r_{12}^B)$. That is, the manager pays agent B the equivalent of its expected utility as if it had observed θ_2^B . However,

Agent A's payments			
A's choice	B's choice		
	q_1^B	q_2^B	Refuse
q_1^A	r_{11}^A	r_{12}^A	r_{11}^A
$q_1^A(\varepsilon)$	$r_{11}^A + s(\varepsilon)$	$r_{12}^A - t(\varepsilon)$	$r_{11}^A + s(\varepsilon)$
q_2^A	r_2^A	r_2^A	$r_2^A - \gamma$
Refuse	-	-	-

Agent B's payments			
A's choice	B's choice		
	q_1^B	q_2^B	Refuse
q_1^A	r_{11}^B	r_2^B	-
$q_1^A(\varepsilon)$	\bar{r}_1^B	$r_2^B + \gamma$	-
q_2^A	r_{12}^B	r_2^B	-
Refuse	r_{11}^B	$r_2^B - \gamma$	-

Fig. 2. The contractors' payments according to Ma et al.'s mechanism for making the contractors choose the equilibrium preferred by the manager. $\gamma > 0$, $s(\varepsilon)$ and $t(\varepsilon)$ are continuous functions that are both strictly positive for $0 < \varepsilon < 1 - s_1^A$.

if agent *B* actually chooses q_2^B , and agent *A* signals some $\varepsilon > 0$ by choosing $q_1^A(\varepsilon)$, then agent *B* is compensated by receiving a higher payment ($r_2^B + \gamma$).³³ The details of the payments to the agents described in [63] are specified in Fig. 2. The continuous functions $s(\varepsilon)$ and $t(\varepsilon)$ that appear in *A*'s payments in Fig. 2 are both strictly positive for $0 < \varepsilon \leq 1 - s_1^A$, and satisfy

$$\begin{aligned}
 &(s_1^A + \varepsilon)(q_{11}^A - r_{11}^A + s(\varepsilon)) + (1 - s_1^A - \varepsilon)(q_{12}^A - r_{11}^A - t(\varepsilon)) \\
 &= (s_1^A + \varepsilon)(q_{11}^A - r_{11}^A) + (1 - s_1^A - \varepsilon)(q_{12}^A - r_{11}^A).
 \end{aligned}
 \tag{33}$$

The idea of this "reward scheme" is as follows. Consider contractor *A* which has observed θ_1^A . Suppose it assesses that agent *B* is choosing q_1^B more than *B* would be choosing q_1^B in the equilibrium preferred by the manager that is described above, e.g., with probability $\bar{s} > s_1^A$. Using construction (33), together with the fact that $s(\varepsilon)$ and $t(\varepsilon)$ are both positive, we can conclude that for all $0 < \varepsilon < (\bar{s} - s_1^A)$ agent *A* prefers to choose $q_1^A(\varepsilon)$ rather than q_1^A . On the other hand, if *B* chooses the output as in the equilibrium preferred by the manager, then *A* does not have an incentive to signal some $\varepsilon > 0$. The proof that this mechanism provides a unique equilibrium that guarantees the manager its second best outcome can be found in [63].

7.1.2. The contractor's effort influences others

In this section we consider situations where the output of a contractor depends both on its level of effort and the other contractors' levels of effort. In addition, there is

³³ The increase $\gamma > 0$ must not be too great, since it turns out that too high a compensation ($R_2^B + \gamma$) might admit unwanted equilibria. See [63] for more details.

symmetrical uncertainty about the state of the world. Suppose there are k possible contractors, and for each contractor i there is a finite set of possible outputs $Outcome^i = \{q_1^i, \dots, q_{n_i}^i\}$ and a finite set of possible effort levels, $Effort_i$. We denote the vector of the possible outcomes by $Outcome$, i.e., $Outcome = \{ \langle q^1, q^2, \dots, q^k \rangle \mid q^i \in Outcome^i \}$. The output of contractor i depends on some unknown (by all agents) features of the world θ_i , in addition to its level of effort and the other contractors' level of effort as we mentioned above. The outcome function is denoted by $f_i(e_1, \dots, e_k, \theta_i)$, and $\theta_1, \dots, \theta_k$ has a joint probability distribution $\phi(\theta_1, \dots, \theta_k)$. This probability distribution induces another probability distribution over vectors of outcomes, for any given vector of actions, as in Section 4.2. This means, that we extend \wp of Section 4.2 to fits the multi-contracted case; $\wp : Effort_1 \times Effort_2 \times \dots \times Effort_k \times Outcome \rightarrow \mathbb{R}$, such that for any e^1, \dots, e^k , $e^i \in Effort_i$, $\sum_{\bar{q} \in Outcome} \wp(e^1, \dots, e^k, \bar{q}) = 1$.

