ELSEVIER

Artificial Intelligence 142 (2002) 239–264

Artificial Intelligence

www.elsevier.com/locate/artint

Surplus equivalence of leveled commitment contracts ☆

Tuomas Sandholm [a,*], Yunhong Zhou [b]

[a] *Computer Science Department, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, USA*
[b] *Compaq Systems Research Center, 130 Lytton Avenue, Palo Alto, CA 94301, USA*

Received 16 March 2001

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 strategically 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 mechanisms lead to different decommitting thresholds, probabilities, and expected social welfare. However, this paper shows that each of these mechanisms leads to the same social welfare when the contract price and penalties are optimized for each mechanism separately. Our derivations allow agents to construct optimal leveled commitment contracts. We show that such integrative bargaining does not hinder distributive bargaining: the surplus can be divided arbitrarily (as long as each agent benefits), e.g., equally, without compromising optimality. Nonuniqueness questions are answered. We also show that surplus equivalence ceases to hold if agents are not risk neutral.

Keywords: Contracting; Automated contracting; Breach; Breach of contract; Mechanism design; Bargaining; Bargaining under uncertainty; Automated negotiation; Game theory; Multiagent systems

☆ Supported by NSF under CAREER Award IRI-9703122, Grant IRI-9610122, Grant IIS-9800994, Grant ITR IIS-0081246, and Grant ITR IIS-0121678. Early versions of this paper appeared at the International Conference on Multi-Agent Systems (ICMAS-2000), Boston, MA, pp. 247–254, and at the AAAI-1999 Workshop on Negotiation: Settling Conflicts and Identifying Opportunities, Orlando, FL, pp. 38–43, and as a Washington University St. Louis, Department of Computer Science Technical Report WUCS-99-03, 1999.

* Corresponding author.
E-mail addresses: sandholm@cs.cmu.edu (T. Sandholm), yzhou@pa.dec.com (Y. Zhou).

0004-3702/02/$ – see front matter
PII: S0004-3702(02)00275-8

1. Introduction

In automated negotiation systems consisting of self-interested agents, contracts have traditionally been binding [5,15,17]. Once an agent agrees to a contract, she 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. Similarly, a contract may have too low expected payoff *ex ante*, but in some realizations of the future events it may be desirable. 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 [24], 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 using a contracting approach [23].

Contingency contracts have been suggested for utilizing the potential provided by future events among self-interested agents [14]. 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, especially as a negotiation instrument among software agents, for several reasons. The space of combinations of future events can be large and it is rare that both agents are cognizant of all possible future worlds *ex ante* and have evaluated their utility in each future world. Even if the real-world parties are cognizant of all possible future worlds, building this information into a software agent can be an error-prone and prohibitively tedious undertaking. Also, to maximize the economic efficiency that a contingency contract can provide, the agents may need to condition the contract on every possible combination of future events, which leads to a combinatorial explosion in the contingency table. Finally, when events are not mutually observable, *ex post*, the observing agent could lie about what transpired. Therefore, contingency contracts generally rely on some nonmanipulable event verification mechanism.

As a response to these practical difficulties associated with contingency contracts, *leveled commitment contracts* were recently introduced as another method for capitalizing on future events [19]. From an AI search perspective, they can be viewed as a backtracking instrument that works even among self-interested agents.[1] 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—that is, wants to be freed from the obligations of the contract—the agent can do so simply by paying the decommitment penalty to the other party. The method requires no explicit conditioning

[1] More conventional backtracking techniques may not be applicable in this setting because a backtrack would only occur if every one of the self-interested agents would gain from it. An agent would not agree to give up gains that it has already obtained through negotiation.

of the contract on future events: each agent can do her own conditioning dynamically. No event verification method against lying is required either.

Principles for assessing decommitment penalties have been studied in the economic analysis of law [3,13], but the purpose has usually 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 is used as an active method for utilizing the potential provided by an uncertain future.[2] By design, breach will occur at times, and this increases the social welfare among the contract parties on an expected value basis.

Another key difference between classic work on the economics of law and this paper is that we do not assume that agents breach sincerely, but rather we take into account the fact that rational agents will breach strategically. Specifically, 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. This also distinguishes our research from work on constructing contracts by combining different option contracts.

Sandholm and Lesser recently showed that despite such strategic 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 [20]. The intuitive reason for this is that in many of the cases where the leveled commitment contract turns out undesirable *ex post*, it will be undone. This paper studies the same setting and the same contract types as they did, but derives new results.

The rest of the paper is organized as follows. Section 2 introduces the contracting setting. In Section 3 we review the different leveled commitment contracting mechanisms (protocols), and how rational agents would decommit in them. This gives rise to the natural question: which mechanism leads to the best results for the agents? In Section 4 we derive the somewhat surprising result that although the optimal contract parameters (price and decommitment penalties) differ among the mechanisms, if the parameters are optimized for each mechanism separately, the mechanisms lead to the *same* social welfare. Section 5 presents positive results regarding the interplay between integrative and distributive bargaining in leveled commitment contracting, and shows how to construct a fair optimal contract. Section 6 discusses nonuniqueness of the optimal contract. Section 7 shows that surplus equivalence ceases to hold if agents are not risk neutral. Finally, Section 8 concludes.

2. Our contracting setting

Consider a contracting setting with two risk neutral agents who attempt to maximize their own expected payoff: the *contractor* who pays to get a task done, and the *contractee*

[2] 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 [4].

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 framework can be interpreted as modeling other types of settings than task allocation also, for example general allocation of rights and obligations where the agents' costs and gains of the rights and obligations may change. In what follows, we word the results in the context of task allocation.

The contractor tries to minimize the contract price ρ that he has to pay the contractee. The contractee tries to maximize the payoff ρ that she receives from the contractor. We study a setting where the future of the agents involves uncertainty. We model this as the agents potentially receiving outside offers (which are full-commitment contracts[3]).[4] The contractor's best (lowest) outside offer v is only probabilistically known *ex ante* by both agents, and is characterized by a probability density function $f(v)$. The contractee's best (highest) outside offer w is also only probabilistically known *ex ante*, and is characterized by a probability density function $g(w)$.[5,6] The variables v and w are assumed statistically independent, and $f(v)$ and $g(w)$ are common knowledge.

The contractor's options are either to make a contract with the contractee or to wait for v. Similarly, the contractee's options are either to make a contract with the contractor or to wait for w. 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.

In this paper, we focus on a setting where each agent has exactly one chance to decommit. The contractor has to decide on decommitting when he knows his outside offer v but does not know the contractee's outside offer w. Similarly, the contractee has to decide on decommitting when she knows her outside offer w but does not know the contractor's. This seems realistic from a practical automated contracting perspective.[7]

As we will show, the contractor will decommit whenever his outside offer is lower than a threshold, v^*, which we call the contractor's decommitting threshold. We denote by p_a the probability that the contractor will decommit, i.e., the chance that his outside offer is

---

[3] Equivalently, the outside offers could be leveled commitment contracts, in which case the agent would use the expected payoff from such a contract as its "value" of the outside offer.

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

[5] If the contractor does not receive an outside offer, v corresponds to its best (lowest) outstanding outside offer. One can also interpret $-v$ as the contractor's fallback payoff, that is, payoff that it receives if no contract is made. Analogously, if the contractee does not receive an outside offer, w corresponds to its best (highest) outstanding outside offer or its fallback payoff.

[6] 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 $f(v)$ and/or $g(w)$ is on one point.

[7] The agents could be forced to make their decommitting decisions at a particular time, or equivalently they could be allowed to make the decisions during an interval (the model would still hold at least as long as the agents' outside options/valuations do not change during the interval). The model could also be used in a setting where the agents do not know their outside offers exactly at the time they have to decide on decommitting. In that case, to evaluate the desirability of decommitting, each agent would use as its best outside offer the *expectation* of its best outside offer based on the probability distribution it has about its best outside offer at that time.

Table 1
Main symbols used in the paper

ρ	Contract price.
a	Contractor's decommitment penalty.
b	Contractee's decommitment penalty.
v	Contractor's best (lowest) outside offer.
w	Contractee's best (highest) outside offer.
$f(v)$	*Ex ante* probability density function over v.
$g(w)$	*Ex ante* probability density function over w.
v^*	Contractor's decommitting threshold.
w^*	Contractee's decommitting threshold.
p_a	Probability that the contractor decommits.
p_b	Probability that the contractee decommits.
π_a	Contractor's expected payoff.
π_b	Contractee's expected payoff.
$H(v^*, w^*)$	Surplus generated by making a contract.

less than v^*. Similarly, the contractee will decommit whenever her outside offer is greater then a threshold, w^*, which we call the contractee's decommitting threshold. We denote by p_b the probability that the contractee will decommit, i.e., the chance that her outside offer is greater than w^*. All of these variables (v^*, p_a, w^*, p_b) are endogenous to the model. That is, their values are not assumed, but rather the values are determined by the model. Table 1 summarizes the main symbols that are used in the paper.

