Games and Economic Behavior **35**, 212–270 (2001)
doi:10.1006/game.2000.0831, available online at http://www.idealibrary.com on IDEAL®

Leveled Commitment Contracts and Strategic Breach

Tuomas W. Sandholm[1]

Computer Science Department, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, Pennsylvania 15213
E-mail: sandholm@cs.cmu.edu

and

Victor R. Lesser

Department of Computer Science, University of Massachusetts, Amherst, Massachusetts 01003
E-mail: lesser@cs.umass.edu

Received July 14, 1999

In (automated) negotiation systems for self-interested agents, contracts have traditionally been binding. They do not accommodate future events. Contingency contracts address this, but are often impractical. As an alternative, we propose *leveled commitment contracts*. The level of commitment is set by breach penalties. To be freed from the contract, an agent simply pays the penalty to the other party. A self-interested agent will be reluctant to breach because the other party might breach, in which case the former agent is freed from the contract, does not incur a penalty, and collects a penalty from the breacher. We show that, despite such strategic breach, leveled commitment increases the expected payoff to both contract parties and can enable deals that are impossible under full commitment. Asymmetric beliefs are also discussed. Different decommitting mechanisms are introduced and compared. Practical prescriptions for market designers are provided. A contract optimizer is provided on the web. *Journal of Economic Literature* Classification Numbers: C72, C78, D82, D83, K12, L14. © 2001 Academic Press

Key Words: contract; breach; search; market design; automated negotiation; bargaining; economics of uncertainty; multiagent systems; artificial intelligence.

1. INTRODUCTION

The importance of automated negotiation systems with self-interested agents is increasing. One reason for this is the technological push of a

A short early version of this paper appeared at the National Conference on Artificial Intelligence (AAAI) (Sandholm and Lesser, 1996).

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

0899-8256/01 $35.00

growing standardized communication infrastructure (the Internet, WWW, EDI, HTML, KQML, FIPA, XML, Java, Odyssey, Voyager, Concordia, Aglets, etc., over which separately designed agents belonging to different organizations can interact in an open environment in real time and safely carry out transactions (Sandholm, 1997; Low *et al*., 1996). The second reason is the strong application pull for computer support for contracting, especially at the operative decision-making level. For example, we are witnessing the advent of small transaction business-to-consumer and consumer-to-consumer commerce on the Internet for purchasing goods, services, information, communication bandwidth, etc. (Choi *et al*., 1997). Hundreds of electronic business-to-business trading sites have also emerged, some of which already incorporate automated negotiation capability. There is also an industrial trend toward virtual enterprises; dynamic alliances of small, agile enterprises which together can take advantage of economies of scale when available (e.g., by being able to respond to larger and more diverse orders than they could individually), but do not suffer from diseconomies of scale.

Multiagent technology facilitates the automated formation of such dynamic alliances on a per order basis by automated contracting. Such automation can save labor time of human negotiators, and in addition other savings are possible because computational agents are often more effective at finding beneficial contracts and contract combinations than humans are in strategically and combinatorially complex settings.

In traditional multiagent negotiation mechanisms among self-interested agents, once a contract is made it is binding, i.e., neither party can back out no matter how future events unravel (Rosenschein and Zlotkin, 1994; Sandholm, 1993; Andersson and Sandholm, 1999; Ephrati and Rosenschein, 1991; Kraus, 1993; Kraus *et al*., 1995; Cheng and Wellman, 1998). Although a contract may be profitable to an agent when viewed *ex ante*, it need not be profitable when viewed *ex post*. Similarly, a contract that is unprofitable *ex ante* may become profitable *ex post*. Full commitment contracts are unable to capitalize on the gains that such probabilistically known future events provide.

On the other hand, many multiagent systems consisting of cooperative agents incorporate some form of decommitment possibility in order to allow the agents to accommodate new events. For example, in the original contract net protocol (Smith, 1980) the agent that had contracted out a task could send a termination message to cancel the contract even when the contractee had already partially fulfilled the contract. This was possible because the agents were not self-interested: the contractee did not mind losing part of its effort without a monetary compensation. Similarly, the role of decommitment possibilities among cooperative agents has been studied in meeting scheduling using a contracting approach (Sen and Durfee, 1994; Sen and Durfee, 1998) and in cooperative coordination protocols (Decker and Lesser, 1995).

Unlike systems with cooperative agents, multiagent systems consisting of self-interested agents require that we consider the case where agents do not follow externally specified strategies, but choose their own strategies. Thus the interaction mechanisms need to be considered from the perspective of noncooperative game theory: given a mechanism (i.e., rules of the game), what is the best strategy from a self-interested viewpoint that each agent can choose, and then what social outcomes will follow?

1.1. *Contingency Contracts*

Some research in noncooperative game theory has focused on utilizing the potential provided by probabilistically known future events by *contingency contracts* among self-interested agents (Raiffa, 1982). The obligations of the contract are made contingent on future events. There are games in which this method provides an expected payoff increase to both parties of the contract compared to any full commitment contract. Also, some deals are enabled by contingency contracts in the sense that there is no full commitment contract that both agents prefer over their fallback positions, but there is a contingency contract that each agent prefers over its fallback.

There are at least three major problems in using contingency contracts among self-interested agents, especially in automated negotiation.

First, the real-world party that an agent represents often does not know all possible future events and cannot therefore use contingency contracts optimally. Furthermore, even if the real-world party does know them, programming that knowledge into the automated agent may be prohibitively laborious and/or error prone.

Second, although contingency contracts can be useful in anticipating a small number of key events, they become cumbersome as the number of relevant future events to monitor increases. In the extreme, all domain events (changes in the domain problem, e.g., new tasks arriving or resources breaking down) and all negotiation events (contracts from other negotiations) can affect the value of the obligations of the original contract to the agent and therefore need to be conditioned on. Furthermore, these future events may not only affect the value of the original contract independently: the value of the original contract may depend on combinations of the future events (Sandholm, 1993; Rosenschein and Zlotkin, 1994; Sandholm and Lesser, 1995). Thus there is a potential combinatorial explosion of possible future worlds, and each of them may need to be associated with a different contingency. This leads to a potential combinatorial explosion of the contract (e.g., the size of the contingency table that represents the contract).

In addition to the two practical difficulties associated with contingency contracts, there is a third, fundamental game-theoretic problem. Sometimes an event is only observable by some of the agents. Those agents may have an incentive to lie to the other contract parties about the event in case the event is associated with a disadvantageous contingency to the observing agents. Thus, to be viable, contingency contracts would require an event verification mechanism that is not manipulable and not prohibitively complicated or costly.

1.2. *Leveled Commitment Contracts*

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

Principles for assessing decommitment penalties have been studied in the economics of law (Calamari and Perillo, 1977; Posner, 1977), 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 abide by 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] Somewhat unintuitively, it turns out that the decommitting possibility in a contract can increase the expected payoff for all contract parties.

[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 (Diamond and Maskin, 1979).

1.3. *Practical Motivations for Leveled Commitment*

The goal of our leveled commitment contracting mechanism is to allow some flexibility as in the case with no commitment while guaranteeing agents some level of security as in the case of full commitment. Full commitment contracts can be viewed as one end of a spectrum where commitment-free contracts are at the other end. Leveled commitment contracts span this entire spectrum based on how the decommitting penalties are chosen.[3] There are several practical reasons why leveled commitment is desirable.

- It allows agents to profitably accommodate new domain events such as new tasks arriving or resources breaking down by allowing an agent to back out of its old contracts that these new events have made unbeneficial or even infeasible.

- It allows agents to profitably accommodate new negotiation events such as new offers or offer-acceptance messages. If these events make some old contracts unbeneficial or infeasible to an agent, that agent can decommit from those old contracts.

- It provides a backtracking instrument for distributed search (in the artificial intelligence sense) that works among self-interested agents, unlike traditional backtracking techniques for distributed search (see, e.g., Yokoo *et al*., 1992). It allows more controlled profitable risk taking. In terms of search this means moving a low-commitment search focus around in the global search space of commitments (because decommitting is not unreasonably expensive), so that more of that space can be explored among self-interested agents which would otherwise avoid risky commitments. For example, in task allocation among agents, an agent can accept a task set and later try to contract out the tasks in that set separately. With full commitment, to avoid risk, an agent needs to have standing offers from the agents it will contract the tasks to, or it has to be able to handle them profitably itself. With the leveled commitment mechanism, the agent can accept the task set even if it is not sure about its chances of getting the tasks handled, because if it does not get them handled it can decommit.

- It allows profitable construction of composite contracts from basic contracts. Often the value of a contract to an agent depends on which other contracts the agent will get. Using leveled commitment, an agent can take on unbeneficial contracts in anticipation of later synergic contracts

[3] The decommitting penalty can also increase with time, decrease as a function of acceptance time of the offer, or be conditioned on events in other negotiations or the environment. It can also be dynamically negotiated over on a per contract or per task set basis (Sandholm, 1996; Sandholm and Lesser, 1995). The analysis of this paper focuses on one contract where the penalties are negotiated up front and they do not change over time.

that will make the sequence beneficial overall. If the later contracts in the sequence do not occur, the agent can backtrack out of the initial parts of the sequence.

- It saves computation and time. In many automated negotiation applications, computing the value of taking on a contract can be intractable and therefore needs to be approximated in practice (Sandholm, 1993; Sandholm, 1996; Sandholm and Lesser, 1995). Leveled commitment allows an agent to bid based on a rough value calculation. If the agent wins the bid, the agent can invest a more thorough value calculation. If the contract no longer looks beneficial in light of this more refined calculation, the agent can decommit. The fact that only the winning bidders carry out a refined calculation can save computation systemwide. Also, the negotiations can be carried out faster because agents can bid based on less computation.

- It makes local feasibility checks unnecessary. When bidding with full commitment, an agent has to make sure that it can handle all of its obligations even if all of its pending bids (and the bid that it might be constructing) are accepted. Such feasibility checks often use a major portion of a contracting software agent's deliberation resources (Sandholm, 1993; Sandholm, 1996). With leveled commitment, agents need not carry out feasibility checks up front because if an agent ends up overcommitted, it can decommit from some of the contracts so as to reobtain feasibility. Avoiding feasibility checks saves computation, and the negotiations can be carried out faster because agents can bid based on less computation.

- It speeds up the negotiation process by increasing parallelism. An agent can make mutually exclusive low-commitment offers to multiple agents. In the case more than one accepts, the agent can backtrack from all but one. This allows the agent to address the other parties in parallel instead of addressing them one at a time and blocking to wait for an answer before addressing the next. For example, if an agent wants one particular contract, it can offer that contract to several parties with meaningful commitment instead of no commitment at all (which would be strategically meaningless).

- It increases Pareto efficiency by reallocating risk. By choosing the contract price and decommitting penalties appropriately, Pareto efficiency can be improved by making the less risk-averse agents carry more of the risk. The more risk-averse agents would be willing to compensate by allowing the former agents a higher expected payoff.

Later in this paper, we substantiate the advantages of leveled commitment contracts more formally. However, before that we discuss some reasons why it is not obvious that leveled commitment contracts are advantageous.

1.4. *Why Are the Advantages of Leveled Commitment Not Obvious*?

Despite their intuitive appeal and the practical motivations for leveled commitment contracts, there are several reasons why it is not obvious that leveled commitment contracts are superior to full commitment contracts.

First, when an agent decommits, his/her profit from decommitting may be smaller than the loss to the victim of the breach; both are computed after the decommitting penalties have been paid. Therefore, decommitting sometimes decreases the sum of the contract parties' payoffs when viewed *ex post*.

Second, one might think that full commitment contracts can never have a higher sum of *expected* payoffs to the contract parties than leveled commitment contracts because the latter incorporate new information (new events) and, according to decision theory, the expected value of information is always nonnegative. However, this result from single-agent decision theory does not hold in games where more than one party can gain new information. In multiagent systems, information can have negative expected value. The prisoner's dilemma provides a simple example. Say that there are two players in separate rooms, and each one can press one of two buttons, cooperate or defect. Based on what buttons the agents press, they receive payoffs according to Table I. Each agent's dominant strategy is to defect, so the sum of the agents' payoffs will be $1 + 1 = 2$. This is also the only outcome which is not Pareto efficient. Now, let us remove some of the information, namely the labels of the buttons. The agents will have to press at random, so the expected sum of payoffs is $\frac{1}{4}(3 + 3) + \frac{1}{4}(0 + 5) + \frac{1}{4}(5 + 0) + \frac{1}{4}(1 + 1) = 4.5$. So, the expected value of the information is $2 - 4.5 = -2.5 < 0$. Therefore, it is not obvious that leveled commitment contracts, which incorporate more information, have higher (or even equal) sums of expected payoffs to the contract parties

TABLE I
Prisoner's Dilemma Game

	Column player	
Row player	Cooperate	Defect
Cooperate	3, 3	0, 5
Defect	5, 0	1, 1

Note. In each square, the row player's payoff is listed first.

than full commitment contracts. Similarly, it is not obvious that leveled commitment contracts have higher (or even equal) Pareto efficiency than full commitment contracts.

Third, agents might decommit insincerely. A nonstrategic agent would decommit whenever its best outside offer plus the decommitting penalty is better than the current contract. However, a rational self-interested agent would be more reluctant to decommit because there is a chance that the other party will decommit, in which case the former agent is freed from the contract obligations, does not have to pay a decommitting penalty, and collects a decommitting penalty from the other party. Similarly, the other contract party will be reluctant to decommit for the same reason. Due to such reluctant decommitting, a contract may end up being inefficiently kept even though each party would be better off by decommitting and paying the penalty (and therefore, the sum of the contract parties' payoffs would be higher if either agent alone, or both agents, would decommit).

In the rest of this paper, we formally show that leveled commitment contracts are superior to full commitment contracts despite such strategic breaching. In Section 2 we present our main models and an analysis of leveled commitment contracts. We show that leveled commitment contracts enable contracts and improve the expected payoffs of both contract parties. We show this for games where both agents' futures involve uncertain events, and we do the analysis both for sequential (Section 2.1) and simultaneous (Section 2.2) decommitting mechanisms. Section 3 compares the equilibria of the different decommitting mechanisms. Section 4 discusses settings where only one agent's future involves uncertain events. In Section 5 we discuss games where the fall-back positions change for the worse before the decommitting game. Section 6 presents what happens when an agent thinks that the distributions of future events are common knowledge and acts according to the equilibrium derived under that assumption, but in reality at least one of the agents has biased beliefs. In Section 7, practical prescriptions are given for market designers. Section 8 concludes and presents future research directions.