If the manager can observe the actions chosen by the contractors then, as in Section 4.1, it can offer the contractors a *forcing contract*. If the manager cannot observe the effort levels, then the contract it should offer will specify for any vector of outcomes $\langle q_{i_1}^1, \dots, q_{i_k}^k \rangle$, a vector of k rewards denoted by $\langle r_{i_1, \dots, i_k}^1, r_{i_1, \dots, i_k}^2, \dots, r_{i_1, \dots, i_k}^k \rangle$. Similar to the maximization problem in the case of one contractor, the manager should maximize its expected utility given similar constraints to (IC)(3) and (IR)(2). A three-step procedure, similar to the one contractor case of Section 4.2, can then be formalized. Given any effort level's vector e^1, \dots, e^k , the manager should find the rewards, r^1, \dots, r^k , that minimize the expected rewards it should pay the contractors, subject to the reservation utility constraint (IR)(5) and *participation* constraint (IC)(6) meaning, that given r^1, \dots, r^k , the contractors will prefer e^1, \dots, e^k over their other options. In some situations, depending on the probability function \wp (e.g., if there is perfect correlation between the θ_i), and the contractors' utility functions³⁴ the manager may gain similar expected utility as in the case where it can observe the agents' effort levels (i.e., as in a first best contract) [75].

In some situations, however, the contracts found by the above maximization problem may fail to uniquely implement the manager's preferred actions, as in the previous section. There may be other actions according to the contract that are better to the contractors, as in the previous section, where the agent's effort does not influence the others.

The main question is how the manager can make the contractors choose the set of actions it prefers. One approach is to try to strengthen the constraints that are related to the contractors, but this, of course, is costly for the manager. Another possibility, as in the previous section, is to construct a sophisticated contract. We may distinguish between two situations:

- (1) Actions are mutually observed by the contractors (but not by the manager).
- (2) Actions are only privately observed.

In the first case, the contractors pick an effort level simultaneously, and afterwards they (but not the manager) can observe each other's actions. There is some delay after the observation and the realization of the outcome, which is then used for message exchange. The manager can try to extract information about the effort levels from the agents and

³⁴ The exact restrictions on the contractors' utility functions and the environment can be found in [75].

although the contractor can provide false information, the accuracy of this information is known to the other contractor. The manager may then appeal to the other agents for verification.

We will consider the case where there are only two contractors [62], denoted by A and B , with the same utility function $U^c(e, r) = v(r) - c(e)$. Suppose by using the techniques of previous sections, and assuming the manager can observe the agent's actions, the manager would then like the two contractors to choose effort levels e_a^* and e_b^* , respectively, in order to maximize its own expected utility, taking into consideration their reservation prices. r_a^* can be the payments that will be awarded to contractor A if the manager can observe efforts, i.e., $U^c(e_a^*, r_a^*) = \hat{u}$, and similarly, r_b^* can be the reward for the second contractor. Note, that since $U^c(e, r) = v(r) - c(e)$, $v(r_a^*) = \hat{u} + c(e_a^*)$. The aim of the manager is to make sure the agents find (e_a^*, e_b^*) attractive and then the above utilities will be awarded in a unique equilibrium. The main idea is to ask A to report the effort levels chosen by the agents, and then ask B to confirm the report or to declare that A did not report honestly. We assume that for all $e_{a_k}, e_{a_m} \in Effort_A$, $e_{b_l}, e_{b_n} \in Effort_B$, $q^i \in Outcome^A$ and $q^j \in Outcome^B$, the following holds:³⁵

$$(\wp(e_{a_k}, e_{b_l}, \langle q^i, q^j \rangle))_{i,j} \neq (\wp(e_{a_m}, e_{b_n}, \langle q^i, q^j \rangle))_{i,j} \text{ whenever } (e_{a_k}, e_{b_l}) \neq (e_{a_m}, e_{b_n}). \tag{34}$$

Note that by condition (34), for any pair of effort levels (e_{a_k}, e_{b_l}) , where $e_{a_k} \in Effort_A$ and $e_{b_l} \in Effort_B$, are chosen by the agents, A cannot announce $(\hat{e}_a, \hat{e}_b) \neq (e_{a_k}, e_{b_l})$ with $(\wp(e_{a_k}, e_{b_l}, \langle q^i, q^j \rangle))_{i,j} = (\wp(\hat{e}_a, \hat{e}_b, \langle q^i, q^j \rangle))_{i,j}$. To see how the manager can use one agent's report to examine the truthfulness of another's, we will suppose A reports (\hat{e}_a, \hat{e}_b) where $\hat{e}_a \in Effort_A$ and $\hat{e}_b \in Effort_B$ concerning the pair of effort levels chosen by the contractors. Subsequently, B is allowed an opportunity to "challenge" A 's report. If B challenges, then it announces an alternative pair of effort levels. On reporting $(\tilde{e}_a, \tilde{e}_b)$ where $\tilde{e}_a \in Effort_A$ and $\tilde{e}_b \in Effort_B$ the manager uses the following function ε to give B an incentive to tell the truth. Let ε be a function such that $\varepsilon : Outcome \times Effort_A \times Effort_B \times Effort_A \times Effort_B \rightarrow \mathbb{R}$ and ε satisfies

$$\begin{aligned} \sum_{\langle q^i, q^j \rangle \in Outcome} \varepsilon(\langle q^i, q^j \rangle, \hat{e}_a, \hat{e}_b, \tilde{e}_a, \tilde{e}_b) \wp(\hat{e}_a, \hat{e}_b, \langle q^i, q^j \rangle) &< 0, \\ \sum_{\langle q^i, q^j \rangle \in Outcome} \varepsilon(\langle q^i, q^j \rangle, \hat{e}_a, \hat{e}_b, \tilde{e}_a, \tilde{e}_b) \wp(\tilde{e}_a, \tilde{e}_b, \langle q^i, q^j \rangle) &> 0. \end{aligned} \tag{35}$$

