Using Option Pricing to Value Commitment Flexibility in Multi-agent Systems[1]

Katia Sycara

September 1997

CMU-CS-97-169

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Author affiliation: Robotics Institute, Carnegie Mellon University, katia@ri.cmu.edu

[1]This material is based upon work supported in part by ONR grant N-00014-96-1-1222 and by NSF grant IRI-9612131.

Abstract

With the explosive growth of internet activity, there will be an increasing reliance on intelligent software agents for electronic commerce and information retrieval. Such multi-agents systems will be comprised of self-motivated agents that interact with each other though negotiation and task delegation. Multi-agent technology models and facilitates these interactions through automated contracting. We develop a domain independent computational model to study in a uniform manner many complex issues that arise in multi-agent contracting, such as modeling commitment flexibility in a contract, valuing a contract under assumptions of uncertainty, risk reduction, making decisions in situations of asymmetric information, or situations of sequential subcontracting where each agent must decide to subcontract part of its current contract to others. Our model is based on *financial option pricing* theory. We believe that modeling contracts as options provides a natural unified framework for taking into account contracting flexibility and complex forms of environmental uncertainty. In addition, option pricing provides a computationally tractable formalism for calculating optimal values of various contracting decision parameters, that to date have not been rigorously modeled. Such parameters include the *value* of a flexible/contingent contract, *when* to give out a contract to a contractee, *when* to break a contract, and *which* contract to accept out of a set of offered contracts. Under our model these aspects of contracting can be explored analytically and experimentally. Moreover, there are some aspects of contracting that have no analogues in financial options. These include contract quality guarantees and multiple sequential sub-contracting. We extend option pricing theory in interesting ways to model such contracts.

Keywords: Software agents, Multi agent systems, Automated negotiation, Contingent contracts, Decommitment, Financial options

1 Introduction

The area of automated negotiation and contracting has been of particular interest due to the important role it can play in facilitating understanding and the achievement of mutually-acceptable deals among entities with differing interests, whether they be individuals, companies, governments, or automated agents. Recent growing interest in autonomous software agents [32, 41, 43] and their potential application in areas such as supply contracting[1], filtering news [22], intelligent information-retrieval [29, 6], investment portfolio management [43, 44, 46], performing secretarial functions [25, 27, 45], electronic commerce [17, 39], electricity transport management, and telecommunication [16], has given increased importance to automated negotiation and contracting. Our long term objective is to develop a computational model of contracting and build automated agents that enter into contracting agreements that demonstrate a variety of complicated characteristics, such as flexibility of contractual commitments, dealing with incomplete information, complex forms of uncertainty and risk.

Software agents contract for services to other agents (or human users). Each agent may handle requests from several other agents and may be in a position to choose which requests it will honor in order to use its local resources most effectively. Correspondingly, an agent chooses to give a contract to the agent who offers the most attractive deal. Thus, in the most general case there is an *electronic marketplace* consisting of *self-interested* agents that have their own goals and resources, and follow their own strategies (e.g. [21], [37]) to maximize their payoffs. As agent populations grow and increased volume of business is transacted in cyberspace, it seems inevitable [28] that there will be increasing use of agents that automatically seek out offers of goods and services, negotiate prices, and make purchases. The design and analysis of interaction protocols for such agents is part of the growing field of *automated negotiation systems*[30, 35]. A major component of automated negotiation is *contracting:* an agent, the *contractor,* gives a job to another agent, the *contractee,* with certain provisions. Traditionally, in the DAI/multi-agent systems literature, the contracts considered are binding. A contract is *binding* for an agent if the agent cannot get out of its contractual obligations. When an agent can get out of a contractual obligation, the contract is called *non-binding* or *contingent*. Contingent contracts allow agents increased flexibility. It has been shown [38] that in many situations non-binding contracts are superior to binding ones. Introducing contingent contracts has two main advantages: (1) The space of possible contracts is enhanced, so the expected utility can be higher, and (2) Contingent contracts can reduce the variability of an agent's payoff, since an agent can postpone a decision for the future when more information could be available.

In this paper we present a model that will be useful for studying contingent contracts involving multiple contractors and contractees in an uncertain environment. Consider the problem of maximizing the expected utility when contingent contracts in addition to binding ones are allowed. One immediate question is: What should be the price of these contingent contracts? When an agent is seeking other agents to perform tasks, or supply goods and services, there may be several candidates available, offering *potential contracts* with various conditions and at various prices. Similarly, each contractee agent may bid for several different contracts. In order to select the most profitable contract, (or bid profitably for contracts), contractors and contractees need to evaluate how much each contract is worth. In game theory, the value of the contract is assumed known and used as input to the game-theoretic solution concepts such as Nash equilibrium and its extensions (e.g., sequential equilibrium, perfect Bayesian equilibrium [12]). However, this approach do not address issues such as contract valuation, contract flexibility, or the nonstationary nature of the underlying uncertainty. More sophisticated computational mechanisms are needed.

Besides the issue of estimating the price of a contingent contract, other important questions are raised, such as *when* is it optimal for an agent to decommit? It must be noted that decommitment is just one kind of flexibility. One can envisage contracts where other features of a contract (e.g., quality of contract results) are allowed to depend on future events. We address these questions in our models.

Our models are based on *financial option pricing* theory. We believe that modeling contingent contracts under time-dependent uncertainty and risk as options provides a natural unified framework for taking into account contracting flexibility and complex forms of environmental uncertainty. In addition, option pricing provides a computationally tractable formalism for calculating optimal values of various contracting decision parameters, that to date have not been rigorously modeled. Such parameters include the value of a flexible/contingent contract, when to give out a contract to a contractee, when to break a contract, and which contract to accept out of a set of offered contracts. In our models these aspects of contracting can be explored analytically and experimentally. There are some aspects of contracting that have no analogues in financial options. These include contracts with quality guarantees, and how to handle multiple sequential sub-contracting. We extend option pricing theory in interesting ways to model such contracts.