2. COMPARING LEVELED AND FULL COMMITMENT CONTRACTS

We analyze a contracting situation from the perspective of two risk-neutral agents, each of which attempts to maximize his own expected payoff. We call one of the agents the *contractor* and the other agent the *contractee*. We define a *full commitment contract* as follows:

DEFINITION 2.1. A *full commitment contract* between a contractor and a contractee is defined by the contract obligations (Γ, ρ_F), which include two parts;

- a description, Γ, of what each of the two agents has to perform (handling tasks, contributing goods, lending resources, etc.), and
- a contract price, $\rho_F \in \Re$, that the contractor has to pay to the contractee.

In a full commitment contract, the contractor's payoff is $-\rho_F$ (because the contractor has to pay the contract price) and the contractee's payoff is ρ_F (because the contractee is paid the contract price).

We focus on a setting where the value of the contract changes for each agent independently, so at contract time each agent only has probabilistic information about the value of the contract to that agent. The change in value can stem from changes in the agent's own characteristics, such as resources failing or becoming available. The change can also be due to computation: based on a rough computation a contract's value may have looked different than it looks after a more refined computation. Also, the value of a contract may change based on changes in outside options, such as offers from third parties. Our framework and results are not specific to any particular source of change. However, we present our results in the setting where the change stems from outside options.

Specifically, the agents might receive outside offers. For simplicity, we assume in the model that all offers have the same description Γ, so price is the only concern. We also assume that each of the two agents only wants to be involved in one contract (e.g., the contractor gets its task handled and does not need to have it handled more than once, and the contractee has limited resources and can only take on one contract).

The contractor's best (lowest) outside offer $\breve{a}$ is only probabilistically known *ex ante* by both agents and is characterized by a probability density function $f(\breve{a})$. If the contractor does not receive an outside offer, $\breve{a}$ corresponds to its best outstanding outside offer or its fallback payoff, i.e., a payoff that it receives if no contract is made. The contractee's best (highest) outside offer $\breve{b}$ is also only probabilistically known *ex ante* and is characterized by a probability density function $g(\breve{b})$. If the contractee does not receive an outside offer, $\breve{b}$ corresponds to its best outstanding outside offer or its fallback payoff. The variables $\breve{a}$ and $\breve{b}$ are assumed to be statistically independent. The distributions f and g are assumed to be common knowledge between the contractor and the contractee.

The contractor and the contractee can either make a contract or wait for their outside offers. We call the latter option the *null deal*.

DEFINITION 2.2. In the *null deal*, the agents do not make a contract. The contractor waits for its best (lowest) outside offer, $\breve{a}$, with an expected payoff of $E[-\breve{a}] = \int_{-\infty}^{\infty} f(\breve{a})[-\breve{a}]\,d\breve{a}$. The contractee waits for its best (highest) outside offer, $\breve{b}$, with an expected payoff of $E[\breve{b}] = \int_{-\infty}^{\infty} g(\breve{b})\breve{b}\,d\breve{b}$.

If the contractor and the contractee do make a contract, they can choose to use some full commitment contract or a contract instrument which we introduce here and call a *leveled commitment contract*.

DEFINITION 2.3. A *leveled commitment contract* between a contractor and a contractee is a 4-tuple (Γ, ρ, a, b), where

- Γ is a description of what each of the two agents has to perform (handling tasks, contributing goods, lending resources, etc.).
- $\rho \in \Re$ is a contract price that the contractor has to pay to the contractee.
- $a \geq 0$ is the contractor's decommitting penalty. If the contractor pays this penalty to the contractee, the contract obligations (Γ, ρ) are canceled: neither party is bound by the contract description Γ, and the contractor does not have to pay the contract price ρ to the contractee.
- $b \geq 0$ is the contractee's decommitting penalty. If the contractee pays this penalty to the contractor the contract obligations (Γ, ρ) are canceled: neither party is bound by the contract description Γ, and the contractor does not have to pay the contract price ρ to the contractee.

We say that the contractor has to decide on decommitting when it knows its outside offer, $\breve{a}$, but does not know the contractee's outside offer, $\breve{b}$. Similarly, the contractee has to decide on decommitting when it knows its outside offer, $\breve{b}$, but does not know the contractor's outside offer, $\breve{a}$. This seems realistic from a practical (automated) contracting perspective. (See Table II for a summary of the variables.)

An easy way to think about the outside offers is to consider them to be full commitment contracts. Alternatively, $\breve{a}$ and $\breve{b}$ can be interpreted to be expected payments of outside offers which themselves are leveled commitment contracts.

Our contracting setting consists of two stages. In the first stage, which we call the *contracting game*, the agents choose a contract (or the null deal) before any future events have unraveled. In the second stage, which we call the *decommitting game*, the agents decide on whether to decommit or not, after the future events have unraveled. Clearly, the equilibrium of the decommitting game affects the agents' preferences over contracts in the contracting game. The decommitting game will be analyzed using the Nash equilibrium or the iterated dominance solution concept (Mas-Colell *et al.*,

TABLE II
Symbols Used in This Paper

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

1995). The contracting game will be analyzed with respect to *individual rationality*.

DEFINITION 2.4. A contract is *individually rational* (IR) for an agent if the agent's expected payoff under the contract is higher (or equal) than that agent's expected payoff under the null deal. A contract is individually rational (IR) if it is individually rational for both agents.

Often there is either no contract that is IR for both agents or there are many such contracts. When there is no time discounting and many IR contracts to choose from, there are uncountably many Nash equilibria in the contracting game. Even with just full commitment contracts, if the contractor's strategy is to offer a contract for price ρ and no more (and that contract is IR for both agents), the contractee's best response is to take the offer as opposed to the null deal. Now the contractor's best response to this is to offer ρ and no more. Thus, a Nash equilibrium exists for any ρ that defines a contract that is IR for both agents. Bargaining theory addresses the choice among IR deals, usually by incorporating some form of time discounting (Rubinstein, 1982), deadlines (Sandholm and Vulkan, 1999), or bargaining costs into the bargaining process or by asserting desirable properties that the chosen contract should fulfill compared to other contracts (Nash, 1950; Osborne and Rubinstein, 1990). In this paper, we will not address the choice among IR contracts. In another paper we present algorithms for computing the IR contracts (contract price and decommitting penalties) that maximize the sum of the contract parties' payoffs (Sandholm *et al.*, 1999b).

We restrict our attention to contracts where $a \geq 0$ and $b \geq 0$, i.e., agents do not get paid for decommitting. This is intuitively appealing. Furthermore, in a recent paper we show that this restriction is not crucial because if there exists an IR contract that maximizes the sum of the two agents' expected payoffs, then there exists one with $a \geq 0$ and $b \geq 0$ (Sandholm and Zhou, 2000).

We do *not* assume that the agents decommit nonstrategically, by which we mean the following.

DEFINITION 2.5. A *nonstrategic* contractee decommits if $\breve{b} - b > \rho$, i.e., $\breve{b} > \rho + b$. A nonstrategic contractor decommits if $-\breve{a} - a > -\rho$, i.e., $\breve{a} < \rho - a$.

Instead of decommitting nonstrategically, an agent may not decommit although its outside offer is better (even after paying the decommitment penalty) for itself than the contract because the agent believes that there is a high probability that the other party will decommit. This would save the agent its decommitment penalty and make the agent receive a decommitment penalty from the other party. Games of this type differ significantly based on whether the agents decommit sequentially or simultaneously. These cases are analyzed in Sections 2.1 and 2.2 respectively. The key distinction is not whether decommitting happens at one point in real time, but whether or not an agent knows the other agent's decommitting decision by the time it has to reveal its own decommitting decision. If neither agent knows the other's decision at that stage, the decommitting is in essence simultaneous.

2.1. *Sequential Decommitting*

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

Let us now analyze the decommitting game using iterated dominance as the solution concept. Specifically, we start reasoning about the agents' actions at the leaves of the game tree and proceed backward to the beginning of the game. In the subgame where the contractee has decommitted, the contractor's best move is not to decommit because $-\breve{a} - a + b \leq -\breve{a} + b$ (because $a \geq 0$).[4] This also holds for a contract where neither agent has to pay a decommitment penalty if both decommit, because

[4] For ease of presentation, throughout the paper we make the *ad hoc* assumption that in the case of equality an agent does not bother to decommit. This assumption is not essential: The results do not hinge on it.

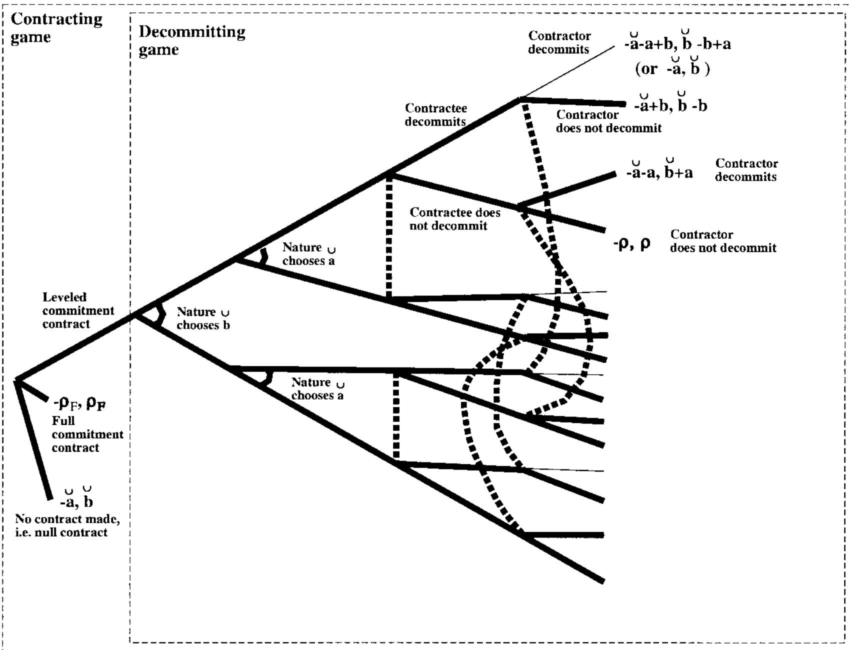

FIG. 1. The "sequential decommitting" (SEQD) game. The game tree represents two alternative mechanisms, i.e., two different games. In the first, both agents have to pay the decommitment penalties to each other if both decommit. In the second, neither agent has to pay if both decommit. The payoffs of the latter mechanism are in parentheses when they differ from the former. Dotted lines show information sets. They model the aspect that the contractor does not know the contractee's outside offer and vice versa. The contractor's payoffs are usually negative because it has to pay for having the task handled.

$-\breve{a} \leq -\breve{a} + b$ (Fig. 1, parenthesized payoffs). In the subgame where the contractee has not decommitted, the contractor's best move is to decommit if $-\breve{a} - a > -\rho$. This happens with probability

$$p_a = p_a(\rho, a) = \int_{-\infty}^{\rho - a} f(\breve{a})\, d\breve{a}. \tag{1}$$

Put together, the contractee gets $\breve{b} - b$ if it decommits, $\breve{b} + a$ if it does not but the contractor does, and ρ if neither decommits. Thus the contractee decommits if

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

If $p_a = 1$, i.e., $\int_{\rho - a}^{\infty} f(\breve{a})\, d\breve{a} = 0$, the above inequality is equivalent to $-b > a$, which is false because a and b are nonnegative. In other words, if the contractee surely decommits, the contractor does not. Clearly, entering into such a contract cannot be strictly IR for the contractee because its decommitting penalty is nonnegative. On the other hand, the above inequality is equivalent to

$$\breve{b} > \rho + \frac{b + p_a a}{1 - p_a} \stackrel{\text{def}}{=} \breve{b}^*(\rho, a, b) \qquad \text{when } p_a < 1. \tag{2}$$

Now the contractee's IR constraint states that the expected payoff from the contract is no less than the expected payoff from the outside offer:

$$\begin{aligned}\pi_b = \pi_b(\rho, a, b) = &\int_{\breve{b}^*(\rho, a, b)}^{\infty} g(\breve{b})[\breve{b} - b]\, d\breve{b} \\ &+ \int_{-\infty}^{\breve{b}^*(\rho, a, b)} g(\breve{b})\left[p_a \cdot (\breve{b} + a) + (1 - p_a) \cdot \rho\right] d\breve{b} \\ \geq E[\breve{b}] = &\int_{-\infty}^{\infty} g(\breve{b})\breve{b}\, d\breve{b}.\end{aligned} \tag{3}$$

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

$$\begin{aligned}\pi_a &= \pi_a(\rho, a, b) \\ &= \int_{\breve{b}^*(\rho, a, b)}^{\infty} g(\breve{b}) \int_{-\infty}^{\infty} f(\breve{a})[-\breve{a} + b]\, d\breve{a}\, d\breve{b} \\ &\quad + \int_{-\infty}^{\breve{b}^*(\rho, a, b)} g(\breve{b}) \left[\int_{-\infty}^{\rho - a} f(\breve{a})[-\breve{a} - a]\, d\breve{a} + \int_{\rho - a}^{\infty} f(\breve{a})[-\rho]\, d\breve{a}\right] d\breve{b} \\ &\geq E[-\breve{a}] = \int_{-\infty}^{\infty} f(\breve{a})[-\breve{a}]\, d\breve{a}.\end{aligned} \tag{4}$$

Now, do SEQD games exist where some full commitment contract is possible but no leveled commitment contract is? Because the contractor can want to decommit only if $-\breve{a} - a > -\rho$, its decommitment penalty can be chosen so high that it will surely not decommit—assuming that $\breve{a}$ is bounded from below. In this case the contractee will decommit whenever $\rho < \breve{b} - b$. If $\breve{b}$ is bounded from above, the contractee's decommitment penalty can be chosen so high that it will surely not decommit. Thus, assuming that $\breve{a}$ is bounded from below and $\breve{b}$ from above, full commitment contracts are a subset of leveled commitment ones. This reasoning holds for contracts where both agents have to pay the penalties if both decommit, and for contracts where neither agent has to pay a penalty if both decommit. Because full commitment contracts are a subset of leveled commitment contracts, the former can be no better in the sense of Pareto efficiency or maximizing the sum of the contract parties' expected payoffs than the latter. It follows that if there exists an IR full commitment contract, then there also exists at least one IR leveled commitment contract.

