Advantages of a Leveled Commitment Contracting Protocol

Tuomas W. Sandholm and Victor R. Lesser *
{sandholm, lesser}@cs.umass.edu
University of Massachusetts at Amherst
Department of Computer Science
Amherst, MA 01003

Abstract

In automated negotiation systems consisting of self-interested agents, contracts have traditionally been binding. Such contracts do not allow agents to efficiently accommodate future events. Game theory has proposed contingency contracts to solve this problem. Among computational agents, contingency contracts are often impractical due to large numbers of interdependent and unanticipated future events to be conditioned on, and because some events are not mutually observable. This paper proposes a leveled commitment contracting protocol that allows self-interested agents to efficiently accommodate future events by having the possibility of unilaterally decommitting from a contract based on local reasoning. A decommitment penalty is assigned to both agents in a contract: to be freed from the contract, an agent only pays this penalty to the other party. It is shown through formal analysis of several contracting settings that this leveled commitment feature in a contracting protocol increases Pareto efficiency of deals and can make contracts individually rational when no full commitment contract can. This advantage holds even if the agents decommit manipulatively.

1 Introduction

The importance of automated negotiation systems is likely to increase as a result of three developments. One is the growth of a standardized communication infrastructure—EDI, NII, KQML, Telescript etc—over which separately designed agents belonging to different organizations can interact in an open environment and safely carry out transactions [8; 16]. The second is the advent of small transaction commerce on the Internet for purchasing goods, information, and communication bandwidth. The third is an industrial trend toward virtual enterprises: dynamic alliances of small enterprises which together can take advantage of economies of scale.

In such multiagent systems consisting of self-interested agents, contracts have traditionally been binding [12; 13; 4; 6]. Once an agent agrees to a contract, it has to follow through with it no matter how future events unravel. Although a contract may be profitable to an agent when viewed *ex ante*, it need not be profitable when viewed after some future events have occurred, i.e. *ex post*. Similarly, a contract may have too low expected payoff *ex ante*, but in some realizations of the future events, the same contract may be desirable when viewed *ex post*. Normal full commitment contracts are unable to efficiently take advantage of the possibilities that such—probabilistically known—future events provide.

On the other hand, many multiagent systems consisting of cooperative agents incorporate some form of decommitment possibility in order to allow the agents to accommodate new events. For example, in the original Contract Net Protocol [19], the agent that had contracted out a task could send a termination message to cancel the contract even when the contractee had already partially fulfilled the contract. This was possible because the agents were not self-interested: the contractee did not mind losing part of its effort without a monetary compensation. Similarly, the role of decommitment among cooperative agents has been studied in meeting scheduling [18] and in cooperative coordination [2].

Game theory has suggested utilizing the potential provided by probabilistically known future events via *contingency contracts* among self-interested agents [11]. The contract obligations are made contingent on future events. There are games in which this method increases the expected payoff to both parties of the contract compared to any full commitment contract. Also, some deals are enabled by contingency contracts in the sense that there is no full commitment contract that both agents prefer over their fall-back positions, but there is a contingency contract that each agent prefers over its fall-back.

There are at least three problems regarding the use of contingency contracts in automated negotiation among self-interested agents. First, it is often impossible to enumerate all possible relevant future events in advance. Second, contingency contracts get cumbersome as the number of relevant events to monitor increases. In the limit, all domain events (e.g. new tasks arriving or resources breaking down) and all negotiation events (other contracts) can affect the value of the obligations of the original contract, and should therefore be conditioned on. Furthermore, these future events may not only affect the value of the original contract independently: the value may depend on *combinations* of the future events [17; 13; 12]. The third problem is that of verifying the unraveling of the events. Sometimes an event is only observable by one of the agents. This agent may have an incentive to lie to the other party of the contract about the event in case the event is associated with an unadvantageous contingency to the directly observing agent. Thus, to be viable, contingency contracts would require an event verification mechanism that is not manipulable and not prohibitively complicated.

*Supported by NSF under Grant No. IRI-9523419.

We propose another method for taking advantage of the possibilities provided by probabilistically known future events. Instead of conditioning the contract on future events, a mechanism is built into the contract that allows unilateral decommitting at any point in time. This is achieved by specifying in the contract decommitment penalties, one for each agent. If an agent wants to decommit—i.e. to be freed from the obligations of the contract—it can do so simply by paying the decommitment penalty to the other party. We will call such contracts *leveled commitment contracts* because the decommitment penalties can be used to choose a level of commitment. The method requires no explicit conditioning on future events: each agent can do its own conditioning dynamically. Therefore no event verification mechanism is required either. This paper presents formal justifications for adding this decommitment feature into a contracting protocol.

Principles for assessing decommitment penalties have been studied in law [1; 10], but the purpose has been to assess a penalty on the agent that has breached the contract *after the breach has occurred.* Similarly, penalty clauses for partial failure—such as not meeting a deadline—are commonly used in contracts, but the purpose is usually to motivate the agents to follow the contract. To our knowledge, the possibility of explicitly allowing decommitment from the whole contract for a predetermined price has not been studied as an active method for utilizing the potential provided by an uncertain future.[1] Somewhat unintuitively, it turns out that the mere existence of a decommitment possibility in a contract can increase each agent's expected payoff.

Key microeconomic concepts are now introduced. *Social welfare* is the sum of the payoffs of the agents under consideration [7; 5]. It does not address distribution. *Pareto efficiency* measures both societal good and distribution [7; 5]. A vector of payoffs to the agents Pareto dominates another vector if each agent's payoff in the first vector is no less than in the second, and there exists an agent whose payoff in the first vector is greater than in the second. Social welfare and Pareto efficiency can be measured either *ex ante* as expected values or *ex post* as realizations. Strategies (mappings from observed history of the game to actions) S_a of the contractor and S_b of the contractee are in *Nash equilibrium* if S_a is a best—expected payoff maximizing—response to S_b, and S_b is a best response to S_a [9; 7; 5]. Finally, a strategy is a *dominant strategy* if it is a best response to *any* strategy of the other agent [7].

We analyze contracting situations from the perspective of two agents: the *contractor* who pays to get a task done, and the *contractee* who gets paid for handling the task. Handling a task can mean taking on any types of constraints. The method is not specific to classical task allocation. The contractor tries to minimize the contract price ρ that it has to pay. The contractee tries to maximize the payoff ρ that it receives from the contractor. Outside offers from third parties will be explicitly discussed. The contracting setting consists of two games. First, the *contracting game* involves the agents choosing a contract—or no contract, i.e. the *null deal*—before any future events have unraveled. Secondly, the *decommitting game* involves the agents deciding on whether to decommit or to carry out the obligations of the contract—after the future events have unraveled. The decommitment game is a subgame of the contracting game: the expected outcomes of the decommitting game affect the agents' preferences over contracts in the contracting game. The decommitting game will be analyzed using the Nash equilibrium and the dominant strategy concepts. The contracting game will be analyzed with respect to *individual rationality* (IR): is the contract better for an agent than the null deal?