1.1 Related Work

In the DAI literature, contracting has been used as a metaphor for task allocation. In the Contract Net protocol,[42], 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). In the contract net, the agents are not self-interested, thus no fee is specified as an inducement to the agent to submit a bid for an advertised contract. In addition, if the manager breaks the contract, the contractor simply abandons the task without monetary compensation. In studies of formalizing commitments in joint multi-agent plans [32, 14], in studies of decommitment in meeting scheduling [40], or in cooperative coordination [7], although the agents do not explicitly form contracts, they are assumed to be not self-interested. In these settings, there is no need for more complex mechanisms, e.g mechanisms to value non-binding contracts or to calculate optimal timing for decommiting from contractual obligations.

Most existing game-theoretic work on automated negotiation (e.g., [33, 35, 47]) has focused on the design of protocols and strategies for agents to arrive at mutually agreeable deals, or socially/globally desirable states. Such research assumes that the value of each potential deal in the space of possible deals is either known or easily calculated. For example in [36] potential contracts are valued based on the *marginal cost* of taking on an additional set of deliveries in a transportation delivery domain. In these works, future uncertainty is usually not taken into account, thus considerably simplifying the valuation problem.

Most contracts in the multi-agent literature have been *binding*, i.e., neither party can abort a contract once it has been entered into [20, 35, 36]. The lack of flexibility of such contracts is a serious limitation: the profitability of a contract may increase or decrease considerably depending on uncertain future events. Game theory has proposed *contingency contracts* to take into consideration the potential afforded by probabilistically known future events. Contingency contracts [1] specify contractual obligations that are made contingent on future events. Such contingency contracts, though potentially more beneficial than binding contracts, suffer from a number of disadvantages, e.g. impossible to enumerate all possible relevant future events [38, 36]. One type of flexibility is afforded by the ability to *decommit* from the obligations of the contract. Sandholm and Lesser [38] present arguments for the usefulness of such contracts. They consider allowing unilateral decommitment by either party of a contract at any point in time. The party that aborts the contract must pay a *decommitment penalty* to the other party. We believe that our valuation methods, based on option pricing theory, will provide a rigorous and equitable way of calculating such penalties.

In [31] the future uncertainties are modelled as absorbing Markov chains. Using this Markov chain an agent computes the expected payoff from a contract. This expected payoff is used by the agents to decide between contracts. In this work, the contracts are binding and uncertainty is modelled by a stationary process, a Markov chain. In our work, we would like to model contingent contracts under non-stationary uncertainty processes. In [21] efficient ways to reach contracting are considered. The focus of the paper is more on negotiation rather than modeling and valuing contingent contracts.

The idea of applying option pricing theory to "real-world" investment and pricing decisions is not new. Indeed, there is a growing interest in *real option theory,* [4, 5, 9, 8, 11, 26] which applies financial options theory to the optimal timing and valuation of irreversible investment decisions. The investor is viewed as having the "option" to invest, which he may "exercise" at any time. However, the use of option pricing theory to value contingent contracts is new, and gives rise to new problems that have not been addressed by previous research in real option theory. For example, prior research does not consider decommitment in contracts, contract quality guarantees, or sequential subcontracting.

To put our research in context, we mention that the economic perspective [24, 48]is gaining popularity in computer science in general and in AI in particular. For example,

[1] We use the word *contingent* or *flexible* contracts in our document to distinguish from the contingency contracts of game theory. In contingent contracts, the particular future contingencies do not need to be explicitly spelled out.

in [49] the argument is made that the economic paradigm should be adopted in situations where (a) there is a problem of allocation of limited resources, (b) rational behavior can be assumed, and (c) authority and activity are decentralized. All three of these assumptions are valid for communities of intelligent agents. In particular, the rationality assumption is perhaps more valid for computerized agents than humans since presumably the former have much more computational powers than the latter. This point has been made by Varian [47], and by Rosenschein and Zlotkin [35].

2 Background

Negotiation of self-interested agents has been extensively studied in game theory. Game theoretic models make the following restrictive assumptions: [2] (1) Both the number of players and their identity are assumed to be fixed and known to everyone. (2) All the players are assumed to be fully rational, and each player knows that the others are rational (common knowledge). Each player's alternative set is fixed and known. (3) Each player's risk-taking attitude and expected-utility calculations are also fixed and known to each and every individual involved in decision making. These assumptions limit the applicability of game theoretic frameworks for solving realistic problems. The search for determinate rational decisions within the framework of game theory has not led to a general model governing rational choice in interdependent situations.

Moreover, the evaluation of an outcome in game theory is one-shot, i.e. it is the agent's expected (final) payoff, that is then propagated backwards to calculate the expected payoff at each stage of the game. In addition, in game theoretic frameworks, dynamic risk attitudes are hard to handle computationally and for simplicity, an agent is typically assumed risk neutral. Furthermore, although game theory does not explicitly mention contract execution, it is assumed that contracts are binding.

Decision making under uncertainty has been extensively discussed in game theoretic[12] and decision analysis[33] literature. Most of these models assume that the decisions to be made are *now or never* propositions, that is, if the decision maker does not undertake the decision now, he will not be able to in the future. Furthermore, if the decision maker commits himself to certain decisions, he is not able to change his commitment regardless of how unfavorable the future might turn out to be. In traditional decision analysis under uncertainty and game theory, the possible outcomes of an action depend on the uncertainty in the environment and on other players' possible actions. In the face of such uncertainty, it may be beneficial for an agent, and possibly for the agent society to allow flexibility in agent commitments.

