CS-TR-3254 UMIACS-TR-94-44 **April 1994**

Contracting Tasks in Multi-Agent Environments*

Sarit Kraus
Institute for Advanced Computer Studies,
University of Maryland, College Park and
Department of Mathematics and Computer Science
Bar Ilan University Ramat Gan, 52900 Israel
sarit@umiacs.umd.edu

Abstract

Agents may contract some of their tasks to other agent 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 to reach contracting in varied situations: the agents have full information about the environment and each other or subcontracting when the 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 where there is competition among possible contracted agents or possible contracting agents are also considered. In all situations we would like the contracted agent to carry out the task efficiently without the need of close supervision by the contracting agent.

*This paper is based upon work supported by the National Science Foundation under Grant No. IRI-9123460 and is an extension of [Kraus, 1993]. I would like to thank Jonathan Wilkenfeld and Barbara Grosz for their comments.

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 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 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 main question addressed in this paper is how one agent may 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 would like the contracted agent to carry out the task efficiently without the need of close supervision of the contracting agent, enabling the contracting agent 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 tasks: by threatening to interfere with the agent carrying out its own tasks, or by promising rewards [Kraus *et al.*, 1993]. This paper concentrates on subcontracting by rewards which may be in two forms. The first approach is a bartering system, where one agent may promise to help the other with its future tasks in return for current help. However, as was long ago observed in economics, barter is not an efficient basis for cooperation. In particular, in a multi-agent environment, an agent that wants 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 utilized for other purposes.

In this paper we present an automated multiagent model where contracting is beneficial. We propose to use a monetary system in the 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 multiagent. 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 every agent to fulfill its goals more efficiently as opposed to working on its own. One of the advantageous of contracting is that there usually is no need for negotiation among the agents and therefore communication is limited.

The issue of contracting by rewards has been investigated in economics and game-theory for the last two decades (e.g., [Arrow, 1985; Ross, 1973; Raizsmusen, 1989;

Grossman and Hart, 1983; Hirshleifer and Riley, 1992]). These works consider different types of contracts for different applications. Examples of these are contracts between: a firm, and an employer or employers (e.g., [Nalebuff and Stiglitz, 1983a; Baiman and Demski, 1980; Banerjee and Beggs, 1989; Macho-stadler and Pérez-Castrillo, 1991]); a government and taxpayers (e.g., [Caillaud *et al.*, 1988]); a landlord and a tenant (e.g., [Arrow, 1985]); an insurance company and a policy holder (e.g., [Rubinstein and Yaari, 1983; Harris and Raviv, 1978; Spence and Zeckhauser, 1971]); a buyer and a seller (e.g., [Matthews, 1983; Myerson, 1983]); a government and firms (e.g., [McAfee and McMillan, 1986]); stockholders and managements (e.g., [Arrow, 1985]); a professional and a client [Shavell, 1979], etc. In these situations two parties are usually be found. The first party (called "the agent" in the economics literature), must choose an action 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 multiagent 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 found appropriate economics mechanism and techniques that can be used for contracting in the automated agents environment. 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. The contracting agent does not supervise the contracted agents' performance and by using the appropriate techniques, the agent that designs the contract maximizes its personal expected utilities, given the constraints of the other agent(s). Throughout the paper, we use a robotics domain to demonstrate the contracting techniques introduced above.

2 Related work in DAI

Research in DAI is divided into two basic classes: *Distributed Problem Solving (DPS)* and *Multi-Agent Systems (MA)* [Bond and Gasser, 1988; Gasser, 1991]. Research in DPS (e.g.,

[Lesser and Erman, 1980; Lesser, 1991; Conry *et al.*, 1990; Smith and Davis, 1983; Durfee, 1988]) considers how the work involved in solving a particular problem can be divided among a number of modules or "nodes." The modules in a DPS system are centrally designed to improve performance, stability, modularity, and/or reliability. They include the development of cooperation mechanisms designed to find a solution to a given problem.

Research in MA (e.g., [Sycara, 1987; Zlotkin and Rosenschein, 1991; Kraus and Lehmann, 1994; Ephrati and Rosenschein, 1991]) 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, no globally consistent knowledge, and no globally shared goals or success criteria. There is, however, a possibility for *real competition* among the agents.

In DPS, on the other hand, there are no implicit conflicts among the agents, and it is assumed that it is in the agents' interest to help one another. This help can be in the form of sharing tasks, results, or information [Durfee, 1992]. 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 [Smith and Davis, 1983]. This approach assumes that agents not otherwise occupied will readily take on the task. Similarly, results and information are shared among agents in such environments with no expectation of reciprocation [Lesser and Erman, 1980; Lesser, 1991; Conry *et al.*, 1990]. 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 as it can, thereby giving the agents shared, often implicit, global goals that they are all unselfishly committed to achieving.

Contracting, in particular, was previously used in the Distributed Problem Solvers framework for tasks allocation. In the Contract Net protocol [Smith and Davis, 1981; Smith and Davis, 1983], 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 advertises the task's existence to other agents. Available agents (potential contractors) then evaluate the task announcements made by several managers and submit bids for which they are suited. 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. In contrast, since agents are self motivated in our work, a contract must specify the fee to be paid to the contracted agent as a function of the contracting agent's observations.

The "benevolent" agents are also taken into account in Malone's refinement of the

contract-net protocol. The refined protocol is based on a more sophisticated economic model [Malone *et al.*, 1988], proving optimal behavior under certain conditions. Similar to the original Contract Net protocol, also in this model there is no need to motivate the agents to bid or to make decisions in order to maximize the global expected utility of the system.

A modified version of the Contract Net protocol for competitive agents in the transportation domain is presented in [Sandholm, 1993]. It provides a formalization of the bidding and awarding decision process based on marginal cost calculating on local agent criteria. In particular, an agent will send a bid for a delivery task only if the maximum price mentioned in the task 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. In contrast to our framework, it is assumed in [Sandholm, 1993] that all agents centrally designed will carry out their contracted tasks efficiently and without monitoring. Furthermore, there are situations where an agent in [Sandholm, 1993] is forced to accept a non beneficial contract. We assume on the other hand, that the agents are heterogeneous and self motivated.

In [Pattison *et al.*, 1987] a language for specification of complex relations among agents in DPS is described. 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. suggested an additional mechanism of contracting to the one presented in the Contract Net protocol using *focused addressing*. This would mean, in addition to broadcasting requests for bids, an agent in [Pattison *et al.*, 1987] 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 knows the other agent's address. In this paper, we also allow both of these addressing methods.

Subcontracting in 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 [Rosenschein, 1982; Corkill, 1979; Lesser, 1990]. Werner [Werner, 1988] presents a formal logical model for a master-slave relationship by one-way communication. The main problem for a master is not to convince other agents to carry out the plan appropriately without its supervision, but rather finding the best plan and synchronizing the agent's actions. The simple master/slaves model was extended by Ephrati and Rosenschein [Ephrati and Rosenschein, 1993] to allow the "slaves" more freedom in

carrying out the plans. However, the slaves' main goal is still to satisfy their master's wishes.

Wellman [Wellman, 1992] proposes the use of 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 small in respect to the entire economy. We consider contracting when there are usually a small number of agents in the environment. We also deal with situations where agents are uncertain about the world, and the contracted agents (the producers in Wellman's terminology) may not carry out the tasks as promised.

Negotiation is the main paradigm that is used for coordination and task distribution in multi-agent systems (MA) where agents are self motivated. While contracting is most appropriate for a hierarchical relationship, negotiation is most appropriate when all agents are on the same level. For example, Sycara [Sycara, 1990; Sycara, 1987] presents a model of negotiation that combines case-based reasoning and optimization of the multi-attributed utilities. This model is used in labor management negotiations where two agents need to agree on an acceptable agreement. In [Kraus *et al.*, 1994; Kraus and Wilkenfeld, 1993] a set of self-motivated autonomous agents have a 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. Zlotkin and Rosenschein [Zlotkin and Rosenschein, 1993] present a general theoretical negotiation model for rational agents who are in a similar status.

Negotiation is communication consuming, therefore, the above models are appropriate when communication is not expensive. The contracting model that we present in this paper usually requires only one round of exchanging messages, but requires intensive computation, and therefore, is more appropriate when communication is expensive and computation resources are available.

Contracting in multi-agent systems was previously studied in [Grosz and Kraus, 1993]. 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 believe 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 contracted agent is not presented in [Grosz and Kraus, 1993]. The main contribution of this paper is the presentation of techniques for identifying possible contracted agents and to drafting beneficial contracts (i.e., the "motivations" action).

3 Multiagent Framework for Contracting

In the environments that we discuss, there are two types of agents. We will refer to the agent (or agents) that subcontracts one of its tasks to another agent or agents as the *contracting agent(s)*, and to the agent(s) that may agree to carry out the task as the *contracted agent(s)*. In order to convince the contracted agent to do the task and motivate it to do well, the contracting agent needs to provide the contracted agent with a beneficial contract. The contracted agent's success in carrying out the task depends on the time and work intensity which the contracted agent puts into fulfilling the task, which is referred to as the *effort* level. We propose constructing a monetary system in the multi-agent environment, which will provide a way for allocating rewards and evaluating outcomes.

What are the conditions that a *contracting multiagent* (CMA) framework should satisfy (for any specific distributed multi-agent domain), such that it should be accepted by all the designers of agents (for that specific domain)?

Simplicity: That the contract be simple and that there be an algorithm to compute it.

Pareto-Optimality: That there be no other contracted arrangement that is preferred by both sides over the one they have reached. Meaning that there be no other contract where the utilities of *both* agents are greater than their utilities in the contract agreed upon.

Stability: That the results be in equilibrium and that the contracts be reached and executed without delay.

3.1 Agents' Utility Function

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. Numeric representations of these preferences offer distinct advantages in compactness and analytic manipulation [Wellman and Doyle, 1992]. We therefore propose that each designer of autonomous agents will develop a numerical utility function that it would like its agent to maximize.

This is especially important in situations where there is uncertainty in the situation and the agents need to make decisions under risk considerations. 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 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 formalism 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, and a linear utility function yields risk neutral behavior [French, 1986].

We propose that a utility function of an automated agent in our contracting multiagent (CMA) environment depends on the agent's monetary gain and effort. 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 contracting agent's strategy in our CMA environment specifies which contract to offer to the contracted agent. The contracted agent's strategy specifies how it should respond to a given offer. Our desire is to obtain strategies which are in an equilibrium, since if the agents use these strategies, the environment becomes more stable. Since we consider different situations, we use different concepts of equilibria to gain stability.

In simple situations, with complete information, we use the **Nash equilibrium** concept. If there are n agents in the environment, a set of strategies $(s_1, s_2, ..., s_n)$ 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, if (s_{cing}, s_{ced}) are pair of strategies for the contracting and contracted agents respectively that are in Nash equilibrium, then if s_{cing} specifies a contract that the contracting agent should offer the contracted agent, the contracted agent does not have a better response but to act according to s_{ced}. On the other hand, given the possible responses of the contracted agent according to s_{ced}, the contracting agent's best strategy is to offer the contract indicated in s_{cing}.[1]

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

[1] As we see in Section 7.1 there are situations where there are 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.

set of beliefs, and the agents update their beliefs, according to Bayes's rule [Harsanyi, 1967 1968].

When there are several stages of the interaction among the agents, we use the concept of **perfect equilibrium**. 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 [Selten, 1975].

There are two approaches for finding equilibria in 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. 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 it is difficult to develop a computer program that will find the Nash equilibrium strategies.[2] The maximization approach, on the other hand, is much easier to implement. The problem with the maximization approach in our context is that the players must solve their optimization problems together: the contracted agent's strategy affects the contracting agent'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 contracting agent) will include, as a constraint, its opponent's (usually the contracted agent) maximization problem. The maximization problem of the contract's designer agent can be solved automatically by the agent. That is, the contracts that we provide maximize the expected utility of the designer of the contract (i.e., the contracting agent). 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 Notations

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

Effort level: Given a task, there are several effort levels that the contracted agent may take when trying to fulfill the task. We denote the set of these efforts by *Effort*. We use $e, e_i \in$ *Effort* to denote specific effort levels. In all cases, the contracted agent will decide how much effort to expend, but its decision may be influenced by the contract offered by the contracting agent.