As we mentioned above, by condition (34), for any pair of effort levels (e_{a_k}, e_{b_l}) , A cannot announce $(\hat{e}_a, \hat{e}_b) \neq (e_{a_k}, e_{b_l})$ with

$$(\wp(e_{a_k}, e_{b_l}, \langle q^i, q^j \rangle))_{i,j} = (\wp(\hat{e}_a, \hat{e}_b, \langle q^i, q^j \rangle))_{i,j}.$$

Suppose the actual effort levels pair is (e_{a_k}, e_{b_l}) . B 's behavior will then depend on A 's report (\hat{e}_a, \hat{e}_b) . First, if A reports $(\hat{e}_a, \hat{e}_b) \neq (e_{a_k}, e_{b_l})$, then B prefers to challenge A . If B reports $(\tilde{e}_a, \tilde{e}_b) = (e_{a_k}, e_{b_l})$ then it gets additional expected reward (as described below) of $\sum_{\langle q^i, q^j \rangle \in Outcome} \varepsilon(\langle q^i, q^j \rangle, \hat{e}_a, \hat{e}_b, e_{a_k}, e_{b_l}) \wp(e_{a_k}, e_{b_l}, \langle q^i, q^j \rangle)$ which by (35) is

³⁵ Below, $(\wp(e_{a_k}, e_{b_l}, \langle q^i, q^j \rangle))_{i,j}$ denotes the vector of probabilities for any $\langle q^i, q^j \rangle \in Outcome$.

A's announces	$\hat{e}_a = a^*$	$\hat{e}_a \neq a^*$
B "agrees"	r_a^*	$\underline{r}^a - \delta$
B "challenges"	$r_a^* - \gamma$	$\underline{r}^a - \delta$

Fig. 3. A's rewards; $\delta > 0, \gamma > r_a^* - \underline{r}^a + \delta > 0$.

A's announces	$\hat{e}_b = b^*$	$\hat{e}_b \neq b^*$
B "agrees"	r_b^*	$\underline{r}^b - \delta$
B "challenges"	$r_a^* + \varepsilon(\langle q^j, q^j \rangle, \hat{e}_a, \hat{e}_b, \tilde{e}_a, \tilde{e}_b)$	$\underline{r}^b - \delta + \varepsilon(\langle q^j, q^j \rangle, \hat{e}_a, \hat{e}_b, \tilde{e}_a, \tilde{e}_b)$

Fig. 4. B's rewards; $\delta > 0, \gamma > r_a^* - \underline{r}^a + \delta > 0$.

positive. Second, if A has reported truthfully that $(\hat{e}_a, \hat{e}_b) = (e_{a_k}, e_{b_l})$, then by (35) the expected value for B from $\sum_{\langle q^j, q^j \rangle \in Outcome} \varepsilon(\langle q^j, q^j \rangle, e_{a_k}, e_{b_l}, \tilde{e}_a, \tilde{e}_b) \wp(e_{a_k}, e_{b_l}, \langle q^j, q^j \rangle)$ is negative. Hence, B will avoid (falsely) accusing A, and B's "challenge" is a signal to the contractor that A has been lying. Using these ideas the manager should offer the following mechanism:

Stage 1: Both contractors take actions simultaneously.

Stage 1+: Contractors observe each other's action.

Stage 2: Agent A announces a pair of effort levels: (\hat{e}_a, \hat{e}_b) where $(\hat{e}_a, \hat{e}_b) \in Effort_A \times Effort_B$.

Stage 3: Agent B can either "agree" or "challenge". If B "challenges" A's announcement then it announces $(\tilde{e}_a, \tilde{e}_b)$ where $(\tilde{e}_a, \tilde{e}_b) \in Effort_A \times Effort_B$ but $(\tilde{e}_a, \tilde{e}_b) \neq (\hat{e}_a, \hat{e}_b)$.

The rewards, as a function of the outputs q^i and q^j , are described in Figs. 3 and 4. We denote the reward that satisfies $v(\underline{r}^a) = \hat{u} + \min_{e_a} \{c(e_a) \mid e_a \in Effort_A\}$ by \underline{r}^a , and similarly for B we denote the reward that satisfies $v(\underline{r}^b) = \hat{u} + \min_{e_b} \{c(e_b) \mid e_b \in Effort_b\}$ by \underline{r}^b .

It was shown in [62] that the following strategies form a unique perfect equilibrium of the described mechanism: Agent A chooses e_a^* at Stage 1, and reports honestly at Stage 2 which action pair was chosen at Stage 1. Agent B chooses e_b^* at Stage 1 and "agrees" at Stage 3 if and only if A is honest at Stage 2. The intuition behind this proof is as follows. The manager elicits information from agent A and uses B's reaction as a policing device as we explained above. If B accuses A of lying, then its outcome depends on ε . However, due to assumption (35), the expected outcome from ε to B is valuable if and only if A has lied. In addition, given that the contractors report honestly, the rewards will motivate them to choose the required actions. These results can easily be extended to the case of more than two contractors [62].

