

Revenue Equivalence of Leveled Commitment Contracts

Tuomas Sandholm* and Yunhong Zhou

{sandholm, yzhou}@cs.wustl.edu

Department of Computer Science

Washington University

St. Louis, MO 63130

Abstract

In automated negotiation systems consisting of self-interested agents, contracts have traditionally been binding. Leveled commitment contracts—i.e. contracts where each party can decommit by paying a predetermined penalty—were recently shown to improve expected social welfare even if agents decommit insincerely in Nash equilibrium. Such contracts differ based on whether agents have to declare their decommitting decisions sequentially or simultaneously, and whether or not agents have to pay the penalties if both decommit. For a given contract, these protocols lead to different decommitting thresholds and probabilities. However, this paper shows that, surprisingly, each protocol leads to the same expected social welfare when the contract price and penalties are optimized for each protocol. Our derivations allow agents to construct optimal leveled commitment contracts. We also show that such integrative bargaining does not hinder distributive bargaining: the excess can be divided arbitrarily (as long as each agent benefits), e.g. equally, without compromising optimality. Revenue equivalence ceases to hold if agents are not risk neutral. A contract optimization service is offered on the web as part of *eMediator*, our next generation electronic commerce server.

1 Introduction

In multiagent systems consisting of self-interested agents, contracts have traditionally been binding (Rosenschein & Zlotkin 1994; Sandholm 1993; Kraus 1993). Once an agent agrees to a contract, she has to follow through 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. Similarly, a contract may have too low expected payoff *ex ante*, but in some realizations of the future events, it may be desirable when viewed *ex post*. Normal full commitment contracts are unable to take advantage of the possibilities that such future events provide.

On the other hand, many multiagent systems consisting of cooperative agents incorporate some form of

decommitment in order to allow agents to accommodate new events. For example, in the original Contract Net Protocol (Smith 1980), the agent that contracts out a task could send a termination message to cancel the contract even when the contractee had partially fulfilled it. 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 (Sen 1993).

Contingency contracts have been suggested for utilizing the potential provided by future events among self-interested agents (Raiffa 1982). The contract obligations are made contingent on future events. In some games this increases the expected payoff to both parties compared to any full commitment contract. However, contingency contracts are often impractical. The space of combinations of future events can be large and it is rare that both agents are cognizant of all possible future worlds. Also, when events are not mutually observable, the observing agent can lie about what transpired.

Leveled commitment contracts are another method for capitalizing on future events (Sandholm & Lesser 1996). Instead of conditioning the contract on future events, a mechanism is built into the contract that allows unilateral decommitting. This is achieved by specifying the level of commitment by 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. The method requires no explicit conditioning on future events: each agent can do her own conditioning dynamically. No event verification mechanism against lying is required either.

Principles for assessing decommitment penalties have been studied in law (Calamari & Perillo 1977; Posner 1977), 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. Instead, in leveled commitment contracts, explicitly allowing decommitting from the contract for a predetermined price

* This material is based upon work supported by the National Science Foundation under CAREER Award IRI-9703122, Grant IRI-9610122, and Grant IIS-9800994.

is used as an active method for utilizing the potential provided by an uncertain future.¹ The decommitment possibility increases each agent's expected payoff under very general assumptions (Sandholm & Lesser 1996). This paper studies the same setting and the same contract types as they did, but derives new results.

We analyze contracting situations from the perspective of two risk neutral agents who attempt to maximize their own expected payoff: 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 he has to pay. The contractee tries to maximize the payoff ρ that she receives. The future of the agents involves uncertainty. Specifically, the agents might receive outside offers.² The contractor's best outside offer \check{a} is only probabilistically known *ex ante* by both agents, and is characterized by a probability density function $f(\check{a})$. If the contractor does not receive an outside offer, \check{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 \check{b} is also only probabilistically known *ex ante*, and is characterized by a probability density function $g(\check{b})$. If the contractee does not receive an outside offer, \check{b} corresponds to its best outstanding outside offer or its fall-back payoff.³ The variables \check{a} and \check{b} are assumed statistically independent.