[2] 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 [Kraus and Wilkenfeld, 1991a; Kraus and Wilkenfeld, 1991b; Kraus *et al.*, 1994].

Notation	Meaning	Comments
Effort	Set of efforts of the contracted agent	$e, e_1, ..., e_i \in$ *Effort*
Outcome	Set of possible monetary outcomes of carrying out a task.	$q, q_1, ...q_j \in$ *Outcome*. $q(e) \in$ *Outcome* when q is a function of $e \in$ *Effort*
Rewards	Set of possible monetary rewards to the contracted agent	$r, r_1, .., r_i \in$ *Rewards* $r(q)$ when r is a function of $q \in$ *Outcome*
U^{ced}	The contracted agent's utility function	
U^{cing}	The contracting agent's utility function	
$\hat{u} \in I\!R$	Contracted agent's utility from outside options.	(reservation price)
$e^* \in$ *Effort*	Efficient effort level for the contracting agent	Given contracted agent constraints
$q^* \in$ *Outcome*	Efficient outcome for the contracting agent	Given contracted agent constraints

Figure 1: Notations used in the paper

Outcome: While the contracted agent's expected utility depends on its effort level in performing a task, the expected utility of the contracting agent depends heavily of the outcome of performing a 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 contracted agent and that it can be expressed using the monetary system. We denote the monetary value of performing a task by $q \in$ *Outcome*. Given an effort level $e \in$ *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 time and effort put in by the contracted agent, the better the outcome.

Rewards: In order to convince the contracted agent to carry out a task, the contracting agent offers to pay the contracted agent a reward using the CMA monetary system. We denote the set of possible rewards by *Rewards* and its elements by r. The reward $r \in$ *Rewards* may be a function of the outcome from carrying on the task (i.e., $q \in$ *Outcome*).

Utility functions: We denote the contracted agent's utility function by U^{ced} : *Effort* $\times$ *Rewards* $\rightarrow I\!R$. We assume that in the CMA environment the contracted agent prefers to do as little as possible and gain higher rewards, therefore U^{ced} is a decreasing function in effort and an increasing function in rewards.

We denote the contracting agent's utility function by U^{cing} : *Outcome* $\times$ *Rewards* $\rightarrow I\!R$. The contracting agent prefers to give lower rewards and obtain larger outcomes. Thus, U^{cing} is an increasing function with the outcome and decreasing function with the reward being paid to the contracted agent.

Outside options: If the contracted agent does not accept the contract from the contracting agent and does not carry out the task it can then either perform another task (its own or others) or just be idle and not do anything. Its expected utility in such a situation is its reservation price and we refer to it as $\hat{u}$.

In our system we assume that the contracting agent rewards the contracted agent *after* the task is carried out. In such situations there should be a technique for enforcing these rewards. So that 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 contracting agent's owner for its contracts toward the contracted agent's owner.

Our last definitions are concerned with the value of the contracts to the contracting agent. The *first-best* contract will provide the contracting agent with a profit that is equal to a profit it could get when there is complete information and the contracting agent can supervise and observe the contracted agent(s)' actions. The *second-best* contract is Pareto-optimal given information asymmetry and constraints on writing contracts, e.g., the contracting agent does not observe the contracted agent(s)' actions.

4 Full Information

At first we would assume that all the relevant information about the environment and the situation is known to both agents. In the simplest case the contracting agent can observe and supervise the contracted agent's effort and actions and force it to make the effort level preferred by the contracting agent by paying only when the required effort is made. The amount of effort required from the contracted agent will be the one that maximizes the contracting agent's outcome, taking into account the task fulfillment and the rewards that need to be made to the contracted agent.

However, in most situations it is either not possible or too costly for the contracting agent to supervise the contracted agent's actions and observe its level of effort. In some cases, the contracting agent may be trying to carry out another task at the same time, or it cannot reach the site of the action (and that is indeed the reason for subcontracting).

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 contracted agent's actions.
- In Section 4.2 there is uncertainty concerning the outcome of an action taken by the contracted agent.

4.1 Contracts under Certainty

Suppose both agents have full information about the world and about each other, but the contracting agent does not supervise the contracted agent's actions. There is no uncertainty however, concerning the results of the contracted agent's actions, i.e., the outcome is a function of the contracted agent's effort. If this function is known to both agents, the contracting agent can offer the contracted agent a *forcing contract* [Harris and Raviv, 1978; Raizsmusen, 1989; Douglas, 1989]. This contract means that the contracting agent will pay the contracted agent only if it provides the outcome required by the contracting agent. If the contracted agent accepts the contract, it has to perform the task with the effort level that the contracting agent finds to be most profitable to itself, even without supervision. Note, the outcome won't necessarily be a result of the highest effort on the part of the contracted agent, but rather a result of the effort which provides the contracting agent with the desired outcome.

The contracting agent should pick an effort level $e^* \in \mathit{Effort}$ that will generate the efficient output level $q^* \in \mathit{Outcome}$. Since we assume that there are several possible agents available for contracting in equilibrium, the contract must provide the contracted agent at least with the utility $\hat{u}$.[3] The contracting agent needs to choose a reward function such that $U^{ced}(e^*, r(q^*)) = \hat{u}$ and $U^{ced}(e, r(q)) < \hat{u}$ for $e \neq e^*$. $\hat{u}$ is the minimal reward that will make the contracted agent accept the contract. Since the contracting agent would like to pay the contracted agent as little as possible, but wants the contracted agent to accept the offer, then if the outcome reveals that the contracted agent provided the required effort level, the contracting agent will pay the contracted agent $\hat{u}$. If the contracted agent accepts the contract but doesn't 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

The US and Germany have sent several mobile robots independently to Mars to collect minerals and ground samples and to conduct experiments. One of the US robots has to dig some minerals on Mars far from the other US robots. There are several German robots in that area and the US robot would like to subcontract some of its digging. The US robot approaches one of the German robots that can dig in three levels of effort (e): Low, Medium and High respectively denoted by 1,2 and 3. The US agent cannot supervise the German robot's effort since it wants to carry out another task simultaneously. The value of digging is $q(e) = \sqrt{100e}$. *The US robot's utility function, if a contract is reached, is* $U^{cing}(q, r) = q - r$ *and the German*

[3] We assume that if the contracted agent is indifferent between two actions, it will choose the one preferred by the contracting agent.

robot's utility function in the case that it accepts the contract is $U^{ced}(e,r) = 17 - \frac{10}{r} - 2e$*, where* r *is the reward to the German robot. If the German 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 the US robot's point of view is* 2*, in which there will be an outcome of* $\sqrt{200}$*. The contract that the US robot offers to the German robot is* $3\frac{1}{3}$ *if the outcome is* $\sqrt{200}$ *and 0 otherwise. This contract will be accepted by the German robot and its effort level will be Medium.*

There are two additional issues of concern, the first being how the contracting agent 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 contracting agents, but only one possible contracted agent. In such cases, there should be information about the utilities of the contracting agents in the event that they do not sign a contract, i.e., the contracting agents' reservation price. The outcome to the contracting agent in this case should be equal to its reservation price. Each contracting agent should offer a contract that maximizes the utility of the contracted agent, rather than its own, as when there are several contracted agents and one contracting. This maximization process should be done under the constraint that the contracting agent 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 contracting agent does not observe the contracted agent's behavior. However, in most subcontracting situations, there is uncertainty concerning the outcome of an action. If the contracted agent chooses some effort level, there are several possibilities for an outcome. For example, suppose an agent on Mars subcontracts digging for samples of a given mineral and suppose that there is uncertainty about the depth of the given mineral at the site. If the contracted agent chooses a high effort level and the mineral level is deep underground the outcome may be similar to the case where the contracted agent chooses a low level of effort and the mineral is located near the surface. However, if the contracted agent chooses a high effort level when the mineral is located near the surface, the outcome may be higher and thus, better to the contracting agent. In such situations the outcome of performing a task does not reveal the exact effort level of the contracted agent and choosing a stable and maximal contract is much more difficult.

Assuming that the world may be in one of several states, neither the contracting agent

nor the contracted agent knows the exact state of the world when agreeing on the contract. There is the possibility that the contracted agent may gain more information about the world during or after completing the task, but only *after* signing the contract and choosing the effort level. The contracting agent can not gain more information about the world.

Following [Harris and Raviv, 1978], we also assume that there is a set of possible outcomes to the contracted agent carrying out the task $Outcome = \{q_1, ..., q_n\}$ such that $q_1 < q_2 < ... < q_n$ depends upon the state of the world and upon the effort level of the contracted agent. Furthermore, we assume that, given a level of effort, there is a probability distribution attached to the outcomes that is known to both agents.[4] Formally, we assume that there is a probability function $\wp : Effort \times Outcome \rightarrow I\!R$, such that for any $e \in Effort$, $\sum_1^n \wp(e, q_i) = 1$ and for all $q_i \in Outcome, \wp(e, q_i) > 0$.[5] This characterizes the situations where the contracting agent is not able to use the outcome to determine the contracted agent's effort level unambiguously.

The contracting agent's problem is to find a contract that will maximize the contracting agent's expected utility, knowing that the contracted agent may reject the contract or, even if it accepts the contract the effort level will be chosen later [Raizsmusen, 1989]. The contracting agent's reward to the contracted agent can be based only on the outcome. Let us assume that in the contract that will be offered by the contracting agent, for any q_i $i = 1, ..., n$, the contracting agent will pay the contracted agent the reward r_i. The maximization problem can be constructed as follows (see also [Raizsmusen, 1989]).

$$Maximize_{r_1,...r_n} \sum_1^n \wp(\hat{e}, q_i) U^{cing}(q_i, r_i) \tag{1}$$

with the constraints:

$$\textbf{(IR)} \qquad \sum_1^n \wp(\hat{e}, q_i) U^{ced}(\hat{e}, r_i) \geq \hat{u} \tag{2}$$

$$\textbf{(IC)} \qquad \hat{e} = argmax_{e \in Effort} \sum_1^n \wp(e, q_i) U^{ced}(e, r_i) \tag{3}$$

[4] A practical question is how the agents find the probability distribution. It may be that they have preliminary information about the world, e.g., what the possibility is that a given mineral will be in that area of Mars. 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 [Page, 1987].

[5] The formal model in which the outcome is a function of the state of the world and the contracted agent's effort level, and in which the probabilistic function gives the probability of the state of the world which is independent of the contracted agent's effort level is a special case of the model described here. [Page, 1987; Ross, 1973; Harris and Raviv, 1978].

Equation (1) states that the contracting agent tries to choose the reward to the contracted agent so as to maximize its expected utility subject to two constraints. First, the rewards to the contracted agent must be large enough to motivate the contracted agent 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 contracted agent will be at least as much as its reservation price ($\hat{u}$). The second constraint (3), which is called *participation* constraint (IC), provides the contracted agent with a motivation to choose the effort level that the contracting agent prefers, given the contract it is offered. This means, given the agreed rewards, $\hat{e}$ will provide the contracted agent 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 to 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 contracting agent and the contracted agent are risk neutral, then solving the maximization problem can be done using any linear programming technique (e.g, simplex, see for example [Pfaffenberger and Walker, 1976; Spivey and Thrall, 1970].) Furthermore, in most situations, the solution will be very simple: the contracting agent will receive a fixed amount out of the outcome and the rest will go to the contracted agent. That is, $r_i = q_i - C$ for $1 \leq i \leq n$, where the constant C is determined by constraint (**IR**:2) [Shavell, 1979].

Example 4.2 Risk Neutral Agents Under Uncertainty
Suppose the utility function of the German robot from Example 4.1 is $U^{ced}(r, e) = r - e$ *and that it can choose between two effort levels, Low (e=1) and High (e=2), and that its reservation price is* $\hat{u} = 1$. *There are two possible monetary outcomes to the digging:* $q_1 = 8$ *and* $q_2 = 10$, *the US robot's utility function remains as it was in the previous example, i.e.,* $U^{cing}(q, r) = q - r$.

If the German robot chooses the Lower level 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, the US 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 German robot will choose the High level effort.*

4.2.2 The Contracted Agent is Risk Averse

When the agents are not neutral toward risk, the problem of solving the contracting agent's maximization problem is much more difficult. However, if the utility functions for the agents are carefully chosen, an algorithm does exist.