ρ	Contract price.
$a \geq 0$	Contractor's decommitment penalty.
$b \geq 0$	Contractee's decommitment penalty.
$\breve{a}$	Price of contractor's best (full commitment) outside offer.
$\breve{b}$	Price of contractee's best (full commitment) outside offer.
$f(\breve{a})$	*Ex ante* probability density function of $\breve{a}$.
$g(\breve{b})$	*Ex ante* probability density function of $\breve{b}$.
p_b	Probability that the contractee decommits.

Table 1: *Symbols used in the paper. We restrict our analysis to contracts where $a \geq 0$ and $b \geq 0$, i.e. we rule out contracts that specify that the decommitting agent receives a payment from the victim of the decommitment.*

In our contracting settings, the future of both agents involves uncertainty. Specifically, the agents might receive outside offers. The contractor's best outside offer $\breve{a}$ is only probabilistically known *ex ante* by both agents, and is characterized by a probability density function $f(\breve{a})$. If the contractor does not receive an outside offer, $\breve{a}$ corresponds to its best outstanding outside offer or its fall-back payoff, i.e. payoff that it receives if no contract is made. The contractee's best outside offer $\breve{b}$ is also only probabilistically known *ex ante*, and is characterized by a probability density function $g(\breve{b})$.[2] If the contractee does not receive an outside offer, $\breve{b}$ corresponds to its best outstanding outside offer or its fall-back payoff. The variables $\breve{a}$ and $\breve{b}$ are assumed statistically independent. The contractor's options are either to make a contract with the contractee or to wait for $\breve{a}$. Similarly, the contractee's options are either to make a contract with the contractor or to wait for $\breve{b}$. The two agents have many mutual contracts to choose from. A leveled commitment contract is specified by the contract price ρ, the contractor's decommitment penalty a, and the contractee's decommitment penalty b. The agents also have the possibility to make a full commitment contract. The contractor has to decide on decommitting when it knows its outside offer $\breve{a}$ but does not know the contractee's outside offer $\breve{b}$. Similarly, the contractee has to decide on decommitting when it knows its outside offer $\breve{b}$ but does not know the contractor's. This seems realistic from a

[1] Decommitting has been studied in other settings, e.g. where there is a constant inflow of agents, and they have a time cost for searching partners of two types: good or bad [3].

[2] Games where at least one agent's future is certain, are a subset of these games. In such games all of the probability mass of $f(\breve{a})$ and/or $g(\breve{b})$ is on one point.

practical automated contracting perspective.

We do not assume that the agents decommit truthfully. An agent may not decommit although its outside offer accompanied by a penalty is better than the contract because the agent believes that there is a high probability that the opponent will decommit. This would save the agent its decommitment penalty and give the agent a decommitment penalty from the opponent. Games of this type differ significantly based on whether the agents have to decommit sequentially or simultaneously. The next two sections analyze these cases. Finally, Section 4 gives some practical prescriptions for building automated negotiation systems, and Section 5 concludes.

2 Sequential decommitting (SEQD)

In our sequential decommitting (SEQD) game, one agent has to declare decommitment before the other. We study the case where the contractee has to decommit first. The case where the contractor has to go first is analogous. Figure 1 presents the game tree. There are two alternative types of leveled commitment contracts that differ on what happens if both agents decommit. In the first, both agents have to pay the decommitment penalties to each other. In the second, neither agent has to pay.

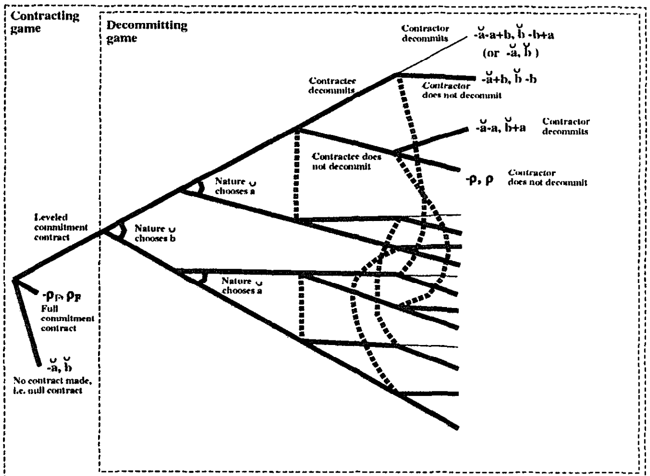

Figure 1: *The "SEQuential Decommitting" (SEQD) game. The game tree of the figure represents two alternative protocols, i.e. two different games. In the first, both agents have to pay the decommitment penalties to each other if both decommit. In the second, neither agent has to pay if both decommit. The payoffs of the latter protocol are in parentheses when they differ from the former. The dotted lines represent information sets: the contractor does not know the contractee's outside offer and vice versa. The contractor's payoffs are usually negative because it has to pay for having the task handled.*

We now analyze the decommitting game using dominance in subgames as the solution concept. Reasoning about the agents' actions starts at the leaves of the tree and proceeds backwards to the beginning of the game. In the subgame where the contractee has decommitted, the contractor's best move is not to decommit because $-\breve{a}-a+b \leq -\breve{a}+b$ (because $a \geq 0$). This also holds for a contract where neither agent has to pay a decommitment penalty if both decommit—because $-\breve{a} \leq -\breve{a}+b$. In the subgame where the contractee has not decommitted, the contractor's best move is to decommit if $-\breve{a} - a > -\rho$. This happens with probability $\int_{-\infty}^{\rho-a} f(\breve{a})d\breve{a}$. Put together, the contractee gets $\breve{b} - b$ if it decommits, $\breve{b} + a$ if it does not but the contractor does, and ρ if neither decommits. Thus the contractee decommits if

$$\breve{b} - b > \int_{-\infty}^{\rho-a} f(\breve{a})d\breve{a}[\breve{b} + a] + \int_{\rho-a}^{\infty} f(\breve{a})d\breve{a}[\rho]$$