The contractor's options are either to make a contract with the contractee or to wait for \check{a} . Similarly, the contractee's options are either to make a contract with the contractor or to wait for \check{b} . The two agents could make a full commitment contract at some price. Alternatively, they can make a leveled commitment contract which is specified by the contract price, ρ , the contractor's decommitment penalty, a , and the contractee's decommitment penalty, b . The contractor has to decide on decommitting when he knows his outside offer \check{a} but does not know the contractee's outside offer \check{b} . Similarly, the contractee has to decide on decommitting when she knows her outside offer \check{b} but does not know the contractor's. This seems realistic from a practical automated contracting perspective.

Section 2 reviews the leveled commitment contracting protocols and how rational agents decommit in them. The question arises: which protocol leads to the best re-

sults for the agents? Section 3 shows that, surprisingly, each protocol leads to the same expected social welfare when the contract price and penalties are optimized for each protocol. Section 4 analyzes the interplay between integrative and distributive bargaining in leveled commitment contracting, and shows how to construct a fair optimal contract. Section 5 discusses nonuniqueness. Section 6 shows that revenue equivalence ceases to hold if agents are not risk neutral. Section 7 concludes.

2 Leveled commitment contracts

One concern with leveled commitment contracts is that a rational agent is reluctant to decommit because there is a chance that the other party will decommit, in which case the former agent gets freed from the contract, does not have to pay a penalty, and collects a penalty from the breacher. (Sandholm & Lesser 1996) showed that despite such insincere decommitting the leveled commitment feature increases each contract party's expected payoff, and enables contracts in settings where no full commitment contract is beneficial to all parties. We derive the Nash equilibrium (Nash 1950b) where each agent's decommitting strategy is a best response to the other agent's decommitting strategy. The results of the paper take into account the fact that agents decommit insincerely in this way.

2.1 Sequential decommitting (SEQD)

In a sequential decommitting (SEQD) game, one agent has to declare her decommitting decision before the other. We assume that the contractee has to decommit first. The case where the contractor has to go first is analogous. There are two alternative types of leveled decommitment contracts based on whether or not the agents have to pay the penalties if both decommit.

If the contractee has decommitted, the contractor's best move is not to decommit because $-\check{a} - a + b \leq -\check{a} + b$ (unless $a < 0$, which would mean—absurdly—that the contractor gets paid for decommitting). This also holds for a contract where neither agent has to pay a decommitment penalty if both decommit since $-\check{a} \leq -\check{a} + b$. In the subgame where the contractee has not decommitted, the contractor's best move is to decommit if $-\check{a} - a > -\rho$, i.e. the contractor decommits if his outside offer, \check{a} , is below a threshold $\check{a}^* = \rho - a$. So, the probability that he decommits is $p_a = \int_{-\infty}^{\check{a}^*} f(\check{a}) d\check{a}$.

The contractee gets $\check{b} - b$ if she decommits, $\check{b} + a$ if she does not but the contractor does, and ρ if neither decommits. Thus the contractee decommits if $\check{b} - b > p_a(\check{b} + a) + (1 - p_a)\rho$. A contract where $p_a = 1$ cannot be strictly individually rational to both agents since breach will occur for sure. On the other hand, when $p_a < 1$ the inequality above shows that the contractee decommits if her outside offer exceeds a threshold $\check{b}^* = \rho + \frac{b + ap_a}{1 - p_a}$. So, the probability that she decommits is $p_b = \int_{\check{b}^*}^{\infty} g(\check{b}) d\check{b}$.

¹Decommitting 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 (Diamond & Maskin 1979).

²The framework can also be interpreted to model situations where the agents' cost structures for handling tasks and for getting tasks handled change e.g. due to resources going off-line or becoming back on-line.

³Games where at least one agent's future is certain, are a subset of these games. In such games all of the probability mass of $f(\check{a})$ and/or $g(\check{b})$ is on one point.