Suppose the contracted agent is risk averse and the contracting agent is risk neutral (the methods are also applicable when both are risk averse). Grossman and Hart [Grossman and Hart, 1983] presented a three-steps procedure 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 contracted agent 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 contracting agent. The third step of the procedure is to choose the effort level that maximizes profits, given the necessity to support that effort with a costly reward contract. Formally, step one and two are as follows: Suppose the contracting agent wants the contracted agent to choose the effort level $e' \in \mathit{Effort}$ then it needs to solve the following:

$$C(e') = Minimize_{r_1,...,r_n} \sum_{i=1}^{n} \wp(e', q_i) r_i \tag{4}$$

with the constraints:

$$\textbf{(IR)} \qquad \sum_{i=1}^{n} \wp(e', q_i) U^{ced}(e', r_i) \geq \hat{u} \tag{5}$$

$$\textbf{(IC)} \qquad \sum_{i=1}^{n} \wp(e', q_i) U^{ced}(e', r_i) \geq \sum_{i=1}^{n} \wp(e, q_i) U^{ced}(e, r_i) \text{for all } e \in \mathit{Effort} \tag{6}$$

The first constraint (5) requires that the expected utility for the contracted agent will be at least as good as its outside options (its reservation price). The second constraint (6) requires that given the contract, the contracted agent will prefer to take the effort level e'. The minimization problem states that the contracting agent is looking for a contract where it can pay as little as possible to induce the contracted agent to choose e'. For this minimization problem there is an algorithm given that U^{ced} satisfies several properties, including the property that the preference of the contracted agent over entering uncertain situations are independent of its actions [Grossman and Hart, 1983; Pfaffenberger and Walker, 1976; Rogerson, 1985].[6]

After finding a set of possible values, $r_1, ..., r_n$ for every $e \in \mathit{Effort}$ (where the set may be empty since there could be effort levels which the contracting agent cannot make the contracted agent choose), and the minimum expected reward $C(e)$ for any effort level, the contracting agent is ready to move to the third step, which is easy to compute, and to choose

[6] In [Rogerson, 1985] the problem of finding a contract when the contracting agent 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 contracted agent. In such situations a solution can be calculated using the Kuhn-Tucker Theorem.

the effort level that will provide it with the maximal outcome:

$$Maximize_{e \in Effort} U^{cing}(\sum_{1}^{n} \wp(e, q_i) q_i, C(e)) \tag{7}$$

The contracted agents computational task is easier. After being offered a contract, the contracted agent needs only to check the validity of the inequalities that appear as constraints in the contracting agent's maximization problem. In particular, the contracted agent needs to check the validity of the individual-rationality constraint (IR) to decide whether to accept the contract. When the contracted agent needs to decide which effort level to provide, it should consider its expected utility from its effort level, similar to 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 4.3 Risk Averse Contracted Agent Under Uncertainty. *Suppose the situation is exactly as in Example 4.2 but the designer of the robot determines that the contracted agent will be risk averse and its utility function is as in Example 4.1:* $U^{ced}(r, e) = 17 - \frac{10}{r} - 2e$ *and* $\hat{u} = 1$.

The maximization problem that the contracting agent should solve is:

$$Maximize_{r_1, \ldots r_n} \sum_{1}^{n} \wp(\hat{e}, q_i)(q_i - r_i) \tag{8}$$

with the constraints:

$$\textbf{(IR)} \qquad \sum_{1}^{2} \wp(\hat{e}, q_i)(17 - \frac{10}{r_i} - 2e) \geq 1 \tag{9}$$

$$\textbf{(IC)} \qquad \hat{e} = argmax_{e \in \{1,2\}} \sum_{1}^{2} \wp(e, q_i)(17 - \frac{10}{r_i} - 2e) \tag{10}$$

Using Grossman and Hart's three-steps procedure [Grossman and Hart, 1983] requires that the contracting agent first determine the minimal reward needed to make the contracted agent choose $e_1 = 1$ *and what the minimal rewards is that will make it choose* $e_2 = 2$*:*

$$C(e_1) = Minimize_{r_1, r_2} \frac{3}{4} r_1 + \frac{1}{4} r_2 \tag{11}$$

with the constraints:

$$\textbf{(IR)} \qquad \frac{3}{4}(17 - \frac{10}{r_1} - 2) + \frac{1}{4}(17 - \frac{10}{r_2} - 2) \geq 1 \tag{12}$$

$$\textbf{(IC)} \qquad \frac{3}{4}(17-\frac{10}{r_1}-2)+\frac{1}{4}(17-\frac{10}{r_2}-2) \geq \frac{1}{8}(17-\frac{10}{r_1}-4)+\frac{7}{8}(17-\frac{10}{r_2}-4) \qquad (13)$$

The results of solving this minimization problem using Lagrangian multipliers is that the minimal reward to make the contracted agent 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 contracted agent choose effort level* $e_2 = 2$ *is* $r'_1 = 1$ *and* $r'_2 = 1\frac{8}{17}$.

Finally, the contracting agent should check which effort level it prefers, given the above rewards, i.e., it should compare between $\wp(e_1, q_1)(q_1 - r_1) + \wp(e_1, q_2)(q_2 - r_2)$ *and* $\wp(e_2, q_1)(q_1 - r'_1) + \wp(e_2, q_2)(q_2 - r'_2)$. *The conclusion is that the contracting agent can obtain the largest expected utilities by offering* $r'_1 = 1$ *and* $r'_2 = 1\frac{8}{17}$.

The contracted agent will then compute its expected utility from choosing effort level e_1 *(i.e.,* $\frac{3}{4}(17-\frac{10}{r'_1}-2)+\frac{1}{4}(17-\frac{10}{r'_2}-2)$*) and from choosing effort level* e_2 *(i.e.,* $\frac{1}{8}(17-\frac{10}{r'_1}-4)+\frac{7}{8}(17-\frac{10}{r'_2}-4)$*), and will realize that its expected utility from both is the same. It will then verify that its expected utility from the offered contract is greater than* $\hat{u}$ *(i.e.,* $\frac{1}{8}(17-\frac{10}{r'_1}-4)+\frac{7}{8}(17-\frac{10}{r'_2}-4) \geq 1$ *), will 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 contracting agent.*[7]

4.2.3 Obtaining Imperfect Information about the Contracted Agent Behavior

Even in situations where the contracting agent cannot observe the actions of the contracted agent, it may be able to gain some information about its behavior. For example, it can gain information by setting up a camera in the digging site. This information may be imperfect, and the process of getting this information is called an imperfect (noisy) monitoring process. In particular, if the contracted agent 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 contracting agent to obtain some estimation of the contracted agent's effort level. The main question is however, whether using such monitoring is beneficial.

It has been shown that if the contracted agent is risk neutral, there are no gains (to either agent) from the use of any monitoring mechanism [Harris and Raviv, 1979]. However,

[7] In the rest of the paper we won't specify the contracted agent's computation procedures, since in most of the situations, given a contract, the contracted agent needs only to check the validity of the inequalities that appear as constraints in the contracting agent's maximization problem, similar to the check done in this example. Since all variables are known, based on the suggested contract, this check is straight forward.

if the contracted agent is risk averse, there are potential gains to monitoring. In particular, this is the case if a contract of the following form is an optimal monitoring contract: If the contracted agent's action is judged acceptable on the basis of the monitored outcome, the contracted agent will then be paid according to a prespecified schedule. Otherwise, it will receive less preferred, fixed rewards [Harris and Raviv, 1979].

To demonstrate this idea we use a modification of an example that appears in [Harris and Raviv, 1979].

Example 4.4 *Suppose the utility function of the German robots from the previous examples is* $U^{ced}(e,r) = r^{0.25} - \frac{4e^{1.25}}{5}$*, its reservation price* $\hat{u} = 0$ *and the utility function of the US robot is, as in previous examples* $U^{cing}(q,r) = q - r$*. Suppose the world is in situation* 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* $[e-\epsilon, e+\epsilon]$ *for some* $\epsilon > 0$*. Meaning that if the contracted agent chooses effort level* e*, the monitor will provide an equal probability number* α*, between* $e-\epsilon$ *and* $e+\epsilon$*.*

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

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

The effort level chosen by the German robot depends on ϵ*. If* $\epsilon < 2^{-1.25}$ *then it will choose* $2e + 2^{-6}e^{-3}$*. In such situations the German robot will always get the reward* $\frac{5}{4}\epsilon$ *and its expected utility is* 0*. The expected utility of the US robot is* $\frac{1}{2} + 2^{-5} + \epsilon^{-3} + \frac{3}{4}\epsilon$*.*

If $\epsilon \geq 2^{-1.25}$ *then the German robot will not choose the required level of effort, but rather will take a lower level effort* $5 * 2^{-6}\epsilon^{-3}$*. It may be that the monitoring value* α *will be lower than* $2e + 2^{-6}e^{-3} - \epsilon$ *and the German robot won't get any reward.*

The probability of this happening is $1 - 2^{-5}\epsilon^{-4}$*, and the German robot's expected utility is still* 0*, while the expected utility of the US robot in this case is* $\frac{1}{2} + 5 * 2^{-7}\epsilon^{-3}$*.*

In both cases, the US' expected utility is more than $\frac{1}{2}$*, which is what it can expect if it does not use monitoring mechanism.*

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

5 Asymmetric and Incomplete Information

There are some situations where the contracted agent may have more information than the contracting agent. The contracted agent may have obtained more information concerning the environment, e.g., the German robot may know the mineral level since it is at the digging site, while the US robot may have only some probabilistic beliefs on the level of the minerals. In other situations the contracting agent may not know the utility function of the contracted agent. The contracted agent may be one of several types that reflect the contracted agent's ability to carry out the task, its efficiency or the cost of its effort. However, we assume that given the contracted agent's type, its utility function would be known to its party. For example, suppose Germany builds robots of two types. The specifications of the robots are known to the German robots and to the US robots; however, the US robots do not know the specific type of the German robots they will encounter.

The contracting agent could simply ask the contracted agent for the additional information, i.e., its type or the state of the world, although the contracted agent will not tell the truth unless the contracting agent provides it with a monetary incentive to do so. This will often cause inefficiency from the contracting agent's point of view.

A useful technique in such situations is for the contracting agent to search for an optimal mechanism [Demougin, 1989] as follows: the contracting agent offers the contracted agent a menu of contracts that are functions of its type (or the state of the world) and the outcome. If the contracted agent accepts the offer, it chooses a contract and announces it to the contracting agent. Given this contract, the contracted agent chooses an effort level which maximizes its own expected utility. In each of the menu's contracts, the contracted agent's expected utility should be at least as high as its expected utility if it does not sign the contract.

One of the useful results in this area is that without loss of generality it is enough to consider only contracts in which it is in the interest of the contracted agent to honestly report its type [Myerson, 1982].

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 contracted agent, but not to the contracting agent.

- In Section 5.2 neither agent knows the state of the world *before* signing the contract, but the contracted agent finds out that information *after* signing the contract, but *before* choosing its effort level.

- In Section 5.3 the contracted agent's information is initially better than that of the

contracting agent, 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 contracted agent cannot predict the outcome based on its private information both before and after signing the contract.
- In Section 5.5 both agents have some private information, e.g., they have some private information on their types.

5.1 Asymmetric Information about the State of the World

Supposing the world can be in one of several states, $\theta_1, ..., \theta_n$. If the contracted agent chooses a level of effort e and the state of the world is θ, then the outcome will be $f(e, \theta)$ [Harris and Townsend, 1981]. As in previous cases the contracted agent's utility function ($U^{ced}(e, r)$) increases with the reward it gets from the contracting agent (r) and decreasing with its effort (e). The contracting agent's utility function ($U^{cing}(q, r)$) increases with the outcome and decreases with its reward to the contracted agent.

We assume that the contracted agent knows the state of the world θ but the contracting agent 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 $i = 1, ..., n$ by ϕ_i and assume that $\sum_i^n \phi_i = 1$.

As we described above, in the first step of the agents' interaction the contracting agent will offer the contracted agent n pairs (one for each state) of an outcome and a payoff (q_i, r_i). The contracted agent then will report its private information, i.e., the state of the world, to the contracting agent. According to this message, the corresponding contract is implemented. In the third step the contracted agent chooses its effort level, and is paid according to the chosen contract and the outcome.