In the case that actions are only privately observed, it is not possible to implement the results of perfect observation (i.e., the first best contract, where the result is that the manager observes the contractors' actions). However, even the implementation of the second best is not so simple. The rewards that were suggested in the beginning of the section are appropriate *only if* the agents follow the actions prescribed by the manager. It is possible, however, that the contractors may be better off (given the

suggested rewards) if they *all* deviated from the required actions. In [62] a multi-stage mechanism is presented that makes the contractors choose the appropriate actions of the second best contract.

7.2. Individual outcome is not observed

There are other situations in which the manager cannot observe the individual outcome (or such an outcome does not exist), but rather can only observe the overall outcome of all the agents' efforts [42, 87]. Even in the case of certainty, i.e., the state of the world is known, there is a problem in making the contractors take the preferred level of action, since there is no way for the manager to find out the effort level of each of the individual agents, given the overall output. For example, suppose two robots have agreed to collect garbage, but they both put the garbage in the same truck; it is not possible to then figure out who collected what. If the manager wants the contractors to take the vector of effort level e^* , then it can search for a contract such that if the outcome is $q \geq q(e^*)$, then $r_i(q) = b_i$ and otherwise 0, such that $U^c(e_i^*, b_i) \geq \hat{u}_i$. That is, if all agents choose the appropriate effort level, each of them gets b_i , and if any of them does not, none of them gets anything.

In some cases the contractors take sequential actions. That is, agent 1 chooses its effort level and performs its part of the task which is observed by the other contractors, but not the manager. The second contractor then, chooses its effort level, based on the first agent's actions, and its effort level is observed by the other contractors, and so on. After the last agent finishes its part, the outcome of the whole vector is figured out and observed by all agents (including the manager). If, in addition, there is also some uncertainty in the environment, the outcome function may be similar to the one presented in Section 7.1.1: $f(e^1, \dots, e^n) = z(e^1, \dots, e^2) + \varepsilon$. If, no matter how low the effort levels exerted by contractors $1, \dots, i$ are, it is possible for the rest of the agents $i + 1, \dots, n$ to compensate for the slack and if for fixed effort levels e^1, \dots, e^i , z is a monotonic function of the effort level of the rest of the contractors, then the manager can construct a contract in which it can obtain its first best outcome [7]. The contract enables agent i , whose choice of effort level is a function of the effort levels of agents $1, \dots, i - 1$, to use its monitoring capability effectively.

Another interesting situation is when a group of contractors can commit themselves to cooperate. Although they can still be individually motivated if they can agree upon a cooperation level, the outcome (under appropriate conditions) can be better for all of them. An even more efficient result may be obtained if the contractors work as a team and share the outcome. Such a situation may occur, for example, if all the contractors are robots of CompC, that have the same general task to maximize CompC's profits [64].

8. Conclusions

In this paper we presented techniques that can be used in different cases where incentive contracting of a task by an agent to another agent or a set of agents in non-

collaborative environments is beneficial. These techniques are useful when the contractor can choose an effort level to carry out the task, and the manager tries to find an incentive to convince the contractor to choose the effort level that the manager prefers. We considered several such situations and described the maximization problems that should be solved by the manager in order to design a beneficial contract for itself. The contractor's computational task is easier than that of the manager. In most of the situations, given a contract, the contractor needs only to check the validity of the inequalities that appear as constraints in the manager's maximization problem. The contractor needs to check the validity of the individual rationality constraint (IR) in order to decide whether to accept the contract, and since all variables are known, based on the suggested contract, this check is very easy. When the contractor needs to decide which effort level to provide, it should consider its expected utility from its effort level, similar to the maximization problem described in the *participation* constraints (IC).

The maximization problems the manager needs to solve are much more difficult. In most of the situations we presented procedures that can be used as the basis for solving these maximization problems. In general, solutions of optimization problems by a single, all purpose, method is cumbersome and inefficient. Optimization problems are therefore classified into particular categories, where each category is defined by the objective function of the maximization and the constraints; special purpose procedures were developed for each case. Currently, there are several computer optimization packages available using a variety of practical optimization methods [22] that can be used for automating those procedures. The designer of the automated agent should build an interface between the chosen package and its agent's software.

The agents' utility functions influence the efficiency of the subcontracting itself and the computation time required for finding efficient contracts and solving the maximization problems. It is clear, that when the agents are risk neutral, all the maximization problems presented in this paper are much easier to solve. In this case the objective function of the maximization problem, as well as its constraints, are linear, and there is a polynomial algorithms to solve the maximization problem. Furthermore, more efficient results are obtained in such situations.

However, if the designer would like its agent to be risk averse, then not all utility functions are appropriate for incentive contracting. In order to support most of the results presented in this paper, the contractor's utility function should be additively separable in rewards and efforts in the form $U^c(e, r) = v(r) - c(e)$ where $v' > 0$, $v'' \leq 0$, $c' > 0$ and $c'' \geq 0$. However, a large set of utility functions satisfies these requirements, and these properties of the utility functions seem useful also in other settings, therefore, it seems reasonable that agents' utility functions will satisfy these conditions. In these cases, the objective function of the maximization function may be nonlinear, as well as the constraints. The library routines in available computer packages for solving such maximization problems, generate an iterative sequence that converges in to the solution in the limit.³⁶ The agent that uses the routines may stage the convergence conditions

³⁶ In all the cases that we considered, if the utility functions satisfy the above conditions, there exist solutions to the maximization problems described in the paper.

that will fit its computation and time limitation. Below we present a summary of the results for the different situations considered in this paper. The results of contracting with symmetric information situations are as follows:

- (1) If the manager can observe the contractor's actions (Section 4), then it can force the contractor to provide the effort level preferred by the manager, and thus the manager maximizes its utility, and the contractor obtains its reservation price.
- (2) If the manager does not observe the contractor's actions, but there is full information and no uncertainty concerning the outcome of the contractor's actions (Section 4.1), then the expected utility to both agents is as in the previous case. That is, in this situation, there is no need for the manager's observation.
- (3) If there is uncertainty in the environment but the contractor is risk neutral (Section 4.2.1), then the manager's utility will be as in the previous two cases (i.e., the agents reach a *first best* contract). The *expected* utility of the contractor will be equal to its reservation price; however, its actual outcome may be less or more than its reservation price.
- (4) If there is uncertainty as in the previous case, but the contractor is risk averse (Section 4.2.2), then the manager's expected utility will be lower than in the previous case (i.e., the agents reach a *second best* contract). The contractor's expected utility is higher than its reservation price.
- (5) Monitoring (Section 4.2.3) cannot improve the manager's utility in case (3) above, but it may increase its utility in case (4), when the contractor is risk averse.

If there is asymmetric information then the contracts should include a menu of options and there is a need for the exchange of messages. However, in all the situations, the agents can consider only contracts in which it is in the interest of the contractor to honestly report its private information. Below is a summary of the results of the main cases in asymmetric information situations:

- (1) If the contractor knows the state of the world but the manager does not (Section 5.1), then the manager's expected utility is lower than if they have symmetric beliefs and the contractor's expected utility is higher.
- (2) If the contractor is able to collect more information *before* it performs the agreed upon task but only after signing the contract, and the contractor cannot opt out after signing an agreement (Section 5.2), then the manager can get its second best utility if the contractor is risk neutral.
- (3) If the manager also has private information (Section 5.5), but its private information does not directly influence the contractor's utilities, then in most of the situations, there exists a mechanism in which all types of managers do strictly better than the fully informed contractor (i.e., even better than in the first best contract).
- (4) If there are several agents in the environment (Section 5.7), then in most situations, the manager can design a *second best* contract.

When there is more than one encounter between the agents (Section 6), then they can reach either short term contracts or enforceable long term contracts. The contracts in the first case are similar to those of one encounter; however, the strategies used by the agents are more complicated.

- (1) If the agents agreed upon short term contracts, and the number of encounters is large enough, even in asymmetric information situations, they can reach first best contracts.
- (2) If the number of encounters is small, then enforceable long term encounters are more beneficial to the manager. However, it is still difficult to design an efficient contract.

The last set of situations considered in this paper are of contracting to a group. The type of contracts that are used depend on the following factors: Whether the individual outcome of each contractor is observed by the manager, whether the effort level of one contractor influences the other agents' outcome, and whether each of the contractors possesses private information. In some of these situations an efficient contract for the manager may be quite complicated and may require two rounds of message exchanges.

We are now in the process of applying the techniques presented in this paper to the performance of robots in a simulated environment.

References

- [1] D. Abreure, P. Dutta and L. Smith, The folk theorem for repeated games: a new condition, *Econometrica* **62** (1994) 939–948.
- [2] K.J. Arrow, The economics of agency, in: J. Pratt and R. Zeckhauser, eds., *Principals and Agents: The Structure of Business* (Harvard Business School Press, Cambridge, MA, 1985) 37–51.
- [3] K.J. Arrow and F.H. Hahn, *General Competitive Analysis* (Holden-Day, San Francisco, CA, 1971).
- [4] K.J. Arrow, L. Hurwicz and H. Uzawa, *Studies in Linear and Non-Linear Programming* (Stanford University Press, Stanford, CA, 1958).
- [5] R.J. Aumann, *Survey of Repeated Games. Essays in Game Theory and Mathematical Economics in Honor of Oskar Morgenstern* (Wissenschaftsverlag, Bibliographisches Institut, Mannheim, 1981).
- [6] S. Baiman and J. Demski, Economically optimal performance evaluation and control systems, *J. Accounting Res.* **18** (1980) 184–220.
- [7] A. Banerjee and A. Beggs, Efficiency in hierarchies: implementing the first-best solution by sequential actions, *Rand J. Econ.* **20** (1989) 637–645.
- [8] A.H. Bond and L. Gasser, An analysis of problems and research in DAI, in: A.H. Bond and L. Gasser, eds., *Readings in Distributed Artificial Intelligence* (Morgan Kaufmann, San Mateo, CA, 1988) 3–35.
- [9] B. Caillaud, R. Guesnerie, P. Rey and J. Tirole, Government intervention in production and incentives theory: a review of recent contributions, *Rand J. Econ.* **19** (1988) 1–26.
- [10] J. Christensen, Communication in agencies, *Bell J. Econ.* **12** (1981) 661–674.
- [11] C. Congdon, M. Huber, D. Kortenkamp, K. Konolige, K. Myres, A. Saffiotti and E. Ruspini, Carmel versus Flakey a comparison of two winners, *AI Magazine* **14** (1993) 49–57.
- [12] S. Conry, D.J. MacIntosh and R. Meyer, DARES: a Distributed Automated REasoning System, in: *Proceedings AAAI-90*, Boston, MA (1990) 78–85.
- [13] D. Corkill, Hierarchical planning in a distributed environment, in: *Proceedings IJCAI-79*, Tokyo (1979) 168–175.
- [14] D. Demougis, A renegotiation-proof mechanism for a principal-agent model with moral hazard and adverse selection, *Rand J. Econ.* **20** (1989) 256–267.
- [15] J.S. Demski and D. Sappington, Optimal incentive contracts with multiple agents, *J. Econ. Theory* **33** (1984) 152–171.
- [16] E.J. Douglas, The simple analytics of the principal-agent incentive contract, *J. Econ. Educ.* **20** (1989) 39–51.
- [17] J. Doyle, Rationality and its role in reasoning, *Comput. Intell.* **8** (1992) 376–409.
- [18] E.H. Durfee, *Coordination of Distributed Problem Solvers* (Kluwer Academic Publishers, Boston, MA, 1988).