[2] It should be noted that some of the very recent game theoretic models are directly motivated by considerations of dropping or relaxing some of these assumptions. Although there has been interesting progress reported in the literature (e.g., [18]), the fundamental framework and methodology of game theory remains almost the same and it might be too early to tell whether these new results will reshape the current game theoretic framework.

In a single agent environment, we may view commitment to a decision as a contract of the decision maker with himself [13, 34]. In a multi-agent environment, most of the time decisions can be viewed as inter-agent contracts. In traditional models, decisions made by a decision maker are binding. Binding decisions are appropriate for some types of tasks and environments (e.g., a static world); in most realistic situations they are not. Recent years have seen tremendous interest in models which are able to deal with more flexible types of decisions. One of the major types of flexible decisions is decisions which could be *deferred* for the future when more information could be obtained. The notion of a *contingent contract* can be used to describe this type of flexible decisions. The basic idea is parallel to the concept of "financial call option" – the *right but not the obligation* to make an investment at some future time of the investor's choosing. Intuitively, by introducing contingent contracts into the decision space available to the decision maker, we naturally enlarge the space of possible decisions from which the decision maker chooses his action. Assuming everything else being equal, we know that the decision maker's payoff can benefit given these flexibilities. This is from the individual decision makers' point of view. If we consider a multi-agent scenario, we might as well improve the overall payoffs (social welfare) by allowing individual decision makers to make more flexible contracts. This is significant for developing computational multi-agent coordination/contracting mechanism in a DAI multi-agent setting.

Option pricing theory is motivated to answer questions regarding how to make decisions given contingent contracts. For some simple scenarios, where the uncertainty can be captured in simple probabilistic models, the traditional decision making under uncertainty techniques can be adapted to deal with contingent contracts. However, traditional decision making theory cannot cope with more complicated situations. By contrast, option pricing theory offers much more powerful representational and computational mechanisms which can be used to address the more difficult and general situations, such as the ones involving non-stationary random processes and/or multiple uncertainties which are interwoven. There are many technical difficulties when allowing contingent contracts that make traditional methods obsolete. For instance, in decision analysis with contingent contracts, we allow the decision maker to have the opportunity to exercise his option at any time point. Mathematically, this entails the introduction of time-dependent processes into the model. Traditional models don't provide computationally tractable methods to address general time-dependent (non-stationary) random processes. Another example is how to deal with continuous information gathering/updating which can happen before the decision maker makes any commitment. Option pricing theory whose baseline mathematical model is based on a fairly general stochastic optimal control framework, provides satisfactory answers for modeling and evaluating these complicated phenomena.

The model proposed in this research is intended to handle contracting flexibility among self-interested agents in situations that the game theoretic model cannot cover, such as settings with complex underlying environmental uncertainty (non-stationary stochastic processes), where expected payoff calculations cannot fully model and take into consideration time dependent probabilistic interactions, and where risk reduction is desirable.

3 The Many Faces of Contingent Contracts

We illustrate contingent contracts and motivate the need for formally modeling multi-agent contingent contracting by means of a few examples.

Example 3.1 (Text Processing.) Imagine that agent A, as part of some larger task, needs to translate N pages of Japanese text into English. Since agent A lacks translation capabilities, it seeks other agents that specialize in translation. Various translating agents offer contracts to agent A. We assume the following model for each such translating agent. The agent measures time in integral "cycles" (say 10 second intervals), numbered $1, 2, 3, \ldots$, time 1 denoting the cycle just after starting to work on the contract. At time k the agent processes an integer number $A_k \geq 0$ pages of text for agent A. Since the agents may be interleaving their cycles among several contractors, A_k may be 0. Thus the total number of A's pages processed by time k by the translating agent is

$$M_k = \sum_{i=0}^{k} A_i.$$

Let X denote the (random) time the translating agent takes to translate all N pages of A's text:

$$X = \min\{k : M_k = N\}.$$

Processing one page costs c dollars for any agent, so by time k, the cost of processing A's pages is cB_k. We assume symmetrically that c is also the worth of a translated page of text to agent A. Assume that whenever a page is translated, it is sent immediately to agent A. Thus if agent A aborts the contract at time k, the payoff to agent A is $G_k = cM_k$, which is also the translator's cost of processing A's pages.

Now suppose translating agents B, C, D, E offer four different kinds of contracts to agent A:

- Contract B: For each k, with probability 0.8, $A_k = 1$, and with probability 0.2, $A_k = 0$; A is not allowed to abort, i.e., this is a *binding* contract.
- Contract C: For each k, $A_k = 1$ or $A_k = 0$ with probability 0.5. However, unlike contract B, this one is a *contingent* contract where agent A can decommit at *any* time.
- Contract D: Same as contract C, except that agent A can abort the job if the throughput M_k/k of the job drops below a certain threshold θ.
- Contract E: Same as contract C, except that agent A will abort the job if it is not finished by a certain deadline d.