3. Leveled commitment contracting mechanisms

A key 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*. Despite such strategic 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 [20]. We derive the Nash equilibrium [10] 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 strategically in this way. The equilibrium depends on the interaction mechanism. We study six natural leveled commitment contracting mechanisms, which cover the space of sensible mechanisms that we conceived. The mechanisms differ based on the order in which the agents have to reveal their decisions (decommit/not), and based on whether or not the agents have to pay the penalties to each other if both decommit:

(1) Contractee has to reveal its decision first, both pay if both decommit (SEQD).
(2) Contractee has to reveal its decision first, neither pays if both decommit. If the contractee decommits, the contractor can only lose by decommitting (assuming the contractor's penalty is positive). Therefore, the situation where both decommit never occurs. Thus this mechanism is analogous to SEQD, and will not be discussed further.

(3) Contractor has to reveal its decision first, both pay if both decommit. This case is mathematically analogous to SEQD, and will not be discussed further.
(4) Contractor has to reveal its decision first, neither pays if both decommit. If the contractor decommits, the contractee can only lose by decommitting (assuming the contractee's penalty is positive). Therefore, the situation where both decommit never occurs. Thus this mechanism is analogous to the one above, and will not be discussed further.
(5) The agents reveal their decisions simultaneously (not necessarily at the same time, as long as the second one to reveal does not learn what the first agent decided before the second agent has to decide), both pay if both decommit (SIMUDBP).
(6) The agents reveal their decisions simultaneously (not necessarily at the same time, as long as the second one to reveal does not learn what the first agent decided before the second agent has to decide), neither pays if both decommit (SIMUDNP).

Put together, in the rest of the paper we can focus on three mechanisms (SEQD, SIMUDBP, SIMUDNP), and the results will apply to all six.

3.1. Sequential decommitting, contractee first (SEQD)

In a sequential decommitting (SEQD) game, one agent has to declare her decommitting decision before the other. We focus on the case where the contractee has to decommit first. In the subgame where the contractee has not decommitted, the contractor's best move is to decommit if $-v - a > -\rho$, i.e., the contractor decommits if his outside offer, v, is below a threshold $v^*_{\mathrm{SEQD}} = \rho - a$. So, the probability that he decommits is

$$p_a = \int_{-\infty}^{v^*_{\mathrm{SEQD}}} f(v)\,\mathrm{d}v.$$

The contractee gets $w - b$ if she decommits, $w + a$ if she does not but the contractor does, and ρ if neither decommits. Thus the contractee decommits if $w - b > p_a(w + a) + (1 - p_a)\rho$. 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 $w^*_{\mathrm{SEQD}} = \rho + (b + ap_a)/(1 - p_a)$. So, the probability that she decommits is

$$p_b = \int_{w^*_{\mathrm{SEQD}}}^{\infty} g(w)\,\mathrm{d}w.$$

The rest of the paper uses the following six shorthand notations:

$$E(v) \equiv \int_{-\infty}^{\infty} vf(v)\,\mathrm{d}v, \qquad E(w) \equiv \int_{-\infty}^{\infty} wg(w)\,\mathrm{d}w,$$

$$E(v, v^*) \equiv \int_{-\infty}^{v^*} v f(v)\,\mathrm{d}v, \qquad E(v^*, v) \equiv \int_{v^*}^{\infty} v f(v)\,\mathrm{d}v,$$

$$E(w, w^*) \equiv \int_{-\infty}^{w^*} w g(w)\,\mathrm{d}w, \qquad E(w^*, w) \equiv \int_{w^*}^{\infty} w g(w)\,\mathrm{d}w.$$

We are now ready to study the value generated by a contract. The contractor's expected payoff under the contract is

$$\begin{aligned}
\pi_a = {} & p_b\left[\int_{-\infty}^{\infty} (-v+b) f(v)\,\mathrm{d}v\right] \\
& + (1-p_b)\left[\int_{-\infty}^{v^*_{\mathrm{SEQD}}} (-v-a) f(v)\,\mathrm{d}v + \int_{v^*_{\mathrm{SEQD}}}^{\infty} (-\rho) f(v)\,\mathrm{d}v\right] \\
= {} & p_b\big[b - E(v)\big] + (1-p_b)\big[-E(v, v^*_{\mathrm{SEQD}}) - a p_a - \rho(1-p_a)\big] \\
= {} & -\big[p_a(1-p_b)a - p_b b + (1-p_a)(1-p_b)\rho\big] - E(v) + (1-p_b)E(v^*_{\mathrm{SEQD}}, v) \\
= {} & -E(v) - \phi_{\mathrm{SEQD}}(\rho, a, b) + (1-p_b)E(v^*_{\mathrm{SEQD}}, v),
\end{aligned}$$

where

$$\phi_{\mathrm{SEQD}}(\rho, a, b) = p_a(1-p_b)a - p_b b + (1-p_a)(1-p_b)\rho.$$

The contractee's expected payoff under the contract is

$$\begin{aligned}
\pi_b = {} & \int_{w^*_{\mathrm{SEQD}}}^{\infty} g(w)(w-b)\,\mathrm{d}w + \int_{-\infty}^{w^*_{\mathrm{SEQD}}} g(w)\big[p_a(w+a) + (1-p_a)\rho\big]\,\mathrm{d}w \\
= {} & -p_b b + (1-p_b)\big[p_a a + (1-p_a)\rho\big] + E(w^*_{\mathrm{SEQD}}, w) + p_a E(w, w^*_{\mathrm{SEQD}}) \\
= {} & \big[p_a(1-p_b)a - p_b b + (1-p_a)(1-p_b)\rho\big] + E(w) - (1-p_a)E(w, w^*_{\mathrm{SEQD}}) \\
= {} & E(w) + \phi_{\mathrm{SEQD}}(\rho, a, b) - (1-p_a)E(w, w^*_{\mathrm{SEQD}}).
\end{aligned}$$

The expected social welfare under the contract is

$$\begin{aligned}
\pi = \pi_a + \pi_b = {} & E(w) - E(v) + (1-p_b)E(v^*_{\mathrm{SEQD}}, v) - (1-p_a)E(w, w^*_{\mathrm{SEQD}}) \\
= {} & \pi^{\mathrm{fallback}} + H(v^*_{\mathrm{SEQD}}, w^*_{\mathrm{SEQD}}),
\end{aligned}$$

where $\pi^{\mathrm{fallback}} = E(w) - E(v)$ is the expected social welfare that would prevail without the contract (i.e., expected welfare from the outside offers), and the *surplus* created by the contract is $H(v^*_{\mathrm{SEQD}}, w^*_{\mathrm{SEQD}})$ where

$$H(x, y) = \int_{-\infty}^{y} g(w)\,\mathrm{d}w \int_{x}^{\infty} v f(v)\,\mathrm{d}v - \int_{x}^{\infty} f(v)\,\mathrm{d}v \int_{-\infty}^{y} w g(w)\,\mathrm{d}w.$$

This notion of surplus is a key concept, and will be used throughout the rest of the paper.

3.2. 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 $p_b \cdot (-v+b-a) + (1-p_b)(-v-a) > p_b \cdot (-v+b) + (1-p_b)(-\rho)$. A contract where $p_b = 1$ cannot be strictly individually rational to both agents since breach will occur for sure. On the other hand, when $p_b < 1$ the inequality above shows that the contractor decommits if his outside offer is less than a threshold $v^*_{\mathrm{SIMUDBP}} = \rho - a/(1-p_b)$. So, the probability that he decommits is

$$p_a = \int_{-\infty}^{v^*_{\mathrm{SIMUDBP}}} f(v)\,\mathrm{d}v.$$

The contractee decommits if $(1-p_a)(w-b) + p_a(w-b+a) > (1-p_a)\rho + p_a(w+a)$. 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 $w^*_{\mathrm{SIMUDBP}} = \rho + b/(1-p_a)$. So,

$$p_b = \int_{w^*_{\mathrm{SIMUDBP}}}^{\infty} g(w)\,\mathrm{d}w.$$

The contractor's expected payoff under the contract is

$$\begin{aligned}
\pi_a &= p_b\left[\int_{-\infty}^{v^*_{\mathrm{SIMUDBP}}} (-v+b-a) f(v)\,\mathrm{d}v + \int_{v^*_{\mathrm{SIMUDBP}}}^{\infty} (-v+b) f(v)\,\mathrm{d}v\right] \\
&\quad + (1-p_b)\left[\int_{-\infty}^{v^*_{\mathrm{SIMUDBP}}} (-v-a) f(v)\,\mathrm{d}v + \int_{v^*_{\mathrm{SIMUDBP}}}^{\infty} (-\rho) f(v)\,\mathrm{d}v\right] \\
&= p_b\big[-E(v, v^*_{\mathrm{SIMUDBP}}) + (b-a)p_a - E(v^*_{\mathrm{SIMUDBP}}, v) + b(1-p_a)\big] \\
&\quad + (1-p_b)\big[-E(v, v^*_{\mathrm{SIMUDBP}}) - a p_a - \rho(1-p_a)\big] \\
&= -\big[p_a a - p_b b + \rho(1-p_a)(1-p_b)\big] - E(v) + (1-p_b)E(v^*_{\mathrm{SIMUDBP}}, v) \\
&= -E(v) - \phi_{\mathrm{SIMUDBP}}(\rho, a, b) + (1-p_b)E(v^*_{\mathrm{SIMUDBP}}, v),
\end{aligned}$$