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

THEOREM 2.1 (Enabling in SEQD Games). *There are SEQD games (defined by f and g) where no full commitment contract is IR for both agents but a leveled commitment contract is.*

Proof. Let

$$f(\breve{a}) = \begin{cases} \frac{1}{100} & \text{if } 0 \le \breve{a} \le 100 \\ 0 & \text{otherwise} \end{cases} \quad \text{and} \quad g(\breve{b}) = \begin{cases} \frac{1}{110} & \text{if } 0 \le \breve{b} \le 110 \\ 0 & \text{otherwise.} \end{cases}$$

Now a full commitment contract F does not satisfy both IR constraints since that would require $E[\breve{b}] \le \rho_F \le E[\breve{a}]$, which is impossible because $55 = E[\breve{b}] > E[\breve{a}] = 50$. Let us choose a leveled commitment contract where $\rho = 52.5$, $a = 30$, and $b = 20$. Now

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

$$= 52.5 + \frac{20 + 0.225 \cdot 30}{0.775} \approx 87.0.$$

The contractor's IR constraint becomes

$$\int_{\breve{b}^*(\rho,a,b)}^{\infty} g(\breve{b}) \int_{-\infty}^{\infty} f(\breve{a})[-\breve{a} + b]\, d\breve{a}\, d\breve{b}$$
$$+ \int_{-\infty}^{\breve{b}^*(\rho,a,b)} g(\breve{b}) \left[\int_{-\infty}^{\rho - a} f(\breve{a})[-\breve{a} - a]\, d\breve{a} + \int_{\rho - a}^{\infty} f(\breve{a})[-\rho]\, d\breve{a} \right] d\breve{b}$$
$$\ge \int_{-\infty}^{\infty} f(\breve{a})[-\breve{a}]$$
$$\Leftrightarrow \int_{\breve{b}^*(\rho,a,b)}^{110} g(\breve{b}) \int_0^{100} f(\breve{a})[-\breve{a} + 20]\, d\breve{a}\, d\breve{b}$$
$$+ \int_0^{\breve{b}^*(\rho,a,b)} g(\breve{b}) \left[\int_0^{52.5 - 30} f(\breve{a})[-\breve{a} - 30]\, d\breve{a} \right.$$
$$\left. + \int_{52.5-30}^{100} f(\breve{a})[-52.5]\, d\breve{a} \right] d\breve{b}$$
$$\ge \int_0^{100} f(\breve{a})[-\breve{a}]$$

$$\Leftrightarrow \int_{\breve{b}^*(\rho,a,b)}^{110} \frac{1}{110}\frac{1}{100}\left[\frac{-(100)^2}{2} + 100\cdot 20\right] d\breve{b}$$

$$+ \int_0^{\breve{b}^*(\rho,a,b)} \frac{1}{110}\left[\frac{1}{100}\left[\frac{-(22.5)^2}{2} - 22.5\cdot 30\right] + \frac{1}{100}\left[-52.5\cdot(100-22.5)\right]\right] d\breve{b}$$

$$\geq -50$$

$$\Leftrightarrow \int_{\breve{b}^*(\rho,a,b)}^{110} \frac{15.5}{110}\, d\breve{b} + \int_0^{\breve{b}^*(\rho,a,b)} - \frac{49.96875}{110}\, d\breve{b} \geq -50.$$

Substituting $\breve{b}^*(\rho, a, b) = 87.0$ gives approximately $-6.3 - 39.5 \geq -50$ for the above inequality. Thus the contractor's IR constraint is satisfied.

The contractee's IR constraint becomes

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

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

$$\geq \int_{-\infty}^{\infty} g(\breve{b})\,\breve{b}\, d\breve{b}$$

$$\Leftrightarrow \int_{\breve{b}^*(\rho,a,b)}^{110} \frac{1}{110}[\breve{b} - 20]\, d\breve{b}$$

$$+ \int_0^{\breve{b}^*(\rho,a,b)} \frac{1}{110}\left[\int_0^{52.5-30} \frac{1}{100}[\breve{b} + 30]\, d\breve{a} + \int_{52.5-30}^{100} \frac{1}{100} 52.5\, d\breve{a}\right] d\breve{b}$$

$$\geq 55$$

$$\Leftrightarrow \int_{\breve{b}^*(\rho,a,b)}^{110} \frac{1}{110}[\breve{b} - 20]\, d\breve{b}$$

$$+ \int_0^{\breve{b}^*(\rho,a,b)} \frac{1}{110}\left[\frac{22.5}{100}[\breve{b} + 30] + \frac{77.5}{100} 52.5\right] d\breve{b}$$

$$\geq 55$$

$$\Leftrightarrow \int_{\breve{b}^*(\rho,a,b)}^{110} \frac{1}{110}[\breve{b} - 20]\, d\breve{b} + \int_0^{\breve{b}^*(\rho,a,b)} 0.43125 + \frac{0.225}{110}\breve{b}\, d\breve{b} \geq 55.$$

Substituting $\breve{b}^*(\rho, a, b) = 87.0$ gives approximately $16.4 + 45.3 \geq 55$ for the above inequality. Thus the contractee's IR constraint is satisfied. ■

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

Case 1 (Some Chance that Either Agent is Going to Decommit). In the case where $a < \rho$ there is some chance that the contractor will decommit (it may happen that $-\breve{a} > -\rho + a$). Now

$$\breve{b}^*(\rho, a, b) = \rho + \frac{b + \int_{-\infty}^{\rho-a} f(\breve{a})\, d\breve{a}[a]}{\int_{\rho-a}^{\infty} f(\breve{a})\, d\breve{a}} = \rho + \frac{b + \frac{1}{100}[\rho - a]a}{\frac{1}{100}[100 - (\rho - a)]}.$$

If $\breve{b}^*(\rho, a, b) < 110$ (i.e., less than the maximum possible $\breve{b}$), there is some chance that the contractee will decommit. This occurs if $\breve{b} > \rho + b$. We programmed a model of the IR constraints (Eqs. (4) and (3)) for this case. To make the algebra tractable (constant $f(\breve{a})$ and $g(\breve{b})$), versions of these IR constraint equations were used that assumed $0 \leq a < \rho$ and $0 < \breve{b}^* < 110$, without loss of generality. The corresponding decommitment penalties a and b that satisfy the IR constraints are plotted in Fig. 2 (left). Furthermore, the boundaries of the programmed model need to be checked. The boundaries $a = 0$, $a = \rho$, and $\breve{b}^* = 110$ are plotted in Fig. 2 (left). The constraint $\breve{b}^* > 0$ is always satisfied in this case and is thus not

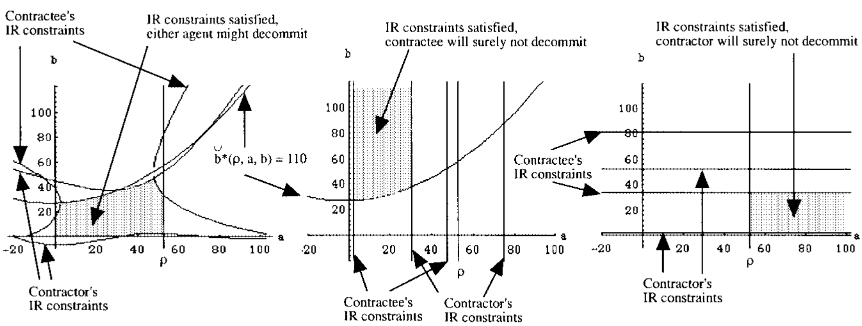

FIG. 2. The decommitment penalties a and b that satisfy both agents' IR constraints in the example SEQD game when $\rho = 52.5$. (Left) The case where either agent might decommit ($a < \rho$, and $\breve{b}^*(\rho, a, b) < 110$). (Middle) The case where the contractor might decommit but the contractee will not ($a < \rho$ and $\breve{b}^*(\rho, a, b) \geq 110$). (Right) The case where $a \geq \rho$, i.e., the contractor will surely not decommit but the contractee might.

plotted. To summarize, in the gray area of Fig. 2 (left) the contracts are IR for both agents, given that the agents decommit according to the iterated dominance equilibrium.

Case 2 (Contractor Will Surely Not Decommit). When $a \geq \rho$, the contractor will surely not decommit because its best possible outside offer is $\breve{a} = 0$. Note that a can be arbitrarily high. The corresponding $\breve{b}^*(\rho, a, b) = \rho + (b + \int_{-\infty}^{\rho-a} f(\breve{a})\, d\breve{a}[a]) / \int_{\rho-a}^{\infty} f(\breve{a})\, d\breve{a} = \rho + b$, i.e., the contractee decommits nonstrategically. Now the contractor's IR constraint (Eq. (4)) becomes

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

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

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

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

$$\Leftrightarrow 2.5 \leq b \leq 52.5$$

by the quadratic equation solution formula.

Similarly, the contractee's IR constraint (Eq. (3)) becomes

$$\int_{\rho+b}^{110} g(\breve{b}) \int_0^{100} f(\breve{a})[\breve{b} - b]\, d\breve{a}\, d\breve{b} + \int_0^{\rho+b} g(\breve{b}) \int_0^{100} f(\breve{a})[\rho]\, d\breve{a}\, d\breve{b} \geq E[\breve{b}]. \tag{6}$$

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

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

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

$$\Leftrightarrow b \leq \sim 34.05 \text{ or } b \geq \sim 80.95$$

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

Put together, the open region $2.5 \leq b \leq 34.05$, $a \geq \rho$, is where this type of contract is IR for both agents, given that the agents decommit according to the iterated dominance equilibrium. This region is colored gray in Fig. 2 (right).

Case 3 (Contractee Will Surely Not Decommit). If b is so high that $\breve{b}^*(\rho, a, b) \geq 110$, the contractee will surely not decommit. Now the contractor will decommit whenever $-\breve{a} - a > -\rho \Leftrightarrow \breve{a} < \rho - a$. In other words, the decommitting threshold $\breve{a}^* = \rho - a$. The contractor's IR constraint becomes

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

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

$$\Leftrightarrow \int_0^{\rho - a} f(\breve{a})[-\breve{a} - a]\, d\breve{a} + \int_{\rho - a}^{100} f(\breve{a})[-\rho]\, d\breve{a} \geq -50. \tag{7}$$

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

$$\frac{1}{100} \left[\int_0^{\rho - a} [-\breve{a} - a]\, d\breve{a} + \int_{\rho - a}^{100} [-\rho]\, d\breve{a} \right] \geq -50$$

$$\Leftrightarrow \frac{1}{100} \left[\left(\frac{-(\rho - a)^2}{2} + (\rho - a)(-a) \right) + \left((100 - (\rho - a)) \cdot (-\rho) \right) \right] \geq -50$$

$$\Leftrightarrow a \leq \sim 30.14 \text{ or } a \geq \sim 74.86$$

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

Similarly, the contractee's IR constraint becomes

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

$$\Leftrightarrow \int_0^{110} \frac{1}{110} \left[[\breve{b} + a] \int_0^{\rho - a} f(\breve{a})\, d\breve{a} + [\rho] \int_{\rho - a}^{100} f(\breve{a})\, d\breve{a} \right] d\breve{b} \geq 55$$

$$\Leftrightarrow \frac{1}{110}\left[\left[\frac{110^2}{2} + 110a\right]\int_0^{\rho - a} f(\breve{a})\, d\breve{a} + 110\rho\int_{\rho - a}^{100} f(\breve{a})\, d\breve{a}\right] \geq 55$$

$$\Leftrightarrow [55 + a]\int_0^{\rho - a} f(\breve{a})\, d\breve{a} + \rho\int_{\rho - a}^{100} f(\breve{a})\, d\breve{a} \geq 55. \tag{8}$$

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

$$[55 + a](\rho - a)\frac{1}{100} + \rho[100 - (\rho - a)]\frac{1}{100} \geq 55$$
$$\Leftrightarrow 2.5 \leq a \leq 47.5$$

by the quadratic equation solution formula. Thus the open region $2.5 \leq a \leq 30.14$, $\breve{b}^* \geq 110$, is where this type of contracts are IR for both agents, given that the agents decommit according to the iterated dominance equilibrium. This region is colored gray in Fig. 2 (middle).

In addition to the above three cases that can occur in this instance of the game, in general there is a fourth, trivial case that can occur. In that case, at least one of the contract parties will decommit for sure. For such a contract to be IR for the breacher, the breacher's decommitting penalty will have to be 0 (and in this case the IR constraint is satisfied with equality rather than strictly). For example, if the contract price is set higher than the contractor's highest possible outside offer, and the contractor's penalty is 0, the contractor will surely decommit. Similarly, if the contract price is lower than the contractee's lowest possible outside offer, and the contractee's decommitting penalty is 0, the contractee will surely decommit. In either case, the contract is trivial in the sense that it will always be breached and zero penalty will be paid. Such a contract does not lead to any benefit: the payoff of each contract party is the same as under the null deal. In the game instance of the proof, the contract price, 52.5, is between the contractee's lowest possible outside offer, 0, and the contractor's highest possible outside offer, 100, so this trivial case cannot occur.

In addition to enabling deals that are impossible using full commitment contracts, leveled commitment contracts can increase the efficiency of a deal even if a full commitment contract were possible. The reverse cannot occur because leveled commitment contracts can emulate full commitment contracts by setting the penalties high enough (assuming that $\breve{a}$ is bounded from below and $\breve{b}$ from above). Leveled commitment improves the efficiency of a deal if there is some chance that the contractor's outside offer is lower than the contractee's expected outside offer or some chance that the contractee's outside offer is higher than the contractor's expected outside offer.

THEOREM 2.2 (Pareto Efficiency Improvement in SEQD Games). *Let F be any full commitment contract in a SEQD game (defined by f and g). If*

1. *$\breve{b}$ is bounded from above, and* $\Pr(\breve{a} < E[\breve{b}]) > 0$, *or*
2. *$\breve{a}$ is bounded from below, and* $\Pr(\breve{b} > E[\breve{a}]) > 0$,

then the game has a leveled commitment contract that increases the contractor's expected payoff as well as the contractee's expected payoff over F. It follows that if F is IR, then so is this leveled commitment contract.