The rest of the paper uses the following shorthand:

$$E(\check{a}) \equiv \int_{-\infty}^{\infty} \check{a} f(\check{a}) d\check{a}, \quad E(\check{b}) \equiv \int_{-\infty}^{\infty} \check{b} g(\check{b}) d\check{b}$$

$$E(\check{a}, \check{a}^*) \equiv \int_{-\infty}^{\check{a}^*} \check{a} f(\check{a}) d\check{a}, \quad E(\check{a}^*, \check{a}) \equiv \int_{\check{a}^*}^{\infty} \check{a} f(\check{a}) d\check{a}$$

$$E(\check{b}, \check{b}^*) \equiv \int_{-\infty}^{\check{b}^*} \check{b} g(\check{b}) d\check{b}, \quad E(\check{b}^*, \check{b}) \equiv \int_{\check{b}^*}^{\infty} \check{b} g(\check{b}) d\check{b}.$$

The contractor's expected payoff under the contract is

$$\begin{aligned} \pi_a &= p_b \int_{-\infty}^{\infty} (-\check{a} + b) f(\check{a}) d\check{a} + (1 - p_b) \cdot \\ &\quad \left[\int_{-\infty}^{\check{a}^*} (-\check{a} - a) f(\check{a}) d\check{a} + \int_{\check{a}^*}^{\infty} (-\rho) f(\check{a}) d\check{a} \right] \\ &= p_b [b - E(\check{a})] + (1 - p_b) [-E(\check{a}, \check{a}^*) - a p_a - \rho(1 - p_a)] \\ &= -[p_a(1 - p_b)a - p_b b + (1 - p_a)(1 - p_b)\rho] \\ &\quad - E(\check{a}) + (1 - p_b)E(\check{a}^*, \check{a}) \\ &= -E(\check{a}) - \phi(\rho, a, b) + (1 - p_b)E(\check{a}^*, \check{a}). \end{aligned}$$

where $\phi(\rho, a, b) = p_a(1 - p_b)a - p_b b + (1 - p_a)(1 - p_b)\rho$. The contractee's expected payoff under the contract is

$$\begin{aligned} \pi_b &= \int_{\check{b}^*}^{\infty} g(\check{b})(\check{b} - b) d\check{b} + \int_{-\infty}^{\check{b}^*} g(\check{b})[p_a(\check{b} + a) + (1 - p_a)\rho] d\check{b} \\ &= -p_b b + (1 - p_b)(p_a a + (1 - p_a)\rho) + E(\check{b}^*, \check{b}) + p_a E(\check{b}, \check{b}^*) \\ &= [p_a(1 - p_b)a - p_b b + (1 - p_a)(1 - p_b)\rho] \\ &\quad + E(\check{b}) - (1 - p_a)E(\check{b}, \check{b}^*) \\ &= E(\check{b}) + \phi(\rho, a, b) - (1 - p_a)E(\check{b}, \check{b}^*) \end{aligned}$$

The expected social welfare under the contract is

$$\begin{aligned} \pi &= \pi_a + \pi_b \\ &= E(\check{b}) - E(\check{a}) + (1 - p_b)E(\check{a}^*, \check{a}) - (1 - p_a)E(\check{b}, \check{b}^*) \\ &= \pi^{fallback} + H(\check{a}^*, \check{b}^*), \end{aligned}$$

where $\pi^{fallback} = E(\check{b}) - E(\check{a})$ is the expected social welfare that would prevail without the contract (i.e. expected welfare from the outside offers), and the *excess*

$$H(x, y) = \int_{-\infty}^y g(\check{b}) d\check{b} \int_x^{\infty} \check{a} f(\check{a}) d\check{a} - \int_x^{\infty} f(\check{a}) d\check{a} \int_{-\infty}^y \check{b} g(\check{b}) d\check{b}.$$

The contractor's individual rationality (IR) constraint states that he will participate in the contract only if that gives him higher expected payoff than waiting for the outside offer:

$$\pi_a \geq -E(\check{a}) \Leftrightarrow \phi(\rho, a, b) \leq (1 - p_b)E(\check{a}^*, \check{a}).$$

Similarly, the contractee's IR constraint is

$$\pi_b \geq E(\check{b}) \Leftrightarrow (1 - p_a)E(\check{b}, \check{b}^*) \leq \phi(\rho, a, b).$$

2.2 Simultaneous decommitting, both pay if both decommit (SIMUDBP)

In our simultaneous decommitting games, agents have to reveal their decommitment decisions simultaneously. We first discuss the SIMUDBP variant where both have to pay the penalties if both decommit. The contractor decommits if $p_b \cdot (-\check{a} + b - a) + (1 - p_b)(-\check{a} - a) > p_b \cdot (-\check{a} + b) + (1 - p_b)(-\rho)$. A contract where $p_b = 1$ cannot be strictly individually rational to both agents since breach will occur for sure. If $p_b < 1$ the inequality above shows that the contractor decommits if his outside offer is less than a threshold $\check{a}^* = \rho - \frac{a}{1 - p_b}$. So, $p_a = \int_{-\infty}^{\check{a}^*} f(\check{a}) d\check{a}$.