where

$$\phi_{\mathrm{SIMUDBP}}(\rho, a, b) = p_a a - p_b b + \rho(1-p_a)(1-p_b).$$

Similarly, the contractee's expected payoff under the contract is

$$\pi_b = p_a\left[\int_{w^*_{\mathrm{SIMUDBP}}}^{\infty} g(w)(w-b+a)\,\mathrm{d}w + \int_{-\infty}^{w^*_{\mathrm{SIMUDBP}}} (w+a)g(w)\,\mathrm{d}w\right]$$
$$+ (1-p_a)\left[\int_{w^*_{\mathrm{SIMUDBP}}}^{\infty} g(w)(w-b)\,\mathrm{d}w + \int_{-\infty}^{w^*_{\mathrm{SIMUDBP}}} \rho g(w)\,\mathrm{d}w\right]$$
$$= p_a\left[E(w^*_{\mathrm{SIMUDBP}}, w) + E(w, w^*_{\mathrm{SIMUDBP}}) + p_b(a-b) + (1-p_b)a\right]$$
$$+ (1-p_a)\left[E(w^*_{\mathrm{SIMUDBP}}, w) - p_b b + \rho(1-p_b)\right]$$
$$= \left[p_a a - p_b b + \rho(1-p_a)(1-p_b)\right] + E(w) - (1-p_a)E(w, w^*_{\mathrm{SIMUDBP}})$$
$$= E(w) + \phi_{\mathrm{SIMUDBP}}(\rho, a, b) - (1-p_a)E(w, w^*_{\mathrm{SIMUDBP}}).$$

The expected social welfare under the contract is

$$\pi = \pi_a + \pi_b$$
$$= E(w) - E(v) + (1-p_b)E(v^*_{\mathrm{SIMUDBP}}, v) - (1-p_a)E(w, w^*_{\mathrm{SIMUDBP}})$$
$$= \pi^{\mathrm{fallback}} + H(v^*_{\mathrm{SIMUDBP}}, w^*_{\mathrm{SIMUDBP}}),$$

where π^{fallback} and $H(x, y)$ are defined as in Section 3.1.

3.3. 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 $p_b \cdot (-v) + (1-p_b)(-v-a) > p_b \cdot (-v+b) + (1-p_b)(-\rho)$. A contract where $p_b = 1$ cannot be strictly individually rational to both agents since breach will occur for sure. On the other hand, when $p_b < 1$, the inequality above shows that the contractor decommits if his outside offer is less than a threshold $v^*_{\mathrm{SIMUDNP}} = \rho - a - bp_b/(1-p_b)$. So,

$$p_a = \int_{-\infty}^{v^*_{\mathrm{SIMUDNP}}} f(v)\,\mathrm{d}v.$$

The contractee decommits if $(1-p_a)(w-b) + p_a w > (1-p_a)\rho + p_a(w+a)$. 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 $w^*_{\mathrm{SIMUDNP}} = \rho + b + ap_a/(1-p_a)$. So,

$$p_b = \int_{w^*_{\mathrm{SIMUDNP}}}^{\infty} g(w)\,\mathrm{d}w.$$

The contractor's expected payoff under the contract is

$$\begin{aligned}
\pi_a &= p_b\left[\int_{-\infty}^{v^*_{\text{SIMUDNP}}} (-v) f(v)\,\mathrm{d}v + \int_{v^*_{\text{SIMUDNP}}}^{\infty} (-v+b) f(v)\,\mathrm{d}w\right] \\
&\quad + (1-p_b)\left[\int_{-\infty}^{v^*_{\text{SIMUDNP}}} (-v-a) f(v)\,\mathrm{d}v + \int_{v^*_{\text{SIMUDNP}}}^{\infty} (-\rho) f(v)\,\mathrm{d}v\right] \\
&= p_b\big[-E(v, v^*_{\text{SIMUDNP}}) + b(1-p_a) - E(v^*_{\text{SIMUDNP}}, v)\big] \\
&\quad + (1-p_b)\big[-E(v, v^*_{\text{SIMUDNP}}) - ap_a - \rho(1-p_a)\big] \\
&= -\big[p_a(1-p_b)a - (1-p_a)p_b b + \rho(1-p_a)(1-p_b)\big] \\
&\quad - E(v) + (1-p_b)E(v^*_{\text{SIMUDNP}}, v) \\
&= -E(v) - \phi_{\text{SIMUDNP}}(\rho, a, b) + (1-p_b)E(v^*_{\text{SIMUDNP}}, v),
\end{aligned}$$

where

$$\phi_{\text{SIMUDNP}}(\rho, a, b) = p_a(1-p_b)a - (1-p_a)p_b b + \rho(1-p_a)(1-p_b).$$

Similarly, the contractee's expected payoff under the contract is

$$\begin{aligned}
\pi_b &= p_a\left[\int_{w^*_{\text{SIMUDNP}}}^{\infty} g(w)w\,\mathrm{d}w + \int_{-\infty}^{w^*_{\text{SIMUDNP}}} (w+a)g(w)\,\mathrm{d}w\right] \\
&\quad + (1-p_a)\left[\int_{w^*_{\text{SIMUDNP}}}^{\infty} g(w)(w-b)\,\mathrm{d}w + \int_{-\infty}^{w^*_{\text{SIMUDNP}}} \rho g(w)\,\mathrm{d}w\right] \\
&= p_a\big[E(w^*_{\text{SIMUDNP}}, w) + E(w, w^*_{\text{SIMUDNP}}) + (1-p_b)a\big] \\
&\quad + (1-p_a)\big[E(w^*_{\text{SIMUDNP}}, w) - p_b b + \rho(1-p_b)\big] \\
&= \big[p_a(1-p_b)a - (1-p_a)p_b b + \rho(1-p_a)(1-p_b)\big] \\
&\quad + E(w) - (1-p_a)E(w, w^*_{\text{SIMUDNP}}) \\
&= E(w) + \phi_{\text{SIMUDNP}}(\rho, a, b) - (1-p_a)E(w, w^*_{\text{SIMUDNP}}).
\end{aligned}$$

The expected social welfare under the contract is

$$\begin{aligned}
\pi &= \pi_a + \pi_b \\
&= E(w) - E(v) + (1-p_b)E(v^*_{\text{SIMUDNP}}, v) - (1-p_a)E(w, w^*_{\text{SIMUDNP}}) \\
&= \pi^{\text{fallback}} + H(v^*_{\text{SIMUDNP}}, w^*_{\text{SIMUDNP}}),
\end{aligned}$$

where π^{fallback} and $H(x, y)$ are defined as in Section 3.1.

3.4. Individual rationality (IR) constraints

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

$$\pi_a \geqslant -E(v).$$

For each of the mechanisms (SEQD, SIMUDBP, SIMUDNP), it follows from the formula for π_a that

$$\pi_a \geqslant -E(v) \ \Leftrightarrow\ \phi(\rho, a, b) \leqslant (1 - p_b)E(v^*, v).$$

For example, for the SEQD mechanism, this means

$$\phi_{\text{SEQD}}(\rho, a, b) \leqslant (1 - p_b)E(v^*_{\text{SEQD}}, v).$$

Similarly, the contractee's IR constraint is

$$\pi_b \geqslant E(w).$$

For each of the mechanisms (SEQD, SIMUDBP, SIMUDNP), it follows from the formula for π_b that

$$\pi_b \geqslant E(w) \ \Leftrightarrow\ (1 - p_a)E(w, w^*) \leqslant \phi(\rho, a, b).$$

For example, for the SEQD mechanism, this means

$$(1 - p_a)E(w, w^*_{\text{SEQD}}) \leqslant \phi_{\text{SEQD}}(\rho, a, b).$$

3.5. (Non)uniqueness of the decommitting equilibrium

Given the distributions of outside offers f and g, and a contract (ρ, a, b), in the sequential decommitting games the Nash equilibrium pair of decommitting thresholds (v^*, w^*) is unique. This is because the formulas from Section 3.1 give a closed form solution for v^* and then for w^*. Specifically, $v^*_{\text{SEQD}} = \rho - a$. It follows that $p_a = \int_{-\infty}^{v^*_{\text{SEQD}}} f(v)\,\mathrm{d}v$ is determined. From this it follows that $w^*_{\text{SEQD}} = \rho + (b + ap_a)/(1 - p_a)$ is determined.