If $\int_{\rho-a}^{\infty} f(\breve{a})d\breve{a} = 0$, this is equivalent to $-b > a$ which is false because a and b are nonnegative. In other words, if the contractee surely decommits, the contractor does not. On the other hand, the above is equivalent to

$$\breve{b} > \rho + \frac{b + \int_{-\infty}^{\rho-a} f(\breve{a})d\breve{a}[a]}{\int_{\rho-a}^{\infty} f(\breve{a})d\breve{a}} \stackrel{def}{=} \breve{b}^*(\rho, a, b) \quad (1)$$

when $\int_{\rho-a}^{\infty} f(\breve{a})d\breve{a} > 0$. Now the contractee's IR constraint states that the expected payoff from the contract is no less than the expected payoff from the outside offer:

$$\int_{\breve{b}^*(\rho,a,b)}^{\infty} g(\breve{b})[\breve{b} - b]d\breve{b} + \int_{-\infty}^{\breve{b}^*(\rho,a,b)} g(\breve{b})[\int_{-\infty}^{\rho-a} f(\breve{a})[\breve{b} + a]d\breve{a} + \int_{\rho-a}^{\infty} f(\breve{a})\rho d\breve{a}]d\breve{b}$$
$$\geq E[\breve{b}] = \int_{-\infty}^{\infty} g(\breve{b})\breve{b}d\breve{b} \quad (2)$$

Similarly, the contractor's IR constraint states that the expected payoff from the contract is no less than that from the outside offer:

$$\int_{\breve{b}^*(\rho,a,b)}^{\infty} g(\breve{b}) \int_{-\infty}^{\infty} f(\breve{a})[-\breve{a} + b]d\breve{a}d\breve{b} +$$
$$\int_{-\infty}^{\breve{b}^*(\rho,a,b)} g(\breve{b})[\int_{-\infty}^{\rho-a} f(\breve{a})[-\breve{a} - a]d\breve{a} + \int_{\rho-a}^{\infty} f(\breve{a})[-\rho]d\breve{a}]d\breve{b}$$
$$\geq E[-\breve{a}] = \int_{-\infty}^{\infty} f(\breve{a})[-\breve{a}]d\breve{a} \quad (3)$$

Because the contractor can want to decommit only if $-\breve{a} - a > -\rho$, its decommitment penalty can be chosen so high that it will surely not decommit (assuming that $\breve{a}$ is bounded from below). In this case the contractee will decommit whenever $\rho < \breve{b} - b$. If $\breve{b}$ is bounded from above, the contractee's decommitment penalty can be chosen so high that it will surely not decommit. Thus, full commitment contracts are a subset of leveled commitment ones. This reasoning holds for contracts where both agents have to pay the penalties if both decommit, and for contracts where neither agent has to pay a penalty if both decommit. Because full commitment contracts are a subset of leveled commitment contracts, the former can be no better in the sense of Pareto efficiency or social welfare than the latter. It follows that if there exists an IR full commitment contract, then there also exist IR leveled commitment contracts. However, leveled commitment contracts can enable deals that are impossible via full commitment contracts:

Theorem 2.1 ***Enabling in a SEQD game.*** *There are SEQD games (defined by $f(\breve{a})$ and $g(\breve{b})$) where no full commitment contract satisfies the IR constraints but a leveled commitment contract does.*

Proof. Let $f(\breve{a}) = \begin{cases} \frac{1}{100} & \text{if } 0 \le \breve{a} \le 100 \\ 0 & \text{otherwise} \end{cases}$ and $g(\breve{b}) = \begin{cases} \frac{1}{110} & \text{if } 0 \le \breve{b} \le 110 \\ 0 & \text{otherwise.} \end{cases}$ A full commitment contract F cannot satisfy both IR constraints because that would require $E[\breve{b}] \le \rho_F \le E[\breve{a}]$ which is impossible because $55 = E[\breve{b}] > E[\breve{a}] = 50$. Choose a leveled commitment contract where $\rho = 52.5$, $a = 30$, and $b = 20$. By substituting these in Equations 1, 2, and 3, it turns out that both agents' IR constraints are strictly satisfied. The substitutions are straightforward but tedious [14]. □

In the game of the above proof, both IR constraints are satisfied by a wide range of leveled commitment contracts—and for no full commitment contract. Which leveled commitment contracts defined by ρ, a, and b satisfy the constraints? There are many values of ρ for which some a and b exist such that the constraints are satisfied. We analyze contracts where $\rho = 52.5$ as an example. Now which values of a and b satisfy both IR constraints? There are three qualitatively different cases.

Case 1. Either agent might decommit. In the case where $a < \rho$ there is some chance that the contractor will decommit (it may happen that $-\breve{a} > -\rho + a$). Now $\breve{b}^*(\rho, a, b) = \rho + \frac{b + \int_{-\infty}^{\rho - a} f(\breve{a}) d\breve{a}[a]}{\int_{\rho-a}^{\infty} f(\breve{a}) d\breve{a}} = \rho + \frac{b + \frac{1}{100}[\rho - a]a}{\frac{1}{100}[100 - (\rho - a)]}$. If $\breve{b}^*(\rho, a, b) < 110$ (i.e. less than maximum possible $\breve{b}$), there is some chance that the contractee will decommit. This occurs if $\breve{b} > \rho + b$. We programmed a model of the IR constraints (Equations 3 and 2) for this case. To make the algebra tractable (constant $f(\breve{a})$ and $g(\breve{b})$), versions of these IR constraint equations were used that assumed $0 \le a < \rho$, and $0 < \breve{b}^* < 110$, without loss of generality. The corresponding decommitment penalties a and b that satisfy the IR constraints are plotted in Figure 2 left. Furthermore, the boundaries of the programmed model need to be checked. The boundaries $a = 0$, $a = \rho$, and $\breve{b}^* = 110$ are plotted in Figure 2 left. The constraint $\breve{b}^* > 0$ is always satisfied in this case and is not plotted.

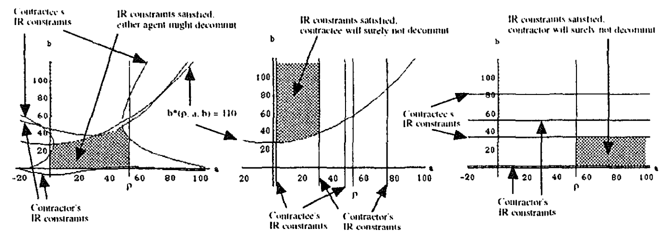

Figure 2: *Decommitment penalties a and b that satisfy both agents' IR constraints in the example SEQD game. Right: either agent might decommit ($a < \rho$, and $\breve{b}^*(\rho, a, b) < 110$). Middle: only contractor might decommit ($a < \rho$, and $\breve{b}^*(\rho, a, b) \ge 110$). Left: $a \ge \rho$, i.e. only the contractee might decommit.*