Proof. We prove this under Condition 1. The proof under Condition 2 is analogous. Under the full commitment contract, the contractor's payoff is $-\rho_F$, and the contractee's is ρ_F. We now construct a leveled commitment contract such that the expected payoff of each of the two agents increases.

Let the contract price be $\rho = \rho_F + \lambda$, and let the contractor's decommitting penalty be $a = \rho_F + \lambda - E[\breve{b}]$. Later in the proof we show how to set the parameter λ, $\lambda > 0$.

We set the contractee's decommitting penalty, b, so high that the contractee will surely not decommit. This can be done because $\breve{b}$ is bounded from above.

If the contractor does not decommit, the contractee's payoff is $\rho = \rho_F + \lambda$. If the contractor decommits, the contractee's expected payoff is $E[\breve{b}] + a = \rho_F + \lambda$. So, in either scenario, the contractee's expected payoff increases by λ.

Clearly, $E[\breve{a}|\breve{a} < E[\breve{b}]] < E[\breve{b}]$ because $\Pr(\breve{a} < E[\breve{b}]) > 0$. Let us call the difference $\Delta = E[\breve{b}] - E[\breve{a}|\breve{a} < E[\breve{b}]] > 0$.

If the contractor does not decommit, its payoff is $-\rho = -\rho_F - \lambda$ (i.e., less than under the full commitment contract). If the contractor decommits, its payoff is $-\breve{a} - a$. The contractor decommits whenever $-\breve{a} - a > -\rho \Leftrightarrow \breve{a} < \rho - a = E[\breve{b}]$. Therefore, the contractor's expected payoff in the scenario where it decommits is $-E[\breve{a}|\breve{a} < E[\breve{b}]] - a = -E[\breve{a}|\breve{a} < E[\breve{b}]] - \rho_F - \lambda + E[\breve{b}] = -\rho_F - \lambda + \Delta$ (i.e., greater than under the full commitment contract as long as $\lambda < \Delta$).

The breach by the contractor occurs with probability $p_a = \Pr(\breve{a} < E[\breve{b}]) > 0$. What remains to be done is to tailor λ so that the contractor's expected payoff increases compared to the full commitment contract. Formally, we want to set λ such that

$$
\begin{aligned}
&(1 - p_a)(-\rho_F - \lambda) + p_a \cdot (-\rho_F - \lambda + \Delta) > -\rho_F \\
&\quad \Leftrightarrow \lambda < p_a \Delta.
\end{aligned}
$$

Put all together, choosing $\lambda > 0$ increases the contractee's expected payoff and choosing $\lambda < p_a\Delta$ increases the contractor's expected payoff. Since $p_a > 0$ and $\Delta > 0$, we have $p_a\Delta > 0$. Therefore, λ can be chosen $(0 < \lambda < p_a\Delta)$ so that the expected payoff increases for each of the two parties. ■

It follows that under the conditions of the theorem, no full commitment contract is Pareto efficient.

2.2. *Simultaneous Decommitting*

In our simultaneous decommitting games, both agents have to declare their decommitting decisions simultaneously. Again, at decommitting time, the contractor knows its outside offer, $\breve{a}$, but not the contractee's outside offer, $\breve{b}$. Similarly, the contractee knows its outside offer, $\breve{b}$, but not the contractor's outside offer, $\breve{a}$. There are two alternative types of leveled commitment contracts that differ based on what happens if both agents decommit. In the first, both agents have to pay the decommitment penalties to each other if both decommit. In the second, neither agent has to pay if both decommit. Figure 3 presents the game trees corresponding to these contract types. Next, these two game types are discussed separately.

2.2.1. *Both Pay if Both Decommit* (*SIMUDBP*). This section discusses simultaneous decommitting games where a mechanism is used where both agents have to pay the decommitting penalties to each other if both decommit. Such settings will be called SIMUDBP games (Fig. 3). Let p_b be the probability that the contractee decommits. The value of this variable in equilibrium depends on f, g, ρ, a, and b. The contractor will decommit if

$$p_b[-\breve{a} + b - a] + (1 - p_b)[-\breve{a} - a] > p_b[-\breve{a} + b] + (1 - p_b)[-\rho].$$

If $p_b = 1$, this is equivalent to $a < 0$. But we already ruled out this type of contract where either one of the agents gets paid for decommitting. On the other hand, the above inequality is equivalent to

$$\breve{a} < \rho - \frac{a}{1 - p_b} \stackrel{\text{def}}{=} \breve{a}^*\big(\rho, a, b, \breve{b}^*\big) \qquad \text{when } p_b < 1. \tag{9}$$

Thus we have characterized a decommitting threshold $\breve{a}^*$ for the contractor. If the contractor's outside offer $\breve{a}$ is less than $\breve{a}^*$, the contractor is best off by decommitting. The probability that the contractor decommits is thus

$$p_a = \int_{-\infty}^{\breve{a}^*(\rho, a, b, \breve{b}^*)} f(\breve{a})\, d\breve{a}. \tag{10}$$

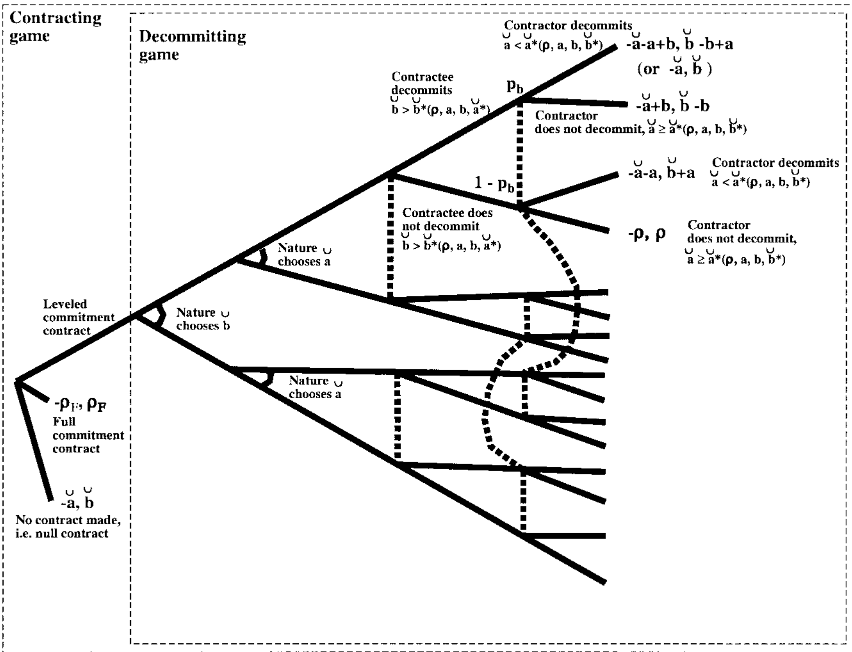

FIG. 3. The "simultaneous decommit—both pay if both decommit" (SIMUDBP) game. The parenthesized payoffs represent the "simultaneous decommit—neither pays if both decommit" (SIMUDNP) game. The dashed lines represent the agents' information sets. They model the aspect that when decommitting the contractor does not know the contractee's outside offer, and vice versa. They also model the aspect that the contractor has to reveal its decommitting decision before it has observed the contractee's decommitting decision, and vice versa.

The contractee decommits if

$$p_a[\breve{b} - b + a] + (1 - p_a)[\breve{b} - b] > p_a[\breve{b} + a] + (1 - p_a)[\rho].$$

If $p_a = 1$, this is equivalent to $b < 0$. But we already ruled out this type of contract where either one of the agents gets paid for decommitting. On the other hand, the above inequality is equivalent to

$$\breve{b} > \rho + \frac{b}{1 - p_a} \stackrel{\text{def}}{=} \breve{b}^*(\rho, a, b, \breve{a}^*) \qquad \text{when } p_a < 1. \tag{11}$$

Now we have characterized a decommitting threshold $\breve{b}^*$ for the contractee. If the contractee's outside offer $\breve{b} > \breve{b}^*$, the contractee is best off by decommitting. The probability that the contractee will decommit is thus

$$p_b = \int_{\breve{b}^*(\rho, a, b, \breve{a}^*)}^{\infty} g(\breve{b})\, d\breve{b}. \tag{12}$$

Inequality (9) defines the contractor's best response (characterized by $\breve{a}^*$) to the contractee's strategy that is characterized by $\breve{b}^*$. Condition (11) defines the contractee's best response (characterized by $\breve{b}^*$) to the contractor's strategy that is characterized by $\breve{a}^*$. Condition (9) uses the variable p_b which is defined by Eq. (12). Condition (11) uses the variable p_a which is defined by Eq. (10). So, together, Eqs. (9), (10), (11), and (12) define the Nash equilibria of the decommitting game.

Now the contractor's IR constraint becomes

$$\begin{aligned}
& p_b\left[\int_{-\infty}^{\breve{a}^*(\rho, a, b, \breve{b}^*)} f(\breve{a})[-\breve{a} + b - a]\, d\breve{a} + \int_{\breve{a}^*(\rho, a, b, \breve{b}^*)}^{\infty} f(\breve{a})[-\breve{a} + b]\, d\breve{a}\right] \\
&\quad + (1 - p_b)\left[\int_{-\infty}^{\breve{a}^*(\rho, a, b, \breve{b}^*)} f(\breve{a})[-\breve{a} - a]\, d\breve{a} \right.\\
&\qquad\qquad \left. + \int_{\breve{a}^*(\rho, a, b, \breve{b}^*)}^{\infty} f(\breve{a})[-\rho]\, d\breve{a}\right] \geq E[-\breve{a}] \\
&\Leftrightarrow \int_{\breve{b}^*(\rho, a, b, \breve{a}^*)}^{\infty} g(\breve{b})\left[\int_{-\infty}^{\breve{a}^*(\rho, a, b, \breve{b}^*)} f(\breve{a})[-\breve{a} + b - a]\, d\breve{a} \right.\\
&\qquad\qquad \left. + \int_{\breve{a}^*(\rho, a, b, \breve{b}^*)}^{\infty} f(\breve{a})[-\breve{a} + b]\, d\breve{a}\right] d\breve{b} \\
&\quad + \int_{-\infty}^{\breve{b}^*(\rho, a, b, \breve{a}^*)} g(\breve{b})\left[\int_{-\infty}^{\breve{a}^*(\rho, a, b, \breve{b}^*)} f(\breve{a})[-\breve{a} - a]\, d\breve{a} \right.\\
&\qquad\qquad \left. + \int_{\breve{a}^*(\rho, a, b, \breve{b}^*)}^{\infty} f(\breve{a})[-\rho]\, d\breve{a}\right] d\breve{b} \geq E[-\breve{a}].
\end{aligned}$$

The first half of the left-hand side of the inequality corresponds to the contractee decommitting, while the second half corresponds to the contractee not decommitting. The second integral in each half corresponds to the contractor decommitting, while the third integral in each half corresponds to the contractor not decommitting.

Using the same logic, the contractee's IR constraint becomes

$$\int_{\breve{b}^*(\rho,a,b,\breve{a}^*)}^{\infty} g(\breve{b})\big[p_a[\breve{b}-b+a]+(1-p_a)[\breve{b}-b]\big]\,d\breve{b}$$

$$+\int_{-\infty}^{\breve{b}^*(\rho,a,b,\breve{a}^*)} g(\breve{b})\big[p_a[\breve{b}+a]+(1-p_a)\rho\big]\,d\breve{b} \geq E[\breve{b}]$$

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

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

If $\breve{a}$ is bounded from below, the contractor's decommitment penalty, a, can be chosen so high that the contractor's decommitment threshold, $\breve{a}^*(\rho, a, b, \breve{b}^*)$, becomes lower than any possible realization of $\breve{a}$. In that case the contractor will surely not decommit. Similarly, if $\breve{b}$ is bounded from above, the contractee's decommitment penalty, b, can be chosen so high that the contractee's decommitment threshold, $\breve{b}^*(\rho, a, b, \breve{a}^*)$, is greater than any possible realization of $\breve{b}$. In that case the contractee will surely not decommit. Thus, assuming that $\breve{a}$ is bounded from below and $\breve{b}$ from above, full commitment contracts are a subset of leveled commitment ones. Therefore, the former can be no better in the sense of Pareto efficiency or maximizing the sum of the contract parties' expected payoffs than the latter. It follows that if there exists an IR full commitment contract, then there also exists at least one IR leveled commitment contract. In addition to leveled commitment contracts never being worse than full commitment ones, they can enable a deal that is impossible via full commitment contracts:

THEOREM 2.3 (Enabling in SIMUDBP Games). *There are SIMUDBP games* (*defined by f and g*) *where no full commitment contract is individually rational for both agents but a leveled commitment contract is*.

Proof. Let

$$f(\breve{a}) = \begin{cases} \frac{1}{100} & \text{if } 0 \leq \breve{a} \leq 100 \\ 0 & \text{otherwise} \end{cases} \quad \text{and} \quad g(\breve{b}) = \begin{cases} \frac{1}{110} & \text{if } 0 \leq \breve{b} \leq 110 \\ 0 & \text{otherwise.} \end{cases}$$

No full commitment contract F satisfies both IR constraints since that would require $E[\breve{b}] \leq \rho_F \leq E[\breve{a}]$, which is impossible because $55 = E[\breve{b}] > E[\breve{a}] = 50$. Let us analyze a leveled commitment contract where $\rho = 52.5$; some other choices would work as well. There are four qualitatively different cases.