On the other hand, as shown by Sandholm, Sikka and Norden [21], given the distributions of outside offers f and g, and a contract (ρ, a, b), the simultaneous decommitting games can sometimes have multiple Nash equilibrium pairs of decommitting thresholds (v^*, w^*). Whether or not such multiplicity occurs depends on the distributions of outside offers f and g and the contract parameters ρ, a, and b. The potential multiplicity of equilibria in the SIMUDBP game can be understood from the following group of four equations from Section 3.2 that defines the Nash equilibria for the SIMUDBP game:

$$v^*_{\text{SIMUDBP}} = \rho - \frac{a}{1 - p_b}, \qquad w^*_{\text{SIMUDBP}} = \rho + \frac{b}{1 - p_a},$$

$$p_a = \int_{-\infty}^{v^*_{\text{SIMUDBP}}} f(v)\,\mathrm{d}v, \qquad p_b = \int_{w^*_{\text{SIMUDBP}}}^{\infty} g(w)\,\mathrm{d}w.$$

This group can have multiple solutions $(v^*_{\text{SIMUDBP}}, w^*_{\text{SIMUDBP}})$.

Similarly, the potential multiplicity of equilibria in the SIMUDNP game can be understood from the following group of four equations from Section 3.3 that defines the Nash equilibria for the SIMUDNP game:

$$v^*_{\text{SIMUDNP}} = \rho - a - \frac{bp_b}{1-p_b}, \qquad w^*_{\text{SIMUDNP}} = \rho + b + \frac{ap_a}{1-p_a},$$

$$p_a = \int_{-\infty}^{v^*_{\text{SIMUDNP}}} f(v)\,\mathrm{d}v, \qquad p_b = \int_{w^*_{\text{SIMUDNP}}}^{\infty} g(w)\,\mathrm{d}w.$$

This group can have multiple solutions $(v^*_{\text{SIMUDNP}}, w^*_{\text{SIMUDNP}})$.

If a unique equilibrium is desired, the sequential mechanism can be used to guarantee uniqueness.

4. Surplus equivalence

Now, which of the leveled commitment contracting mechanisms would be best for the agents? In this section we prove the main result of the paper, i.e., 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 social welfare (in the game's social welfare maximizing equilibria). In other words, each of the games leads to the same surplus over what the agents would get without a contract (by taking their outside offers/fallbacks). This is surprising since the optimal contracts differ for the games. Also, for a given contract, the decommitting thresholds, decommitting probabilities, and expected social welfare generally differ across the games.

We start by proving that if a leveled commitment contract can generate positive surplus, 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*

$$H(x,y) = \int_{-\infty}^{y} g(w)\,\mathrm{d}w \int_{x}^{\infty} vf(v)\,\mathrm{d}v - \int_{x}^{\infty} f(v)\,\mathrm{d}v \int_{-\infty}^{y} wg(w)\,\mathrm{d}w. \tag{†}$$

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

$$a^* = \frac{\int_{-\infty}^{b^*} wg(w)\,\mathrm{d}w}{\int_{-\infty}^{b^*} g(w)\,\mathrm{d}w}, \qquad b^* = \frac{\int_{a^*}^{\infty} vf(v)\,\mathrm{d}v}{\int_{a^*}^{\infty} f(v)\,\mathrm{d}v}. \tag{‡}$$

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

$$p_x = \int_{-\infty}^{a^*} f(v)\,\mathrm{d}v, \qquad p_y = \int_{b^*}^{\infty} g(w)\,\mathrm{d}w.$$

The proof of Lemma 1 is technical, and is presented in Appendix A.

Now we are ready to prove the main result:

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 mechanisms* (*SEQD, SIMUDBP, and SIMUDNP*) *separately* (*and an expected social welfare maximizing equilibrium is played in the SIMUDBP and SIMUDNP games*), each mechanism yields the same social welfare, i.e., the same sum of payoffs to the agents.[8] *The mechanisms have at least one expected social welfare maximizing equilibrium pair (v^*, w^*) of decommitting thresholds* (*and the associated decommitting probabilities*) *in common. The optimal contracts can differ across the mechanisms, but each mechanism has an optimal contract where the decommitment penalties are nonnegative.*

Proof. As shown earlier in the paper, for each mechanism,

$$\pi = \pi^{\text{fallback}} + H(v^*, w^*),$$
$$\pi_a = -E(v) - \phi(\rho, a, b) + (1 - p_b)E(v^*, v),$$
$$\pi_b = E(w) + \phi(\rho, a, b) - (1 - p_a)E(w, w^*).$$

Therefore, the IR constraints reduce to

$$(1 - p_a)E(w, w^*) \leqslant \phi(\rho, a, b) \leqslant (1 - p_b)E(v^*, v).$$

If $\max_{x,y} H(x, y) \leqslant 0$, then $\pi = \pi^{\text{fallback}} + H(v^*, w^*) \leqslant \pi^{\text{fallback}}$, i.e., there exists no contract that is (strictly) IR for both agents. In other words, the agents will wait for the outside offers. Thus all three mechanisms have the same payoffs.

If $\max_{x,y} H(x, y) > 0$, then by Lemma 1, there exists a finite, global optimum (a^*, b^*) that satisfies (‡). For the three mechanisms, v^*, w^*, and $\phi(\rho, a, b)$ are determined differently based on $\rho, a, b, f,$ and g. If we can prove that for each game there exist ρ, a, and b such that the threshold values v^*, and w^* determined by them are identical to a^* and b^*, then the social welfare is $\pi^{\text{fallback}} + H(a^*, b^*)$, i.e., a maximal value of $\pi^{\text{fallback}} + H(x, y)$. We also have to guarantee that this configuration satisfies the IR constraints. These facts would mean that all three mechanisms lead to the same social welfare. We also show that at an optimum, $a \geqslant 0$ and $b \geqslant 0$.

For shorthand, let $\lambda(z) \equiv zb^* + (1 - z)a^*$ for $0 \leqslant z \leqslant 1$. Now, $a^* \leqslant \lambda(z) \leqslant b^*$, and $\lambda(z)$ increases monotonically. The IR constraints can be simplified to

$$(1 - p_x)(1 - p_y)a^* \leqslant \phi(\rho, a, b) \leqslant (1 - p_x)(1 - p_y)b^*. \quad (\S)$$

From the formula for $H(a^*, b^*)$ in Lemma 1 and the fact that $H(a^*, b^*) > 0$ we get $p_x < 1$ and $p_y < 1$.

Case 1. *SEQD*: Here $a^* = v^*_{\text{SEQD}} = \rho - a$, so $a = \rho - a^*$. Because $b^* = w^*_{\text{SEQD}} = \rho + (b + p_x a)/(1 - p_x)$, $b = p_x a^* + (1 - p_x)b^* - \rho = \lambda(1 - p_x) - \rho$. Substituting the expressions for a and b into the expression for $\phi_{\text{SEQD}}(\rho, a, b)$ gives

[8] Note that this equivalence holds not only on an expected value basis, but also *ex post* when the agents have found out their outside offers, have made their decommitting decisions, and have paid the associated decommitting penalties.

$$\begin{aligned}\phi_{\mathrm{SEQD}}(\rho,a,b) &= p_x(1-p_y)a - p_y b + (1-p_x)(1-p_y)\rho \\ &= p_x(1-p_y)[\rho - a^*] - p_y\big[p_x a^* + (1-p_x)b^* - \rho\big] \\ &\quad + (1-p_x)(1-p_y)\rho \\ &= \rho - p_x a^* - (1-p_x)p_y b^*.\end{aligned}$$

So the IR constraints (§) become

$$\begin{aligned}&(1-p_x)(1-p_y)a^* \leqslant \rho - p_x a^* - (1-p_x)p_y b^* \leqslant (1-p_x)(1-p_y)b^* \\ \Leftrightarrow\ &(1-p_x)p_y b^* + (1-p_y+p_x p_y)a^* \leqslant \rho \leqslant (1-p_x)b^* + p_x a^* \\ \Leftrightarrow\ &\lambda\big((1-p_x)p_y\big) \leqslant \rho \leqslant \lambda(1-p_x).\end{aligned}$$

Because λ is increasing and $p_y \leqslant 1$, there exists a ρ that satisfies the IR constraints. For such ρ values, $a = \rho - a^* \geqslant 0$ and $b = \lambda(1-p_x) - \rho \geqslant 0$.