Let us assume that the fees for these contracts are C_B, C_C, C_D, C_E respectively. In the absence of a rigorous valuation procedure, it is difficult to compare these contracts. For instance let us compare contracts B and C. On the one hand, contract B has the disadvantage of being binding, but there is a high probability (0.8) that each time cycle is devoted to processing a page of agent A. On the other hand, agent C's contract is more flexible but only works on a page of agent A with probability 0.5 Which contract should agent A choose? Let V_B, V_C, V_D, V_E denote the respective values (from A's perspective) of these contracts. A rational policy for agent A is to pick the contract with the largest value of V_x/C_x where $x \in \{B, C, D, E\}$. How should one determine V_B, V_C, V_D, V_E? Questions such as this are the focus of our proposal. One difficulty in valuing these contracts is that the payoff to the contractor depends on how events play out in the future, e.g., how many time cycles are devoted to agent A, etc. However, the future is *uncertain*, and agent A must make a decision *now* in the absence of information about the future. In Section 4 we show one way to compute the values of these contracts. ∎

Example 3.2 (Buying Chips) Consider the example of an electronic market place. In general, an agent who wants to purchase a certain quantity of an item is faced with many options. The agent would want to choose the best deal. Suppose an agent wants to buy 5000 micro-processor chips in 2 months. Let us consider two choices for the agent:

- *Choice A:* The agent can buy 5000 chips in two months at the prevailing price.
- *Choice B:* A supplier can give the agent a license that will cost the agent 50 dollars for the right to buy the chips in two months for 100 per chip.

The agent is faced with the question of which choice to make. Notice that this is not as trivial as it seems because the prices of micro-processor chips might have random fluctuations.

Example 3.3 (Information Retrieval) Consider two Web sources A and B with different performance characteristics. Suppose an agent sends a request to source A. Suppose x units of time pass and the agent hasn't received an answer. Further suppose that the agent can have only one active request at a time. Should the agent abort the request on source A and restart the request on Source B?

The above examples motivate our view that several features of contracts, and particularly contingent contracts, are analogous to those of financial derivative securities, such as options. There is a well-established and widely-used mathematical theory for pricing such securities (see [15]). We believe that this theory is highly applicable to modeling various aspects of inter-agent contracts. In addition, some types of contracts have no analogues in the financial world. To value such contracts we will develop new models and extend option pricing theory in interesting ways.

4 Overview of Options

We now give a very brief introduction to options and their valuation. For more details the reader is referred to Hull's [15] excellent introductory text. For simplicity we confine ourselves to discrete time, although we intend to consider continuous time in our research. An *option* on a stock is a contract that confers upon the holder the *right* but not the obligation to *exercise* it at certain specified times. Once exercised, the option ceases to exist. When exercised at time k, the option yields a nonegative payoff G_k. The payoff in general may depend on the history of stock prices $S_0, S_1, S_2, \ldots, S_k$ up to time k. The values S_i are random, so the payoff G_k is also random. An option is *bought* from the *seller* of the option for a certain *price*. An *American* option can be exercised at any time k before the option expiration date n. *European* options can only be exercised at the expiration date n. Intuitively, since an American option has more flexibility, it is more valuable than an otherwise equivalent European option. As we show later, contingent contracts where an agent can decommit at any time can be viewed as American options.

The value, or "fair" price, of an option is defined as the price that prevents arbitrage opportunities, which are opportunities for unlimited riskless profit. This is the basis for the Arbitrage Pricing Theory (APT) [10, 15] of option valuation. In this pricing theory, it is assumed that one can lend/borrow money at a risk-less rate of r, so that one dollar at time k is worth $R = 1 + r$ dollars at time $k+1$. Thus one dollar at time k is worth $1/R^k$ dollars at time 0.

Assume that we have a stochastic process $\mathcal{P} = \{P_k\}_{k=0}^{\infty}$. P_k is the random variable which gives the value of the process at time k. Intuitively, a *measure* is the distribution of the process. The reader is referred to [2] for a formal definition. A measure lets us reason about the process. For example, using the measure we can compute the probability that the stochastic process will ever go above a certain pre-specified bound k. A filtration $\{\mathcal{F}_k\}_{k=0}^{\infty}$ lets us talk about information known at a certain time k. Informally, $\mathcal{F}_k$ is the information known upto time k. Formally, $\mathcal{F}_k$ is a σ-algebra (see [2]). A process $\mathcal{P}$ is a *martingale* if and only if the value of the process at time $k+1$ conditioned back to time k is the value of the process at time k. Formally,

$$E(P_{k+1} \mid \mathcal{F}_k) \quad = \quad P_k$$

Next, we give an example of a martingale. Consider an infinite sequence of independent random variables $X_1, X_2, \cdots$. Assume that $X_i = 1$ with probability p_i and $X_i = -1$ with probability $1 - p_i$. Also, let $Y_i = X_i - (2p_i - 1)$. Let $P_k = \sum_{i=1}^{k} Y_i$. Let $\mathcal{F}_k$ be the information known by watching the variables $X_1, \cdots, X_k$. In this case $\{P_k\}_{k=1}^{\infty}$ is a martingale. This is shown in the equations given below:

$$\begin{aligned} E(P_{k+1} \mid X_1, \cdots, X_k) &= E(\sum_{i=1}^{k+1} Y_i \mid X_1, \cdots, X_k) \\ &= E(Y_{k+1}) + \sum_{i=1}^{k} Y_i \\ &= P_k \end{aligned}$$

Notice that if $p_i \neq p_j$ (for $i \neq j$), then $\{P_k\}_{k=1}^{\infty}$ is not a stationary process. Hence, a martingale need not be a stationary process. Intuitively, a martingale represents a process with no drift.

Under APT, the value V_0 of a European option at time 0 is given by

$$V_0 = \mathbf{E}(G_n)/R^n, \tag{1}$$

where the expectation **E** is taken with respect to the the martingale measure **P** for the stock price process S_k. In words, the value of a European option at time 0 is the expectation of the present value of the terminal payoff G_n.