- [19] E.H. Durfee, What your computer really needs to know, you learned in kindergarten, in: *Proceedings AAAI-92*, San Jose, CA (1992) 858–864.
- [20] E. Ephrati and J.S. Rosenschein, The clarke tax as a consensus mechanism among automated agents, in: *Proceedings AAAI-91*, Anaheim, CA (1991) 173–178.
- [21] E. Ephrati and J.S. Rosenschein, Planning to please: following another agent's intended plan, *J. Group Decision Negotiation* **2** (1993) 219–235.
- [22] R. Fletcher, *Practical Methods of Optimization* (Wiley, New York, 1987).
- [23] S. French, *Decision Theory: An Introduction to the Mathematics of Rationality* (Ellis Horwood, Chichester, 1986).
- [24] D. Fudenberg and D. Levine, Limit games and limit equilibrium, *J. Econ. Theory* **38** (1986) 261–279.
- [25] D. Fudenberg and E. Maskin, The folk theorem in repeated games with discounting or incomplete information, *Econometrica* **54** (1986) 533–554.
- [26] D. Fudenberg and J. Tirole, Moral hazard and renegotiation in agency contracts, *Econometrica* **58** (1990) 1279–1319.
- [27] D. Fudenberg, J. Tirole and P. Milgrom, Short-term contracts and long-term relationships, *J. Econ. Theory* **51** (1990) 1–3.
- [28] L. Gasser, Social concepts of knowledge and action: DAI foundations and open systems semantics, *Artif. Intell.* **47** (1991) 107–138.
- [29] L. Gasser, Social knowledge and social action, in: *Proceedings IJCAI-93*, Washington, DC (1993) 751–757.
- [30] F. Gjesdal, Information and incentives: the agency information problem, *Rev. Econ. Stud.* **49** (1982) 373–390.
- [31] S. Grossman and O. Hart, An analysis of the principal-agent problem, *Econometrica* **51** (1983) 7–45.
- [32] B.J. Grosz and S. Kraus, Collaborative plans for complex activities, *Artif. Intell.* (1996); also: Tech. Report TR-20-95, Harvard University, Center for Research and Computing Technology (1995).
- [33] P. Haddaway and S. Hanks, Issues in decision-theoretic planning: symbolic goals and numeric utilities, in: *Proceedings Workshop on Innovative Approaches to Planning, Scheduling and Control*, San Diego, CA (1990) 48–58.
- [34] M. Harris and A. Raviv, Some results on incentive contracts with applications to education and employment, health insurance, and law enforcement, *Am. Econ. Rev.* **68** (1978) 20–30.
- [35] M. Harris and A. Raviv, Optimal incentive contracts with imperfect information, *J. Econ. Theory* **20** (1979) 231–259.
- [36] M. Harris and R. Townsend, Resource allocation under asymmetric information, *Econometrica* **49** (1981) 33–64.
- [37] J. Harsanyi, Games with incomplete information played by bayesian players, *Manage. Sci.* **14** (1967/1968) 159–182, 320–334, 486–502.
- [38] O.D. Hart and J. Tirole, Contract renegotiation and Coasian dynamics, *Rev. Econ. Stud.* **55** (1988) 509–540.
- [39] G. Heal, Planning without prices, *Rev. Econ. Stud.* **36** (1969) 346–362.
- [40] J. Hirshleifer and J. Riley, *The Analytics of Uncertainty and Information* (Cambridge University Press, Cambridge, 1992).
- [41] Y. Ho, L. Servi and R. Suri, A class of center-free resource allocation algorithms, *Large Scale Syst.* **1** (1980) 51–62.
- [42] B. Holmstrom, Moral hazard in teams, *Bell J. Econ.* **13** (1982) 324–340.
- [43] B. Holmstrom and P. Milgrom, Aggregation and linearity in the provision of intertemporal incentives, *Econometrica* **55** (1987) 303–328.
- [44] I. Horswill, Polly: A vision-based artificial agent, in: *Proceedings AAAI-93*, Washington, DC (1993) 824–829.
- [45] L. Hurwicz, The design mechanisms for resource allocation, *Am. Econ. Rev.* **63** (1963) 1–30.
- [46] R. Keeney and H. Raiffa, *Decisions with Multiple Objectives: Preferences and Value Tradeoffs* (Wiley, New York, 1976).
- [47] S. Kraus, Agents contracting tasks in non-collaborative environments, in: *Proceedings AAAI-93*, Washington, DC (1993) 243–248.