Case 2. *SIMUDBP*: Here $a^* = v^*_{\mathrm{SIMUDBP}} = \rho - a/(1-p_y)$, so $a = (1-p_y)(\rho - a^*)$. Also, $b^* = w^*_{\mathrm{SIMUDBP}} = \rho + b/(1-p_x)$, so $b = (1-p_x)(b^* - \rho)$. Substituting the expressions for a and b into the expression for $\phi_{\mathrm{SIMUDBP}}(\rho,a,b)$ gives

$$\begin{aligned}\phi_{\mathrm{SIMUDBP}}(\rho,a,b) &= p_x a - p_y b + (1-p_x)(1-p_y)\rho \\ &= p_x\big[(1-p_y)(\rho - a^*)\big] - p_y\big[(1-p_x)(b^* - \rho)\big] \\ &\quad + (1-p_x)(1-p_y)\rho \\ &= (1-p_x p_y)\rho - p_x(1-p_y)a^* - p_y(1-p_x)b^*.\end{aligned}$$

So the IR constraints (§) become

$$\begin{aligned}&(1-p_x)(1-p_y)a^* \leqslant \phi_{\mathrm{SIMUDBP}}(\rho,a,b) \leqslant (1-p_x)(1-p_y)b^* \\ \Leftrightarrow\ &(1-p_x)(1-p_y)a^* \leqslant (1-p_x p_y)\rho - p_x(1-p_y)a^* - p_y(1-p_x)b^* \\ &\qquad \leqslant (1-p_x)(1-p_y)b^* \\ \Leftrightarrow\ &\frac{p_y(1-p_x)}{1-p_x p_y}b^* + \frac{1-p_y}{1-p_x p_y}a^* \leqslant \rho \leqslant \frac{(1-p_x)}{1-p_x p_y}b^* + \frac{p_x(1-p_y)}{1-p_x p_y}a^* \\ \Leftrightarrow\ &\lambda\left(\frac{p_y(1-p_x)}{1-p_x p_y}\right) \leqslant \rho \leqslant \lambda\left(\frac{(1-p_x)}{1-p_x p_y}\right).\end{aligned}$$

Because λ is increasing and

$$0 \leqslant \frac{p_y(1-p_x)}{1-p_x p_y} \leqslant \frac{(1-p_x)}{1-p_x p_y} \leqslant 1,$$

there exists a ρ that satisfies the IR constraints. For such ρ values, $a = (1-p_y)(\rho - a^*) \geqslant 0$, and $b = (1-p_x)(b^* - \rho) \geqslant 0$ because $a^* \leqslant \lambda(z) \leqslant b^*$.

Case 3. *SIMUDNP*: Recall that

$$a^* = v^*_{\mathrm{SIMUDNP}} = \rho - a - \frac{p_y b}{1-p_y}, \quad b^* = w^*_{\mathrm{SIMUDNP}} = \rho + b + \frac{p_x a}{1-p_x}.$$

Since $p_x < 1$ and $p_y < 1$, these formulas can be converted into linear equations:

$$(1-p_y)a + p_y b = (1-p_y)(\rho - a^*), \quad p_x a + (1-p_x)b = (1-p_x)(b^* - \rho).$$

There are two subcases based on the value of $p_x + p_y$.

In the subcase where $p_x + p_y = 1$, the linear equation group has no solution or infinitely many solutions depending on ρ. For a solution to exist, ρ must satisfy

$$(1 - p_y)(\rho - a^*) = (1 - p_x)(b^* - \rho), \quad \text{i.e.,} \quad \rho = p_x a^* + p_y b^* = \lambda(p_y)$$

and a and b must satisfy

$$p_x a + p_y b = p_x p_y (b^* - a^*).$$

If so, we can compute b as $b = p_x(b^* - a^*) - (p_x/p_y)a$. Substituting the formulas for ρ and b into $\phi_{\mathrm{SIMUDNP}}(\rho, a, b)$ gives

$$\begin{aligned}\phi_{\mathrm{SIMUDNP}}(\rho, a, b) &= p_x(1 - p_y)a - (1 - p_x)p_y b + (1 - p_x)(1 - p_y)\rho \\ &= p_x a + p_x p_y a^*.\end{aligned}$$

The IR constraints (§) become

$$\begin{aligned}&(1 - p_x)(1 - p_y)a^* \leqslant p_x a + p_x p_y a^* \leqslant (1 - p_x)(1 - p_y)b^* \\ \Leftrightarrow\ & 0 \leqslant a \leqslant p_y(b^* - a^*).\end{aligned}$$

Given this restriction on a, and the above relationship between a and b, we get $0 \leqslant b \leqslant p_x(b^* - a^*)$. So, a solution of the desired type exists for this subcase.

In the subcase where $p_x + p_y \neq 1$, we can solve a and b directly as a function of ρ:

$$\begin{aligned}a &= \frac{1 - p_x}{1 - p_x - p_y}\big[\rho - (1 - p_y)a^* - p_y b^*\big] = \frac{1 - p_x}{1 - p_x - p_y}\big[\rho - \lambda(p_y)\big], \\ b &= \frac{1 - p_y}{1 - p_x - p_y}\big[p_x a^* + (1 - p_x)b^* - \rho\big] = \frac{1 - p_y}{1 - p_x - p_y}\big[\lambda(1 - p_x) - \rho\big].\end{aligned}$$

Substituting these into $\phi_{\mathrm{SIMUDNP}}(\rho, a, b)$ gives

$$\begin{aligned}\phi_{\mathrm{SIMUDNP}}(\rho, a, b) &= p_x(1 - p_y)a - (1 - p_x)p_y b + (1 - p_x)(1 - p_y)\rho \\ &= \frac{(1 - p_x)(1 - p_y)(\rho - p_x a^* - p_y b^*)}{1 - p_x - p_y}.\end{aligned}$$

Then, the IR constraints (§) become

$$a^* \leqslant \frac{\rho - a^* p_x - b^* p_y}{1 - p_x - p_y} \leqslant b^*.$$

In the subsubcase where $p_x + p_y < 1$, this is equivalent to $p_y b^* + (1 - p_y)a^* \leqslant \rho \leqslant (1 - p_x)b^* + p_x a^*$, i.e., $\lambda(p_y) \leqslant \rho \leqslant \lambda(1 - p_x)$, so a solution of the desired type exists. Furthermore,

$$a = \frac{1 - p_x}{1 - p_x - p_y}\big[\rho - \lambda(p_y)\big] \geqslant 0, \quad b = \frac{1 - p_y}{1 - p_x - p_y}\big[\lambda(1 - p_x) - \rho\big] \geqslant 0.$$

In the subsubcase where $p_x + p_y > 1$, the IR constraints become $(1 - p_x)b^* + p_x a^* \leqslant \rho \leqslant p_y b^* + (1 - p_y)a^*$, i.e., $\lambda(1 - p_x) \leqslant \rho \leqslant \lambda(p_y)$, so again a solution of the desired type exists. Furthermore,

$$a = \frac{(1 - p_x)(\lambda(p_y) - \rho)}{p_x + p_y - 1} \geqslant 0, \quad b = \frac{(1 - p_y)(\rho - \lambda(1 - p_x))}{p_x + p_y - 1} \geqslant 0. \qquad \square$$

4.1. Existence of optimal IR contracts

It follows from the proof of Theorem 1 that if some leveled commitment contract generates positive surplus to the agents in the aggregate, then there exists an optimal leveled commitment contract that generates positive surplus 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 individually rational (IR) for both agents.*

Proof. Immediate from the proof of Theorem 1. □

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

5. Integrative vs. distributive bargaining

Proposition 1 showed that among optimal contracts there exist ones that are beneficial for both parties. However, the question of how to divide the surplus between the agents remains, that is, how to choose among the expected social welfare maximizing, individually rational contracts. Each agent's expected surplus is her expected payoff under the contract minus the expected fallback payoff: $e_a = \pi_a - \pi_a^{\text{fallback}} = \pi_a + E(v)$ and $e_b = \pi_b - \pi_b^{\text{fallback}} = \pi_b - E(w)$. 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 surplus between the agents). It could be that some splits cannot be supported by an optimal contract. However, we show that this problem cannot occur. 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 $e_b = (1-\beta) H(a^*, b^*)$.*

Proof. Follows from the proof of Theorem 1. The split of surplus, fixed by the value of β, is controlled by choosing ρ in the contract. The decommitting penalties, a and b, are then chosen based on ρ using the formulas in the proof of Theorem 1 to maximize expected social welfare. □

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.

The fact that distributive bargaining does not hinder integrative bargaining in leveled commitment contracts makes the leveled commitment mechanism a *modular* component technology for automated negotiation. The leveled commitment feature can be used to increase expected social welfare when there is uncertainty about the future. The fact that in leveled commitment contracts, distributive bargaining does not hinder integrative bargaining means that the leveled commitment contract technique can be directly used in conjunction with any standard distributive bargaining method. These methods include noncooperative bargaining mechanisms (such as alternating offers bargaining with time discounting [6,16], alternating offers bargaining with a constant cost per bargaining round [16], and unrestricted bargaining under deadlines [22]), as well as axiomatic bargaining solutions [7,9,11,12]. See [8] for an overview of the bargaining literature.

5.1. 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 surplus 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 surpluses [9,12,15]. It turns out that in leveled commitment contracts, such product maximization is equivalent to choosing an expected social welfare maximizing contract that splits surplus 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.*

Proof. By Proposition 2 we know that there exists a contract (ρ, a, b) that satisfies $e_a = e_b$ and maximizes expected social welfare. Next we calculate an upper bound on the product of surpluses:

$$
\begin{aligned}
e_a e_b &\leqslant \left(\frac{e_a + e_b}{2}\right)^2 = \left(\frac{\pi - \pi^{\text{fallback}}}{2}\right)^2 \\
&= \left(\frac{H(v^*(\rho, a, b), w^*(\rho, a, b))}{2}\right)^2 \leqslant \left(\frac{\max_{x,y} H(x, y)}{2}\right)^2.
\end{aligned}
$$

We proceed to show that this upper bound is reached—implying that the product is maximized—when surplus is equally split and expected social welfare maximized. The first inequality holds with equality iff $e_a = e_b$, i.e., surplus is equally split. The second inequality holds with equality at the optimum, and there only. □

We now give the closed form formulas for determining such a contract that maximizes expected social welfare, maximizes $e_a e_b$, and splits the surplus equally. Now,

$$\begin{aligned} & e_a = e_b \\ \Leftrightarrow\ & -\phi(\rho,a,b) + (1-p_x)(1-p_y)b^* = \phi(\rho,a,b) - (1-p_x)(1-p_y)a^* \\ \Leftrightarrow\ & \phi(\rho,a,b) = (1-p_x)(1-p_y)(a^*+b^*)/2. \end{aligned}$$

We use the formulas from the proof of Theorem 1.