The contractee decommits if $(1 - p_a)(\check{b} - b) + p_a(\check{b} - b + a) > (1 - p_a)\rho + p_a(\check{b} + a)$. A contract where $p_a = 1$ cannot be strictly individually rational to both agents since breach will occur for sure. If $p_a < 1$ the inequality shows that the contractee decommits if her outside offer exceeds a threshold $\check{b}^* = \rho + \frac{b}{1 - p_a}$. So, $p_b = \int_{\check{b}^*}^{\infty} g(\check{b}) d\check{b}$.

The contractor's expected payoff under the contract is

$$\begin{aligned} \pi_a &= p_b \left[\int_{-\infty}^{\check{a}^*} (-\check{a} + b - a) f(\check{a}) d\check{a} + \int_{\check{a}^*}^{\infty} (-\check{a} + b) f(\check{a}) d\check{a} \right] \\ &\quad + (1 - p_b) \left[\int_{-\infty}^{\check{a}^*} (-\check{a} - a) f(\check{a}) d\check{a} + \int_{\check{a}^*}^{\infty} (-\rho) f(\check{a}) d\check{a} \right] \\ &= p_b [-E(\check{a}, \check{a}^*) + (b - a)p_a + b(1 - p_a) - E(\check{a}^*, \check{a})] \\ &\quad + (1 - p_b) [-E(\check{a}, \check{a}^*) - a p_a - \rho(1 - p_a)] \\ &= -[p_a a - p_b b + \rho(1 - p_a)(1 - p_b)] \\ &\quad - E(\check{a}) + (1 - p_b)E(\check{a}^*, \check{a}) \\ &= -E(\check{a}) - \phi(\rho, a, b) + (1 - p_b)E(\check{a}^*, \check{a}), \text{ where} \\ &\quad \phi(\rho, a, b) = p_a a - p_b b + \rho(1 - p_a)(1 - p_b). \end{aligned}$$

The contractee's expected payoff under the contract is

$$\begin{aligned} \pi_b &= p_a \left[\int_{\check{b}^*}^{\infty} g(\check{b})(\check{b} - b + a) d\check{b} + \int_{-\infty}^{\check{b}^*} (\check{b} + a) g(\check{b}) d\check{b} \right] \\ &\quad + (1 - p_a) \left[\int_{\check{b}^*}^{\infty} g(\check{b})(\check{b} - b) d\check{b} + \int_{-\infty}^{\check{b}^*} \rho g(\check{b}) d\check{b} \right] \\ &= p_a [E(\check{b}^*, \check{b}) + p_b(a - b) + E(\check{b}, \check{b}^*) + (1 - p_b)a] \\ &\quad + (1 - p_a) [E(\check{b}^*, \check{b}) - p_b b + \rho(1 - p_b)] \\ &= [p_a a - p_b b + \rho(1 - p_a)(1 - p_b)] \\ &\quad + E(\check{b}) - (1 - p_a)E(\check{b}, \check{b}^*) \\ &= E(\check{b}) + \phi(\rho, a, b) - (1 - p_a)E(\check{b}, \check{b}^*) \end{aligned}$$

The expected social welfare under the contract is

$$\begin{aligned} \pi &= \pi_a + \pi_b \\ &= E(\check{b}) - E(\check{a}) + (1 - p_b)E(\check{a}^*, \check{a}) - (1 - p_a)E(\check{b}, \check{b}^*) \\ &= \pi^{fallback} + H(\check{a}^*, \check{b}^*), \end{aligned}$$

where $\pi^{fallback}$ and $H(x, y)$ are defined as in Sec. 2.1.