The valuation of an American option is complicated by the fact that it can be exercised at any time before expiration. The holder of the option may follow an arbitrary exercise strategy specified by a *stopping time* τ. A stopping time can be thought of as a non-clairvoyant decision rule that says when to exercise; for any evolution of the world ω, (i.e. sequence of stock prices from S_0 to S_k) $\tau(\omega)$ is the exercise time for that evolution. For a random process $X_0, X_1, \ldots, X_n$, the symbol X_τ denotes a new random variable whose value on a path ω equals $X_{\tau(\omega)}(\omega)$. Under APT, value V_0 of an American option is the *maximum,* over all possible exercise strategies τ, of the expectation of the discounted payoff G_τ/R^τ:

$$V_0 = \max_{0 \leq \tau \leq n} \mathbf{E}(G_\tau/R^\tau). \tag{2}$$

More generally, the value of the option at time k is

$$V_k = R^k \max_{k \leq \tau \leq n} \mathbf{E}(G_\tau/R^\tau).$$

Notice that since $\tau = k$ is a valid stopping time, $V_k \geq G_k$ holds for all k. However, since we have the option of *waiting* to exercise our option, V_k may in fact be greater than the immediate payoff G_k. It can also be shown [10] that a stopping time τ^* that achieves the $\max$ in the expression (2) for V_0 is:

$$\tau^* = \min\{k \geq 0;\ G_k = V_k\}. \tag{3}$$

Thus a holder of an American option wishing to maximize the expectation (with respect to the martingale measure) of the time-discounted payoff from his option would use the following rule to exercise his option: exercise the option as soon as the value V_k equals the immediate payoff G_k.

4.1 Contracts as options

Now let us consider Contract B, the binding contract in Example 3.1. What is the value V_B of this contract, from A's point of view? If the time taken by agent B to process all N pages of A's text is X, the payoff to agent A is cN. Let's denote by $\mathbf{E}(X)$ the expectation of the completion time X. Thus in analogy with European options, the value of the contract (at time 0) can be defined as

$$V_0 = \mathbf{E}(cN/R^X) = cN\mathbf{E}(1/R^X),$$

where R is an appropriate discounting factor.

The similarities between options and contracts are particularly evident in the case of contingent contracts. Returning to Example 3.1, consider the contingent contract offered by agent C, where agent A has the right to abort at any time of its choosing. In other words, agent A has the *option* to abort the contract, and it may *exercise* this option at any time. Thus contract C is analogous to an American option and its value can be computed by equation 2.

The decommitment flexibility of A is modeled by a stopping time τ. For example, for contract E (the one with the deadline d), if X denotes the random time agent E will take to finish the job, the stopping time τ is $\min(d, X)$. If agent A aborts the contract at time k, the payoff is $G_k = cM_k$. Thus in analogy with the pricing of an American option, the value of contract E can be defined as:

$$V_0^\tau = \mathbf{E}(G_\tau/R^\tau).$$

We now show how to compute the value of contract E in a very simple scenario.

Example 4.1 Consider the text processing Example 3.1, and suppose agent A wants $N = 2$ pages of text translated. Consider contract E in that example, which has the following features: (a) with probability 0.5, $A_k = 1$ (i.e. one page of A's text will be translated in cycle k), and with probability 0.5, $A_k = 0$, (b) agent A can abort the contract if the job is not finished by 3 cycles, (c) processing each page of text costs agent E $c = 1$ dollar, and so $G_k = M_k$, where it will be recalled that M_k is the number of pages of A processed by time k.

How much is this contract worth to agent A? The possible evolution of events is depicted in figure 1. The right/left branches represent times when a cycle is/isn't alloted to A. The square nodes represent completion of the job. Note that we did not expand the tree beyond three levels because A aborts the job after 3 cycles. The value of this contract is computed using backward recursion on the tree. The value at the leaves is shown in the figure. The values at nodes (3,1) and (3,2) are 0 and 1 respectively. Assume that the interest rate is $r = \frac{1}{4}$. Now we compute the value at node (2,1). This value is the expected value of the value at (3,1) and (3,2) (the two children of (2,1)) discounted by $R = 1 + r$. Recall that T dollars at time 3 is worth $\frac{T}{1+r}$ dollars at time 2. This gives the value of $\frac{2}{5}$ at node (2,1).

Continuing this way, we get the value $\frac{96}{125}$ at the root node (0,1). This is the fair price of this contract. ■

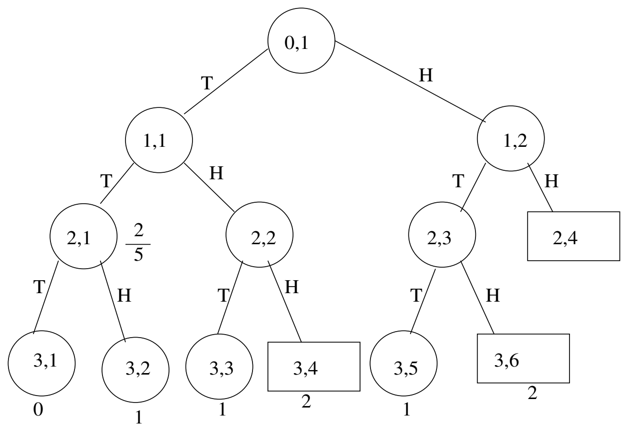

Figure 1: Value Tree

We now show how an agent can view the situations of buying chips in the future (example 3.2) in an options theoretic framework.

Let P represent the random price of chips in 2 months.

Case A: In this case the agent pays $5000 \times P$.

Case B: In this case the agent pays 5000×100.