As was mentioned above, we will restrict our attention to *direct mechanisms* under which the contracting agent reports the situation of the world honestly, motivated by the contract. That is, if the state of the world is θ_i , (q_i, r_i) is the best contract among the ones offered by the contracting agent. This constraint is called "self-selection".

Formally,

$$(\mathbf{SS})\ \forall i \in \{1, ..., n\}\ \ U^{ced}(e_i, r_i) \geq U^{ced}(e_j, r_j)\ \ where\ 1 \leq j \leq n, \quad f(\theta_i, e_i) = q_i, \ \ f(\theta_i, e_j) = q_j \tag{14}$$

In addition, in each of the n contracts offered by the contracting agent to contracted agent's utility should be higher than its reservation price. The contracting agent should find

a set of such self-selection contracts that will maximize its expected utility, based on its probabilistic beliefs. Formally:

$$Maximize_{(q_1,r_1),\ldots(q_n,r_n)} \sum_{i=1}^{n} \phi_i U^{cing}(q_i, r_i) \tag{15}$$

subject to: (SS:14)

$$(\mathbf{IR}) \qquad U^{ced}(e_i, r_i) \geq \hat{u}, \quad where\ f(\theta_i, e_i) = q_i \tag{16}$$

We demonstrate this maximization problem in the next example.

Example 5.1 Contracting Under Asymmetric Information *Suppose the digging site of the German/US example may be in two states* $\theta_1 = 1$ *and* $\theta_2 = 2$. *The outcome function is* $f(e, \theta) = e\theta$, *the US robot's utility function is* $U^{cing}(q, r) = q - r$ *and the German robot's utility function is* $U^{ced}(e, r) = r - e^2$. *Hence, with* $f(e, \theta) = e\theta$, *the German's utility function as a function of the output, reward and the state of the world is* $U^{ced}(q, r, \theta) = r - (q/\theta)^2$. *We also assume that the contracted agent's reservation price is* $\hat{u} = 1$ *and the contracting agent believes with probability* 0.25 *that the state is* θ_1 *(i.e.,* $\phi_1 = 0.25$*) and with probability* 0.75 *that the situation is* θ_2.

In such a situation the contracting agent should solve the following maximization problem:

$$Maximize_{(r_i,q_i),i=1,2} 0.25(q_1 - r_1) + 0.75(q_2 - r_2) \tag{17}$$

subject to:

$$\begin{array}{lcl} r_1 - q_1^2 & \geq & r_2 - q_2^2 \\ r_2 - (q_2/2)^2 & \geq & r_1 - (q_1/2)^2 \\ r_1 - q_1^2 & \geq & 1 \\ r_2 - (q_2/2)^2 & \geq & 1 \\ 0 \leq r_i & \leq & q_1, i = 1, 2 \end{array}$$

If the output function f is twice differentiable in e, with $f_e > 0$ and $f_{ee} < 0$[8] for all θ, then there is an interesting result concerning the contracting agent's preference over the information available to the contracted agent. If the contracted agent has full information about the state of the world *before* signing the contract, then the contracting agent's expected utility is *lower* than in the case where it and the contracted agent have symmetric beliefs (either perfect or imperfect) about the state of the world *before* signing the contract [Baiman and Demski, 1980; Demski and Sappington, 1984]. This finding is a result of the fact that when they share the same (perfect or imperfect) state of information the contracted agent can be held to its reservation level of expected utility.

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

5.2 Asymmetric Information After Reaching an Agreement

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

If agreements are enforced, i.e., if the contracted agent 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 contracted agent, rather than its eventual utilities, since at the time of the contract, the exact utility is not known to the contracted agent. 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:

$$(\mathbf{IR}) \qquad \sum_{i=1}^{n} \phi_i U^{ced}(e_i, r_i) \geq \hat{u}, \quad where\ f(\theta_i, e_i) = q_i \tag{18}$$

We demonstrate this in the following example.

Example 5.2 Risk Neutral Agents Under Asymmetric Information *Suppose the situation is exactly as in Example 4.2, but the German robot can find out more information after the agents have reached a contract, but before choosing its level of effort. As in Example 4.2 the contracted agent can choose between two effort levels Low (e=1) and High (e=2) and its reservation price is* $\hat{u} = 1$*, and there are two possible monetary outcomes to the digging:* $q_1 = 8$ *and* $q_2 = 10$*. The agents' utility functions are as in Example 4.2.*

The world can be in one of eight possible states $\theta_1, ..., \theta_8$ *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* $2 \leq i \leq 8$*,* $f(2, \theta_1) = q_2$*. Note that this yields the same probabilistic outcome as in Example 4.2.*

There are two possibilities for constructing the contracts, depending on which effort level the contracted agent will choose if the state of the world is either $\theta_2, ..., \theta_6$*. It is clear that if the state is* θ_1*,*θ_7 *or* θ_8 *the contracted agent will choose the Low effort level.*

If the contracting agent would like the contracted agent to choose effort level High in these states, then the contracting agent should solve the following minimization problem (we list only the binding constraints):

$$Minimize_{r_1, r_2} \frac{1}{8} r_1 + \frac{7}{8} r_2 \tag{19}$$

subject to:

$$(\mathbf{IR}) \qquad \frac{1}{8}(r_1 - 1) + \frac{5}{8}(r_2 - 2) + \frac{2}{8}(r_2 - 1) \geq 1 \tag{20}$$

$$\textbf{(IC)} \qquad r_2 - r_1 \geq 1 \tag{21}$$

By solving this problem we can conclude that the contracting agent can always keep $7\frac{1}{8}$ *of the outcome and pay the contracted agent* $r_1 = \frac{7}{8}$ *and* $r_2 = 2\frac{7}{8}$.

Similarly, we can formalize the problem where the contracted agent 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 contracting agent is* $6\frac{1}{2}$.

In order for the contracting agent to maximize its expected utility, the first option is better since it yields the contracting agent an expected outcome of $7\frac{1}{8}$. *This is higher than in Example 4.2, where its expected outcome there is* $6\frac{3}{4}$.

We would like to consider the option of monitoring in such situations. It was proved in [Harris and Raviv, 1979] that if the contracted agent 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 contracted agent is risk averse, monitoring may be beneficial as we will explain in Section 5.6.

The contracting agent can design a contract that will make the contracted agent choose the Pareto efficient effort level for the real state of the world.

If it is possible for the contracted agent 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 [Shappington, 1983].

When the contracted agent can opt out of an agreement, the question is what are its alternatives at that point. It may be that it can still gets 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 can gain less from a new option. Let us denote the contracted agent's new reservation price by $\hat{u}^{new}$.

In such situations, the contracting agent 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:

$$U^{ced}(e_i, r_i) \geq \hat{u}^{new} \qquad 1 \leq i \leq n \ where \ f(\theta_i, e_i) = q_i \tag{22}$$

This constraint verifies that even when the contracted agent find out more information about the environment *before* it chooses its level of effort, it will benefit from choosing the level e_i and will keep the agreement.

Of course, these constraints reduce the contracting agent's expected utility, and it will need to suggest to the contracted agent higher payments to make sure it won't opt out. We

will demonstrate this in the case that the contracted agent is risk neutral as in Example 4.2.

Example 5.3 Risk Neutral Agents Under Asymmetric Information with Opting Out

Suppose the situation is exactly as in Example 5.2, but before choosing its level of effort, the German robot can opt out of the agreement and get its original reservation price (i.e., $\hat{u}^{new} = \hat{u} = 1$). Therefore, instead of constraint 20, the following should be stated:

$$r_1 - 1 \geq 1 \qquad r_2 - 2 \geq 1 \tag{23}$$

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

5.3 Asymmetric and Imperfect Information Before Contracting

We would like to consider the situation where the contracted agent's information is initially better than that of the contracting agent, but it knows the exact state of the world only after a contract is signed. For example, the German robot may initially have better information about the level of the minerals than the US robot. However, it does not have full information about the state of the world. Only after reaching the digging site (*after* signing an agreement), does it find the real level of the minerals. Note that in the previous section, 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 contracted agent 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 the previous Section 5.2.

As in previous situations, we assume that the outcome is a function of the contracted agent's effort level and the state of the world, i.e., $q = f(e, \theta)$. At no time can the contracting agent observe either e or θ.

Suppose that the possible states of the world are $\theta_1, \theta_2, ...\theta_n$, such that $\theta_i < \theta_{i+1}$ for $1 \leq i \leq n$. Furthermore, the contracting agent does not know the exact probabilistic distribution of θ, but rather knows that there are D possible probabilistic distributions $\wp^d$, and it believes with probability ϕ^d that the real distribution is $\wp^d$.

Before signing the contract, the contracted agent does not know the actual state of the world either, but it knows which probabilistic distribution function is the correct one.

We assume that the utility function of the contracted agent can be written as a function of q and r as follows: $U^{ced}(q, r) = r - e(q, \theta)$ where $f(e(q, \theta), \theta) = q$. In such situations the optimal strategy for the contracting agent [Harris and Townsend, 1981] is to design at most

D distinct contracts from which the contracted agent can make a binding choice by sending a message to the contracting agent.

Thus the maximization problem of the contracting agent is as follows [Sappington, 1984]:

$$Maximize_{(q_i^d, r_i^d),...(q_n^d, r_n^d)} \sum_{i=1}^{D} \phi^d \sum_{j=1}^{n} \wp^d(\theta_i) U^{cing}(q_i, r_i) \tag{24}$$

subject to:

$$\textbf{(IR)} \qquad \sum_{i=1}^{n} \wp^d(\theta_i)(r_i^d - e(q_i^d, \theta_i)) \geq \hat{u} \quad \forall d = 1, ..., D \tag{25}$$

$$\textbf{(SS)} \qquad \sum_{i=1}^{n} \wp^d(\theta_i)(r_i^d - e(q_i^d, \theta_i)) \geq \sum_{i=1}^{n} \wp^d(\theta_i)(r_i^r - e(q_i^r, \theta_i)) \quad \forall r, d = 1, ..., D \tag{26}$$

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

where $\wp^d(\theta_i)$ is the probability that the state of the world is θ_i according to distribution d ($\wp^d(\theta_i) > 0 \; \forall i, d$), q_i^d is the output produced by the contracted agent in state θ_i under contract $\{(q_i^d, r_i^d)\}$ and r_i^d is the reward to the contracted agent 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 contracted agent will report honestly about the actual distribution (i.e., will choose contract $\{(q_i^d, r_i^d)\}$ when $\wp^d$ 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 contracted agent 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 [Christensen, 1981]. In the first stage of the interaction, the contracting agent offers the contracted agent a menu of contracts based on a message it will send and the observed outcome. The contracted agent may reject the offer or agree to it and sign a contract. In the second stage, the contracted agent 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 it to better predict what the outcome will be given its level of effort. For example, when the German robot reaches the area where

it needs to dig, it determines the structure of this area (i.e., it collects information about the world state). This information may not be complete, but it is not known to the US robot at all. In the third stage, the contracted agent sends a message to the contracting agent and chooses a level of effort. In the fourth stage the outcome is observed by both agents and the contracted agent is paid according to the outcome and its earlier message.

Note that in such situations, the contracted agent has committed itself not to leave the agreement once it has observed ξ.[9] Also in this case [Christensen, 1981], the agents can concentrate on the class of contracts that induce the contracted agent to send a truthful message to the contracting agent. This is due to the fact that it has been shown [Christensen, 1981] 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 contracting agent is similar to the one in Sections 5.2, where, under the constraints, the contracted agent's utility is replaced by its expected utility given ξ.

5.5 Both Parties Have Private Information

There are some situations where both the contracting agent and the contracted agent have private information, for example, the contracting agent has private information about its type, and the contracted agent has private information about the world. To put it simply, we assume that the actions taken by the contracted agent are observable by the contracting agent, but there is uncertainty about the outcome. That is, we assume that, given a level of effort, there is a probability distribution $\wp$ which is attached to the possible 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 contracted agent 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 contracting agent to do better than an uninformed one, it must actively participate in the contract selection and not only in the mechanism design. One possibility is as follows [Maskin and Tirole, 1990]: there are up to four possible stages in an interaction.

1. In the first stage of the interactions, the contracting agent offers a mechanism to the contracted agent which specifies:

[9] In most of the situations the contracting agent is better off making such a commitment. But in some situations, both agents can be made better off through reconstructing [Fudenberg and Tirole, 1990; Laffont and Tirole, 1990; Demougin, 1989; Hart and Tirole, 1988].