Case 2, Contractor will surely not decommit. When $a \ge \rho$, the contractor will surely not decommit because its best possible outside offer is $\breve{a} = 0$. Note that a can be arbitrarily high. The corresponding $\breve{b}^*(\rho, a, b) = \rho + \frac{b + \int_{-\infty}^{\rho - a} f(\breve{a}) d\breve{a}[a]}{\int_{\rho-a}^{\infty} f(\breve{a}) d\breve{a}} = \rho + b$, i.e. the contractee decommits truthfully. The contractor's IR constraint (Eq. 3) becomes

$$\int_{\rho+b}^{110} g(\breve{b}) \int_0^{100} f(\breve{a})[-\breve{a} + b] d\breve{a} d\breve{b} + \int_0^{\rho+b} g(\breve{b}) \int_0^{100} f(\breve{a})[-\rho] d\breve{a} d\breve{b} \ge E[-\breve{a}] \tag{4}$$

If $\rho + b \ge 110$, this is equivalent to $-\rho \ge E[-\breve{a}]$ which is false. If $0 < \rho + b < 110$, this is equivalent to

$$\frac{1}{110}\frac{1}{100}[(110 - (\rho + b)) \cdot (\frac{-(100)^2}{2} + 100b) + (\rho + b) \cdot (-100\rho)] \ge E[-\breve{a}]$$

$$\Leftrightarrow \frac{1}{110}\frac{1}{100}[(57.5 - b) \cdot (-5000 + 100b) + (52.5 + b) \cdot (-5250)] \ge -50$$

$$\Leftrightarrow 2.5 \le b \le 52.5$$

by the quadratic equation solution formula.

Similarly the contractee's IR constraint (Eq. 2) becomes

$$\int_{\rho+b}^{110} g(\breve{b}) \int_0^{100} f(\breve{a})[\breve{b} - b] d\breve{a} d\breve{b} + \int_0^{\rho+b} g(\breve{b}) \int_0^{100} f(\breve{a})[\rho] d\breve{a} d\breve{b} \ge E[\breve{b}] \tag{5}$$

If $\rho + b \ge 110$, this is equivalent to $\rho \ge E[\breve{b}]$ which is false. If $0 < \rho + b < 110$, this is equivalent to

$$\int_{\rho+b}^{110} g(\breve{b})[\breve{b} - b] \int_0^{100} f(\breve{a}) d\breve{a} d\breve{b} + \int_0^{\rho+b} g(\breve{b})[\rho] \int_0^{100} f(\breve{a}) d\breve{a} d\breve{b} \ge E[\breve{b}]$$

$$\Leftrightarrow \frac{1}{110}\frac{1}{100}[(\frac{110^2}{2} - 110b - (\frac{(\rho + b)^2}{2} - (\rho + b)b)) \cdot 100 + (\rho + b)\rho \cdot 100] \ge 55$$

$$\Leftrightarrow b \le \text{ approximately } 34.05 \text{ or } b \ge \text{ approximately } 80.95$$

by the quadratic equation solution formula. The latter violates $\rho + b < 110$.

Put together, in the open region $2.5 \le b \le 34.05$, $a \ge \rho$ (Fig. 2 right) this type of contracts are IR for both agents even though the agents decommit insincerely.

Case 3, Contractee will surely not decommit. If b is so high that $\breve{b}^*(\rho, a, b) \ge 110$, the contractee will surely not decommit. The contractor will decommit whenever $-\breve{a} - a > -\rho \Leftrightarrow \breve{a} < \rho - a$, i.e. the decommitting threshold $\breve{a}^* = \rho - a$. The contractor's IR constraint becomes

$$\int_{-\infty}^{\breve{b}^*(\rho,a,b)} g(\breve{b})[\int_{-\infty}^{\rho-a} f(\breve{a})[-\breve{a} - a] d\breve{a} + \int_{\rho-a}^{\infty} f(\breve{a})[-\rho] d\breve{a}] d\breve{b} \ge E[-\breve{a}] \tag{6}$$

$$\Leftrightarrow \int_0^{110} \frac{1}{110}[\int_0^{\rho-a} f(\breve{a})[-\breve{a} - a] d\breve{a} + \int_{\rho-a}^{100} f(\breve{a})[-\rho] d\breve{a}] d\breve{b} \ge -50$$

$$\Leftrightarrow \int_0^{\rho-a} f(\breve{a})[-\breve{a} - a] d\breve{a} + \int_{\rho-a}^{100} f(\breve{a})[-\rho] d\breve{a} \ge -50$$

If $a \ge \rho$, this is equivalent to $-\rho \ge -50$ which is false. If $0 \le a < \rho$, this is equivalent to

$$\frac{1}{100}[\int_0^{\rho-a} [-\breve{a} - a] d\breve{a} + \int_{\rho-a}^{100} [-\rho] d\breve{a}] \ge -50$$

$$\Leftrightarrow \quad \tfrac{1}{100}[(\tfrac{-(\rho-a)^2}{2} + (\rho - a)(-a)) + ((100 - (\rho - a)) \cdot (-\rho)) \ge -50$$

$$\Leftrightarrow \quad a \le \text{ approximately } 30.14 \text{ or } a \ge \text{ approximately } 74.86$$

by the quadratic equation solution formula. The latter violates $a < \rho$.

Similarly, the contractee's IR constraint becomes

$$\int_{-\infty}^{\breve{b}^*(\rho,a,b)} g(\breve{b})[\int_{-\infty}^{\rho-a} f(\breve{a})[\breve{b}+a]d\breve{a} + \int_{\rho-a}^{\infty} f(\breve{a})[\rho]d\breve{a}]d\breve{b} \geq E[\breve{b}] \quad (7)$$

$$\Leftrightarrow \quad \int_{0}^{110} \frac{1}{110}[[\breve{b}+a]\int_{0}^{\rho-a} f(\breve{a})d\breve{a} + [\rho]\int_{\rho-a}^{100} f(\breve{a})d\breve{a}]d\breve{b} \geq 55$$

$$\Leftrightarrow \quad \frac{1}{110}[[\frac{110^2}{2}+110a]\int_{0}^{\rho-a} f(\breve{a})d\breve{a} + 110\rho\int_{\rho-a}^{100} f(\breve{a})d\breve{a}] \geq 55$$

$$\Leftrightarrow \quad [55+a]\int_{0}^{\rho-a} f(\breve{a})d\breve{a} + \rho\int_{\rho-a}^{100} f(\breve{a})d\breve{a} \geq 55$$

If $a \geq \rho$, this is equivalent to $\rho \geq 55$ which is false. If $0 \leq a < \rho$, this is equivalent to

$$[55+a](\rho-a)\frac{1}{100} + \rho[100-(\rho-a)]\frac{1}{100} \geq 55$$

$$\Leftrightarrow \quad 2.5 \leq a \leq 47.5$$

Thus the open region $2.5 \leq a \leq 30.14$, $\breve{b}^* \geq 110$ (Fig. 2 middle) is where this type of contracts are IR for both agents even though the agents decommit insincerely.