Case 1 (Some Chance that Either Agent is Going to Decommit). If $0 < \breve{a}^* < 100$ and $0 < \breve{b}^* < 110$, there is a nonzero probability for each agent to decommit. The unique Nash equilibrium is plotted out for different values of a and b in Fig. 4. The Nash equilibrium decommitment thresholds $\breve{a}^*$ and $\breve{b}^*$ differ from the nonstrategic ones. Yet there exist Nash equilibria within the proper range of $\breve{a}^*$ and $\breve{b}^*$. It is not guaranteed that all of these Nash equilibria satisfy the agents' IR constraints, however. Therefore, we programmed a model of Eqs. (9), (10), (11), and (12) and the IR constraints. To make the algebra tractable (constant $f(\breve{a})$ and $g(\breve{b})$), versions of these equations were used that assumed $0 < \breve{a}^* < 100$ and $0 < \breve{b}^* < 110$, without loss of generality. Therefore the first task was to check the boundaries of the validity of the model. The boundaries $\breve{a}^* = 0$ and $\breve{b}^* = 110$ are plotted in Fig. 5. The boundary $\breve{a}^* = 100$ turns out to be the line $b = 0$. There exists no boundary $\breve{b}^* = 0$ because $\breve{b}^*$ was always greater than zero. After plotting the validity boundaries of the model, the curves at which the IR constraints held with equality were plotted; see Fig. 5. Each agent's IR constraint induced three curves, two of which actually bound the IR region. The third one is just a root of the IR constraint, but at both sides of that curve the IR constraint is satisfied. The dark gray area of Fig. 5 represents the values of the decommitment penalties a and b for which the validity constraints of the programmed model and the IR constraints are satisfied. In other words, for any such a and b, there exist decommitment thresholds $\breve{a}^*$ and $\breve{b}^*$ such that these form a Nash equilib-

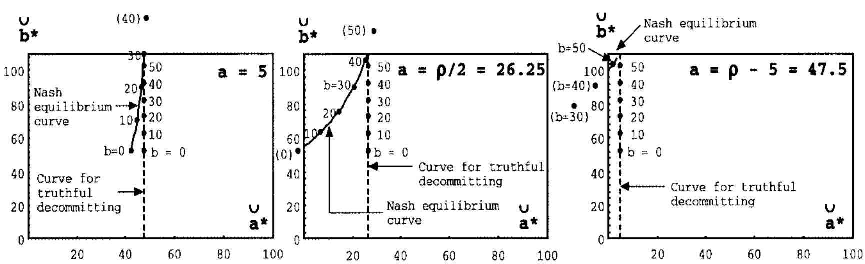

FIG. 4. The Nash equilibrium decommitment thresholds $\breve{a}^*$ and $\breve{b}^*$ of our example SIMUDBP game for different values of the decommitment penalties a and b ($\rho = 52.5$). The Nash equilibrium deviates from nonstrategic ("truthful") decommitting. If $0 < \breve{a}^* < 100$ and $0 < \breve{b}^* < 110$, there is some chance that either agent will decommit.

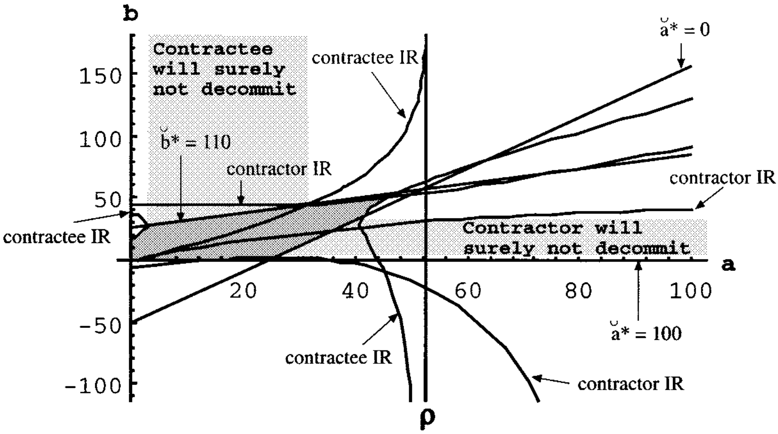

FIG. 5. IR regions in the SIMUDBP decommitting game with $\rho = 52.5$. The gray areas are three qualitatively different regions of contracts that are IR for both agents and allow an equilibrium in the SIMUDBP decommitting game. In the dark gray area either agent might decommit while in the light gray areas only one agent might decommit. The curves represent the IR constraints and validity constraints of the programmed model that requires $0 < \breve{a}^* < 100$ and $0 < \breve{b}^* < 110$. Both agents have one curve from their IR constraint that is just a root of the constraint but is satisfied on both sides.

rium, and there is a nonzero probability for either agent to decommit, and each agent has higher expected payoff with the contract than without it.

As a numerical example, pick a contract where $a = \rho/2 = 26.25$ and $b = 30$. Now in Nash equilibrium, the decommitment thresholds are $\breve{a}^* \approx 20.50$ and $\breve{b}^* \approx 90.24$ (Fig. 4). The contractor's expected payoff is approximately $-44.94 > E[-\breve{a}] = -50$, and the contractee's is approximately $59.81 > E[\breve{b}] = 55$. Thus both agents' expected payoffs are higher than without the contract, i.e., the contract is IR for both agents. This suffices to prove the theorem. Nevertheless, for illustration purposes we present the other types of equilibria that can occur.

Case 2 (Contractor Will Surely Not Decommit). If $\breve{a}^* \leq 0$, the contractor will surely not decommit. Now $\breve{b}^*(\rho, a, b, \breve{a}^*) = \rho + b/(\int_{\breve{a}^*}^{\infty} f(\breve{a})\, d\breve{a}) = \rho + b$, i.e., the contractee decommits nonstrategically. The contractor's IR constraint becomes exactly the same as in Case 2 of the example SEQD game (Eq. (5)). This constraint was proven equivalent to $2.5 \leq b \leq 52.5$. Similarly, the contractee's IR constraint becomes exactly the same as in the SEQD game (Eq. (6)). It was proven equivalent to $b \leq \sim 34.05$. Thus the open region $2.5 \leq b \leq 34.05$, $\breve{a}^* \leq 0$, is where this type of contract is IR for both agents and in equilibrium. This region is colored light gray in Fig. 5.

Case 3 (Contractee Will Surely Not Decommit). If $\breve{b}^* \geq 110$, the contractee will surely not decommit ($p_b = 0$). Now $\breve{a}^*(\rho, a, b, \breve{b}^*) = \rho - a/(1 - p_b) = \rho - a$, i.e., the contractor decommits nonstrategically. The contractor's IR constraint becomes exactly the same as in Case 3 of the example SEQD game (Eq. (8)). This constraint was proven equivalent to $a \leq \sim 30.14$. Similarly, the contractee's IR constraint becomes the same as in the SEQD game (Eq. (9)). It was proven equivalent to $2.5 \leq a \leq 47.5$. Thus the open region $2.5 \leq a \leq 30.14$, $\breve{b}^* \geq 110$ is where this type of contract is IR for both agents and in equilibrium. This region is colored light gray in Fig. 5.

Case 4 (Trivial Case). A contract where at least one agent will surely decommit, i.e., $\breve{a}^* \geq 100$ or $\breve{b}^* \leq 0$, can be IR. For such a contract to be IR for the decommitting agent, its decommitment penalty would have to be zero. Thus the decommitting agent gets the same payoff as without the contract. Similarly, the other agent gets the same payoff as it would get without the contract. Although this contract is IR for both agents (barely because it does not increase either agent's payoff), it is equivalent to no contract at all: decommitment occurs and no payment is transferred. ■

In addition to enabling deals that are impossible using full commitment contracts, leveled commitment contracts can increase the efficiency of a deal even if a full commitment contract were possible. The reverse cannot occur because leveled commitment contracts can emulate full commitment contracts by setting the penalties high enough (assuming that $\breve{a}$ is bounded from below and $\breve{b}$ from above). Leveled commitment contracts improve the efficiency of a deal if there is some chance that the contractor's outside offer is lower than the contractee's expected outside offer, or some chance that the contractee's outside offer is higher than the contractor's expected outside offer.

THEOREM 2.4 (Pareto Efficiency Improvement in SIMUDBP Games). *Let F be any full commitment contract in a SIMUDBP game* (*defined by f and g*). *If*

1. $\breve{b}$ *is bounded from above and* $\Pr(\breve{a} < E[\breve{b}]) > 0$, *or*
2. $\breve{a}$ *is bounded from below and* $\Pr(\breve{b} > E[\breve{a}]) > 0$,

then the game has a leveled commitment contract that increases the contractor's expected payoff as well as the contractee's expected payoff over F. It follows that if F is IR, *then so is this leveled commitment contract*.

Proof. In the proof of Theorem 2.2, a leveled commitment contract was constructed where one agent was sure not to decommit. When one agent is known not to decommit, SIMUDBP games are equivalent to SEQD games. Therefore, the proof of Theorem 2.2 applies. ■

It follows that under the conditions of the theorem, no full commitment contract is Pareto efficient.

2.2.2. *Neither Pays if Both Decommit* (*SIMUDNP*). This section discusses simultaneous decommitting games where a mechanism is used where neither agent has to pay a decommitting penalty if both agents decommit. Such settings will be called SIMUDNP games (Fig. 3). In a game of this type the contractor will decommit if

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

If $p_b = 1$, this is equivalent to $0 > b$. But we already ruled out this type of contract where either one of the agents gets paid for decommitting. On the other hand, the above inequality is equivalent to

$$\breve{a} < \rho - a - \frac{bp_b}{1 - p_b} \stackrel{\text{def}}{=} \breve{a}^*\big(\rho, a, b, \breve{b}^*\big) \qquad \text{when } p_b < 1. \tag{13}$$

The probability that the contractor will decommit is therefore

$$p_a = \int_{-\infty}^{\breve{a}^*(\rho, a, b, \breve{b}^*)} f(\breve{a})\, d\breve{a}. \tag{14}$$

The contractee decommits if

$$p_a \breve{b} + (1 - p_a)[\breve{b} - b] > p_a[\breve{b} + a] + (1 - p_a)\rho.$$

If $p_a = 1$, i.e., $\int_{\breve{a}^*(\rho, a, b, \breve{b}^*)}^{\infty} f(\breve{a})\, d\breve{a} = 0$, this is equivalent to $0 > a$. But we already ruled out this type of contract where either one of the agents gets paid for decommitting. On the other hand, the above inequality is equivalent to

$$\breve{b} > \rho + b - \frac{ap_a}{1 - p_a} \stackrel{\text{def}}{=} \breve{b}^*(\rho, a, b, \breve{a}^*) \qquad \text{when } p_a < 1. \tag{15}$$

The probability that the contractee will decommit is therefore

$$p_b = \int_{\breve{b}^*(\rho, a, b, \breve{a}^*)}^{\infty} g(\breve{b})\, d\breve{b}. \tag{16}$$

Condition (13) defines the contractor's best response (characterized by $\breve{a}^*$) to the contractee's strategy that is characterized by $\breve{b}^*$. Condition (15) defines the contractee's best response (characterized by $\breve{b}^*$) to the contractor's strategy that is characterized by $\breve{a}^*$. Condition (13) uses the variable p_b which is defined by Eq. (16). Condition (15) uses the variable p_a which is defined by Eq. (14). So all together, Eqs. (13), (14), (15), and (16) define the Nash equilibria of the decommitting game.

Now the contractor's IR constraint becomes

$$p_b\left[\int_{-\infty}^{\breve{a}^*(\rho,a,b,\breve{b}^*)} f(\breve{a})[-\breve{a}]\,d\breve{a} + \int_{\breve{a}^*(\rho,a,b,\breve{b}^*)}^{\infty} f(\breve{a})[-\breve{a}+b]\,d\breve{a}\right]$$
$$+(1-p_b)\left[\int_{-\infty}^{\breve{a}^*(\rho,a,b,\breve{b}^*)} f(\breve{a})[-\breve{a}-a]\,d\breve{a} + \int_{\breve{a}^*(\rho,a,b,\breve{b}^*)}^{\infty} f(\breve{a})[-\rho]\,d\breve{a}\right] \geq E[-\breve{a}]$$
$$\Leftrightarrow \int_{\breve{b}^*(\rho,a,b,\breve{a}^*)}^{\infty} g(\breve{b})\left[\int_{-\infty}^{\breve{a}^*(\rho,a,b,\breve{b}^*)} f(\breve{a})[-\breve{a}]\,d\breve{a} + \int_{\breve{a}^*(\rho,a,b,\breve{b}^*)}^{\infty} f(\breve{a})[-\breve{a}+b]\,d\breve{a}\right] d\breve{b}$$
$$+\int_{-\infty}^{\breve{b}^*(\rho,a,b,\breve{a}^*)} g(\breve{b})\left[\int_{-\infty}^{\breve{a}^*(\rho,a,b,\breve{b}^*)} f(\breve{a})[-\breve{a}-a]\,d\breve{a} + \int_{\breve{a}^*(\rho,a,b,\breve{b}^*)}^{\infty} f(\breve{a})[-\rho]\,d\breve{a}\right] d\breve{b} \geq E[-\breve{a}].$$

The first half of the left-hand side of the inequality corresponds to the contractee decommitting, while the second half corresponds to the contractee not decommitting. The second integral in each half corresponds to the contractor decommitting, while the third integral corresponds to the contractor not decommitting.

Using the same logic, the contractee's IR constraint becomes

$$\int_{\breve{b}^*(\rho,a,b,\breve{a}^*)}^{\infty} g(\breve{b})\left[p_a\breve{b} + (1-p_a)[\breve{b}-b]\right] d\breve{b}$$
$$+\int_{-\infty}^{\breve{b}^*(\rho,a,b,\breve{a}^*)} g(\breve{b})\left[p_a[\breve{b}+a] + (1-p_a)\rho\right] d\breve{b} \geq E[\breve{b}]$$
$$\Leftrightarrow \int_{\breve{b}^*(\rho,a,b,\breve{a}^*)}^{\infty} g(\breve{b})\left[\int_{-\infty}^{\breve{a}^*(\rho,a,b,\breve{b}^*)} f(\breve{a})[\breve{b}]\,d\breve{a} + \int_{\breve{a}^*(\rho,a,b,\breve{b}^*)}^{\infty} f(\breve{a})[\breve{b}-b]\,d\breve{a}\right] d\breve{b}$$
$$+\int_{-\infty}^{\breve{b}^*(\rho,a,b,\breve{a}^*)} g(\breve{b})\left[\int_{-\infty}^{\breve{a}^*(\rho,a,b,\breve{b}^*)} f(\breve{a})[\breve{b}+a]\,d\breve{a} + \int_{\breve{a}^*(\rho,a,b,\breve{b}^*)}^{\infty} f(\breve{a})[\rho]\,d\breve{a}\right] d\breve{b} \geq E[\breve{b}].$$

If $\breve{a}$ is bounded from below, the contractor's decommitment penalty a can be chosen so high that the contractor's decommitment threshold $\breve{a}^*(\rho, a, b, \breve{b}^*)$ becomes lower than $\breve{a}$. In that case the contractor will surely not decommit. Similarly, if $\breve{b}$ is bounded from above, the contractee's decommitment penalty b can be chosen so high that the contractee's decommitment threshold $\breve{b}^*(\rho, a, b, \breve{a}^*)$ is greater than $\breve{b}$. In that case the contractee will surely not decommit. Thus, assuming that $\breve{a}$ is bounded from below and $\breve{b}$ from above, full commitment contracts are a subset of leveled commitment ones. Therefore, the former can be no better in the sense of Pareto efficiency or maximizing the sum of the contract parties' expected payoffs than the latter. It follows that if there exists an IR full commitment contract, then there also exists at least one IR leveled commitment contract. In addition to these arguments that state that leveled commitment contracts are never worse than full commitment ones, the following theorem states the positive result that in SIMUDNP games, leveled commitment contracts can enable (via increased efficiency) a deal that is not possible via full commitment contracts.