(a) a set of possible messages that each party can choose

(b) for each pair of messages m_{cing}, m_{ced} can be chosen simultaneously by the contracting and contracted agents respectively, a corresponding probabilistic function of the effort level that will be chosen by the contracted agent (note that the probabilistic choice mechanism and the effort level are observable by the contracting agent).

(c) pairs of outcomes and rewards.

2. In the next stage the contracted agent accepts or refuses the mechanism. If it refuses it gets some of the reservation price $\hat{u}$, and the interaction ends.

3. The agents send each other the messages simultaneously.

4. The contracted agent performs the task at the appropriate effort level and is paid according to the outcome.

As in previous cases, the agents can limit themselves to honest reports.

In situations where the exact type of the contracting agent does not directly influence the contracted agent's utilities, it was shown [Maskin and Tirole, 1990; Myerson, 1983] that the contracting agent can profit from the contracted agent's incomplete information. The intuition behind these results is as follows; When the contracting agent proposes a contract, it does it subject to two types of constraints. The (IR) constraints requires that the expected utility of the contracted agent when accepting the contract will be higher than the contracted agent's reservation price. There are also constraints to ensure that when the contract is carried out, the contracted agent behaves in the appropriate way, given its private information (IC). When the contracting agent does not have private information, the constraints must hold individually for each type of contracting agent. If the contracted agent has incomplete information about the contracting agent, the constraints need to hold only in "expectation" over the suggested contracts which are functions of the contracting agent's type. For this reason, a given type of contracting agent can raise its utility above the case where the contracted agent is fully informed, by violating some constraints, as long as they are offset by other types. Actually, in most of the situations, there exists a mechanism in which all types of contracting agents do strictly better than the fully informed contracted agent.

However, in order to take advantage of the contracted agent's incomplete information, the contracting agent 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 just for the expected types.

This means that if the selection of the mechanism by the contracting agent depends in any way upon the contracting agent's individual type, then the selection of the mechanism itself will convey information about its type to the contracted agent. Therefore, any contracting agent, regardless of its type, should offer the same mechanism.[10]

Cases in which the contracting agent's private information influences the contracted agent's utilities are more complex [Maskin and Tirole, 1992]. In such situations it is no longer true that, without loss of generality, the contracting agent can postpone revealing its type until the third stage of the interaction. The contracting agent may wish to disclose information about itself in order to influence the contracted agent's action, if so then the contracting agent's proposal should balance between total disclosure and complete concealment. Furthermore, the contracting agent's expected utility when it has private information which influences the contracted agent's utility may be even lower than in the case the contracting agent doesn't have any private information at all. This is because the contracted agent's expected utility may be low, given some of the contracting agent's types denoted by "bad" types. Therefore, when the contracted agent's probabilistic belief is that its opponent's "bad" type is high (even if the actual type is not "bad"), the contracted agent must be paid correspondingly high rewards to encourage it to accept the contract. Note that in the first case we considered, where the contracted agent is not directly influenced by the contracting agent's type, its original beliefs do not play an important role, since the contracted agent cares only about how the contracting agent'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 situations of Asymmetric information [Melumad and Reichelstein, 1989; Christensen, 1981]:

1. Will the contracting agent always be better off, the more the contracted agent knows about the world?

2. Is communication beneficial to the contracting agent. Meaning, is it better to the contracting agent to suggest a menu of contracts to the contracted agent and ask it to send a message informing it of the current state of the world, or will it be better off offering only a single contract based only on the jointly observed outcome?

[10] Maskin and Tirole [Maskin and Tirole, 1990] show that any equilibrium of the mechanism design presented here can be computed as a Walrasian equilibrium of fictitious economy. In this economy, the traders are the different types of contracting agent. For more technical and formal details see [Maskin and Tirole, 1990].

The second question is essential when communication is costly to the contracting agent. Intuitively, it seems that both communications and a knowledgeable contracted agent will allow for more efficient contracting. The contracted agent may use its knowledge to choose the correct actions, and with a menu of contracts the contracted agent 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 contracted agents will improve the contracting agents benefits, but rather their effect depends on the exact details of the situation. There are even situations when less information by the contracting agent is preferred to more [Gjesdal, 1982].

As we explained in Section 5.1, when the contracted agent has full private information before signing the contract, the contracting agent's expected utility is lower than if they have symmetric beliefs.

If the contracted agent acquires its information *after* signing the agreement, then its effect on the contracting agent varies. The contracted agent may use its additional information in two ways: it may use its information to shirk, thereby reducing the benefits for the contracting agents, or it may use the information to improve the outcome (see two demonstrating examples in [Christensen, 1981]). Any additional information gained by the contracting agent after a contract is reached is only valuable if it is influenced by the choices made by the contracted agent [Gjesdal, 1982].

The disadvantage of communications is that the "self-selection" constraint can sometimes be very restrictive, so that the information received by the contracting agent is not beneficial. This occurs particularly, if the contracted agent 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 contracting agent can then replicate its benefits, using a single contract. Furthermore, even if the contracted agent does not have perfect information, there are many situations in which there is no value for communication [Melumad and Reichelstein, 1989; Demougin, 1989]. These situations are such that the stochastic outcome is informative.

If the outcome is not informative, however[11], then communication is valuable. It is valuable for two reasons; because it allows the contracting agent to implement a more efficient level of effort choices without having to pay the contracted agent for making it choose correctly. Alternatively, menu contracts can be valuable even though the contracted agent's action choices are unchanged. In such situations, the value of communication results from rewards to the contracted agent.

There are, of course, situations where the contracting agent can use the information gathered in the menu contracts for other purposes (e.g., later contracts with other agents).

[11]See [Melumad and Reichelstein, 1989] for exact conditions.

In such a case, it may prefer the menu of contracts, even if it cannot benefit in the current interaction.

5.7 Several Contracted Agents Compete for the Job

There may be a situation where there are several agents in the environment, and the contracting agent 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 contracting agent does not know the types of the other agents, the following mechanism is appropriate: The contracting agent announces a set of contracts based on the agent's type and asks the potential contracted agents to report their types. On the basis of these reports the contracting agent chooses one agent [McAfee and McMillan, 1987].[12] The agent that is chosen, chooses a level of effort that is not observable by the contracting agent. The rewards to the chosen contracted agent depend upon the contracted agent's reported type and the observed outcome. As in previous cases, the contracting agent can use, with out loss of generality, contracts in which the agents report their types honestly [Myerson, 1982].

An important aspect in the design of the contracts is the marginal return to the contracting agent by increasing the probability that a specific type (e.g., z_i) will be chosen. This marginal return consists of the outcome minus the contracted agent's costs to produced the required effort level, minus the rewards the contracted agent 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 contracting agent makes it more attractive for higher types to pretend to be z_i. To prevent this the contracting agent 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 [McAfee and McMillan, 1987]) related to the above described aspect, and if the highest reported type is chosen, then the contract may be optimal for the contracting agent. However, the contracting agent's benefits will be lower than in the case where it can observe the contracted agent's effort level (i.e., it gets only the "second best" benefits).

[12] There are situations where the agents' types are multidimensional. That is, the contracting agent is uncertain about different aspects of the contracted agent that are independent. For example, its digging capabilities and its disk space. Techniques to formalized the maximization problem in such situations, and methods to solve it can be found in [Laffont *et al.*, 1987; McAfee and McMillan, 1988].

6 Repeated Encounters

Suppose the contracting agent 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 encounter starts, and short term contracts, i.e., in each encounter a new contract is agreed upon by the agents.

6.1 Short Term Contracts

Repetition of the encounters between the contracting and the contracted agents enables the agents to reach efficient short term contracts if the number of encounters is large enough and if the contracted agent can be "punished" sufficiently severely [Radner, 1981; Radner, 1985; Malcomson and Spinnewyn, 1988].

Based on the average outcome, the contracting agent could form an accurate estimate of the contracted agent's effort over time. That is, if the contracting agent wants the contracted agent to make a certain effort level of $\hat{e} \in \mathit{Effort}$ in all the encounters, it can compute the expected outcome over time if the contracted agent actually performs the task with that effort level. The contracting agent can keep track of the cumulative sum of the actual outcomes and compare it with the expected outcome. If after several encounters the contracting agent realizes that the cumulative outcome is below a given function of the expected outcome, it should impose a big "punishment" on the contracted agent. If the function over the expected outcome is chosen carefully [Radner, 1981], the probability of imposing a "punishment" when the contracted agent 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.

In particular, 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 contracted agent's effort level e_t in that encounter and the state of the world θ_t (which may change from one encounter to the other). In each encounter, the contracting agent offers a reward function of $r_t(q^t)$, and the contracted agent 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 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 contracting and the contracted agent by $\hat{v}$ and $\hat{x}$ respectively.

The notion of a perfect *epsilon equilibrium* is used [Radner, 1981], although it is a weaker condition than that of 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.

The first-best strategies can be sustained in perfect epsilon equilibria of the multiple encounters situation by "trigger strategies." The trigger strategy for the contracted agent denoted by ρ is very simple: it uses the effort level function $\hat{e}$ until the first encounter where the contracting agent does not use the reward function $\hat{r}$; at that encounter and in each encounter thereafter the contracted agent will optimize against the reward function announced for each encounter.

The suitable trigger strategy for the contracting agent is more complicated. In each encounter t, based on the history of outcomes through encounters $(t-1)$, the contracting agent 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 contracted agent 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 contracting agent will switch to r^* before the contracted agent ever starts cheating.

Define $C_t = f(e_t(\theta_t), \theta_t)$, i.e., C_t is the outcome in encounter t if the contracted agent 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 $S_n = C_1 + ... + C_n$. We let $\hat{C}$ denote the outcome in period t if the contracted agent uses $\hat{e}$ which is bounded by B, 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 and identically distributed since the θ_t's are. 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:

$$\tilde{N} = min\{n \geq 1 : S_n - n\hat{c} \leq -b_n\}, \quad N = min\{\tilde{N}, T\} \tag{28}$$

The following trigger strategy should be used by the contracting agent: pay the contracted agent $\hat{r}$ in each period through N and thereafter use the reward function r^*. We shall denote this strategy by $\sigma((b_n))$.

The main result of [Radner, 1981] on these strategies is as follows: For any $\epsilon > 0$ there exists a sequence (b_n) in B and T_ϵ such that for all $T \geq T_\epsilon$ the pair of strategies $(\sigma((b_n)), \rho)$ is an ϵ equilibrium, and yields the contracting and contracted agent average expected utilities respectively of at least $(\hat{v} - \epsilon)$ and $(\hat{x} - \epsilon)$.[13]

[13]In [Radner, 1985] 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 infinitely many encounters.

6.2 Long Term Contracts

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

For example, suppose that the agent is evaluated according to its average performance and there is an uncertainty about the state of the world (i.e., each single encounter is as in Section 4.2). If the contracted agent 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 contracted agent therefore has a motivation to adjust its effort over time as a function of its previous performance. As a result of this phenomenon, the optimal contracts in such situations will not be a simple function of the average outcomes [Lambert, 1983] in general. This type of behavior also arises when the number of encounters is very large. However, such behavior will eventually be detected.

The problem of subcontracting when the number of repeated encounters is small is considered in [Lambert, 1983]. It is assumed that the contracting agent 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 contracted agent's effort level (which is unobservable to the contracting agent) 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 [Lambert, 1983] and before the first encounter the contracting agent offers a binding contract. The reward in the first encounter depends upon the outcome of that encounter, but the reward of the second encounter depends upon the outcomes of the first and second encounters. If the contract is accepted by the contracted agent it should then choose the effort level of the first encounter. The outcome is observed by both agents and the contracted agent is paid according to the contract.

In the second encounter, the contracted agent 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 contracting agent chooses the contract, it should solve a maximization problem similar to that of Section 4.2. However, in its expected utility of 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 contracted agent in both encounters.

Subject to these constraints, the contracting agent is able to update the contracted agent's rewards over time in any fashion that it desires. It was shown in [Lambert, 1983] 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 contracting agent wants to contract can be performed by a group of agents. Each of the contracted agents is independent in the sense that it tries to maximize its own utility. The contracting agent offers a contract to each of the possible contracted agents. If one of them rejects the offer, than the contracting agent cannot subcontract the task[14]. Otherwise, the contracted agents can simultaneously choose effort levels.