In addition to enabling deals that are impossible via full commitment contracts, leveled commitment contracts can increase the efficiency of deals which are possible via full commitment contracts (the reverse cannot occur because the former can emulate the latter) if there is enough *ex ante* variance in the outside offers:

Theorem 2.2 ***Pareto improvement.*** *If a SEQD game has at least one IR full commitment contract F and*

1. *$\breve{b}$ is bounded from above, f is bounded, and $\int_{-\infty}^{E[\breve{b}]} f(\breve{a})d\breve{a} > 0$, or*
2. *$\breve{a}$ is bounded from below, g is bounded, and $\int_{E[\breve{a}]}^{\infty} g(\breve{b})d\breve{b} > 0$,*

then that game has a leveled commitment contract that increases both agents' expected payoffs over any full commitment contract. Therefore, the leveled commitment contract is Pareto superior and IR.

Proof. We prove this under condition 1. The proof under 2 is analogous. With F, the contractor's payoff is $-\rho_F$, and the contractee's ρ_F. We construct a leveled commitment contract where the *contractee will surely not decommit* because its penalty is chosen high and $\breve{b}$ is bounded from above. Choose $\rho = \rho_F$, and $a = \rho_F - E[\breve{b}] + \epsilon$. The contractor decommits if $-\breve{a} - a > -\rho \Leftrightarrow \breve{a} < \rho - a = E[\breve{b}] - \epsilon$. This has nonzero probability because bounded f and $\int_{-\infty}^{E[\breve{b}]} f(\breve{a})d\breve{a} > 0$ imply $\exists\epsilon > 0$ s.t. $\int_{-\infty}^{E[\breve{b}]-\epsilon} f(\breve{a})d\breve{a} > 0$. The contractee's expected payoff increased: it is ρ_F if the contractor does not decommit, and $E[\breve{b}] + a = \rho_F + \epsilon > \rho_F$ if the contractor does. The contractor's expected payoff also increased:

$$-\rho_F < \int_{-\infty}^{\rho-a} f(\breve{a})[-\breve{a}-a]d\breve{a} + \int_{\rho-a}^{\infty} f(\breve{a})[-\rho]d\breve{a}$$

$$\Leftrightarrow 0 < \int_{-\infty}^{\rho-a} f(\breve{a})[-\breve{a}-a+\rho_F]d\breve{a}$$

which is implied by $\int_{-\infty}^{E[\breve{b}]-\epsilon} f(\breve{a})d\breve{a} > 0$ (Because f is bounded, all of this probability mass cannot be on a single point $\breve{a} = \rho - a$ $(= E[\breve{b}] - \epsilon)$). □

3 Simultaneous decommitting

In our simultaneous decommitting games, both agents have to declare decommitment simultaneously. There are two alternative leveled commitment protocols that differ on what happens if both agents decommit, Fig. 3. In the first, both agents have to pay the decommitment penalties to each other. In the second, neither agent has to pay. The next two sections analyze them.

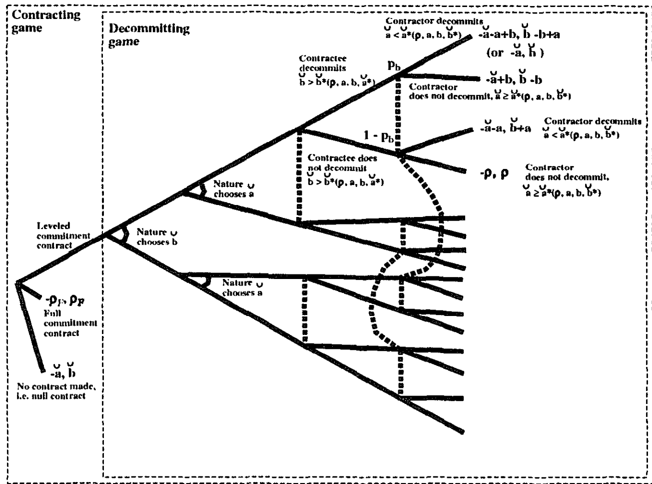

Figure 3: *The "SIMUltaneous Decommit - Both Pay if both decommit" (SIMUDBP) game. The parenthesized payoffs represent the "SIMUltaneous Decommit - Neither Pays if both decommit" (SIMUDNP) game. The dashed lines represent the agents' information sets. When decommitting, the contractor does not know the contractee's outside offer and vice versa. Furthermore, the contractor has to decide on decommitting before it has observed the contractee's decommitting decision, and vice versa.*

3.1 Both pay if both decommit

Simultaneous decommitting games where both agents have to pay the penalties if both decommit will be called SIMUDBP games, Fig. 3. Let p_b be the probability that the contractee decommits, which depends on $f(\breve{a})$, $g(\breve{b})$, ρ, a, and b. The contractor decommits if

$$p_b \cdot (-\breve{a}+b-a) + (1-p_b)(-\breve{a}-a) > p_b \cdot (-\breve{a}+b) + (1-p_b)(-\rho)$$

If $p_b = 1$, this equates to $a < 0$, but we already ruled out contracts where an agent gets paid for decommitting. On the other hand, the above inequality is equivalent to

$$\breve{a} < \rho - \frac{a}{1-p_b} \stackrel{def}{=} \breve{a}^*(\rho,a,b,\breve{b}^*) \text{ when } p_b < 1 \quad (8)$$

If the contractor's outside offer is below the threshold ($\breve{a} < \breve{a}^*$), the contractor is best off by decommitting.

The contractee decommits if

$$\int_{\breve{a}^*(\rho,a,b,\breve{b}^*)}^{\infty} f(\breve{a})d\breve{a}[\breve{b}-b] + \int_{-\infty}^{\breve{a}^*(\rho,a,b,\breve{b}^*)} f(\breve{a})d\breve{a}[\breve{b}-b+a]$$

$$> \int_{\breve{a}^*(\rho,a,b,\breve{b}^*)}^{\infty} f(\breve{a})d\breve{a}[\rho] + \int_{-\infty}^{\breve{a}^*(\rho,a,b,\breve{b}^*)} f(\breve{a})d\breve{a}[\breve{b}+a]$$

If $\int_{\breve{a}^*(\rho,a,b,\breve{b}^*)}^{\infty} f(\breve{a})d\breve{a} = 0$, this equates to $b < 0$, but we ruled out contracts where an agent gets paid for decommitting. However, the above inequality equates to

$$\breve{b} > \rho + \frac{b}{\int_{\breve{a}^*}^{\infty} f(\breve{a})d\breve{a}} \stackrel{def}{=} \breve{b}^*(\rho,a,b,\breve{a}^*) \text{ when } \int_{\breve{a}^*(\rho,a,b,\breve{b}^*)}^{\infty} f(\breve{a})d\breve{a} > 0 \quad (9)$$

If the contractee's outside offer exceeds the threshold ($\breve{b} > \breve{b}^*$), the contractee is best off by decommitting. The probability that the contractee will decommit is

$$p_b = \int_{\breve{b}^*(\rho,a,b,\breve{a}^*)}^{\infty} g(\breve{b})d\breve{b} \tag{10}$$

Condition 8 states the contractor's best response (defined by $\breve{a}^*$) to the contractee's strategy that is defined by $\breve{b}^*$. Condition 9 states the contractee's best response $\breve{b}^*$ to the contractor's strategy that is defined by $\breve{a}^*$. Condition 8 uses the variable p_b which is defined by Equation 10. So together, Equations 8, 9, and 10 define the Nash equilibria of the decommitting game.