2.3 Simultaneous decommitting, neither pays if both decommit (SIMUDNP)

In a simultaneous decommitting game where neither agent has to pay the penalty if both decommit (SIMUDNP), the contractor decommits if $p_b \cdot (-\check{a}) + (1 - p_b)(-\check{a} - a) > p_b \cdot (-\check{a} + b) + (1 - p_b)(-\rho)$. A contract where $p_b = 1$ cannot be strictly individually rational to both agents since breach will occur for sure. When $p_b < 1$ the inequality above shows that the contractor decommits if his outside offer is less than a threshold $\check{a}^* = \rho - a - \frac{bp_b}{1-p_b}$. So, $p_a = \int_{-\infty}^{\check{a}^*} f(\check{a})d\check{a}$.

The contractee decommits if $(1 - p_a)(\check{b} - b) + p_a\check{b} > (1 - p_a)\rho + p_a(\check{b} + a)$. A contract where $p_a = 1$ cannot be strictly individually rational to both agents since breach will occur for sure. If $p_a < 1$ the inequality above shows that the contractee decommits if her outside offer exceeds a threshold $\check{b}^* = \rho + b - \frac{ap_a}{1-p_a}$. So, $p_b = \int_{\check{b}^*}^{\infty} g(\check{b})d\check{b}$.

The contractor's expected payoff under the contract is

$$\begin{aligned} \pi_a &= p_b \left[\int_{-\infty}^{\check{a}^*} (-\check{a})f(\check{a})d\check{a} + \int_{\check{a}^*}^{\infty} (-\check{a} + b)f(\check{a})d\check{a} \right] \\ &\quad + (1 - p_b) \left[\int_{-\infty}^{\check{a}^*} (-\check{a} - a)f(\check{a})d\check{a} + \int_{\check{a}^*}^{\infty} (-\rho)f(\check{a})d\check{a} \right] \\ &= p_b [-E(\check{a}, \check{a}^*) + b(1 - p_a) - E(\check{a}^*, \check{a})] \\ &\quad + (1 - p_b) [-E(\check{a}, \check{a}^*) - ap_a - \rho(1 - p_a)] \\ &= -[p_a(1 - p_b)a - (1 - p_a)p_b b + \rho(1 - p_a)(1 - p_b)] \\ &\quad - E(\check{a}) + (1 - p_b)E(\check{a}^*, \check{a}) \\ &= -E(\check{a}) - \phi(\rho, a, b) + (1 - p_b)E(\check{a}^*, \check{a}), \text{ where} \end{aligned}$$

$\phi(\rho, a, b) = p_a(1 - p_b)a - (1 - p_a)p_b b + \rho(1 - p_a)(1 - p_b)$. The contractee's expected payoff under the contract is

$$\begin{aligned} \pi_b &= p_a \left[\int_{\check{b}^*}^{\infty} g(\check{b})\check{b}d\check{b} + \int_{-\infty}^{\check{b}^*} (\check{b} + a)g(\check{b})d\check{b} \right] \\ &\quad + (1 - p_a) \left[\int_{\check{b}^*}^{\infty} g(\check{b})(\check{b} - b)d\check{b} + \int_{-\infty}^{\check{b}^*} \rho g(\check{b})d\check{b} \right] \\ &= p_a [E(\check{b}^*, \check{b}) + E(\check{b}, \check{b}^*) + (1 - p_b)a] \\ &\quad + (1 - p_a) [E(\check{b}^*, \check{b}) - p_b b + \rho(1 - p_b)] \\ &= [p_a(1 - p_b)a - (1 - p_a)p_b b + \rho(1 - p_a)(1 - p_b)] \\ &\quad + E(\check{b}) - (1 - p_a)E(\check{b}, \check{b}^*) \\ &= E(\check{b}) + \phi(\rho, a, b) - (1 - p_a)E(\check{b}, \check{b}^*) \end{aligned}$$

The expected social welfare under the contract is

$$\begin{aligned} \pi &= \pi_a + \pi_b \\ &= E(\check{b}) - E(\check{a}) + (1 - p_b)E(\check{a}^*, \check{a}) - (1 - p_a)E(\check{b}, \check{b}^*) \\ &= \pi^{fallback} + H(\check{a}^*, \check{b}^*), \end{aligned}$$

where $\pi^{fallback}$ and $H(x, y)$ are defined as in Sec. 2.1.