Case 1. *SEQD*:

$$\phi_{\text{SEQD}}(\rho,a,b) = \rho - p_x a^* - (1-p_x)p_y b^*.$$

Combining this with the formula for $\phi(\rho,a,b)$, we get

$$\begin{aligned} \rho &= p_x a^* + (1-p_x)p_y b^* + (1-p_x)(1-p_y)(b^*+a^*)/2 \\ &= \frac{1+p_x-p_y+p_x p_y}{2}a^* + \frac{(1-p_x)(1+p_y)}{2}b^* = \lambda\left(\frac{(1-p_x)(1+p_y)}{2}\right). \end{aligned}$$

The optimal penalties are then determined by ρ:

$$\begin{aligned} a &= \rho - a^* = \lambda\left(\frac{(1-p_x)(1+p_y)}{2}\right) - \lambda(0) = \frac{(1-p_x)(1+p_y)}{2}(b^*-a^*), \\ b &= \lambda(1-p_x) - \rho = \left(1-p_x-\frac{(1-p_x)(1+p_y)}{2}\right)(b^*-a^*) \\ &= \frac{(1-p_x)(1-p_y)}{2}(b^*-a^*). \end{aligned}$$

Case 2. *SIMUDBP*:

$$\phi_{\text{SIMUDBP}}(\rho,a,b) = (1-p_x p_y)\rho - p_x(1-p_y)a^* - (1-p_x)p_y b^*.$$

Combining this with the formula for $\phi(\rho,a,b)$, we get

$$\begin{aligned} & (1-p_x p_y)\rho = \frac{(1+p_x)(1-p_y)}{2}a^* + \frac{(1-p_x)(1+p_y)}{2}b^* \\ \Leftrightarrow\ & \rho = \lambda\left(\frac{(1-p_x)(1+p_y)}{2(1-p_x p_y)}\right). \end{aligned}$$

The optimal penalties are then determined by ρ:

$$\begin{aligned} a &= (1-p_y)(\rho-a^*) = (1-p_y)\frac{(1-p_x)(1+p_y)}{2(1-p_x p_y)}(b^*-a^*) \\ &= \frac{(1-p_x)(1-p_y^2)}{2(1-p_x p_y)}(b^*-a^*), \\ b &= (1-p_x)(b^*-\rho) = (1-p_x)\frac{(1+p_x)(1-p_y)}{2(1-p_x p_y)}(b^*-a^*) \\ &= \frac{(1-p_x^2)(1-p_y)}{2(1-p_x p_y)}(b^*-a^*). \end{aligned}$$

Case 3. *SIMUDNP*: There are two subcases based on $p_x + p_y$.

If $p_x + p_y = 1$, then $\rho = \lambda(p_y)$. By using the formula for $\phi(\rho, a, b)$ together with $\phi_{\mathrm{SIMUDNP}}(\rho, a, b) = p_x a + p_x p_y a^*$, we get

$$a = \frac{p_y}{2}(b^* - a^*), \qquad b = \frac{p_x}{2}(b^* - a^*).$$

If $p_x + p_y \neq 1$, then

$$\begin{aligned}\phi_{\mathrm{SIMUDNP}}(\rho, a, b) &= \frac{(1-p_x)(1-p_y)(\rho - p_x a^* - p_y b^*)}{1 - p_x - p_y}\\ &= (1-p_x)(1-p_y)\frac{b^* + a^*}{2}.\end{aligned}$$

We solve for ρ from the equality above to get

$$\begin{aligned}\rho &= p_x a^* + p_y b^* + \frac{1 - p_x - p_y}{2}(b^* + a^*)\\ &= \frac{1 + p_x - p_y}{2}a^* + \frac{1 - p_x + p_y}{2}b^* = \lambda\left(\frac{1 - p_x + p_y}{2}\right).\end{aligned}$$

The optimal penalties are then determined by ρ:

$$\begin{aligned}a &= \frac{1-p_x}{1-p_x-p_y}\bigl[\rho - \lambda(p_y)\bigr] = \frac{1-p_x}{1-p_x-p_y}\frac{1-p_x-p_y}{2}(b^* - a^*)\\ &= \frac{1-p_x}{2}(b^* - a^*),\\ b &= \frac{1-p_y}{1-p_x-p_y}\bigl[\lambda(1-p_x) - \rho\bigr] = \frac{1-p_y}{1-p_x-p_y}\frac{1-p_x-p_y}{2}(b^* - a^*)\\ &= \frac{1-p_y}{2}(b^* - a^*).\end{aligned}$$

The two subcases can be combined independent of $p_x + p_y$:

$$\rho = \lambda\left(\frac{1 - p_x + p_y}{2}\right), \qquad a = \frac{1-p_x}{2}(b^* - a^*), \qquad b = \frac{1-p_y}{2}(b^* - a^*).$$

Finally, it is worth noting that if the surplus is divided equally, and the conditions of Theorem 1 are met, then each of the decommitting mechanisms gives each agent the same payoff.

6. Nonuniqueness of the social welfare maximizing equilibrium and the optimal contract

In Section 3.5 we discussed the fact that *for a given contract*, the equilibrium is unique in the sequential mechanisms, but there can be multiple equilibria in the simultaneous mechanisms.

In this section we study uniqueness issues when we allow the *contract* itself to vary. The main findings are:

- The social welfare maximizing equilibrium is not always unique. This holds for each of the decommitting mechanisms (SEQD, SIMUDBP, and SIMUDNP).
- For any one of the decommitting mechanisms in turn, any given equilibrium is supported by a continuum of contracts—specifically, those contracts that are on a line in the 3-dimensional space (ρ, a, b).

We will now discuss these issues in more detail.

The surplus, $H(x, y)$, can have multiple global maxima (even truly distinct ones that are not in the same neighborhood, as the following example shows). In particular, the social welfare maximizing decommitment threshold 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 global maxima. Let the distributions of the outside offers be

$$f(v) = \begin{cases} 1/10 & \text{if } 0 \leqslant v \leqslant 10, \\ 0 & \text{otherwise,} \end{cases}$$

$$g(w) = \begin{cases} 117/3520 & \text{if } 0 \leqslant w < 320/47, \\ 8/33 & \text{if } 320/47 \leqslant w \leqslant 10, \\ 0 & \text{otherwise.} \end{cases}$$

We use Lemma 1 to find all local maxima. Because $f(v) > 0$ and $g(w) > 0$ for all $v, w \in [0, 10]$, each local maximum (x, y) must satisfy the pair of equations (‡). Substituting our particular f into the second equation of the pair gives $y = (10 + x)/2$. After substituting our particular g into the first equation of the pair, we treat the cases $y < 320/47$ and $y \geqslant 320/47$ separately. The former case gives a line, and the latter case gives a quadratic curve. For each case, the pair of equations is solved to find (x, y). The former case leads to one solution, namely $(10/3, 20/3)$. The latter case leads to two solutions: $(4, 7)$, and $(5, 15/2)$. The corresponding surplus values are $H(10/3, 20/3) = H(5, 15/2) = 65/132$, and $H(4, 7) = 27/55$. Since $27/55 < 65/132$, both $(10/3, 20/3)$ and $(5, 15/2)$ are global maxima, while $(4, 7)$ is only a local maximum.

Now, it is easy to work backwards from any maximum (x, y), be it local or global, and determine the contracts that support it. For example, consider the sequential decommitting (SEQD) mechanism. The probability that the contractor decommits is now a constant $p_a = \int_{-\infty}^{x} f(v)\,\mathrm{d}v$. Then, the contracts that support the specific equilibrium (x, y) are again (as in Section 3.1) defined by the equations

$$x = \rho - a,$$
$$y = \rho + \frac{b + a p_a}{1 - p_a},$$

which are now linear equations. So, these two linear equations in three variables (ρ, a, b) define the SEQD contracts that support the equilibrium (x, y). Since there are only two equations and three variables, there is a continuum of contracts that support the equilibrium. These contracts are on a line in the 3-dimensional space (ρ, a, b).

It is easy to work backwards from any maximum (x, y), be it local or global, and determine the contracts that support it even for the simultaneous decommitting mechanisms. The decommitting probabilities, $p_a = \int_{-\infty}^{x} f(v)\,\mathrm{d}v$ and $p_b = \int_{y}^{\infty} g(w) w$, are now constants since x and y are constants. For the SIMUDBP mechanism, the contracts that support the equilibrium (x, y) are defined (as in Section 3.2) by the two equations

$$x = \rho - \frac{a}{1 - p_b},$$
$$y = \rho + \frac{b}{1 - p_a},$$

which are now linear equations. Similarly, for the SIMUDNP mechanism, the contracts that support the equilibrium (x, y) are defined (as in Section 3.3) by the two equations

$$x = \rho - a - \frac{bp_b}{1 - p_b},$$
$$y = \rho + b + \frac{ap_a}{1 - p_a},$$

which are now linear equations.