Now the contractor's IR constraint becomes

$$\int_{\breve{b}^*(\rho,a,b,\breve{a}^*)}^{\infty} g(\breve{b})[\int_{-\infty}^{\breve{a}^*(\rho,a,b,\breve{b}^*)} f(\breve{a})[-\breve{a}+b-a]d\breve{a} + \int_{\breve{a}^*(\rho,a,b,\breve{b}^*)}^{\infty} f(\breve{a})[-\breve{a}+b]d\breve{a}]d\breve{b}$$
$$+\int_{-\infty}^{\breve{b}^*(\rho,a,b,\breve{a}^*)} g(\breve{b})[\int_{-\infty}^{\breve{a}^*(\rho,a,b,\breve{b}^*)} f(\breve{a})[-\breve{a}-a]d\breve{a} + \int_{\breve{a}^*(\rho,a,b,\breve{b}^*)}^{\infty} f(\breve{a})[-\rho]d\breve{a}]d\breve{b} \geq E[-\breve{a}]$$

The first row corresponds to the contractee decommitting, while the second corresponds to the contractee not decommitting. The second integral in each row corresponds to the contractor decommitting, while the third integral corresponds to the contractor not decommitting. Similarly, the contractee's IR constraint becomes

$$\int_{\breve{b}^*(\rho,a,b,\breve{a}^*)}^{\infty} g(\breve{b})[\int_{-\infty}^{\breve{a}^*(\rho,a,b,\breve{b}^*)} f(\breve{a})[\breve{b}-b+a]d\breve{a} + \int_{\breve{a}^*(\rho,a,b,\breve{b}^*)}^{\infty} f(\breve{a})[\breve{b}-b]d\breve{a}]d\breve{b}$$
$$+\int_{-\infty}^{\breve{b}^*(\rho,a,b,\breve{a}^*)} g(\breve{b})[\int_{-\infty}^{\breve{a}^*(\rho,a,b,\breve{b}^*)} f(\breve{a})[\breve{b}+a]d\breve{a} + \int_{\breve{a}^*(\rho,a,b,\breve{b}^*)}^{\infty} f(\breve{a})[\rho]d\breve{a}]d\breve{b} \geq E[\breve{b}]$$

If $\breve{a}$ is bounded from below, the contractor's decommitment penalty a can be chosen so high that the contractor's decommitment threshold $\breve{a}^*(\rho, a, b, \breve{b}^*)$ becomes lower than any $\breve{a}$. In that case the contractor will surely not decommit. Similarly, if $\breve{b}$ is bounded from above, the contractee's decommitment penalty b can be chosen so high that the contractee's decommitment threshold $\breve{b}^*(\rho, a, b, \breve{a}^*)$ is greater than any $\breve{b}$. In that case the contractee will surely not decommit. Thus, full commitment contracts are a subset of leveled commitment ones. Therefore, the former can be no better in the sense of Pareto efficiency or social welfare than the latter.

In addition to leveled commitment contracts never being worse than full commitment ones, the following theorem shows that in SIMUDBP games they can enable a deal that is impossible via full commitment contracts.

Theorem 3.1 ***Enabling in a SIMUDBP game.*** *There are SIMUDBP games (defined by $f(\breve{a})$ and $g(\breve{b})$) where no full commitment contract satisfies the IR constraints but a leveled commitment contract does.*

Proof. Let $f(\breve{a}) = \begin{cases} \frac{1}{100} & \text{if } 0 \leq \breve{a} \leq 100 \\ 0 & \text{otherwise} \end{cases}$ and $g(\breve{b}) = \begin{cases} \frac{1}{110} & \text{if } 0 \leq \breve{b} \leq 110 \\ 0 & \text{otherwise.} \end{cases}$ No full commitment contract F satisfies both IR constraints because that would require $E[\breve{b}] \leq \rho_F \leq E[\breve{a}]$ which is impossible because $55 = E[\breve{b}] > E[\breve{a}] = 50$. We build a leveled commitment contract with $\rho = 52.5$ as an example. Four cases result:

Case 1. Either agent might decommit. If $0 < \breve{a}^* < 100$, and $0 < \breve{b}^* < 110$, there is a nonzero probability for each agent to decommit. The unique Nash equilibrium is plotted out for different values of a and b in Figure 4. The equilibrium decommitment thresholds $\breve{a}^*$ and $\breve{b}^*$ differ from the truthful ones. Yet there exist equilibria in the proper range of $\breve{a}^*$ and $\breve{b}^*$.

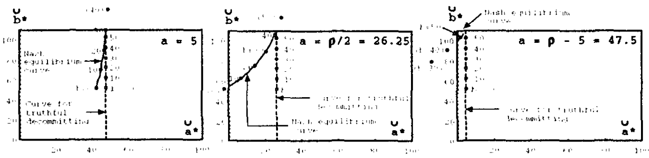

Figure 4: *The Nash equilibrium decommitment thresholds $\breve{a}^*$ and $\breve{b}^*$ of our example SIMUDBP game for different values of the decommitment penalties a and b. The Nash equilibrium deviates from truthful decommitting. If $0 < \breve{a}^* < 100$, and $0 < \breve{b}^* < 110$, there is some chance that either agent will decommit.*

It is not guaranteed that all of these Nash equilibria satisfy the agents' IR constraints however. We programmed a model of Equations 8, 9, and 10 and the IR constraints. To make the algebra tractable (constant $f(\breve{a})$ and $g(\breve{b})$), versions of these equations were used that assumed $0 < \breve{a}^* < 100$, and $0 < \breve{b}^* < 110$, without loss of generality. Therefore the first task was to check the boundaries of the validity of the model. The boundaries $\breve{a}^* = 0$ and $\breve{b}^* = 110$ are plotted in Figure 5. The boundary $\breve{a}^* = 100$ turns out to be the line $b = 0$. There exists no boundary $\breve{b}^* = 0$ because $\breve{b}^*$ was always positive.