3 Revenue equivalence

Now, which of the leveled commitment contracting protocols would be best for the agents? In this sec-

tion we show that if the contract price and the decommitting penalties are optimized for each game (SEQD, SIMUDBP, or SIMUDNP) separately, each of the games leads to the same expected social welfare. This is surprising since the optimal contracts differ for the games. Also, for a given suboptimal contract, the decommitting thresholds, decommitting probabilities, and expected welfare generally differ across the games.

We start by showing that if a leveled commitment contract can generate positive excess, H , i.e. it can lead to higher expected social welfare than making no contract and waiting for the outside offers, then an unconstrained optimum exists.

Lemma 1 *Let f and g be probability distributions on $(-\infty, \infty)$ with finite expectations. Let*

$$H(x, y) = \int_{-\infty}^y g(\check{b})d\check{b} \int_x^{\infty} \check{a}f(\check{a})d\check{a} - \int_x^{\infty} f(\check{a})d\check{a} \int_{-\infty}^y \check{b}g(\check{b})d\check{b} \quad (\dagger)$$

If $\max_{x,y} H(x, y) > 0$, then there exists a global maximal point (a^, b^*) of H that satisfies*

$$a^* = \frac{\int_{-\infty}^{b^*} \check{b}g(\check{b})d\check{b}}{\int_{-\infty}^{b^*} g(\check{b})d\check{b}}, \quad b^* = \frac{\int_{a^*}^{\infty} \check{a}f(\check{a})d\check{a}}{\int_{a^*}^{\infty} f(\check{a})d\check{a}}. \quad (\ddagger)$$

Specifically,

$$H(a^*, b^*) = \max_{x,y} H(x, y) = (b^* - a^*)(1 - p_x)(1 - p_y) \text{ where}$$

$$p_x = \int_{-\infty}^{a^*} f(\check{a})d\check{a}, \quad p_y = \int_{b^*}^{\infty} g(\check{b})d\check{b}.$$

The proofs are highly nontrivial. They are omitted due to the space limitation, but they can be found in an extended technical report (Sandholm & Zhou 1999).

Now we are ready to present the main result (its proof uses Lemma 1):

Theorem 1 *Let f and g have finite expectations. If an expected social welfare maximizing IR leveled commitment contract is chosen for each of the protocols (SEQD, SIMUDBP, and SIMUDNP) separately, each protocol yields the same expected social welfare. The pairs (possibly multiple per protocol) of decommitting thresholds and the associated decommitting probabilities will also be the same. The optimal contract may differ for the different protocols, but in each protocol the optimal decommitment penalties are nonnegative.*

3.1 Existence of optimal IR contracts

It turns out that if some leveled commitment contract generates positive excess to the agents in the aggregate, then there exists an optimal leveled commitment contract that generates positive excess to each agent, i.e. the contract is agreeable in the sense of individual rationality. More strongly:

Proposition 1 *Let f and g have finite expectations. For SEQD, SIMUDBP, and SIMUDNP, $\max_{x,y} H(x, y) > 0$ iff there exists an expected social welfare maximizing contract (ρ, a, b) that is IR for both agents.*

Based on this result, throughout the rest of the paper we assume that $\max_{x,y} H(x,y) > 0$. Recall that we denote an optimal (x,y) by (a^*, b^*) .

4 Integrative vs. distributive bargaining

Proposition 1 showed that among optimal contracts there are ones that are beneficial for both parties. However, the question of how to divide the excess between the agents remains, i.e. how to choose among the individually rational contracts. Each agent's excess is her expected payoff under the contract minus the expected fallback payoff: $e_a = \pi_a - \pi_a^{\text{fallback}} = \pi_a + E(\check{a})$ and $e_b = \pi_b - \pi_b^{\text{fallback}} = \pi_b - E(\check{b})$. It is conceivable that in leveled commitment contracts there is a tradeoff between integrative bargaining (maximizing the expected social welfare) and distributive bargaining (splitting the excess between the agents). It could be that some splits cannot be supported by an optimal contract. However, it turns out that any individually rational split can be supported by an optimal contract:

Proposition 2 *Let f and g have finite expectations. For each one of the games (SEQD, SIMUDBP, and SIMUDNP), for any given $\beta \in [0, 1]$ there exists an expected social welfare maximizing contract where $e_a = \beta H(a^*, b^*)$, and $e_b = (1 - \beta)H(a^*, b^*)$.*

Since the agents would only agree to individually rational splits anyway, Proposition 2 means that for all practical purposes, integrative and distributive bargaining do not hinder each other in leveled commitment contracts. Of course, the contract has to be chosen carefully. First ρ should be chosen (in the IR range) which determines the distributive part. Then the penalties, a and b , are calculated based on ρ in order to maximize expected social welfare. Choosing the penalties first does not allow the same separation of integrative and distributive bargaining because once a and b are fixed, the choice of ρ is limited if one wants to construct an expected social welfare maximizing contract.

4.1 Fair optimal contracts

Proposition 2 implies that there is no tradeoff between expected social welfare maximization and fairness (aka. symmetry, equality) in leveled commitment contracts since both of these desiderata can be satisfied simultaneously. There exists an expected social welfare maximizing contract where the excess is split equally between the agents ($e_a = e_b$).

Distributive bargaining is a large research field of its own, and a literature review is beyond the scope of this short paper. However, significant support has been given for solutions that maximize the product of the excesses (Nash 1950a; Rosenschein & Zlotkin 1994). It turns out that in leveled commitment contracts, such product maximization is equivalent to choosing an expected social welfare maximizing contract that splits excess equally:

Proposition 3 *Let f and g have finite expectations. For each one of the games (SEQD, SIMUDBP, and*

SIMUDNP), $e_a e_b$ is maximized iff the contract maximizes expected social welfare and $e_a = e_b$. Such a contract always exists.

5 Nonuniqueness

Usually the excess, $H(x,y)$, has a unique global maximum, but not always. Let (x_0, y_0) be a global maximum. If $f(x) = 0$ in some neighborhood of x_0 and $g(y) = 0$ in a neighborhood of y_0 , there exists a neighborhood of (x_0, y_0) in which all (x,y) maximize $H(x,y)$.

The excess, $H(x,y)$, can also have multiple global maxima that are not in the same neighborhood. In particular, the pair (a^*, b^*) determined in Lemma 1 is not always unique. The following example shows a case with 3 local maxima of which 2 are globally maximal.

$$f(x) = 1/10 \text{ if } 0 \leq x \leq 10, \text{ and } 0 \text{ otherwise.}$$

$$g(y) = \begin{cases} 117/3520 & \text{if } 0 \leq y < \frac{320}{47} \\ 42939/165440 & \text{if } \frac{320}{47} \leq y \leq 10 \\ 0 & \text{otherwise.} \end{cases}$$

We use Lemma 1 to find all local maxima. Because $f(x) > 0$ and $g(y) > 0$ for all $x, y \in [0, 10]$, each local maximum (x,y) must satisfy the mutual equations (\dagger). The first of those equations can be reduced to $y = (10 + x)/2$, i.e. $x = 2y - 10$. For the second one, the cases $y \leq y_0$ and $y > y_0$ have to be treated separately. For each case, the mutual equations are solved to find (x,y) . The solutions, i.e. the local maxima, are $(10/3, 20/3)$, $(4, 7)$, and $(5, 15/2)$. The excess values are $H(10/3, 20/3) = H(5, 15/2) = \frac{65}{132}$, $H(4, 7) = \frac{27}{55}$. Since $27/55 < 65/132$, both $(10/3, 20/3)$ and $(5, 15/2)$ are global maxima, $(4, 7)$ is only a local maximum.

Nonuniqueness of the optimal threshold pair—and the associated nonuniqueness of the optimal contract (ρ, a, b) —does not prevent the use of leveled commitment contracts. To maximize expected social welfare, the agents can pick any one of the optimal contracts.