Hence the savings from using Case B is $5000 \times (100 - P)$. Now consider an option which pays $5000 \times (100 - P)$ in two months. If the value of this option is more that 50, the agent should pay the license fee of 50 dollars and choose Case B. In general, the situation is more complicated. The agent could buy x chips in 2 months at the current price and $5000 - x$ chips at 100. In this case the payoff of the option is $5000 \times 100 - x(100 + P)$. Hence the agent should buy the license if there exists a $x < 5000$, such that the value of the option paying off $5000 \times 100 - x(100 + P)$ in two months is ≤ 50.

Now consider the example with two WWW sources A and B. Suppose the agent issues a query to WWW source A. Lets say t units of time has elapsed. Suppose at this time the agent aborts the query on A and restarts it on B. The gain from aborting the query is $G_t - L_t$.

- The agent loses some resources because it has to restart the query on source B. The agent also loses (opportunity cost) because in certain situations once the query is restarted on source B the answer from source A could have arrived before the answer from source B. These two factors contribute to the loss factor L_t.

- By restarting the query on source B there is some gain because the answer might arrive earlier than it would have if the agent had waited for source A to answer. This contributes to the gain G_t.

Now the question is what is the optimal time for the agent to abort the query on the WWW source A. Consider an option which pays off $G_t - L_t$ at time t. The optimal exercise rule for this option exactly corresponds to the optimal aborting rule for the agent (equation 3).

5 Quality in a Contingent Contract

As mentioned in Section 3 there are some aspects of contingent contracts that have no parallels in financial options. The quality-guarantee is one of them. In this section we offer some preliminary ideas for valuing such contracts.

Suppose an agent A wants to give a contingent contract to B to finish a job J with the following conditions:

- Agent A will abort the job if it is not finished by a deadline d, i.e., $\tau = \min\{X, d\}$,
- Agent B provides the quality guarantee that the probability that the random completion time X exceeds d is at most ϵ, i.e., $\mathbf{P}(X > d) \leq \epsilon$.

In order to satisfy the quality guarantee, B will need to devote a sufficient fraction of its time cycles to agent A's job. This corresponds to choosing a measure **P** that satisfies $\mathbf{P}(X > d) \leq \epsilon$, and there might be many such measures. Let the set of such measures be $\mathcal{P}$. In order to offer A the most attractive price, B should compute the contract value using the tightest schedule, i.e., B does just enough to fulfil the quality aspect of the contract but no more. This corresponds to valuing the contract using the measure from $\mathcal{P}$ that gives the smallest value. In other words, the value of the contract is

$$V_0^\tau = \min_{\mathbf{P} \in \mathcal{P}} \mathbf{E}_\mathbf{P} \left[\frac{G_\tau}{R^\tau} \right].$$

where $\mathbf{E}_\mathbf{P}$ denotes the expected value under the measure **P**.

6 The Case of Sequential Sub-contracting

We now consider another type of contract that has no obvious parallels in the options world: sequential sub-contracting. We illustrate one approach to valuing such contracts by means of the text processing scenario of Example 3.1.

Suppose agent A gives an N-page contract to agent B, characterized by measure **P** and cost-per-page c, M_k (the number of pages of A processed by time k). Agent A wants to abort the job if it is unfinished by time d, so that its stopping time is $\tau = \min\{X, d\}$. At each time k, B considers whether or not to subcontract out all the remaining pages to agent C, characterized by measure **P′**, cost-per-page c', and process M'_k. Agent B decides to subcontract all remaining pages to C if

- C's expected page output is at least as much as that of B, i.e.,

$$\mathbf{E}'(M'_\tau / R^\tau | \mathcal{F}_k) \geq \mathbf{E}(M_\tau / R^\tau | \mathcal{F}_k), \qquad (4)$$

- and C's expected cost of processing the remaining pages is smaller than that of B:

$$\mathbf{E}'\left[c'(M'_\tau - M_k)/R^\tau|\mathcal{F}_k\right] < \mathbf{E}\left[c(M_\tau - M_k)/R^\tau|\mathcal{F}_k\right]. \tag{5}$$

Thus, B subcontracts out to agent C only if it gets at least as much processing rate at a smaller cost. The effect of subcontracting at time k is thus to change the probability measure on the random process A_k (the number of pages processed in time k), and also to change the cost (or worth) per page processed. Let us define a new stopping time τ' which we call the "switching time": $\tau' = \min\{k : (4), (5) \text{ hold}\}$. We then define the new measure $\widetilde{\mathbf{P}}$ as follows. For any integer $i \geq 0$,$\widetilde{\mathbf{P}}(A_k = i)$ equals $\mathbf{P}(A_k = i)$ if $k \leq \tau'$, and equals $P'(A_k = i)$ otherwise. The payoff function is $G_k = c\sum_{i=1}^{\tau'} A_i + c'\sum_{i=\tau'}^{k} A_i$. Then the value of the contract can be defined as

$$V_0^\tau = \widetilde{\mathbf{E}}(G_\tau/R^\tau). \tag{6}$$

In general, one has a directed acyclic graph called the *sub-contracting graph* whose nodes represents agents, and a directed arc from node u to node v is drawn when agent u can subcontract a job to agent v. There is a designated root node representing the first contractor agent. The value of the contract at each node of the sub-contracting graph can be expressed as in (6), and this allows the value of the contract at the root node to be computed recursively.