So, the two linear equations in three variables (ρ, a, b) define the contracts that support the equilibrium (x, y). Since there are only two equations and three variables, there is a continuum of contracts that support the equilibrium. These contracts are on a line in the 3-dimensional space (ρ, a, b).

To summarize, in this section we studied uniqueness questions when the contract itself is allowed to vary. There can be multiple welfare maximizing equilibria (as well as other equilibria which are only local maxima of the surplus function). For each of the mechanisms in turn, the contracts that support any given equilibrium form a line in the 3-dimensional space (ρ, a, b). Again, one can interpret the contract price, ρ, as the parameter (which is to be used for distributive bargaining), which determines the values of the decommitting penalties a and b that support the desired equilibrium.[9]

Nonuniqueness of the optimal threshold pair—and the associated nonuniqueness of the optimal contract line in the space (ρ, a, b)—does not prevent the use of leveled commitment contracts. To maximize expected social welfare, any one of the optimal contract lines can be used. Finally, choosing a specific point from that line corresponds to distributive bargaining, and we showed earlier in the paper that this does not hinder the integrative bargaining.

7. 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, $u_i : \pi_i \to \Re$, would be linearly increasing. Risk attitudes are captured in the usual way by making u_i nonlinear. We now show that the surplus equivalence of leveled commitment contracts does not generally hold for agents that are not risk neutral, and in different

[9] The individual rationality constraints can be incorporated directly into this view. Under the interpretation that ρ is the parameter, the contractor's IR constraint gives an upper bound on ρ and the contractee's IR constraint gives a lower bound on ρ.

settings, different leveled commitment mechanisms are best in terms of expected social welfare (that is, sum of expected utilities).[10]

Consider the following distributions of outside offers:

$$f(v) = \begin{cases} 1/100 & \text{if } v \in [0, 100], \\ 0 & \text{otherwise,} \end{cases}$$

$$g(w) = \begin{cases} 1/110 & \text{if } w \in [0, 110], \\ 0 & \text{otherwise.} \end{cases}$$

Now, for risk seeking agents with utility functions of the form $u_{\text{contractor}}(x) = u_{\text{contractee}}(x) = x^3$, the maximal expected social welfares are: $\max_{SEQD} \pi \approx 284192$, $\max_{SIMUDBP} \pi \approx 322522$, $\max_{SIMUDNP} \pi \approx 334194$. So, the simultaneous mechanism where neither pays if both decommit is the best, the simultaneous mechanism where both pay if both decommit is the second, and the sequential mechanism is the worst.

On the other hand, for risk averse agents with utility functions of the form $u_{\text{contractor}}(x) = u_{\text{contractee}}(x) = x^{1/3}$, the maximal expected social welfares are: $\max_{SEQD} \pi \approx 0.914$, $\max_{SIMUDBP} \pi \approx 0.925$, $\max_{SIMUDNP} \pi \approx 0.905$. So, the simultaneous mechanism where both pay if both decommit is better than the one where neither pays if both decommit, and the sequential mechanism is between the two simultaneous mechanisms.

8. Conclusions

In automated negotiation systems consisting of self-interested agents, contracts have traditionally been binding. Such contracts do not allow the contract parties to capitalize on uncertain future events. Contingency contracts can be used to avoid this problem, but as we discussed, they suffer from several practical problems—especially when applied to automated negotiation.

Recently, leveled commitment contracts were introduced as a more practical alternative for addressing this problem [19]. In a leveled commitment contract, the agents can backtrack out of the contract unilaterally by paying a predetermined, agent-specific penalty. In such mechanisms, 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. It was recently shown that despite such strategic decommitting (in Nash equilibrium of the decommitting game), 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 [20].

In this paper we studied six different leveled commitment contracting mechanisms that 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. In

[10] There exist particular distributions of outside offers (f and g) and particular utility functions that go along with these distributions, such that the surplus *happens* to be the same for the different leveled commitment mechanisms.

general, given the distributions of outside offers and the contract parameters (contract price, the contractor's breach penalty, and the contractee's breach penalty), these mechanisms lead to different decommitting thresholds and probabilities for the contract parties, and different levels of expected social welfare.

Leveled commitment contracts are often more practical than contingency contracts. However, they cannot always achieve the same social welfare because the agents decommit strategically: some contracts are inefficiently kept. Our intuitions suggested that sequential decommitting mechanisms would lead to higher social welfare than simultaneous ones since the last agent decommits truthfully. We also thought that mechanisms where neither agent pays a penalty if both decommit would promote decommitting and increase welfare.

However, we showed that, somewhat surprisingly, *all six leveled commitment contract mechanisms lead to the same social welfare when the contract price and decommitting penalties are optimized for each mechanism separately*.

We showed that in leveled commitment contracts, integrative bargaining does not hinder distributive bargaining—as long as the interaction between the two is handled appropriately. Specifically, the surplus from leveled commitment can be divided arbitrarily (as long as each agent benefits), e.g. equally, without compromising optimality, as long as the decommitting penalties are tailored to the contract price so as to maximize expected social welfare. One practical way to do this is to conduct distributive bargaining over the contract price, and to then optimize the decommitting penalties to that price. The fact that in leveled commitment contracts, integrative bargaining does not hinder distributive bargaining, means that leveled commitment can be used as a modular technology in automated negotiation to increase expected social welfare among the contract parties, while using any standard noncooperative distributive bargaining mechanism as the technology for dividing the surplus.

In cooperative bargaining theory, the product maximizing solution has a central role [9, 12,15]. We showed that in each of the leveled commitment contracting mechanisms this solution exists, and it is exactly the solution that divides the expected surplus equally.

We also studied uniqueness questions. Given distributions of outside offers and a contract, the equilibrium is unique in the sequential mechanisms, but not always in the simultaneous mechanisms. When we allow the contract to vary, the social welfare maximizing equilibrium is not always unique. This holds for each of the decommitting mechanisms—even the sequential ones. For any one of the decommitting mechanisms in turn, any given equilibrium is supported by a continuum of contracts—specifically, those contracts that are on a line in the 3-dimensional space spanned by the contract price, the contractor's breach penalty, and the contractee's breach penalty.

Finally, we showed that the surplus equivalence ceases to hold if agents are not risk neutral. The ranking of the mechanisms (in terms of expected social welfare) depends on the utility functions of the contract parties.

Our derivations allow agents to construct optimal leveled commitment contracts, and to divide the gains arbitrarily (as long as each agent benefits), for example equally. Using this theory we developed fast algorithms for contract optimization [21], and provide a free contract optimization service on the web as part of *eMediator*, our next generation electronic commerce server prototype (http://www.cs.cmu.edu/~amem/eMediator/) [18].

Several important avenues for future research remain. Sequences of multiple leveled commitment contracts among multiple parties lead to interesting behavior where one decommit can trigger another, etc. That is beyond the scope of the model in this paper, but we have experimentally studied such cascade effects and different mechanisms and parameterizations of mechanisms among strategic agents [2] and myopic agents [1]. Future research involves analytical work in extending the theory of this paper to the setting of multiple sequential contracts. Even the analysis of a single contract where the penalties and/or distributions of outside offers change over time promises to be interesting.

Appendix A. Proof of Lemma 1

Proof. Because f and g have finite expectations, we can extend the domain of $H(x, y)$ to $[-\infty, \infty] \times [-\infty, \infty]$, i.e., we treat infinities as numbers. Choose an arbitrary global maximum (x_0, y_0). Because $\max_{x,y} H(x, y) > 0$, $H(x_0, y_0) > 0$. Using this and (†), we get

$$\int_{-\infty}^{y_0} g(w)\,\mathrm{d}w > 0, \qquad \int_{x_0}^{\infty} f(v)\,\mathrm{d}v > 0,$$

i.e., there is nonzero probability for each agent to keep the contract. Let

$$p(y) \equiv \frac{\int_{-\infty}^{y} w g(w)\,\mathrm{d}w}{\int_{-\infty}^{y} g(w)\,\mathrm{d}w}, \qquad q(x) \equiv \frac{\int_{x}^{\infty} v f(v)\,\mathrm{d}v}{\int_{x}^{\infty} f(v)\,\mathrm{d}v}.$$

The partial derivatives of H with respect to x and y are

$$H_x(x, y_0) = f(x) \cdot \big(p(y_0) - x\big) \cdot \int_{-\infty}^{y_0} g(w)\,\mathrm{d}w,$$

$$H_y(x_0, y) = g(y) \cdot \big(q(x_0) - y\big) \cdot \int_{x_0}^{\infty} f(v)\,\mathrm{d}v.$$

If $p(y_0) = x_0$ and $q(x_0) = y_0$, then (x_0, y_0) satisfies (‡), i.e., it is the desired point (a^*, b^*), so we are done. Otherwise, we continue as follows. Fix y_0. Then $H(x, y_0)$ gets its maximal value at $x = p(y_0)$ because $H_x(x, y_0) \geqslant 0$ when $x \leqslant p(y_0)$, and $H_x(x, y_0) \leqslant 0$ when $x \geqslant p(y_0)$. Let $x_1 = p(y_0)$. So, $H(x_1, y_0) \geqslant H(x_0, y_0)$. Thus $H(x_1, y_0)$ is also a global maximum.