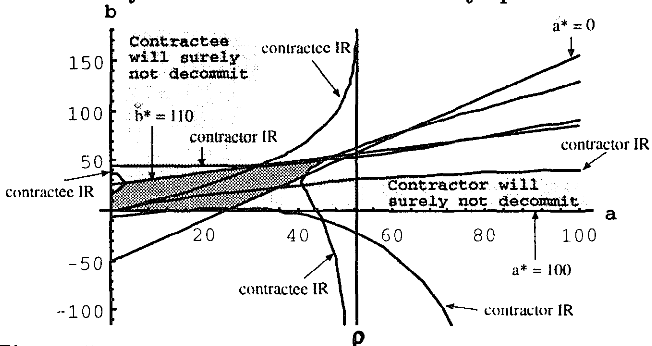

Figure 5: *Three different regions of contracts that are IR for both agents and allow an equilibrium in the SIMUDBP decommitting game. In the dark gray area either agent might decommit but in the light gray areas only one of the agents might. Curves represent the IR constraints and validity constraints of the programmed model that requires $0 < \breve{a}^* < 100$, and $0 < \breve{b}^* < 110$.*

Each agent's IR constraint induces three curves (Fig. 5), two of which actually bound the IR region. The third is also a root, but at both sides of that curve, the IR constraint is satisfied. The dark gray area of Figure 5 represents the values of the decommitment penalties a and b for which the validity constraints of the programmed model and the IR constraints are satisfied. In other words, for any such a and b, there exists decommitment thresholds $\breve{a}^*$ and $\breve{b}^*$ such that these form a Nash equilibrium, and there is a nonzero probability for either agent to decommit, and each agent has higher expected payoff with the contract than without it.

Case 2, Contractor will surely not decommit. If $\breve{a}^* \le 0$, the contractor will surely not decommit. Now $\breve{b}^*(\rho, a, b, \breve{a}^*) = \rho + \frac{b}{\int_{\breve{a}^*}^{\infty} f(\breve{a})d\breve{a}} = \rho + b$, i.e. the contractee decommits truthfully. The contractor's IR constraint becomes the same as in case 2 of the example SEQD game (Eq. 4). This constraint was proven equivalent to $2.5 \le b \le 52.5$. The contractee's IR constraint also equates to that in the SEQD game (Eq. 5). It was proven equivalent to $b \le$ approximately 34.05. Thus these contracts are IR for both agents and in equilibrium in the open region $2.5 \le b \le 34.05$, $\breve{a}^* \le 0$, Fig. 5.

Case 3, Contractee will surely not decommit. If $\breve{b}^* \ge 110$, the contractee will surely not decommit ($p_b = 0$). Now $\breve{a}^*(\rho, a, b, \breve{b}^*) = \rho - \frac{a}{1-p_b} = \rho - a$, i.e. the contractor decommits truthfully. The contractor's IR constraint becomes the same as in case 3 of the example SEQD game (Eq. 6). This constraint was proven equivalent to $a \le$ approximately 30.14. The contractee's IR constraint equates to Eq. 7 of the SEQD game. It was proven equivalent to $2.5 \le a \le 47.5$. Thus the open region $2.5 \le a \le 30.14$, $\breve{b}^* \ge 110$ is where these contracts are IR for both agents, and in equilibrium, Fig. 5.

Case 4, Trivial case. A contract where at least one agent will surely decommit, i.e. $\breve{a}^* \ge 100$ or $\breve{b}^* \le 0$ can be IR—barely because it does not increase either agent's payoff. For it to be IR for the decommitting agent, the decommitment penalty has to be zero: the decommitting agent gets the same payoff as without the contract. Similarly, the other agent gets the same payoff as it would get without the contract. This contract is equivalent to no contract at all: decommitment occurs and no payment is transferred. □

In addition to enabling deals that are impossible using full commitment, leveled commitment contracts can increase the efficiency of a deal even if a full commitment contract were possible (the reverse cannot occur):

Theorem 3.2 ***Pareto efficiency improvement.*** *Theorem 2.2 applies to SIMUDBP games.*

Proof. When one agent is known not to decommit, SIMUDBP games are equivalent to SEQD games. □

3.2 Neither pays if both decommit

Simultaneous decommitting games where a protocol is used where *neither agent has to pay a decommitting penalty if both agents decommit* (SIMUDNP games, Fig. 3) can be analyzed in the same way as SIMUDBP games, but the decommitting thresholds differ [14].

If $\breve{a}$ is bounded from below, and $\breve{b}$ from above, a can be chosen so high that the contractor will surely not decommit, and b so high that the contractee will not. So, full commitment contracts are a subset of leveled commitment ones. Thus the former cannot enable a deal whenever the latter cannot. Also, leveled commitment can enable a deal that is impossible via full commitment:

Theorem 3.3 ***Enabling in a SIMUDNP game.*** *There exist SIMUDNP games (defined by $f(\breve{a})$ and $g(\breve{b})$) where no full commitment contract satisfies the IR constraints but a leveled commitment contract does.*

The proof is like that of Theorem 3.1 except that the formulas for decommitting differ [14]. With the same $f(\breve{a})$, $g(\breve{b})$, and ρ as in the proof of Theorem 3.1, the Nash equilibria of the SIMUDNP game are as shown in Figure 6. The decommitment thresholds $\breve{a}^*$ and $\breve{b}^*$ differ from the truthful ones. They are closer to the truthful ones than what they were with a protocol where both agents pay if both decommit, Figure 4. The shapes of the curves using these two protocols also differ significantly.

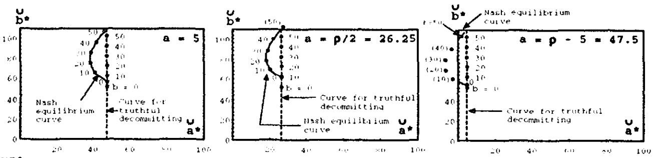

Figure 6: ***The Nash equilibrium decommitment thresholds $\breve{a}^*$ and $\breve{b}^*$ of our example SIMUDNP game for different values of the decommitment penalties a and b. The Nash equilibrium deviates from truthful decommitting. If $0 < \breve{a}^* < 100$, and $0 < \breve{b}^* < 110$, there is some chance that either agent will decommit.***

We programmed a model to check which of the SIMUDNP equilibria allow contracts that are IR for both agents. The results are quantitatively different, but qualitatively same as in SIMUDBP games (Fig. 5) [14].

Leveled commitment contracts can also increase the efficiency of a deal even if a full commitment contract were possible (the reverse cannot occur):

Theorem 3.4 ***Pareto efficiency improvement.*** *Theorem 2.2 applies to SIMUDNP games.*

Proof. When one agent is known not to decommit, SIMUDNP games are equivalent to SEQD games. □

4 Prescriptions for system builders

The results from the above canonical games suggest that it is worthwhile from a contract enabling and a contract Pareto improving perspective to incorporate the decommitment mechanism into automated contracting protocols. The decommitment penalties are best chosen by the agents dynamically at contract time as opposed to statically in the protocol. This allows the tuning of the penalties not only to specific negotiation situations and environmental uncertainties, but also to specific belief structures of the agents. An extended paper analyzes the impact of agents' biased beliefs on the benefits of contracts, and the distribution of these gains [14].