- [48] S. Kraus and D. Lehmann, Designing and building a negotiating automated agent, *Comput. Intell.* **11** (1995) 132–171.
- [49] S. Kraus, M. Nirkhe and K.P. Sycara, Reaching agreements through argumentation: a logical model, in: *Proceedings DAI93* (1993) 233–247; also presented in AAAI-93 workshop on AI theories of Groups and Organizations: Conceptual and Empirical Research.
- [50] S. Kraus and J. Wilkenfeld, The function of time in cooperative negotiations, in: *Proceedings AAAI-91*, Anaheim, CA (1991) 179–184.
- [51] S. Kraus and J. Wilkenfeld, Negotiations over time in a multi agent environment: Preliminary report, in: *Proceedings IJCAI-91*, Sydney, Australia (1991) 56–61.
- [52] S. Kraus, J. Wilkenfeld and G. Zlotkin, Multiagent negotiation under time constraints, *Artif. Intell.* **75** (1995) 297–345.
- [53] J. Kurose and R. Simha, A microeconomic approach to optimal resource allocation in distributed computer systems, *IEEE Trans. Comput.* **38** (1089) 705–717.
- [54] J. Laffont, E. Maskin and J. Rochet, Optimal nonlinear pricing with two-dimensional characteristic, in: T. Groves, R. Radner and S. Reiter, eds., *Information, Incentives, and Economic Mechanisms* (University of Minnesota Press, Minneapolis, MN, 1987) 256–266.
- [55] J. Laffont and J. Tirole, Adverse selection and renegotiation in procurement, *Rev. Econ. Stud.* **57** (1990) 597–625.
- [56] J. Laffont and J. Tirole, *A Theory of Incentives in Procurement and Regulation* (MIT Press, Cambridge, MA, 1993).
- [57] R.A. Lambert, Long term contracts and moral hazard, *Bell J. Econ.* **14** (1983) 441–452.
- [58] M. Landsberger and I. Meilijson, Monopoly insurance under adverse selection when agents differ in risk aversion, *J. Econ. Theory* **63** (1994) 392–407.
- [59] V. Lesser, A retrospective view of FA/C distributed problem solving, *IEEE Trans. Syst. Man Cybern.* **21** (1991) 1347–1362.
- [60] V.R. Lesser, Distributed problem solving, in: S.C. Shapiro, ed., *Encyclopedia of Artificial Intelligence* (Wiley, New York, 1990) 245–251.
- [61] V.R. Lesser and L.D. Erman, Distributed interpretation: a model and experiment, *IEEE Trans. Comput.* **29** (1980) 1144–1163.
- [62] C. Ma, Unique implementation of incentive contracts with many agents, *Rev. Econ. Stud.* **51** (1984) 555–572.
- [63] C. Ma, J. Moore and S. Turnbull, Stopping agents from “cheating”, *J. Econ. Theory* **46** (1988) 355–372.
- [64] I. Macho-Stadler and J. Pérez-Castrillo, Moral hazard and cooperation, *Econ. Lett.* **35** (1991) 17–20.
- [65] J. Malcomson and F. Spinnewyn, The multiperiod principal-agent problem, *Rev. Econ. Stud.* **55** (1988) 391–408.
- [66] E. Malinvaud, Decentralized procedures for planning, in: *Activity Analysis in the Theory of Growth and Planning* (St Martin’s Press, New York, 1967) 170–208.
- [67] T.W. Malone, R.E. Fikes, K.R. Grant and M.T. Howard, Enterprise: a marketlike task schedule for distributed computing environments, in: B.A. Huberman, ed., *The Ecology of Computation* (North-Holland, Amsterdam, 1988) 177–205.
- [68] E. Maskin and J. Tirole, The principal-agent relationship with an informed principal: the case of private values, *Econometrica* **58** (1990) 379–409.
- [69] E. Maskin and J. Tirole, The principal-agent relationship with an informed principal II: common values, *Econometrica* **60** (1992) 1–42.
- [70] S. Matthews, Selling to risk averse buyers with unobservable tastes, *J. Econ. Theory* **30** (1983) 370–400.
- [71] R. McAfee and J. McMillan, Multidimensional incentive compatibility and mechanism design, *J. Econ. Theory* **46** (1988) 335–354.
- [72] R.P. McAfee and J. McMillan, Bidding for contracts: a principal-agent analysis, *Rand J. Econ.* **17** (1986) 326–338.
- [73] R.P. McAfee and J. McMillan, Competition for agency contracts, *Rand J. Econ.* **18** (1987) 296–307.
- [74] N.D. Melumad and S. Reichelstein, Value of communication in agencies, *J. Econ. Theory* **47** (1989) 334–368.
- [75] D. Mookherjee, Optimal incentive schemes with many agents, *Rev. Econ. Stud.* **51** (1984) 433–446.