THEOREM 2.5 (Enabling in SIMUDNP Games). *There are SIMUDNP games* (*defined by f and g*) *where no full commitment contract is individually rational for both agents*, *but a leveled commitment contract is*.

Proof. Let

$$f(\breve{a}) = \begin{cases} \frac{1}{100} & \text{if } 0 \le \breve{a} \le 100 \\ 0 & \text{otherwise} \end{cases} \quad \text{and} \quad g(\breve{b}) = \begin{cases} \frac{1}{110} & \text{if } 0 \le \breve{b} \le 110 \\ 0 & \text{otherwise.} \end{cases}$$

No full commitment contract, F, satisfies both IR constraints since that would require $E[\breve{b}] \le \rho_F \le E[\breve{a}]$, which is impossible because $55 = E[\breve{b}] > E[\breve{a}] = 50$. Let us analyze a leveled commitment contract where $\rho = 52.5$. There are four qualitative different cases.

Case 1 (Some Chance that Either Agent is Going to Decommit). If $0 < \breve{a}^* < 100$, and $0 < \breve{b}^* < 110$, there is a nonzero probability for each agent to decommit. The Nash equilibrium is plotted out for different values of a and b in Fig. 6. Note that the Nash equilibrium decommitment thresholds $\breve{a}^*$ and $\breve{b}^*$ indeed do differ from the nonstrategic ones. It is not guaranteed that all of these Nash equilibria satisfy the agents' IR constraints. Therefore, we programmed a model of Eqs. (13)–(16) and the IR constraints. To make the algebra tractable (constant $f(\breve{a})$ and $g(\breve{b})$), versions of these equations were used that assumed $0 < \breve{a}^* < 100$ and $0 < \breve{b}^* < 110$. Therefore the first task was to check the validity boundaries of the model. The boundaries $\breve{a}^* = 0$, $\breve{a}^* = 100$, $\breve{b}^* = 0$, and $\breve{b}^* = 110$ are plotted with bold lines in Fig. 7. After plotting the validity boundaries of

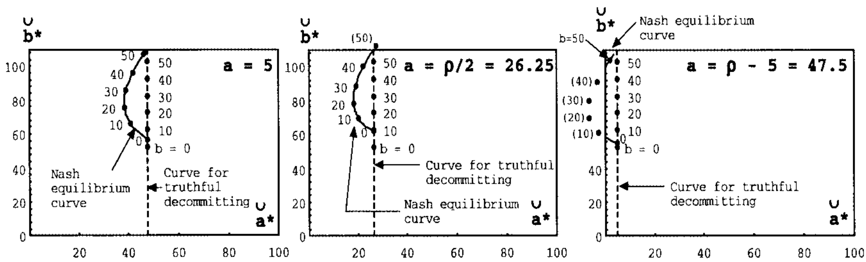

FIG. 6. The Nash equilibrium decommitment thresholds $\breve{a}^*$ and $\breve{b}^*$ of our example SIMUDNP game for different values of the decommitment penalties a and b ($\rho = 52.5$). The Nash equilibrium deviates from nonstrategic ("truthful") decommitting. If $0 < \breve{a}^* < 100$, and $0 < \breve{b}^* < 110$, there is some chance that either agent will decommit.

the model, the curves at which the IR constraints held with equality were plotted in Fig. 7. Note that each agent's IR constraint induced three curves, two of which actually bound the IR region. The third one is just a root of the IR constraint, but at both sides of that curve the IR constraint is satisfied. Now, the dark gray area of Fig. 7 represents the values of the decommitment penalties a and b for which the validity constraints of the programmed model and the IR constraints are satisfied. In other words,

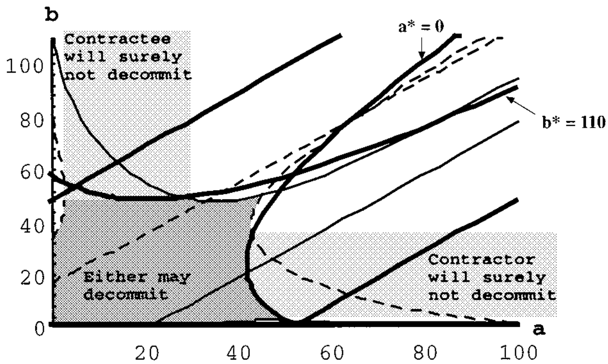

FIG. 7. IR regions in the SIMUDNP decommitting game. The gray areas are three qualitatively different regions of contracts that are IR for both agents and allow an equilibrium in the SIMUDNP decommitting game. The bold lines are the validity constraints for the programmed model that requires $0 < \breve{a}^* < 100$, and $0 < \breve{b}^* < 110$. One of the constraints that slices the "either may decommit" region is just a root of a constraint, but the constraint is satisfied on both sides of the line. The solid lines represent the contractor's IR constraint from the programmed model, and the dashed lines represent the contractee's IR constraint. Both agents have one curve from their constraint that is just a root of the constraint but is satisfied on both sides.

for any such a and b, there exist decommitment thresholds $\breve{a}^*$ and $\breve{b}^*$ such that these form a Nash equilibrium, and there is a nonzero probability for either agent to decommit, and each agent has higher expected payoff with the contract than without it.

As a numerical example, pick a contract where $a = \rho/2 = 26.25$ and $b = 30$. Now, in Nash equilibrium the decommitment thresholds are $\breve{a}^* \approx 19.03$ and $\breve{b}^* \approx 88.67$ (Fig. 6). The contractor's expected payoff is approximately $-44.74 > E[-\breve{a}] = -50$, and the contractee's is approximately $59.65 > E[\breve{b}] = 55$. Thus both agents' expected payoffs are higher than without the contract, i.e., the contract is IR for both agents. This suffices to prove the theorem. Nevertheless, for illustration purposes, we proceed to present the other types of equilibria that can occur.

Case 2 (Contractor Will Surely Not Decommit). If $\breve{a}^* \leq 0$, the contractor will surely not decommit. Now $\breve{b}^*(\rho, a, b, \breve{a}^*) = \rho + b - a\int_{-\infty}^{\breve{a}^*(\rho, a, b, \breve{b}^*)} f(\breve{a})\, d\breve{a} / \int_{\breve{a}^*(\rho, a, b, \breve{b}^*)}^{\infty} f(\breve{a})\, d\breve{a} = \rho + b$, i.e., the contractee decommits nonstrategically. The contractor's IR constraint becomes exactly the same as in Case 2 of the example SEQD game (Eq. (5)). (It is also the same as in Case 2 of the example SIMUDBP game.) This constraint was proven equivalent to $2.5 \leq b \leq 52.5$. Similarly, the contractee's IR constraint becomes exactly the same as in the SEQD game (Eq. (6)). (It is also same as in Case 2 of the example SIMUDBP game.) It was proven equivalent to $b \leq \sim 34.05$. Thus the open region $2.5 \leq b \leq 34.05$, $\breve{a}^* \leq 0$, is where this type of contract is IR for both agents and in equilibrium. This region is colored light gray in Fig. 7.

Case 3 (Contractee Will Surely Not Decommit). If $\breve{b}^* \geq 110$, the contractee will surely not decommit ($p_b = 0$). Now $\breve{a}^*(\rho, a, b, \breve{b}^*) = \rho - a - bp_b/(1 - p_b) = \rho - a$, i.e., the contractor decommits nonstrategically. The contractor's IR constraint becomes exactly the same as in Case 3 of the example SEQD game (Eq. 8). (It is also the same as in Case 3 of the example SIMUDBP game.) This constraint was proven equivalent to $a \leq \sim 30.14$. Similarly, the contractee's IR constraint becomes the same as in the SEQD game (Eq. (9)). (It is also the same as in Case 3 of the example SIMUDBP game.) It was proven equivalent to $2.5 \leq a \leq 47.5$. Thus the open region $2.5 \leq a \leq 30.14$, $\breve{b}^* \geq 110$ is where this type of contract is IR for both agents and in equilibrium. This region is colored light gray in Fig. 7.

Case 4 (Trivial Case). A contract where one agent will surely decommit, i.e., $\breve{a}^* \geq 100$ or $\breve{b}^* \leq 0$, can be IR. In such cases the other agent's dominant strategy is to not decommit, i.e., to collect the decommitment penalty from the first agent. For such a contract to be IR for the decommitting agent, its decommitment penalty would have to be zero.

Thus the decommitting agent gets the same payoff as without the contract. Similarly, the other agent gets the same payoff as it would get without the contract. Although this contract is IR for both agents (barely, because it does not increase either agent's payoff), it is equivalent to no contract at all: decommitment occurs and no payoff is transferred. ■

In addition to the fact that leveled commitment contracts may enable deals that are impossible using full commitment contracts, leveled commitment contracts can increase the efficiency of a deal even if a full commitment contract were possible. The reverse cannot occur because leveled commitment contracts can emulate full commitment contracts by setting the penalties high enough (assuming that $\breve{a}$ is bounded from below and $\breve{b}$ from above).

THEOREM 2.6 (Pareto Efficiency Improvement in SIMUDNP Games). *Let F be any full commitment contract in a SIMUDNP game (defined by f and g). If*

1. *$\breve{b}$ is bounded from above and* $\Pr(\breve{a} < E[\breve{b}]) > 0$, *or*
2. *$\breve{a}$ is bounded from below and* $\Pr(\breve{b} > E[\breve{a}]) > 0$,

then the game has a leveled commitment contract that increases the contractor's expected payoff as well as the contractee's expected payoff over F. It follows that if F is IR, then so is this leveled commitment contract.

Proof. In the proof of Theorem 2.2, a leveled commitment contract was constructed where one agent was sure not to decommit. When one agent is known not to decommit, SIMUDNP games are equivalent to SEQD games. Therefore, the proof of Theorem 2.2 applies. ■

It follows that, under the conditions of the theorem, no full commitment contract is Pareto efficient.

3. COMPARING THE EQUILIBRIA OF THE THREE LEVELED COMMITMENT MECHANISMS

The Nash equilibrium decommitting strategies of the simultaneous decommitting mechanisms differ significantly depending on whether the agents have to pay the penalties to each other (SIMUDBP) or not (SIMUDNP) in the case both decommit.

THEOREM 3.1. *For any f and g, as the contractee's decommitting penalty approaches zero, in the SIMUDBP mechanism the contractee becomes nonstrategic, while in the SIMUDNP mechanism the contractor does. As the contractee's decommitting penalty increases (contractee's decommitting proba-*

bility approaches zero), the contractor becomes nonstrategic in both SIMUDBP and SIMUDNP.

Analogously, as the contractor's penalty approaches zero, in the SIMUDBP mechanism the contractor becomes nonstrategic, while in the SIMUDNP mechanism the contractee does. As the contractor's decommitting penalty increases (contractor's decommitting probability approaches zero), the contractee becomes nonstrategic in both SIMUDBP and SIMUDNP.

Proof. Immediate from the formulas for $\breve{a}^*$ and $\breve{b}^*$ (Eqs. (9), (11), (13), and (15)). ■

The phenomena that the theorem describes can be clearly observed in the context of our example f and g in Figs. 4 and 6.

In the sequential decommitting mechanism (SEQD), if the first mover decommits, then the second mover will never decommit since its decommitting penalty is nonnegative and it has already been freed from the contract obligations. If the first mover does not decommit, the second mover decommits nonstrategically since there are no remaining strategic effects. In the sequential mechanism, the first mover is even more reluctant to decommit than in the simultaneous decommitting mechanisms.

PROPOSITION 3.1. *As before, let p_a in the simultaneous mechanisms be the probability that the contractor decommits. Also as before, let p_a in the sequential mechanism be the probability that the contractor decommits given that the contractee (first mover) did not decommit. For given $p_a < 1$, ρ, $a \geq 0$, and $b \geq 0$, the contractee is most reluctant to decommit (most strategic) in the SEQD mechanism, less reluctant to decommit in the SIMUDBP mechanism, and least reluctant to decommit in the SIMUDNP mechanism. If $a > 0$ and $p_a > 0$, the contractee is strictly more reluctant in SEQD than in SIMUDBP. If $a > 0$ or $b > 0$, and $p_a > 0$, the contractee is strictly more reluctant in SIMUDBP than in SIMUDNP.*

Proof. The higher the decommitting threshold $\breve{b}^*$, the more reluctant the contractee is. From Eqs. (2), (11), and (15) we get $\breve{b}^*$ for each of the mechanisms.

We first prove that the contractee is more reluctant in SEQD ($\breve{b}^* = \rho + (b + p_a a)/(1 - p_a)$) than in SIMUDBP ($\breve{b}^* = \rho + b/(1 - p_a)$).

$$\rho + \frac{b + p_a a}{1 - p_a} \geq \rho + \frac{b}{1 - p_a}$$

$$\Leftrightarrow \frac{p_a a}{1 - p_a} \geq 0,$$

which holds because $0 \le p_a < 1$ and $a \ge 0$. If $a > 0$ and $p_a > 0$, the inequality is strict.

Next we prove that the contractee is more reluctant in SIMUDBP than in SIMUDNP ($\breve{b}^* = \rho + b - ap_a/(1 - p_a)$).

$$\rho + \frac{b}{1 - p_a} \ge \rho + b - \frac{ap_a}{1 - p_a}$$

$$\Leftrightarrow \frac{b}{1 - p_a} \ge \frac{b - bp_a - ap_a}{1 - p_a}$$

$$\Leftrightarrow 0 \ge \frac{-bp_a - ap_a}{1 - p_a},$$

which holds because $0 \le p_a < 1$, $a \ge 0$, and $b \ge 0$. If $a > 0$ or $b > 0$, and $p_a > 0$, the inequality is strict. ■

4. UNCERTAINTY ABOUT ONLY ONE AGENT'S OUTSIDE OFFER

So far we discussed games where both agents receive uncertain outside offers (or analogously, the value of a deal is uncertain and changes independently for both parties). In this section we discuss games where only one agent's outside offer is uncertain. In such games all of the probability mass of $f(\breve{a})$ or $g(\breve{b})$ is on one point.