As in previous sections, the contracting agent cannot observe the effort levels and does not supervise the group while the members carry out the task.

7.1 Individual Outcome is Observed

In this section we assume that each contracted agent yields an observable outcome of q_i and that the overall outcome will be equal to the sum of the q_is. 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_is [Nalebuff and Stiglitz, 1983b], 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 other agents' outcomes

The contracted agents 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., $q_i = f(e_i, \theta, \epsilon_i)$. For example, in the German-US robots case, θ could reflect the level of the mineral in the whole site, while ϵ_i represents the level of minerals in the exact location of contracted agent i. Each of the contracted agents observes θ *before* it chooses its effort level, but it does not observe ϵ_i before making its choice.

We assume that the contracted agents are identical, i.e., have the same utility function $U^{ced}(e, r) = v(r) - c(e)$ and the same abilities. We will assume that $f(e_i, \theta, \epsilon_i) = e_i\theta + \epsilon_i$ and

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

that $\wp_{\epsilon_i}$ 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 can 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 that a contracted agent's reward depend only on its own outcome?

When the contracted agents' outcomes are independent, then observing all the q_is provides no additional information about the contracted agent's effort. In this case, the rewards should depends only on the individual outcome.

Sometimes it is possible to find enough statistics from $q_1, ..., q_n$, denoted by $T(\{q_1, ..., q_n\})$, about the state of the world. The rewards of a specific agent should then depend upon its individual outcome and on $T(\{q_1, ..., q_n\})$ [Nalebuff and Stiglitz, 1983b].

For example, if both θ and ϵ are normally distributed random variables, then the average value of $\{q_1, ..., q_n\}$ provides sufficient statistical information for θ. When the number of contracted agents 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 for designing a contract for a group of contracted agents is to pay the contracted agents 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 contracted agents. 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 contracted agents to choose greater effort levels. A larger prize for the winner, motivates greater effort by all agents and increases the contracting agent's outcome [Nalebuff and Stiglitz, 1983b].

If the first contracted agent chooses effort level e_1, and the second chooses effort level e_2, then the first one will "win" if $\theta e_1 + \epsilon_1 > \theta e_2 + \epsilon_2$. Each of the contracted agents 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 contracted agent i is therefore,

$$\frac{1}{2}[v(r_r) + v(r_l)] - c(e_i) \tag{29}$$

The details of how to compute r_w and r_l in a given situation are described in [Nalebuff and Stiglitz, 1983b]. An interesting result from this is that in some situations it is possible to make the contracted agents choose an effort level, using the above "contest" mechanism,

which is even larger than when the contracting agent 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 contracted agents solely on the basis of the relative position of their outcomes, the contracting agent can rank one contracted agent above another if that agent's outcome is greater than its opponent's by a positive margin. 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 reward for a group: giving a reward only to the agent whose output is the highest, or by punishing the agent that came in last [Nalebuff and Stiglitz, 1983b]. Rewards that are based on relative performance are generally more flexible, and reduce the risk taken by the contracted agents [Nalebuff and Stiglitz, 1983a].

Contracted agents have private information

In this case we assume that each contracted agent's outcome 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 contracting agent cannot observe any of them. For example, if a US robot subcontracts its digging to a German robot and a French robot, then each of them can observe the level of the mineral in its own digging site *before* signing the contract and since they dig in adjacent sites, their mineral levels are correlated. The US robot, however, does not know either levels.

Suppose there are only two agents, A and B, and two output function $f^l(e^l, \theta^l)$, $l = A, B$ [Ma *et al.*, 1988]. Then we also assume that θ^l can be θ^l_1 or θ^l_2 (i.e., the world can be in four different states with two possibilities for each variable). For $l = A, B$ let $\wp(\theta^l_i)$ be the probability that $\theta^l = \theta^l_i$ for $i = 1, 2$. We denote this probability by p^l_i and assume that $p^l_i > 0$ and that $\wp(\theta^l_1) + \wp(\theta^l_2) = 1$.

As in previous sections, the level of effort, e^l is not observable. We do assume however, that for each l $f^l(e^l, \theta^l_1) < f^l(e^l, \theta^l_2)$ for all e^l, therefore, θ^l_2 represents a "good" state and θ^l_1 a bad state.

The state variables are positively but imperfectly correlated. We denote by s^A_i the probability of $\theta^B = \theta^B_1$, given that $\theta^A = \theta^A_i$ and similarly s^B_i denotes the probability of $\theta^A = \theta^A_1$ given that $\theta^B = \theta^B_i$. We assume that $1 > s^l_1 > s^l_2 > 0$.

Agent l $(= A, B)$ *privately* observes θ^l *before* signing a contract with the contracting agent. The contracting agent is risk neutral and the contracted agents are risk averse. Their utilities

functions are similar to that which appears in Section 7.1.1. Given the utility function of the contracted agent l, and the state of the world, one can compute the "disutility" of producing an outcome such as q^l. Therefore, the contracted agent's utility can be expressed as a function of the rewards and the outcome (as we did, for example, in Section 5.3). We will assume that $U^{ced,l}(q^l, r^l) = v(r^l) - d(q^l)$ and that the contracted agent's reservation price is $\hat{u}^l$.

A typical contract that can be offered by the contracting agent to agent A in this case, is of the following form [Ma *et al.*, 1988]:

> 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, 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 [Demski and Sappington, 1984] the maximization problem of the contracting agent was stated. It restricted the contracted agent's output choices to be a Bayes-Nash equilibrium, given that they are guaranteed at least their reservation price (conditional on their private information). This is done for $l = A, B$.

$$Maximize_{q_i^l, r_{ij}^l} p_1^l[s_1^l(q_1^l - r_{11}^l) + (1 - s_1^l)(q_1^l - r_{12}^l)] + p_2^l[s_2^l(q_2^l - r_{11}^2) + (1 - s_2^l)(q_2^l - r_{22}^l)] \quad (30)$$

subject to:

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

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

The result of this maximization provides the contracting agent with rewards that discourage a contracted agent from choosing output q_1^l when it has observed 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 contracting agent the highest *possible* expected outcome. If the contracting agent 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 contracting agent *will* get the maximum outcome *if both agents respond as the contracting agent desires*, i.e., sign their respective contracts and produce output q_i^l when they observe θ_i^l. In the case of a single agent, the constraints ensure that the contracted agent will choose the desired effort level. However, if there are two agents, there exists another pair of equilibrium strategies whose outcome, from the contracted agents' point of view, is better to both agents than the outcome in the equilibrium the contracting agent wants to implement. The outcome for the contracting agent if they choose that level of effort however, is low [Demski and Sappington, 1984]. In particular, there is an equilibrium for both contracted agents 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 contracting agent (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 the contracting agent will definitely be worse off.

It was suggested in [Demski and Sappington, 1984] to strengthen the incentive constraints of *one* contracted agent so that its chosen strategy will provide a better outcome for the contracting agent. But although this method does guarantee a unique equilibrium, it is also costly to the contracting agent.

Another costless method of making the contracted agents choose the "correct" strategies was suggested in [Ma *et al.*, 1988]. This method, however, makes the contracts more complicated. The main idea is that the contracting agent offers one of the contracted agents, e.g., A, a range of extra possible output options $q_1^A(\epsilon)$, indexed by ϵ where $0 < \epsilon \leq 1 - s_1^A$. The additional ϵ is costless to agent A, but it can be used as a signal to the contracting agent, e.g.,

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

A detailed mechanism based on this idea is described in [Ma *et al.*, 1988] and proves that it provides a unique equilibrium that guarantees the contracting agent its second best outcome.

7.1.2 Contracted agent's effort influences others

In this section we consider situations where the output of a contracted agent depends both on its level of effort and the other contracted agents' level of effort, and where there is symmetrical uncertainty about the state of the world.

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

If the contracting agent can observe the actions chosen by the contracted agents then, as in Section 4.1, it can offer the contracted agents a *forcing contract*.

If the contracting agent cannot observe the effort levels then the contract it should offer will specify for any sequence of outcomes $(q_{i_1}^1, ..., q_{i_k}^k)$, a sequence of k rewards denoted by $(r_{i_1,...i_k}^1, r_{i_1,...i_k}^2, ..., r_{i_1,...i_k}^k)$. Similar to the maximization problem in the case of one contracted agent, the contracting agent should maximize its expected utility given similar constraints to (IC:3) and (IR:2). A similar three steps procedure, as in the one contracted agent case of Section 4.2, can then be formalized. Given any effort level's sequence $e^1, ..., e^k$, the contracting agent should find the rewards, $r^1, ..., r^k$, that minimize the expected rewards of the contracting agents subject to the reservation utility constraint (IR:5) and *participation* constraint (IC:6) meaning that given $r^1, ..., r^k$ the contracted agents will prefer $e^1, ..., e^k$ over their other options. In some situations, depending on the probability function $\wp$ (e.g., if there is perfect correlation between the θ_is), and the possible "punishments" the contracting agent can impose on the "shrinking" contracted agents, the contracting agent 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) [Mookherjee, 1984].

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

The main question is how the contracting agent can make the contracted agents choose the set of actions it prefers. One approach is to try and strengthen the constraints that are related to contracted agents, but this of course, is costly for the contracting agent. 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 contracted agents (but not by the contracting agent).

2. actions are only privately observed

In the first case, the contracted agents pick an effort level simultaneously, and afterwards they (but not the contracting agent) 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 contracting agent can try to extract information about the effort levels from the agents and although the contracted agent can provide false information, the accuracy of this information is known to the other contracted agent. The contracting agent may then appeal to the other agents for verification.

We will consider the case where there are only two contracted agents [Ma, 1988] denoted by A and B. Suppose by using the techniques of previous sections, and assuming the contracting agent can observe the agent's actions, the contracting agent would like the two contracted agents 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 contracted agent A if the contracting agent can observe efforts, i.e., $U^{ced}(e_a^*, r_a^*) = \hat{u}$ and similarly r_b^* can be the reward for the second contracted agent. Note, that since $U^{ced}(e, r) = v(r) - c(e)$, $v(r_a^*) = \hat{u} + c(e_a^*)$.

The aim of the contracting agent is to make sure the agents find (e_a^*, e_b^*) attractive and the above utilities will then be awarded in a unique equilibrium.

We denote $\prod(e_{a_k}, e_{b_l}) = (\wp(e_{a_k}, e_{b_l}, q^i, q^j))_{i,j}$. We assume that

$$\prod(e_{a_k}, e_{b_l}) \neq \prod(e_{a_m}, e_{b_n}) \quad whenever \ (e_{a_k}, e_{b_l}) \neq (e_{a_m}, e_{b_n}) \tag{33}$$

Given, $(\tilde{e_a}, \tilde{e_b}) \in \mathit{Effort}_A \times \mathit{Effort}_B$ and $(\hat{e_a}, \hat{e_b}) \in \mathit{Effort}_A \times \mathit{Effort}_B$ let $\epsilon(\hat{d}, \tilde{d})$ be a function of (q^i, q^j), where $\hat{d} = \prod(\hat{e_a}, \hat{e_b})$ and $\tilde{d} = \prod(\tilde{e_a}, \tilde{e_b})$, and $\epsilon(\hat{d}, \tilde{d})$ satisfies

$$\sum_{ij} \epsilon(\hat{d}, \tilde{d})(q^i, q^j)\hat{d}_{ij} < 0 \quad and \quad \sum_{ij} \epsilon(\hat{d}, \tilde{d})(q^i, q^j)\tilde{d}_{ij} > 0 \tag{34}$$

The contracting agent should offer the following mechanism:

Stage 1: Both contracted agents take actions simultaneously.

Stage 1+: Contracted agents observe each other's action.

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

Table 1: A's rewards; $\delta > 0$, $\gamma > r_a^* - \underline{r}^a + \delta > 0$, $\hat{d} = \prod(\hat{e_a}, \hat{e_b})$ and $\tilde{d} = \prod(\tilde{e_a}, \tilde{e_b})$.