6 Agents with risk attitudes

So far we discussed agents that attempt to maximize expected payoff, i.e. they are risk neutral. For a utility maximizing agent, i , to be risk neutral, the utility function, $u_i : \pi_i \rightarrow \mathbb{R}$, would be linearly increasing. Risk attitudes are captured in the usual way by making u_i nonlinear. We now show that the revenue equivalence of leveled commitment contracts does not always hold for agents that are not risk neutral, and in different settings, different leveled commitment protocols are best in terms of expected social welfare. Let $f(\check{a}) = \frac{1}{100}$ if $\check{a} \in [0, 100]$, and $g(\check{b}) = \frac{1}{110}$ if $\check{b} \in [0, 110]$. If $u_a(x) = u_b(x) = x^3$, $\max_{\text{SEQD}} \pi \approx 284192$, $\max_{\text{SIMUDBP}} \pi \approx 322522$, $\max_{\text{SIMUDNP}} \pi \approx 334194$. If $u_a(x) = u_b(x) = x^{1/3}$, $\max_{\text{SEQD}} \pi \approx 0.912$, $\max_{\text{SIMUDBP}} \pi \approx 0.925$, $\max_{\text{SIMUDNP}} \pi \approx 0.905$.

7 Conclusions

Leveled commitment contracts are often more practical than contingency contracts. However, they cannot always achieve the same social welfare because the agents decommit insincerely: some contracts are inefficiently kept. Our intuitions suggested that sequential decommitting protocols would lead to higher social welfare than simultaneous ones since the last agent decommits truthfully. We also thought that protocols where neither agent pays a penalty if both decommit would promote decommitting and increase welfare. However, we showed that, surprisingly, all of the protocols lead to the same expected social welfare when the contract price and decommitting penalties are optimized for each protocol separately.

Our derivations allow agents to construct optimal leveled commitment contracts, and to divide the gains arbitrarily (as long as each agent benefits), e.g. equally. Using this theory we have developed fast algorithms for contract optimization,⁴ and provide a contract optimization service on the web as part of *eMediator*, our next generation electronic commerce server, see <http://ecommerce.cs.wustl.edu/contracts.html>.

References

- Calamari, J. D., and Perillo, J. M. 1977. *The Law of Contracts*. West Publishing Co., 2nd edition.
- Diamond, P. A., and Maskin, E. 1979. An equilibrium analysis of search and breach of contract, I: Steady states. *Bell J. of Economics* 10:282–316.
- Kraus, S. 1993. Agents contracting tasks in non-collaborative environments. In *Proc. Eleventh National Conference on Artificial Intelligence (AAAI-93)*, 243–248.
- Nash, J. 1950a. The bargaining problem. *Econometrica* 18:155–162.
- Nash, J. 1950b. Equilibrium points in n-person games. *Proc. of the National Academy of Sciences* 36:48–49.
- Posner, R. A. 1977. *Economic Analysis of Law*. Little, Brown and Company, 2nd edition.
- Raiffa, H. 1982. *The Art and Science of Negotiation*. Cambridge, Mass.: Harvard Univ. Press.
- Rosenschein, J. S., and Zlotkin, G. 1994. *Rules of Encounter: Designing Conventions for Automated Negotiation among Computers*. MIT Press.
- Sandholm, T. W., and Lesser, V. R. 1996. Advantages of a leveled commitment contracting protocol. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, 126–133.
- Sandholm, T. W., and Zhou, Y. 1999. Revenue equivalence of leveled commitment contracts. Technical Report WUCS-99-03, Washington University, Department of Computer Science.
- Sandholm, T. W.; Sikka, S.; and Norden, S. 1999. Algorithms for optimizing leveled commitment contracts. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence (IJCAI)*. Extended version: Washington University, Department of Computer Science technical report WUCS-99-02.
- Sandholm, T. W. 1993. An implementation of the contract net protocol based on marginal cost calculations. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, 256–262.
- Sen, S. 1993. *Tradeoffs in Contract-Based Distributed Scheduling*. Ph.D. Dissertation, Univ. of Michigan.
- Smith, R. G. 1980. The contract net protocol: High-level communication and control in a distributed problem solver. *IEEE Transactions on Computers* C-29(12):1104–1113.

⁴See (Sandholm, Sikka, & Norden 1999) for the algorithms.