7 Example Application Domains

7.1 Supply Contracting

A rapidly growing application area for agent technologies is supply contracting, which is an emerging area in Operations Management/Management Science[1, 3, 23]. Supply contracting investigates the research question of how manufacturers can make sensible deals with suppliers. Since financial information, commodity prices, inventory information, etc., are easily available on-line and in real-time, agents can be used to monitor this information and respond quickly to changing conditions, such as changes in demand [19]. In addition, agents could be used in negotiating supply contracts. Early research in Management Science and Operations Management has assumed that there is no risk or uncertainty involved in acquiring an adequate amount of raw material and parts supply in a timely manner. This assumption was reasonable since in those times the dominant philosophy and practice was that in order for a manufacturing company to operate efficiently, it should also produce its own parts. This lead to vertically integrated organizations. Since the mid 80's, however, it has been advocated that companies should outsource the supply of parts for their products to better address global competition issues. Much of the recently growing research in supply contracting is devoted to exploring the significant question of how to negotiate an optimal supply contract considering many OM-related issues such as price, parts delivery time, contract length, long-term commitment, etc. Recently, sophisticated forms of supply contracts, such as ones which involve contingent claims, optional decommit, etc., have gained popularity in industries (e.g., fashion industry). The major technical challenges in these lines of research include: (1) how to evaluate a supply contract given incomplete information, (2) how to come up with compromises which offer reasonable risk-sharing and profit-sharing properties and therefore are potentially acceptable to both parties.

8 Advantages of the option pricing approach

In this section we summarize the advantages of the option pricing approach over traditional techniques employed in DAI.

- The traditional scenario in DAI is *now or never*, or in other words, with few exceptions most work concentrates on binding contracts. In contrast,frameworks based on option pricing allow modeling contingent contracts in a very natural way. In addition, option pricing provides a unified framework for computing values for binding and flexible contracts.

- Option pricing provides a computational framework for optimally calculating decisions, such as, whether it is more advantageous to decommit, when to decommit, what contract to accept out of an offered set of contracts,

- Most techniques based on traditional game theory do not provide tractable computational tools to deal with the *risk* of a decision. These techniques are based on maximizing the expected payoff of a decision. In option pricing, risk can be thought of as variability of a decision. Under certain assumptions, option pricing techniques provide mechanisms for reducing decision risk. This can be useful in searching for optimal decisions with low risk or low variability.

- In traditional decision analysis the uncertainty is modelled as a stationary process, i.e., the variation between time t and $t - j$ only depends on the difference of times (j) and not on t. Notice that in general, this assumption is not true. For example, in our text processing example the random process governing the number of pages an agent processes could depend on the number of pages processed so far by the agent, i.e., the past history. In contrast, the option pricing literature can handle generalized time-dependent stochastic processes.

The options pricing framework provides a more general framework to evaluate and compute optimal decisions in the face of uncertainty. However, applying that framework to DAI problems is not a straight forward task because of assumptions inherent in these techniques.

9 Conclusion

We expect our approach to flexibel multi-agent contracting to have an impact on the state of knowledge in three different areas: (1) *Intelligent agents.* Our approach provides, for the first time, a rigorous way for agents to evaluate their contingent contracts and decide between contracts. Many automated negotiation models and protocols assume that the value of contracts is known. Our valuation methods can be integrated into these models to provide a computational mechanism to enable them to compute contract values. (2) *General contracts.* Our model is applicable to “real-world” contracts, in addition to task allocation type contracts. Moreover, we expect our method to impact areas of practical significance, such as electronic commerce and supply contracting management. (3) *Option pricing theory.* In order to handle contracts that have no analogues in the financial world, we start extending option pricing theory in several interesting ways.

References

[1] R. Anupindi. *Supply Management Under Uncertainty*. PhD thesis, Graduate School of Industrial Administration, Carnegie Mellon University, 1993.

[2] P. Billingsley. *Probability and Measure*. John Wiley & Sons, 3rd edition, 1995.

[3] M. A. Cohen and H. L. Lee. Strategic analysis of integrated production-distribution systems: Models and methods. *Operations Research*, 36(2):216–228, 1988.

[4] S. Cox and A. Hogan. Life insurance risk-based capital: an option pricing approach. *J. actuarial practice*, 3(1):5–23, 1995.

[5] G. Daily. Beyond puts and calls: Option pricing as a powerful tool in decision-making. *The Actuary*, 30(3), 1996.

[6] K. Decker, P. A., K. Sycara, and W. M. Designing behaviors for information agents. In *Proceedings of the First International Conference on Autonomous Agents (Agents-97)*, Los Angeles, February 1997.

[7] K. Decker and V. Lesser. Designing a family of coordination algorithms. In *Proceedings of the First International Conference on Multi-Agent Systems (ICMAS-95)*.

[8] A. Dixit and R. Pindyck. *Investment under uncertainty*. Princeton Univ. Press, 1994.

[9] A. Dixit and R. Pindyck. The options approach to capital invesment. *Harvard Business Review*, pages 105–115, May-June 1995.

[10] D. Duffie. *Security Markets: Stochastic Models*. Academic Press, 1988.

[11] J. Flatto. *The application of real options to the information technology valuation process: A benchmark study*. PhD thesis, Univ. of New Haven, 1996.

[12] D. Fudenberg and J. Tirole. *Game Theory*. The MIT Press, 1991.

[13] M. P. Georgeff and F. F. Ingrand. Decision-making in an embedded reasoning system. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence (IJCAI-89)*, pages 972–978, Detroit, MI, 1989.

[14] B. Grosz and S. Kraus. Collaborative plans for group activities. In *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence (IJCAI-93)*.

[15] J. Hull. *Options, Futures, and Other Derivative Securities*. Prentice Hall, 2nd edition, 1993.