Theorems 2.1, 2.3, and 2.5 apply to this setting as well (in the proof it was possible to pick a contract where one party was surely not going to decommit; by such a choice, the proof only capitalizes on one-sided uncertainty). The theorems state that there are instances of the game where no full commitment contract is IR but a leveled commitment contract is.

Theorems 2.2, 2.4, and 2.6 also apply to this setting (the proof was generated by capitalizing on one-sided uncertainty only). Worded in the setting of one-sided uncertainty, these theorems state that if there is some chance that the contractor's outside offer is lower than the contractee's or some chance that the contractee's outside offer is higher than the contractor's, then compared to any full commitment contract there exists a leveled commitment contract that has higher expected payoff for both contract parties.

We now briefly discuss the game of one-sided uncertainty in order to be able to present our later results regarding biased asymmetric beliefs among agents. We present the case where the contractee has the certain outside offer, $\breve{b}$, and, as before, the contractor's (best) outside offer is only known probabilistically by the agents via a probability density function $f(\breve{a})$ which is common knowledge. The case where the contractor instead has the certain outside offer is analogous. To distinguish from another game of one-sided uncertainty which we will present later, let us call this game the *certain offer prevails* (COP) game (Fig. 8).

In a sequential decommitting COP game where the contractee reveals decommitment first (Fig. 8), because the contractee gains no information between the beginning of the contracting game and the decommitting game, it will not find decommitting beneficial (for any $b \geq 0$) if it found the original contract beneficial (better than its outside offer $\breve{b}$) and thus agreed to it.

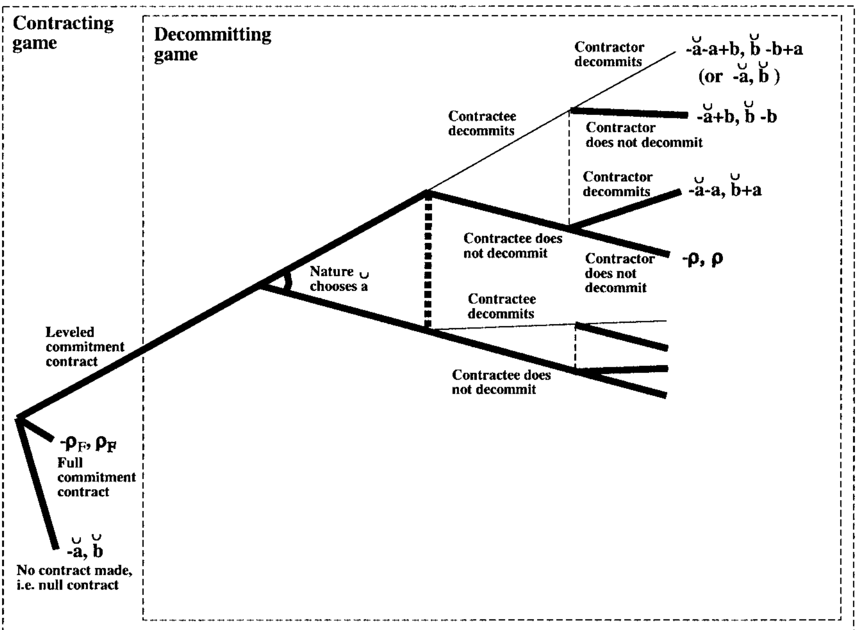

FIG. 8. The "certain offer prevails" (COP) game. If the contractor decommits, the contractee can still accept the outside offer. In the figure, the contractor's payoff is listed before the contractee's. The bold solid lines show choices that may actually occur in any subgame. The bold dashed line represents the contractee's information set. It models the aspect that the contractee does not know $\breve{a}$ in the decommitting game. The thin dashed lines represent the alternative mechanism where both agents have to reveal their decommitment decisions simultaneously: when the contractor has to reveal its decommitting decision, it has not observed the contractee's decommitting decision.

This holds even for a game where the agents reveal decommitting simultaneously as opposed to contractee first (this game is depicted by the information sets denoted by thin dashed lines in Fig. 8).

Even in a sequential decommitting game where the contractor moves first, the contractee will not want to decommit. If the contractor decommits, the contractee can save its decommitment penalty by not decommitting and the contract becomes void anyway. Now let us discuss the branches where the contractor did not decommit. In these branches the contractee's payoff is independent of the contractor's outside offer, and thus all of these branches are equivalent. Now, if the contractee would be better off decommitting in such a branch, the contractor would know that. Therefore the contractor would never decommit (i.e., no matter what its outside offer turns out to be). Thus the decommitting game would always be played by the contractor not decommitting and the contractee decommitting. Clearly this kind of a contracting game cannot be strictly IR for the contractee. Therefore the contractee never decommits in a strictly IR contract.

Even if a mechanism is used that specifies that neither agent has to pay the decommitment penalty if both decommit (payoffs in parentheses in Fig. 8), the contractee wants to decommit in none of the three cases above.

Thus the only agent to possibly decommit is the contractor. In any one of the above three mechanisms, the contractor can reason that the contractee will not decommit. Therefore the three cases become equivalent. This holds for the mechanism that specifies that both have to pay if both decommit and for the mechanism that specifies that neither has to pay if both decommit.

The contractor's payoff is $-\rho$ if it does not decommit and $-\breve{a} - a$ if it does. Therefore, the contractor will decommit if $\breve{a} + a < \rho$. Thus the probability that the contractor will decommit is

$$p_a = \int_{-\infty}^{\rho - a} f(\breve{a})\, d\breve{a}.$$

The contractee's individual rationality (IR) constraint states that the contract has to have higher expected payoff than the fixed outside offer,

$$\breve{b} \leq [1 - p_a]\rho + p_a[\breve{b} + a]$$

The contractor can choose (*ex post*) whether it wants to decommit or stay with the contract. Therefore the contractor's *ex ante* IR constraint is based on the idea that $E[-\breve{a}] \leq E[\max[-\breve{a} - a, -\rho]]$:

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

5. OUTSIDE OFFER THAT BECOMES VOID BEFORE THE DECOMMITTING GAME

This section discusses a setting where the contractee has a fixed outside offer, $\breve{b}$, but this offer has to be accepted before the contractor finds out the price of its (best) outside offer, $\breve{a}$ (in case the contractor receives no outside offer, $\breve{a}$ is its fallback payoff). Otherwise the $\breve{b}$-offer becomes void. Thus, to agree to the contract, the contractee has to get a higher expected payoff by passing on the $\breve{b}$-offer and agreeing to the risky contract than he would gain by accepting the $\breve{b}$-offer. If the contract is made, decommitment happens (if at all) when the contractor's outside offer is valid (and known to the contractor but not to the contractee) but the contractee's is not anymore. In this case the contractee gets its fallback payoff, $\bar{b}$, plus the contractor's decommitment penalty payment, a. The fallback, $\bar{b}$, can be interpreted, for example, as the contractee's second best outside offer (best that is still available) or, in the case no outside offers are outstanding, as the contractee's payoff without any contracts. We will call the setting the *certain offer becomes void* (COBV) game; see Fig. 9.

In a sequential decommitting COBV game where the contractee reveals decommitment first (Fig. 9), because the contractee gains no information between the beginning of the contracting game and the decommitting game, it will not find decommitting beneficial for any $b \geq 0$ if it found the original contract beneficial (better than its outside offer $\breve{b}$) and thus agreed to it.

This holds even for a game where the agents reveal decommitting simultaneously, as opposed to contractee first. This game is depicted by the information sets denoted by thin dashed lines in Fig. 9.

Even in a sequential decommitting game where the contractor moves first, the contractee will not want to decommit. If the contractor decommits, the contractee can save its decommitment penalty by not declaring decommitment and the contract becomes void anyway. Now let us discuss the branches where the contractor did not decommit. In these branches the contractee's payoff is independent of the contractor's outside offer, and thus all of these branches are equivalent. Now, if the contractee would be better off decommitting in such a branch, the contractor would know that. Therefore the contractor would never decommit, no matter what its outside offer turns out to be. Thus the decommitting game would always be played by the contractor not decommitting and the contractee decommitting. Clearly this kind of a contracting game cannot be strictly IR for the contractee. Therefore the contractee never decommits.

Even if a mechanism is used that specifies that neither agent has to pay a decommitment penalty if both decommit (payoffs in parentheses in Fig. 9), the contractee wants to decommit in none of the three cases above.

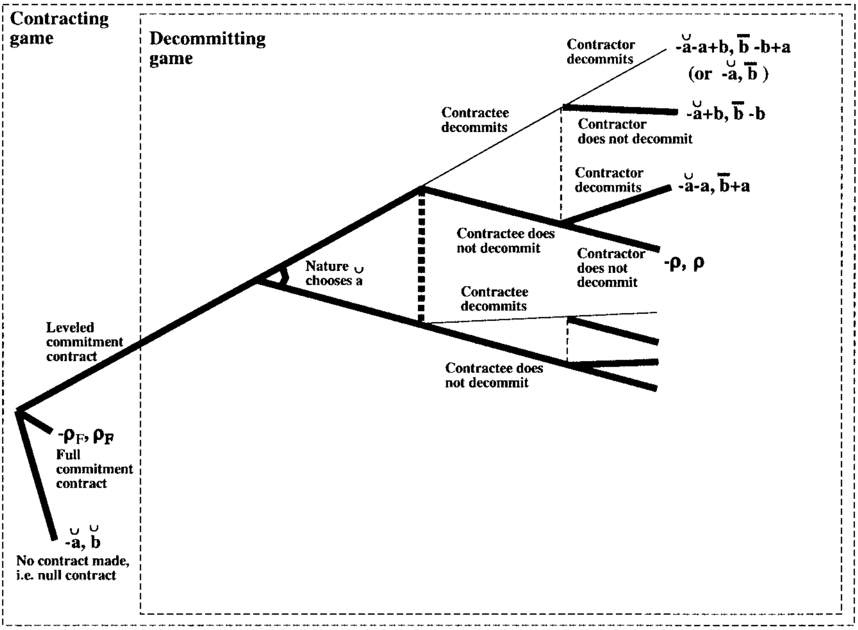

FIG. 9. The "certain offer becomes void" (COBV) game. If at all, the contractee's outside offer $\breve{b}$ has to be accepted before the contractor's outside offer $\breve{a}$ becomes known. The bold solid lines show choices that may actually occur in any subgame. The bold dashed line represents the contractee's information set. It models the aspect that the contractee does not know $\breve{a}$ in the decommitting game. The thin dashed lines represent the alternative mechanism where both agents have to reveal their decommitting decisions simultaneously: when the contractor has to reveal its decommitting decision, it has not yet observed the contractee's decommitting decision.

Thus the only agent to possibly make a move in the decommitting game is the contractor. In any one of the above three settings, the contractor can reason that the contractee will not decommit. Therefore the three cases become equivalent. This holds for the mechanism that specifies that both have to pay if both decommit and for the mechanism that specifies that neither has to pay if both decommit.

The contractor's payoff is $-\rho$ if it does not decommit and $-\breve{a} - a$ if it does. Therefore, the contractor will decommit if $\breve{a} + a < \rho$. Thus,

$$p_a = \int_{-\infty}^{\rho - a} f(\breve{a})\, d\breve{a}.$$

The contractee's IR constraint is

$$\breve{b} \leq [1 - p_a]\rho + p_a[\bar{b} + a],$$

where $\bar{b}$ is the contractee's fallback position, i.e., the payoff it gets if it does not get its outside offer $\breve{b}$ or the contract with the contractor.

The contractor's IR constraint is based on the idea that, *ex post*, the contractor can choose whether it wants to decommit or stay with the contract. *Ex post*, the contractor finds the contract individually rational if $-\breve{a} \leq \max[-\breve{a} - a, -\rho] \Leftrightarrow \breve{a} \geq \rho$. Thus the *ex ante* IR constraint is

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

This is the same constraint as the contractor's IR constraint in the COP game.

Full commitment contracts are a subset of leveled commitment ones in COBV games because the contractor's decommitment penalty can be chosen so high that the contractor will surely not decommit (assuming that $\breve{a}$ is bounded from below). As discussed earlier, the contractee will not decommit for any $b \geq 0$ either. Thus, the class of leveled commitment contracts is no worse than the class of full commitment ones.

Furthermore, the result from earlier in this paper that leveled commitment contracts can enable deals where no full commitment contract is IR to both parties and the result that leveled commitment contracts can improve the expected payoff of both contract parties imply that these two advantages hold even for COBV games. This is easy to see for example when $\bar{b} = \breve{b}$. Under this condition, the contractee's IR constraint becomes the same as in the COP game (Section 4), and the contractor's IR constraint is always the same as in the COP game. Therefore, under this condition the two positive results from the COP game apply to the COBV game. Naturally they also hold if $\bar{b} > \breve{b}$ because this can only make the risky leveled commitment contract more desirable to the contractee without affecting the desirability to the contractor.

Although leveled commitment contracts have these advantages in COBV games when the contractee's fallback is sufficiently high, the following two propositions show the nonsurprising result that if the contractee's fallback is too low, leveled commitment contracts are not helpful in COBV games.