In the presented instance of the simultaneous decommitting game, the Nash equilibrium decommitting strategies were usually closer to truthful ones when a protocol was used where neither pays if both decommit than when a protocol was used where both pay if both decommit. Also, as an agent's opponent's decommitment penalty approaches zero, the agent becomes truthful in the former protocol, but starts to increasingly bias its decommitment decisions in the latter. This suggests using the former protocol in practical systems. It also minimizes the number of payment transfers because it does not require any such transfer if both decommit.

In a web of multiple mutual contracts among several agents, classical full commitment contracts induce one negotiation focus consisting of the obligations of the contracts. Under the protocol proposed in this paper, there are multiple such foci, and any agent involved in a contract can swap from one such focus to another by decommitting from a contract. Such a swap may make it beneficial for another agent to decommit from another contract, and so on. To avoid loops of decommitting and recommitting in practise, recommitting can be disabled. This can be implemented by a protocol that specifies that if a contract offer is accepted and later either agent decommits, the original offer becomes void—as opposed to staying valid according to its original deadline which may not have been reached at decommitment time.

Even if two agents cannot explicitly recommit to a contract, it is hard to specify and monitor in a protocol that they will not make another contract with an identical content. This gives rise to the possibility of the equivalent of useless decommit-recommit loops. Such loops can be avoided by a mechanism where the decommitment penalties increase with real-time or with the number of domain events or negotiation events. This allows a low commitment negotiation focus to be moved in the joint search space while still making the contracts meaningful by some level of commitment. The increasing level of commitment causes the agents to not backtrack deeply in the negotiations, which can also save computation.

The initially low commitment to contracts can also be used as a mechanism to facilitate linking of deals. Often, there is no contract over a single item that is beneficial, but a combination of contracts among two agents would be [13; 17]. Even if explicit clustering of issues into contracts [13; 17] is not used, an agent can agree to an unbeneficial contract in anticipation of synergic future contracts from the other agent that will make the first contract beneficial [17]. If no such contracts appear, the agent can decommit. Similarly, low commitment contracts can be used to facilitate deals among more than two agents. Even without explicit multiagent contract protocols [17], multiagent contracts can be implemented by one agent agreeing to an unbeneficial contract in anticipation of synergic future contracts from third parties that will make the first contract beneficial [17]. If no such contracts appear, the agent can decommit.

In many practical automated contracting settings limited computation resources bound the agents' capability to solve combinatorial problems [17; 15; 13]. The value of a contract may only be probabilistically known to an agent at contract time. The leveled commitment contracting protocol allows the agent to continue deliberation regarding the value of the contract after the contract is made. If the value turns out to be lower than expected, the agent can decommit. However, decommitment penalties which increase quickly in time may be appropriate with computationally limited agents so that the agents do not need to consider the combinatorial number of possible future worlds where alternative combinations of decommitments have occurred [17].

5 Conclusions

A protocol was presented for automated contracting that allows agents to accommodate future events more profitably than traditional full commitment contracts. Each contract specifies a decommitment penalty for both agents involved. To decommit, an agent just pays that penalty to the other agent. This mechanism is better suited for complex computerized contracting settings than contingency contracts. The analysis handled the fact that agents decommit manipulatively. This analysis also serves as a normative tool for agents to decide which contracts to accept based on individual rationality.

Leveled commitment contracts can emulate full commitment ones by setting the decommitting penalties high. Therefore, full commitment contracts cannot be better than leveled commitment ones in the sense of Pareto efficiency or social welfare to the two agents. Neither can they enable a deal that is impossible—based on individual rationality—via a leveled commitment contract. We proved that in these games the new protocol surprisingly enables deals that are impossible via full commitment contracts. It also increases the expected payoff to both agents in settings where a full commitment contract is possible. Obviously one can also construct game instances where the null deal is so profitable to both agents that no contract—even a leveled commitment one—is individually rational to both.

References

[1] J. D. Calamari and J. M. Perillo. The Law of Contracts. West Publishing Co., 2nd edition, 1977.

[2] K. Decker and V. R. Lesser. Designing a family of coordination algorithms. In ICMAS, pages 73–80, San Francisco, CA, June 1995.

[3] P.Diamond and E.Maskin. An equilibrium analysis of search and breach of contract, i: Steady states. Bell J of Economics, 10:282–316, 1979.

[4] E. Ephrati and J. S. Rosenschein. The clarke tax as a consensus mechanism among automated agents. In AAAI, pages 173–178, 1991.

[5] D. Fudenberg and J. Tirole. Game Theory. MIT Press, 1991.

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

[7] D. Kreps. A course in microeconomic theory. Princeton U Press, 1990.

[8] D. M. Kristol, S. H. Low, and N. F. Maxemchuk. Anonymous internet mercantile protocol. 1994. Submitted.

[9] J. Nash. Equilibrium points in n-person games. Proc. of the National Academy of Sciences, 36:48–49, 1950.

[10] R. Posner. Economic Analysis of Law. Little, Brown & Co, 2nd ed, 1977.

[11] H. Raiffa. The Art and Science of Negotiation. Harvard U Press, 1982.

[12] J. S. Rosenschein and G. Zlotkin. Rules of Encounter. MIT Press, 1994.

[13] T. W. Sandholm. An implementation of the contract net protocol based on marginal cost calculations. In AAAI, pages 256–262, July 1993.

[14] T. W. Sandholm and V. R. Lesser. Advantages of a leveled commitment contracting protocol. Umass TR 95-72, 1995.

[15] T. W. Sandholm and V. R. Lesser. Coalition formation among bounded rational agents. In IJCAI, pages 662–669, Montreal, Canada, Aug. 1995. Extended version: Umass TR 95-71.

[16] T. W. Sandholm and V. R. Lesser. Equilibrium analysis of the possibilities of unenforced exchange in multiagent systems. In IJCAI, pages 694–701, Montreal, Canada, Aug. 1995.

[17] T. W. Sandholm and V. R. Lesser. Issues in automated negotiation and electronic commerce: Extending the contract net framework. In ICMAS-95, pages 328–335, San Francisco, CA, June 1995.

[18] S. Sen. Tradeoffs in Contract-Based Distributed Scheduling. PhD thesis, Univ. of Michigan, 1993.

[19] R. G. Smith. The contract net protocol: High-level communication and control in a distributed problem solver. IEEE Transactions on Computers, C-29(12):1104–1113, Dec. 1980.