[16] N. R. Jennings, J. M. Corera, and I. Laresgoiti. Developing industrial multi-agent systems. In *First International Conference on Multi-Agent Systems (ICMAS'95)*, pages 423–430, San Francisco, CA., June 12-14 1995.

[17] D. Jones and D. Navin-Chandra. Industrynet: A model of commerce on the World Wide Web. *IEEE Expert*, pages 54–59, Oct. 1995.

[18] J. S. Jordan. The exponential covergence of bayesian learning in normal form games. *Games and Economic Behavior*, 4:202–217, 1992.

[19] R. Kalakota, J. Stallaert, and A. Whinston. World wide real time decision support systems for electronic commerce applications. *J. organizational computing*, 1996.

[20] S. Kraus. Agents contracting tasks in non-collaborative environments. In *Proc. 11th Nat. Conf. on Artificial Intelligence (AAAI)*, 1993.

[21] S. Kraus. Contracting tasks in multi-agent environments. In *Technical Report CS-TR 3254 UMIACS-TR-94-44*, U of Maryland, 1994.

[22] K. Lang. Newsweeder: Learning to filter netnews. In *Proc. Machine Learning Conf.*, 1995.

[23] H. L. Lee and C. Billington. 1992. *Sloan Management Review*, pages 65–73, Spring 1992.

[24] J. MacKie-Mason and H. Varian. Pricing the internet. In B. Kahin and J. Keller, editors, *Public access to the internet*. MIT Press, 1995.

[25] P. Maes. Agents that reduce work and information overload. *Comm. of the ACM*, 37(7), July 1994.

[26] R. McDonald and D. Siegel. The value of waiting to invest. *Quarterly Journal of Economics*, 101:707–727, 1986.

[27] T. Mitchell, R. Caruana, D. Freitag, and J. McDermott. Experience with a learning personal assistant. *Comm. of the ACM*, 37(7), July 1994.

[28] W. J. Mitchell. *Space, Place and the Infobahn: City of Bits*. MIT Press, 1996.

[29] T. Oates, M. Prasad, and V. Lesser. Cooperative information gathering: a distributed problem solving approach. Technical Report 94-66, Univ. of Mass., Sep 1994.

[30] J. Oliver. *On automation negotiation and electronic commerce*. PhD thesis, The Wharton School, Univ. of Pennsylvania, 1996.

[31] S. Park, E. H. Durfee, and W. P. Birmingham. Advantages of strategic thinking in multiagent contracts. In *Proceedings of the Second International Conference on Multiagent Systems (ICMAS-96)*, Kyoto, Japan., 1996.

[32] P.R.Cohen and H.J.Levesque. Intention = choice + commitment. In *Proc. AAAI*, pages 410–415, 1987.

[33] H. Raiffa. *The art and science of negotiation*. Harvard University Press, 1982.

[34] A. S. Rao and M. P. Georgeff. A model-theoretic approach to the verification of situated reasoning systems. In *Proceedings of IJCAI-93*, pages 318–324, Chambery, France, 28 Aug. - 3 Sept. 1993. IJCAI.

[35] J. Rosenschein and G. Zlotkin. *Rules of encounter: Designing conventions for automated negotiation systems*. MIT Press, 1994.

[36] T. Sandholm. An implementation of the contract net protocol based on marginal costs. In *AAAI*, pages 252–262, Washington, D.C., 1993.

[37] T. Sandholm and V. Lesser. Issues in automated negotation and electronic commerce: extending the contract net protocol. In *Proc. First Int. Conf. on Multiagent Systems (ICMAS-95)*, June 1995.

[38] T. Sandholm and V. Lesser. Advantages of a leveled commitment contracting portocol. In *AAAI*, pages 126–133, Portland, Oregon, 1996.

[39] T. W. Sandholm and V. R. Lesser. Coalition fromation among bounded rational agents. In *Proceedings of 14th International Joint Conference on Artificial Intelligence*, 1995.

[40] S. Sen and E. H. Durfee. The effects of search bias on flexibility in distributed scheduling. In *Proceedings of the Twelfth International Workshop on DAI*, 1993.

[41] Y. Shoham. Agent-oriented programming. *Artificial Intelligence*, 60(1):51–92, 1993.

[42] S. F. Smith. Exploiting temporal knowledge to organize constraints. Technical Report CMU-RI-TR-83-12, CMU Robotics Institute, July 1983.

[43] K. Sycara, K. Decker, A. Pannu, M. Williamson, and D. Zeng. Distributed intelligent agents. *IEEE Expert*, 11(6), December 1996.

[44] K. Sycara, K. Decker, and D. Zeng. Intelligent agents in portfolio management. In M. Wooldridge and N. R. Jennings, editors, *Applications of Intelligent Agents*, 1997.

[45] K. Sycara and D. Zeng. Visitor-hoster: Towards an intelligent electronic secretary. In *Proceedings of the CIKM-94 (International Conference on Information and Knowledge Management) Workshop on Intelligent Information Agents*, National Institute of Standards and Technology, Gaithersburg, Maryland, Dec. 1994.

[46] K. Sycara and D. Zeng. Coordination of multiple intelligent software agents. *International Journal of Cooperative Information Systems*, 5(2 & 3):181–211, 1996.

[47] H. Varian. Economic mechanism design for computerized agents. In *USENIX Workshop on Electronic Conference*, New York, July 1995.

[48] M. Wellman. A market-oriented programming environment and its application to distributed multicommodity flow problems. *J. Artificial Intelligence*, 1:1–23, 1993.

[49] M. Wellman. The economic approach to artificial intelligence. *ACM Computing Surveys Symp. on Artif. Intell.*, 27(3), 1995.