A's announces	$\hat{e_b} = b^*$	$\hat{e_b} \neq b^*$
B "agrees"	r_b^*	$\underline{r}^b - \delta$
B "challenges"	$r_a^* + \epsilon(\hat{d}, \tilde{d}))(q^i, q^j)$	$\underline{r}^b - \delta + \epsilon(\hat{d}, \tilde{d}))(q^i, q^j)$

Table 2: B's rewards; $\delta > 0$, $\gamma > r_a^* - \underline{r}^a + \delta > 0$, $\hat{d} = \prod(\hat{e_a}, \hat{e_b})$ and $\tilde{d} = \prod(\tilde{e_a}, \tilde{e_b})$.

Stage 2: Agent A announces a pair of effort levels: $(\hat{e_a}, \hat{e_b})$ where $(\hat{e_a}, \hat{e_b}) \in \mathit{Effort}^A \times E^B$.

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 $(\tilde{e_a}, \tilde{e_b}) \in \mathit{Effort}^A \times E^B$ 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 Tables 1 and 2. We denote by $\underline{r}^a$ the reward that satisfies $v(\underline{r}^a) = \hat{u} + min_{e_a}\{c(e_a)|e_a \in \mathit{Effort}_A\}$ and similarly for B we denote by $\underline{r}^b$ the reward that satisfies $v(\underline{r}^b) = \hat{u} + min_{e_b}\{c(e_b)|e_b \in \mathit{Effort}_b\}$.

It was shown in [Ma, 1988] that the following strategies form a unique perfect equilibrium of the described mechanism: Agent A chooses e_a^* at Stage 1, and at Stage 2 reports honestly, whatever 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 contracting agent elicits information from agent A and uses B's reaction as a policing device. If B accuses A in lying, its outcome depends on ϵ. However, due to assumption 34, the expected outcome from ϵ to B is valuable, if and only if, A has lied. In addition, given that the contracted agents 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 contracted agents [Ma, 1988].

In the case that actions are only privately observed, it is not possible to implement the results of perfect supervision (i.e., the first best contract, where the result is that the contracting agent observes the contracted agents' 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 contracting agent. It is possible however, that the contracted

agents may be better off (given the suggested rewards) if they *all* deviated from the required actions. In [Ma, 1988] a multi-stage mechanism is presented that makes the contracted agents choose the appropriate actions of the second best contract.

7.2 Individual Outcome is Not Observed

There are other situations in which the contracting agent cannot observe the individual outcome (or such an outcome does not exist), but rather can only observe the overall outcome of all the agents' effort [Holmstrom, 1982; Rasmusen, 1987]. Even in the case of certainty, i.e., the state of the world is known, there is a problem in making the contracted agents take the preferred level of action, since there is no way for the contracting agent to find out the effort level of each of the individual agent, given the overall output. For example, suppose two robots agreed to dig minerals, but they both put the minerals in the same truck, it is not possible then to figure out who dug what. If the contracting agent wants the contracted agents to take the vector of effort level e^* 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 $U^{ced}(e_i^*, b_i) \geq \hat{u}_i$. That is, if all agents choose the appropriate effort level, each of them gets b_i and if any of them does not, all get nothing.

In some cases the contracted agents take sequential actions. That is, agent 1 chooses its effort level and performs its part of the task which is observed by the other contracted agents, but not the contracting agent. The second contracted agent then, chooses its effort level, based on the first agent's actions, and its effort level is observed by the other contracted agents, and so on. After the last agent finishes its part, the outcome of the whole sequence is figured out and observed by all agents (including the contracting agent). 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: $f(e^1, ...e^n) = z(e^1, ..., e^2) + \epsilon$. If, no matter how low the effort levels exerted by contracted agents $1, ..., i$ are, it is possible for the rest of the agents $i+1, ..., n$ to compensate for the slack also and if for fixed effort levels $e^1, ..., e^i$, z is a monotonic function from the effort level of the rest of the contracted agents, then the contracting agent can construct a contract in which it can obtain its first best outcome [Banerjee and Beggs, 1989]. The contract enables agent i, whose choice of effort level is a function of the effort levels of agents $1, ...i-1$, to use its monitoring capability effectively.

Another interesting situation is when a group of contracted agents 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 to all of them. An even more efficient result may be obtained if the contracted agents work as a team and

share the outcome. Such a situation may occur, for example, if all the contracted agents are German robots, that have the same general task to maximize Germany's profits [Machostadler and Pérez-Castrillo, 1991].

8 Conclusions

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

We considered several such situations and described the maximization problems that should be solved by the contracting agent in order to design a beneficial contract to itself. In most of the situations we also presented procedures that can be used for solving these maximization problems by the contracting agent.

Currently, there are several optimization computer packages (e.g., Nag [Nag, 1991]) available using all sorts of practical optimization methods [Fletcher, 1987] 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 contracted agents computational task is easier. In most of the situations, given a contract, the contracted agent needs only to check the validity of the inequalities that appear as constraints in the contracting agent's maximization problem. The contracted agent needs to check the validity of the individual-rationality constraint (**IR**) 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 contracted agent needs to decide which effort level to provide, it should consider its expected utility from its effort level, similar to maximization problem described in the *participation* constraints (IC).

We present below a summary of the results for the different situations considered in this paper.

When two agents have full information about each other, contracts can be signed without a delay. The results of contracting in full information situations are as follows:

1. If the contracting agent can observe and supervise the contracted agent's actions (Section 4), then it can force the contracted agent to provide the effort level preferred by the contracting agent, and thus the contracting agent maximizes its utility. The contracted agent obtains its reservation price.

2. If the contracting agent does not supervise the contracted agent's actions, but there is

full information and no uncertainty concerning the outcome of the contracted agent'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 contracting agent's supervision.

3. If there is uncertainty in the environment but the contracted agent is risk neutral (Section 4.2.1) then the contracting agent's utility will be as in the previous two cases (i.e., the agents reach a *first best* contract). The *expected* utility of the contracted agent will be equal to its reservation price; however, its actual outcome may be less than its reservation price.

4. If there is uncertainty as in the previous case, but the contracted agent is risk averse (Section 4.2.2) then the contracting agent's expected utility will be lower than in the previous case (i.e., the agents reach a *second best* contract). The contracted agent's expected utility is higher than its reservation price.

5. Monitoring (Section 4.2.3) cannot improve the contracting agent's utility in case 3 above, but may increase its utility in the previous case (4) when the contracted agent is neutral toward risk.

If there is asymmetric information 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 contracted agent to honestly report its private information. Below is a summary of the results of main cases in asymmetric information situations:

1. If the contracted agent knows the state of the world but the contracting agent does not (Section 5.1), then the contracting agent's expected utility is lower than if they have symmetric beliefs and the contracted agent's expected utility is higher.

2. If the contracted agent is able to collect more information *before* it performs the agreed-upon task but only after signing the contract, and the contracted agent cannot opt out after signing an agreement (Section 5.2), then the contracting agent can get is second best utility if the contracted agent is risk neutral.

3. If the contracting agent also has private information (Section 5.5), but its private information does not directly influence the contracted agent's utilities, then in most of the situations, there exists a mechanism in which all types of the contracting agents do strictly better than the fully informed contracted agent (i.e., even better than in the first best contract).

4. If there are several agents in the environment (Section 5.7) in most situations the contracting agent can design a *second best* contract.

When there are more than one encounter between the agents (Section 6) 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 are large enough, even in asymmetric information situations, they can reach first-best contracts.

2. If the number of encounters is small, enforceable long term encounters are more beneficial to the contracting agent. However, it is still difficult to design an efficient contract.

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

The most important problem that a designer of an agent faces in a CMA environment, is which utility function to provide its agent with. Of course, the personality of the designer (his/her attitude toward risk) will affect this decision, but computational considerations should also be taken into consideration. It is clear that when the agents are risk neutral, all the maximization problems presented in this paper are much easier to solve. Furthermore, more efficient results are obtained in such situations.

However, if the designer would like its agent to be risk averse, then the utility function should be chosen carefully. In order to support most of the results presented in this paper, the contracted agent's utility function should be additively separable in rewards and effort of the form $U^{ced}(e,r) = v(r) - c(e)$ where $v' > 0, v'' \leq 0, c' > 0$ and $c'' \geq 0$.

If there is more than one possible contracted agent in the environment, the measure of risk aversion should be considered by the designer. A less risk-averse agent will usually have the ability to win over more risk-averse agents in service of *any* risk averse contracting agent [Ross, 1979].

We are now in the process of applying the techniques presented in this paper to the performance of trucks in the Truckworld [Nguyen *et al.*, 1993] simulation environment.

References

[Arrow, 1985] K. J. Arrow. The economics of agency. In J. Pratt and R. Zeckhauser, editors, *Principals and Agents: The Structure of Business*, pages 37–51. Harvard Business School Press, 1985.

[Baiman and Demski, 1980] S. Baiman and J. Demski. Economically optimal performance evaluation and control systems. *Journal of Accounting Research*, 18S:184–220, 1980.

[Banerjee and Beggs, 1989] A. Banerjee and A. Beggs. Efficiency in hierarchies: implementing the first-best solution by sequential actions. *The Rand Journal of Economics*, 20(4):637–645, 1989.

[Bond and Gasser, 1988] A. H. Bond and L. Gasser. An analysis of problems and research in DAI. In A. H. Bond and L. Gasser, editors, *Readings in Distributed Artificial Intelligence*, pages 3–35. Morgan Kaufmann Publishers, Inc., San Mateo, California, 1988.

[Caillaud *et al.*, 1988] B. Caillaud, R. Guesnerie, P. Rey, and J. Tirole. Government intervention in production and incentives theory: a review of recent contributions. *Rand Journal of Economics*, 19(1):1–26, 1988.

[Christensen, 1981] J. Christensen. Communication in agencies. *Bell Journal of Economics*, 12:661—674, 1981.

[Conry *et al.*, 1990] S.E. Conry, D. J. MacIntosh, and R.A. Meyer. DARES: A distributed automated REasoning system. In *Proc. of AAAI-90*, pages 78–85, Boston, MA, 1990.

[Corkill, 1979] D. Corkill. Hierarchical planning in a distributed environment. In *Proceedings of the Sixth International Joint Conference on Artificial Intelligence*, pages 168–175, Tokyo, August 1979.

[Demougin, 1989] D. Demougin. A renegotiation-proof mechanism for a principle-agent model with moral hazard and adverse selection. *The Rand Journal of Economics*, 20:256–267, 1989.

[Demski and Sappington, 1984] J. S. Demski and D. Sappington. Optimal incentive contracts with multiple agents. *Journal of Economic Theory*, 33:152–171, 1984.

[Douglas, 1989] E. J. Douglas. The simple analytics of the principal-agent incentive contract. *The Journal of Economic Education*, 20(1):39–51, 1989.

[Durfee, 1988] E. H. Durfee. *Coordination of Distributed Problem Solvers.* Kluwer Academic Publishers, Boston, 1988.

[Durfee, 1992] E. Durfee. What your computer really needs to know, you learned in kindergarten. In *Proc. of AAAI-92*, pages 858–864, California, 1992.

[Ephrati and Rosenschein, 1991] E. Ephrati and J. Rosenschein. The clarke tax as a consensus mechanism among automated agents. In *Proc. of AAAI-91*, pages 173–178, California, 1991.

[Ephrati and Rosenschein, 1993] E. Ephrati and J. S. Rosenschein. Planning to please: Following another agent's intended plan. *Journal of Group Decision and Negotiation*, 1993. To appear.

[Fletcher, 1987] R. Fletcher. *Practical Methods of Optimization.* John Wiley & Sons, 1987.

[French, 1986] S. French. *Decision Theory: An Introduction to the Mathematics of Rationality.* Ellis Horwood Limited, 1986.

[Fudenberg and Tirole, 1990] D. Fudenberg and J. Tirole. Moral hazard and renegotiation in agency contracts. *Econometrica*, 58(6):1279–1319, 1990.

[Gasser, 1991] L. Gasser. Social concepts of knowledge and action: DAI foundations and open systems semantics. *Artificial Intelligence*, 47(1–3):107–138, 1991.

[Gjesdal, 1982] F. Gjesdal. Information and incentives: The agency information problem. *Review of Economic Studies*, ?:373–390, 1982.

[Grossman and Hart, 1983] S. Grossman and O. Hart. An anaysis of the principal-agent problem. *Econometrica*, 51(1):7–45, 1983.

[Grosz and Kraus, 1993] B. Grosz and S. Kraus. Collaborative plans for group activities. In *IJCAI-93*, pages 367–373, French, 1993.