If $q(x_1) = y_0$, then (x_1, y_0) satisfies (‡), so we are done. Otherwise we continue as follows. Fix x_1. Then $H(x_1, y)$ gets its maximal value at $y = q(x_1)$ because $H_y(x_1, y) \geqslant 0$ when $y \leqslant q(x_1)$, and $H_y(x_1, y) \leqslant 0$ when $y \geqslant q(x_1)$. Let $y_1 = q(x_1)$. So, $H(x_1, y_1) \geqslant H(x_1, y_0)$. Thus $H(x_1, y_1)$ is also a global maximum.

Keep alternating the above two paragraphs by always replacing x_{n-1} by x_n, and y_{n-1} by y_n until (x_n, y_n) satisfies (‡). If this does not occur in a finite number of steps, it gives a sequence $\{(x_n, y_n)\}$ of global maxima for function H. We now prove that this sequence converges to a point (x^*, y^*) that satisfies (‡). The functions $p(y)$ and $q(x)$ are increasing because

$$p'(y) = \frac{\mathrm{d}p(y)}{\mathrm{d}y} = \frac{g(y)\int_{-\infty}^{y}(y-w)g(w)\,\mathrm{d}w}{\left(\int_{-\infty}^{y} g(w)\,\mathrm{d}w\right)^2} \geqslant 0$$

and

$$q'(x) = \frac{\mathrm{d}q(x)}{\mathrm{d}x} = \frac{f(x)\int_{x}^{\infty}(v-x)f(v)\,\mathrm{d}v}{\left(\int_{x}^{\infty} f(v)\,\mathrm{d}v\right)^2} \geqslant 0.$$

Therefore, $p \circ q$ is also increasing. Note that $x_{n+1} = p \circ q \circ x_n$. If $x_1 > x_0$, then $x_0 \leqslant x_1 \leqslant x_2 \leqslant \cdots \leqslant x_n \leqslant \cdots$ and $\{x_n\}$ is increasing. If $x_1 < x_0$, $\{x_n\}$ is decreasing. Similarly, $y_{n+1} = q \circ p \circ y_n$ and $\{y_n\}$ is monotonic. Because $\{x_n\}$ and $\{y_n\}$ are monotonic, we can define $x^* = \lim_n x_n$ and $y^* = \lim_n y_n$. Because the set of global maxima is closed, (x^*, y^*) is a global maximum. Finally,

$$x^* = \lim x_{n+1} = \lim p \circ q \circ x_n = \lim p \circ y_n = p(y^*),$$

and

$$y^* = \lim y_{n+1} = \lim q \circ p \circ y_n = \lim q \circ x_{n+1} = q(x^*),$$

so (x^*, y^*) satisfies (‡).

What remains to be shown is the exact form of $H(a^*, b^*)$. If (a^*, b^*) is a global maximum, then $H(a^*, b^*) > 0$, so from (†) we get $\int_{-\infty}^{b^*} g(w)\,\mathrm{d}w > 0$, and $\int_{a^*}^{\infty} f(v)\,\mathrm{d}v > 0$. Using this and the fact that f and g have finite expectations, we know that a^* and b^* are finite:

$$-\infty < a^* = p(b^*) < \infty \quad \text{and} \quad -\infty < b^* = q(a^*) < \infty,$$

so

$$\begin{aligned} H(a^*, b^*) &= \int_{-\infty}^{b^*} g(w)\,\mathrm{d}w \int_{a^*}^{\infty} vf(v)\,\mathrm{d}v - \int_{a^*}^{\infty} f(v)\,\mathrm{d}v \int_{-\infty}^{b^*} wg(w)\,\mathrm{d}w \\ &= \int_{-\infty}^{b^*} g(w)\,\mathrm{d}w \cdot b^* \int_{a^*}^{\infty} f(v)\,\mathrm{d}v - \int_{a^*}^{\infty} f(v)\,\mathrm{d}v \cdot a^* \int_{-\infty}^{b^*} g(w)\,\mathrm{d}v \\ &= (b^* - a^*)(1 - p_x)(1 - p_y). \qquad \square \end{aligned}$$

References

[1] M. Andersson, T. Sandholm, Leveled commitment contracting among myopic individually rational agents, in: Proc. Third International Conference on Multi-Agent Systems (ICMAS) Paris, France, 1998, pp. 26–33.

[2] M. Andersson, T. Sandholm, Leveled commitment contracts with myopic and strategic agents, J. Economic Dynamics and Control (Special Issue on Agent-Based Computational Economics) 25 (2001) 615–640. Early version: National Conference on Artificial Intelligence (AAAI), Madison, WI, 1998, pp. 38–45.

[3] J.D. Calamari, J.M. Perillo, The Law of Contracts, 2nd Edition, West Publishing, 1977.
[4] P.A. Diamond, E. Maskin, An equilibrium analysis of search and breach of contract, I: Steady states, Bell J. Economics 10 (1979) 282–316.
[5] S. Kraus, Agents contracting tasks in non-collaborative environments, in: Proc. AAAI-93, Washington, DC, 1993, pp. 243–248.
[6] S. Kraus, J. Wilkenfeld, G. Zlotkin, Multiagent negotiation under time constraints, Artificial Intelligence 75 (1995) 297–345.
[7] A. Mas-Colell, M. Whinston, J.R. Green, Microeconomic Theory, Oxford University Press, Oxford, 1995.
[8] A. Muthoo, Bargaining Theory with Applications, Oxford University Press, Oxford, 1999.
[9] J. Nash, The bargaining problem, Econometrica 18 (1950) 155–162.
[10] J. Nash, Equilibrium points in n-person games, Proc. Nat. Acad. Sci. 36 (1950) 48–49.
[11] J. Nash, Two-person cooperative games, Econometrica 21 (1950) 128–140.
[12] M.J. Osborne, A. Rubinstein, Bargaining and Markets, Academic Press, New York, 1990.
[13] R.A. Posner, Economic Analysis of Law, 2nd Edition, Little, Brown and Company, New York, 1977.
[14] H. Raiffa, The Art and Science of Negotiation, Harvard Univ. Press, Cambridge, MA, 1982.
[15] J.S. Rosenschein, G. Zlotkin, Rules of Encounter: Designing Conventions for Automated Negotiation Among Computers, MIT Press, Cambridge, MA, 1994.
[16] A. Rubinstein, Perfect equilibrium in a bargaining model, Econometrica 50 (1982) 97–109.
[17] T. Sandholm, An implementation of the contract net protocol based on marginal cost calculations, in: Proc. AAAI-93, Washington, DC, 1993, pp. 256–262.
[18] T. Sandholm, eMediator: A next generation electronic commerce server, Computational Intelligence (Special issue on Agent Technology for Electronic Commerce), to appear. Early versions appeared in Proc. Fourth International Conference on Autonomous Agents (AGENTS), Barcelona, Spain, 2000, pp. 73–96, and in AAAI-99 Workshop on AI in Electronic Commerce, Orlando, FL, July 1999, pp. 46–55, and Washington University, St. Louis, Dept. of Computer Science Technical Report WU-CS-99-02, January 1999.
[19] T. Sandholm, V.R. Lesser, Issues in automated negotiation and electronic commerce: Extending the contract net framework, in: Proc. First International Conference on Multi-Agent Systems (ICMAS), San Francisco, CA, 1995, pp. 328–335. Reprinted in: Huhns, Singh (Eds.), Readings in Agents, 1997, pp. 66–73.
[20] T. Sandholm, V.R. Lesser, Leveled commitment contracts and strategic breach, Games and Economic Behavior (Special issue on AI and Economics) 35 (2001) 212–270. Early versions appeared as Advantages of a Leveled Commitment Contracting Protocol, in: Proc. National Conference on Artificial Intelligence (AAAI), Portland, OR, 1996, pp. 126–133, and as UMass Amherst Computer Science Department Technical Report 95-72.
[21] T. Sandholm, S. Sikka, S. Norden, Algorithms for optimizing leveled commitment contracts, in: Proc. IJCAI-99, Stockholm, Sweden, 1999, pp. 535–540. Extended version: Washington University, Department of Computer Science Technical Report WUCS-99-04.
[22] T. Sandholm, N. Vulkan, Bargaining with deadlines, in: Proc. AAAI-99, Orlando, FL, 1999, pp. 44–51.
[23] S. Sen, Tradeoffs in contract-based distributed scheduling, PhD Thesis, Univ. of Michigan, 1993.
[24] R.G. Smith, The contract net protocol: High-level communication and control in a distributed problem solver, IEEE Trans. Comput. C-29 (12) (1980) 1104–1113.