- [76] R. Myerson, Optimal coordination mechanisms in generalized principal-agent problem, *J. Math. Econ.* **10** (1982) 67–81.
- [77] R. Myerson, Mechanism design by an informed principal, *Econometrica* **51** (1983) 1767–1798.
- [78] B. Nalebuff and J. Stiglitz, Information, competition, and markets, *Am. Econ. Rev.* **73** (1983) 278–283.
- [79] B. Nalebuff and J. Stiglitz, Prizes and incentives: towards a general theory of compensation and competition, *Bell J. Econ.* **14** (1983) 21–43.
- [80] I. Nourbakhsh, S. Morse, C. Becker, M. Balabanovic, E. Gat, R. Simmons, S. Goodridge, H. Potlapalli, D. Hinkle, K. Jung and D.V. Vactor, The winning robots from the 1993 robot competition, *AI Magazine* **14** (4) (1993) 51–62.
- [81] F. Page, The existence of optimal contracts in the principal-agent model, *J. Math. Econ.* **16** (1987) 157–167.
- [82] H. Pattison, D. Corkill and V. Lesser, Instantiating descriptions of organizational structures, in: M.N. Huhns, ed., *Distributed Artificial Intelligence* (Pitman/Morgan Kaufman Publishers, London/San Mateo, CA, 1987) 59–96.
- [83] R. Pfaffenberger and D. Walker, *Mathematical Programming for Economics and Business* (IOWA State University Press, Ames, IA, 1976).
- [84] R.H. Porter, Optimal cartel trigger price strategies, *J. Econ. Theory* **29** (1983) 313–338.
- [85] R. Radner, Monitoring cooperative agreements in a repeated principal-agent relationship, *Econometrica* **49** (1981) 1127–1148.
- [86] R. Raciner, Repeated principal agents games with discounting, *Econometrica* **53** (1985) 1173–1198.
- [87] E. Rasmusen, Moral hazard in risk-averse teams, *Rand J. Econ.* **18** (1987) 324–340.
- [88] E. Rasmusen, *Games and Information* (Basil Blackwell, Cambridge, MA, 1989).
- [89] W. Rogerson, The first-order approach to the principle-agent problems, *Econometrica* **53** (1985) 1357–1367.
- [90] J.S. Rosenschein, Synchronization of multi-agent plans, in: *Proceedings AAAI-82*, Pittsburgh, PA (1982) 115–119.
- [91] S. Ross, The economic theory of agency: the principal's problem, *Am. Econ. Rev.* **63** (1973) 134–139.
- [92] S. Ross, Equilibrium and agency-inadmissible agents and in the public agency problem, *Am. Econ. Rev.* **69** (1979) 308–312.
- [93] A. Rubinstein and M. Yaari, Repeated insurance contracts and moral hazard, *J. Econ. Theory* **30** (1983) 74–97.
- [94] T. Sandholm, An implementation of the contract net protocol based on marginal cost calculations, in: *Proceedings AAAI-93*, Washington, DC (1993) 256–262.
- [95] D. Sappington, Limited liability contracts between principle and agent, *J. Econ. Theory* **29** (1983) 1–21.
- [96] D. Sappington, Incentive contracting with asymmetric and imperfect precontractual knowledge, *J. Econ. Theory* **34** (1984) 52–70.
- [97] R. Selten, Re-examination of the perfectness concept for equilibrium points in extensive games, *Int. J. Game Theory* **4** (1975) 25–55.
- [98] S. Shavell, Risk sharing and incentives in the principal and agent relationship, *Bell J. Econ.* **10** (1979) 55–79.
- [99] R. Smith, The contract net protocol: high-level communication and control in a distributed problem solver, *IEEE Trans. Comput.* **29** (1980) 1104–1113.
- [100] R. Smith and R. Davis, Framework for cooperation in distributed problem solvers, *IEEE Trans. Syst. Man Cybern.* **29**(12) (1981) 61–70.
- [101] R. Smith and R. Davis, Negotiation as a metaphor for distributed problem solving, *Artif. Intell.* **20** (1983) 63–109.
- [102] M. Spence and R. Zeckhauser, Insurance, information and individual action, *Am. Econ. Rev.* **61** (1971) 380–391.
- [103] W. Spivey and R. Thrall, *Linear Optimization* (Holt, Rinehart and Winston, New York, 1970).
- [104] K. Sycara, Resolving adversarial conflicts: an approach to integrating case-based and analytic methods, Ph.D. Thesis, School of Information and Computer Science, Georgia Institute of Technology, Atlanta, GA (1987).
- [105] K.P. Sycara, Persuasive argumentation in negotiation, *Theory Decis.* **28** (1990) 203–242.

- [106] C. Waldspurger, T. Hogg, A. Huberman, J. Kephart and W. Stornetta, Spawn: a distributed computational economy, *IEEE Trans. Softw. Eng.* **18** (1992) 103–117.
- [107] M.P. Wellman, A general-equilibrium approach to distributed transportation planning, in: *Proceedings AAAI-92*, San Jose, CA (1992) 282–289.
- [108] M.P. Wellman and J. Doyle, Modular utility representation for decision-theoretic planning, in: *Proceedings AI Planning Systems (1992)* 236–242.
- [109] E. Werner, Toward a theory of communication and cooperation for multiagent planning, in: *Proceedings Second Conference on Theoretical Aspects of Reasoning about Knowledge*, Pacific Grove, CA (1988) 129–143.
- [110] G. Zlotkin and J.S. Rosenschein, Incomplete information and deception in multi-agent negotiation, in: *Proceedings IJCAI-91*, Sydney, Australia (1991) 225–231.
- [111] G. Zlotkin and J.S. Rosenschein, A domain theory for task oriented negotiation, in: *Proceedings IJCAI-93*, Chambery (1993) 416–422.