[Harris and Raviv, 1978] M. Harris and A. Raviv. Some results on incentive contracts with applications to education and employment, health insurance, and law enforcement. *The American Economic Review*, 68(1):20–30, 1978.

[Harris and Raviv, 1979] M. Harris and A. Raviv. Optimal incentive contracts with imperfect information. *J. Economics Theory*, 20:231–259, 1979.

[Harris and Townsend, 1981] M. Harris and R. Townsend. Resource allocation under asymmetric information. *Econometrica*, 49:33–64, 1981.

[Harsanyi, 1967 1968] J. Harsanyi. Games with incomplete information played by bayesian players. *Management Science*, 14:159–182,320–334,486—502, 1967-1968.

[Hart and Tirole, 1988] O. D. Hart and J. Tirole. Contract renegotiation and coasian dynamics. *Review of Economic Studies*, LV:509–540, 1988.

[Hirshleifer and Riley, 1992] J. Hirshleifer and J. Riley. *The Analytics of Uncertainty and Information.* Cambridge University Press, Cambridge, 1992.

[Holmstrom, 1982] B. Holmstrom. Moral hazard in teams. *Bell Journal of Economics*, 13(2):324—340, 1982.

[Kraus and Lehmann, 1994] S. Kraus and D. Lehmann. Designing and building a negotiating automated agent. *Computational Intelligence*, 1994. To appear.

[Kraus and Wilkenfeld, 1991a] S. Kraus and J. Wilkenfeld. The function of time in cooperative negotiations. In *Proc. of AAAI-91*, pages 179–184, California, 1991.

[Kraus and Wilkenfeld, 1991b] S. Kraus and J. Wilkenfeld. Negotiations over time in a multi agent environment: Preliminary report. In *Proc. of IJCAI-91*, pages 56–61, Australia, 1991.

[Kraus and Wilkenfeld, 1993] S. Kraus and J. Wilkenfeld. The updating of beliefs in negotiations under time constraints with uncertainty. In *Proc. of IJCAI-93 workshop on Artificial Economics*, pages 57–68, 1993. Also presented in BISFAI93.

[Kraus *et al.*, 1993] S. Kraus, M. Nirkhe, and K. Sycara. Reaching agreements through argumentation: a logical model. In *Proc. of DAI93*, 1993. Also presented in AAAI-93 workshop on AI theories of Groups and Organizations: Conceptual and Empirical Research.

[Kraus *et al.*, 1994] S. Kraus, J. Wilkenfeld, and G. Zlotkin. Multiagent negotiation under time constraints. *Artificial Intelligence*, 1994. To appear.

[Kraus, 1993] S. Kraus. Agents contracting tasks in non-collaborative environments. In *Proc. of AAAI-93*, pages 243–248, Washington, D.C., 1993.

[Laffont and Tirole, 1990] J. Laffont and J. Tirole. Adverse selection and renegotiation in procurement. *Review Economics Studies*, 57:597–625, 1990.

[Laffont *et al.*, 1987] J. Laffont, E. Maskin, and J. Rochet. Optimal nonlinear pricing with two-dimensional characteristic. In T. Groves, R. Radner, and S. Reiter, editors, *Information, Incentives, and Economic Mechanisms*, pages 256–266. University of Minnesota Press, 1987.

[Lambert, 1983] R. A. Lambert. Long term contracts and moral hazard. *Bell Journal of Economics*, 14:441–452, 1983.

[Lesser and Erman, 1980] V. R. Lesser and L. D. Erman. Distributed interpretation:a model and experiment. *IEEE Transaction on Computers*, 29(12):1144–1163, 1980.

[Lesser, 1990] Victor R. Lesser. Distributed problem solving. In Stuart C. Shapiro, editor, *Encyclopedia of Artificial Intelligence*, pages 245–251. John Wiley and Sons, New York, 1990.

[Lesser, 1991] V.R. Lesser. A retrospective view of fa/c distributed problem solving. *IEEE Transactions on Systems, Man, and Cybernetics, Special Issue on Distributed Artificial Intelligence*, 21(6):1347–1362, December 1991.

[Ma *et al.*, 1988] C. Ma, J. Moore, and S. Turnbull. Stopping agents from "cheating". *Journal of Economic Theory*, 46:355–372, 1988.

[Ma, 1988] C. Ma. Unique implementation of incentive contracts with many agents. *Review of Economic Studies*, 51(3):555–572, 1988.

[Macho-stadler and Pérez-Castrillo, 1991] I. Macho-stadler and J. Pérez-Castrillo. Moral hazard and cooperation. *Economics Letters*, 35:17–20, 1991.

[Malcomson and Spinnewyn, 1988] J. Malcomson and F. Spinnewyn. The multiperiod principal-agent problem. *Review of Economic Studies*, LV:391–408, 1988.

[Malone *et al.*, 1988] 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, editor, *The Ecology of Computation*, pages 177–205. North Holland, 1988.

[Maskin and Tirole, 1990] E. Maskin and J. Tirole. The principal-agent relationship with an informed principal: The case of private values. *Econometrica*, 58(2):379–409, 1990.

[Maskin and Tirole, 1992] E. Maskin and J. Tirole. The principal-agent relationship with an informed principal ii:common values. *Econometrica*, 60:1–42, 1992.

[Matthews, 1983] S. Matthews. Selling to risk averse buyers with unobservable tastes. *Journal of Economy Theory*, 30:370–400, 1983.

[McAfee and McMillan, 1986] R. P. McAfee and J. McMillan. Bidding for contracts: a principal-agent analysis. *The Rand Journal of Economics*, 17(3):326–338, 1986.

[McAfee and McMillan, 1987] R. P. McAfee and J. McMillan. Competition for agency contracts. *The Rand Journal of Economics*, 18(2):296–307, 1987.

[McAfee and McMillan, 1988] R. McAfee and J. McMillan. Multidimensional incentive compatibility and mechanism design. *Journal of Economic Theory*, 46:335–354, 1988.

[Melumad and Reichelstein, 1989] N. D. Melumad and S. Reichelstein. Value of communication in agencies. *Journal of Economic Theory*, 47:334–368, 1989.

[Mookherjee, 1984] D. Mookherjee. Optimal incentive schemes with many agents. *Review of Economic Studies*, 51(3):433—446, 1984.

[Myerson, 1982] R. Myerson. Optimal coordination mechanisms in generalized principal-agent problem. *Journal of Mathematical Economics*, 10:67–81, 1982.

[Myerson, 1983] R. Myerson. Mechanism design by an informed principal. *Econometrica*, 51:1767—1798, 1983.

[Nag, 1991] The NAG Fortran library manual, mark 15. Technical report, Nag Ltd., Oxford, United Kingdom, 1991.

[Nalebuff and Stiglitz, 1983a] B. Nalebuff and J. Stiglitz. Information, competition, and markets. *American Economic Review*, 73(2):278–283, 1983.

[Nalebuff and Stiglitz, 1983b] B. Nalebuff and J. Stiglitz. Prizes and incentives: towards a general theory of compensation and competition. *Bell Journal of Economics*, 14(1):21–43, 1983.

[Nguyen *et al.*, 1993] D. Nguyen, S. Hhanks, and C. Thomas. The TRUCKWORLD manual. Technical Report TR 93-09-08, Dept. of Computer Science and Engineering, Univ. of Washington, 1993.

[Page, 1987] F. Page. The existence of optimal contracts in the principal-agent model. *Journal of Mathematical Economics*, 16(2):157–167, 1987.

[Pattison *et al.*, 1987] H. Pattison, D. Corkill, and V. Lesser. Instantiating descriptions of organizational structures. In Michael N. Huhns, editor, *Distributed Artificial Intelligence*, pages 59—96. Pitman Publishing/Morgan Kaufman Publishers, San Matheo, CA, 1987.

[Pfaffenberger and Walker, 1976] R. Pfaffenberger and D. Walker. *Mathematical Programming for Economics and Business.* The IOWA State University Press, Ames, IOWA, 1976.

[Radner, 1981] R. Radner. Monitoring cooperative agreements in a repeated principal-agent relationships. *Econometrica*, 49(5):1127–1148, 1981.

[Radner, 1985] R. Radner. Repeated principal agents games with discounting. *Econometrica*, 53:1173–1198, 1985.

[Raizsmusen, 1989] E. Raizsmusen. *Games and Information.* Basil Blackwell Ltd., Cambridge, Ma, 1989.

[Rasmusen, 1987] E. Rasmusen. Moral hazard in risk-averse teams. *Rand Journal of Economics*, 18(3):324—340, 1987.

[Rogerson, 1985] W. Rogerson. The first-order approach to the principle-agent problems. *Econometrica*, 53(6):1357–1367, 1985.

[Rosenschein, 1982] Jeffrey S. Rosenschein. Synchronization of multi-agent plans. In *Proceedings of the National Conference on Artificial Intelligence*, pages 115–119, Pittsburgh, Pennsylvania, August 1982.

[Ross, 1973] S. Ross. The economic theory of agency: The principal's problem. *The American Economic Review*, 63(2):134–139, 1973.

[Ross, 1979] S. Ross. Equilibrium and agency-inadmissible agents and in the public agency problem. *The American Economic Review*, 69:308–312, 1979.

[Rubinstein and Yaari, 1983] A. Rubinstein and M. Yaari. Repeated insurance contracts and moral hazard. *Journal of Economic Theory*, 30:74–97, 1983.

[Sandholm, 1993] T. Sandholm. An implementation of the contract net protocol based on marginal cost calculations. In *Proc. of AAAI-93*, pages 256–262, Washington D.C., 1993.

[Sappington, 1984] D. Sappington. Incentive contracting with asymmetric and imperfect precontractual knowledge. *Journal of Economic Theory*, 34:52–70, 1984.

[Selten, 1975] R. Selten. Re-examination of the perfectness concept for equilibrium points in extensive games. *International Journal of Game Theory*, 4:25–55, 1975.

[Shappington, 1983] D. Shappington. Limited liability contracts between principle and agent. *Journal of Economic Theory*, 29:1–21, 1983.

[Shavell, 1979] S. Shavell. Risk sharing and incentives in the principal and agent relationship. *Bell Journal of Economics*, 10:55–79, 1979.

[Smith and Davis, 1981] R. Smith and R. Davis. Framework for cooperation in distributed problem solvers. *IEEE Transactions on Systems, Man and Cybernetic*, C-29(12):61–70, 1981.

[Smith and Davis, 1983] R.G. Smith and R. Davis. Negotiation as a metaphor for distributed problem solving. *Artificial Intelligence*, 20:63–109, 1983.

[Spence and Zeckhauser, 1971] M. Spence and R. Zeckhauser. Insurance, information and individual action. *The American Economic Review*, 61(1):380–391, 1971.

[Spivey and Thrall, 1970] W. Spivey and R. Thrall. *Linear Optimization.* Holt, Rinehart and Winston, INC, 1970.

[Sycara, 1987] K.P. Sycara. *Resolving Adversarial Conflicts: An Approach to Integrating Case-Based and Analytic Methods.* PhD thesis, School of Information and Computer Science, Georgia Institute of Technology, 1987.

[Sycara, 1990] K. P. Sycara. Persuasive argumentation in negotiation. *Theory and Decision*, 28:203–242, 1990.

[Wellman and Doyle, 1992] M. Wellman and J. Doyle. Modular utility representation for decision-theoretic planning. In *Proc. of AI planning Systems*, pages 236—242, Maryland, 1992.

[Wellman, 1992] M. Wellman. A general-equilibrium approach to distributed transportation planning. In *Proc. of AAAI-92*, pages 282–289, San Jose, California, 1992.

[Werner, 1988] Eric Werner. Toward a theory of communication and cooperation for multi-agent planning. In *Proceedings of the Second Conference on Theoretical Aspects of Reasoning about Knowledge*, pages 129–143, Pacific Grove, California, March 1988.

[Zlotkin and Rosenschein, 1991] G. Zlotkin and J. Rosenschein. Incomplete information and deception in multi-agent negotiation. In *Proc. IJCAI-91*, pages 225–231, Australia, 1991.

[Zlotkin and Rosenschein, 1993] G. Zlotkin and J. S. Rosenschein. A domain theory for task oriented negotiation. In *Proceedings of IJCAI-93*, pages 416–422, Chambery, France, August 1993.