PROPOSITION 5.1 (No Enabling in a COBV Game). *Let* $\bar{b} \leq (\int_{-\infty}^{\rho-a} f(\breve{a})\breve{a}\, d\breve{a})/(\int_{-\infty}^{\rho-a} f(\breve{a})\, d\breve{a})$ *in a COBV game. If no full commitment contract satisfies the IR constraints, no leveled commitment contract satisfies them either.*

Proof. For a full commitment contract, F, to be IR for both parties, it has to be the case that $\breve{b} \leq \rho_F \leq E[\breve{a}]$. No such ρ_F exists iff $\breve{b} > E[\breve{a}]$. Now say that $\breve{b} > E[\breve{a}]$, and assume (for contradiction) that some leveled commitment contract defined by ρ, a, and b satisfies both IR constraints. Thus,

$$\begin{aligned}
&\left[1 - \left(\int_{-\infty}^{\rho-a} f(\breve{a})\, d\breve{a}\right)\right]\rho + \left(\int_{-\infty}^{\rho-a} f(\breve{a})\, d\breve{a}\right)[\bar{b} + a] \geq \breve{b} > E[\breve{a}] \\
&\quad \geq \int_{-\infty}^{\rho-a} f(\breve{a})[\breve{a} + a]\, d\breve{a} + \int_{\rho-a}^{\infty} f(\breve{a})\,\rho\, d\breve{a} \\
&\quad \Rightarrow \left(\int_{-\infty}^{\rho-a} f(\breve{a})\, d\breve{a}\right)[\bar{b} + a] > \int_{-\infty}^{\rho-a} f(\breve{a})[\breve{a} + a]\, d\breve{a} \\
&\quad \Leftrightarrow \left(\int_{-\infty}^{\rho-a} f(\breve{a})\, d\breve{a}\right)\bar{b} > \int_{-\infty}^{\rho-a} f(\breve{a})\,\breve{a}\, d\breve{a} \\
&\quad \Leftrightarrow \bar{b} > \frac{\int_{-\infty}^{\rho-a} f(\breve{a})\breve{a}\, d\breve{a}}{\int_{-\infty}^{\rho-a} f(\breve{a})\, d\breve{a}} \geq \bar{b}.
\end{aligned}$$

Thus we have a contradiction. Thus no leveled commitment contract satisfies both IR constraints. ■

PROPOSITION 5.2 (No Pareto Efficiency Improvement in a COBV Game). *Let $\bar{b} \leq (\int_{-\infty}^{\rho-a} f(\breve{a})\breve{a}\, d\breve{a})/(\int_{-\infty}^{\rho-a} f(\breve{a})\, d\breve{a})$ in a COBV game. Let F be an arbitrary full commitment contract that satisfies both IR constraints, i.e., $\breve{b} \leq \rho_F \leq E[\breve{a}]$. There exists no leveled commitment contract that increases (over F) at least one agent's expected payoff without decreasing the other agent's expected payoff.*

Proof. Under F, the contractor's payoff is $-\rho_F$ and the contractee's is ρ_F. Assume, for contradiction, that there exists a leveled commitment contract (defined by ρ, a, and b) that increases at least one of these payoffs while not decreasing the other, i.e.,

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

and

$$\left[1 - \left(\int_{-\infty}^{\rho-a} f(\breve{a})\, d\breve{a}\right)\right]\rho + \left(\int_{-\infty}^{\rho-a} f(\breve{a})\, d\breve{a}\right)[\bar{b} + a] \geq \rho_F,$$

and at least one of the above inequalities is strict,

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

$$< \left[1 - \left(\int_{-\infty}^{\rho-a} f(\breve{a})\,d\breve{a}\right)\right]\rho + \left(\int_{-\infty}^{\rho-a} f(\breve{a})\,d\breve{a}\right)[\bar{b}+a]$$

$$\Leftrightarrow \int_{-\infty}^{\rho-a} f(\breve{a})\breve{a}\,d\breve{a} < \left(\int_{-\infty}^{\rho-a} f(\breve{a})\,d\breve{a}\right)\bar{b}$$

$$\Leftrightarrow \frac{\int_{-\infty}^{\rho-a} f(\breve{a})\breve{a}\,d\breve{a}}{\int_{-\infty}^{\rho-a} f(\breve{a})\,d\breve{a}} < \bar{b}$$

$$\Rightarrow \bar{b} < \bar{b}.$$

We have a contradiction. Thus no such leveled commitment contract exists. ■

The constraint $\bar{b} \leq \int_{-\infty}^{\rho-a} f(\breve{a})\breve{a}\,d\breve{a} / \int_{-\infty}^{\rho-a} f(\breve{a})\,d\breve{a}$ is satisfied for example if $\forall \breve{a} \leq 0$, $f(\breve{a}) = 0$, and $\bar{b} \leq 0$. This means that the contractor's outside offer will require some nonnegative payment to do the contractor's task, and that the contractee has a nonpositive fallback. The former requirement does not seem very restrictive, but the latter does. Thus these two propositions with negative results have relatively limited scope.

6. BIASED BELIEFS ABOUT THE DISTRIBUTIONS OF OUTSIDE OFFERS

So far we have discussed games where the distributions, f and g, of the outside offers are common knowledge. In this section we study what happens when an agent thinks that they are common knowledge and acts according to the equilibrium derived under that assumption, but in reality at least one of the agents has biased beliefs about the distributions. The main conclusion will be that in the games with one-sided uncertainty an agent's biased beliefs can only hurt that agent, but in the games with two-sided uncertainty, the biased agent can end up hurting the other agent as well.

6.1. *Effect of Biased Asymmetric Beliefs in SEQD Games*

Recall that in SEQD games, both agents have uncertain futures and they have to reveal their decommitting decisions sequentially. In SEQD games, one agent's expected payoff for a given contract can be affected by the other agent's biased beliefs. For example, the contractee's decision of whether to decommit depends on its subjective distribution, $f_b(\breve{a})$, of the contractor's best upcoming outside offer,

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

That decommitting decision affects the contractor's expected payoff, which really is (the contractor could perceive it differently)

$$\pi_a = \int_{\breve{b}^*(\rho, a, b)}^{\infty} g(\breve{b}) \int_{-\infty}^{\infty} f(\breve{a})[-\breve{a} + b]\, d\breve{a}\, d\breve{b}$$
$$+ \int_{-\infty}^{\breve{b}^*(\rho, a, b)} g(\breve{b}) \left[\int_{-\infty}^{\rho-a} f(\breve{a})[-\breve{a} - a]\, d\breve{a} + \int_{\rho-a}^{\infty} f(\breve{a})[-\rho]\, d\breve{a} \right] d\breve{b}.$$

6.2. *Effect of Biased Asymmetric Beliefs in SIMUDBP Games*

Recall that, in SIMUDBP games, both agents have uncertain futures and they have to reveal their decommitting decisions simultaneously. If both decommit, both have to pay the penalties to each other. In SIMUDBP games (like SEQD games) an agent's expected payoff for a given contract can be affected by the other agent's biased beliefs. For example, the contractee's decision of whether to decommit depends on its subjective distribution, $f_b(\breve{a})$, of the contractor's best upcoming outside offer. If the contractee receives a good outside offer, it would decommit if it acted nonstrategically. But if the contractee believes (according to $f_b(\breve{a})$) that the contractor is likely to get a good outside offer and decommit, then the contractee can save the decommitment penalty by not decommitting. On the other hand, the contractee's decommitting decision affects the contractor's expected payoff because in the case the contractee decommits, the contractor's payoff is either $-\breve{a} + b$ or $-\breve{a} + b - a$, and in case the contractee does not decommit the contractor's payoff is either $-\rho$ or $-\breve{a} - a$. Because of such dependencies, an agent's preference order over potential contracts may depend on the other agent's beliefs. Therefore, in SIMUDBP games with asymmetric biased information, an agent may need to counterspeculate the other agent's beliefs in order to determine a preference order over contracts.

6.3. *Effect of Biased Asymmetric Beliefs in SIMUDNP Games*

Recall that in SIMUDNP games both agents have uncertain futures and they have to reveal their decommitting decisions simultaneously. If both decommit, neither has to pay the decommitment penalty. In SIMUDNP games (like SIMUDNP and SEQD games) an agent's expected payoff for a given contract can be affected by the other agent's biased beliefs. For example, the contractee's decision of whether to decommit depends on its subjective distribution, $f_b(\breve{a})$, of the contractor's best upcoming outside offer. If the contractee receives a good outside offer, it would decommit if it acted nonstrategically. But if the contractee believes (according to $f_b(\breve{a})$) that the contractor is likely to get a good outside offer and decommit, then the contractee can save the decommitment penalty by not decommitting. On the other hand, the contractee's decommitting decision affects the contractor's expected payoff, because in the case the contractee decommits the contractor's payoff is either $-\breve{a} + b$ or $-\breve{a}$, and in case the contractee does not decommit the contractor's payoff is either $-\rho$ or $-\breve{a} - a$. Because of such dependencies, an agent's preference order over potential contracts may depend on the other agent's beliefs. Therefore, in SIMUDNP games with biased asymmetric information, an agent may need to counter-speculate the other agent's beliefs in order to determine a preference order over contracts.

6.4. *Effect of Biased Asymmetric Beliefs in COP Games*

Recall that COP games are like SEQD, SIMUDBP, and SIMUDNP games except that only one agent's future is uncertain. Unlike in the SEQD, SIMUDBP, and SIMUDNP games, in COP games an agent cannot be hurt by the other agents' biased beliefs as the following proposition shows. For any specific contract, an agent with precise information has an expected payoff of what it thinks it has independent of the other agent's reasoning process or information sources. Thus an agent need not counter-speculate its negotiation partner's beliefs. For the analysis, let $f_a(\breve{a})$ be the contractor's subjective distribution of its best outside offer, and let $f_b(\breve{a})$ be the contractee's subjective distribution of the contractor's best outside offer.

PROPOSITION 6.1 (Payoff Unaffected by Opponent's Beliefs in COP Games). *Say that one agent's information is unbiased, i.e., either $f_a = f$ or $f_b = f$. That agent's expected payoffs for contracts are unaffected by the other agent's subjective distribution. Thus the former agent's preference ordering over contracts is unaffected.*

Proof. Say the contractor's information is unbiased, i.e., $f_a = f$. The contractor's expected payoff for not accepting either contract is $\int_{-\infty}^{\infty} f_a(\breve{a})[-\breve{a}]\, d\breve{a}$. The contractor's payoff for the full commitment contract is $-\rho_F$, and its expected payoff for the leveled commitment contract is

$\int_{-\infty}^{\rho-a} f(\breve{a})[-\breve{a} - a]\,d\breve{a} + \int_{\rho-a}^{\infty} f_a(\breve{a})[-\rho]\,d\breve{a}$. None of these depend on the contractee's information.

Now say that the contractee's information is unbiased, i.e., $f_b = f$. The contractee's payoff for not accepting either contract is $\breve{b}$. Its payoff for the full commitment contract is ρ_F, and its expected payoff for the leveled commitment contract is $[1 - p_a]\rho + p_a[\breve{b} + a] = [1 - (\int_{-\infty}^{\rho-a} f_b(\breve{a})\,d\breve{a})]\rho + (\int_{-\infty}^{\rho-a} f_b(\breve{a})\,d\breve{a})[\breve{b} + a]$. None of these depend on the contractor's information. ■

6.5. *Effect of Biased Asymmetric Beliefs in COBV Games*

Recall that, in COBV games, only one agent has an uncertain future outside offer coming, and the other agent's fallback offer changes before the decommitting game. This section discusses COBV games where the agents' beliefs differ. Specifically, let $f_a(\breve{a})$ be the contractor's subjective distribution of its best outside offer, and let $f_b(\breve{a})$ be the contractee's subjective distribution of the contractor's best outside offer. Let everything else be common knowledge. Now the contractee's *perceived individual rationality* (PIR) constraint is

$$\breve{b} \leq \left[1 - \left(\int_{-\infty}^{\rho-a} f_b(\breve{a})\,d\breve{a}\right)\right]\rho + \left(\int_{-\infty}^{\rho-a} f_b(\breve{a})\,d\breve{a}\right)[\bar{b} + a].$$

Similarly, the contractor's PIR constraint is

$$\int_{-\infty}^{\infty} f_a(\breve{a})[-\breve{a}]\,d\breve{a} \leq \int_{-\infty}^{\rho-a} f_a(\breve{a})[-\breve{a} - a]\,d\breve{a} + \int_{\rho-a}^{\infty} f_a(\breve{a})[-\rho]\,d\breve{a}.$$

The following proposition shows that even though no contract is beneficial to the agents and no full commitment contract seems beneficial, both agents may perceive that some leveled commitment contract is beneficial.

PROPOSITION 6.2 (Perceived Enabling in COBV Games). *There are COBV games (defined by f, f_a, f_b, $\bar{b}$, and $\breve{b}$) where no full commitment contract satisfies both IR constraints, no leveled commitment contract satisfies both IR constraints, and no full commitment contract satisfies both PIR constraints, but some leveled commitment contract satisfies both PIR constraints.*

Proof. A full commitment contract can satisfy the IR constraints iff $\int_{-\infty}^{\infty} f_a(\breve{a})\,\breve{a}\,d\breve{a} \geq \breve{b}$. Now say that

$$f(\breve{a}) = f_a(\breve{a}) = \begin{cases} 0.01 & \text{if } 0 \leq \breve{a} \leq 100 \\ 0 & \text{otherwise} \end{cases} \qquad \text{and} \qquad \breve{b} = 55,$$

i.e., no full commitment contract satisfies both IR constraints. Now let $\bar{b} = 0$. It follows by Proposition 5.1 that no leveled commitment contract satisfies both IR constraints. No full commitment contract satisfies both PIR constraints because it would require $55 = \breve{b} \le \rho_F \le E_a[\breve{a}] = 50$. Now we show a leveled commitment contract that satisfies both PIR constraints. Let

$$f_b(\breve{a}) = \begin{cases} 0.01 & \text{if } 50 \le \breve{a} \le 150 \\ 0 & \text{otherwise}. \end{cases}$$

Now the contractee's PIR constraint is

$$55 \le \left[1 - \left(\int_{-\infty}^{\rho - a} f_b(\breve{a})\, d\breve{a}\right)\right]\rho + \left(\int_{-\infty}^{\rho - a} f_b(\breve{a})\, d\breve{a}\right)[0 + a].$$

Substituting $\rho = 60$, $a = 10$ gives

$$55 \le \left[1 - \left(\int_{-\infty}^{60-10} f_b(\breve{a})\, d\breve{a}\right)\right]60 + \left(\int_{-\infty}^{60-10} f_b(\breve{a})\, d\breve{a}\right)10 \Leftrightarrow 55 \le 60 + 0.$$

The contractor's PIR constraint is

$$-50 \le \int_{-\infty}^{\rho - a} f_a(\breve{a})[-\breve{a} - a]\, d\breve{a} + \int_{\rho - a}^{\infty} f_a(\breve{a})[-\rho]\, d\breve{a}$$

and substituting $\rho = 60$, $a = 10$ gives

$$-50 \le \int_{-\infty}^{60-10} f_a(\breve{a})[-\breve{a} - 10]\, d\breve{a} + \int_{60-10}^{\infty} f_a(\breve{a})[-60]\, d\breve{a}$$
$$\Leftrightarrow -50 \le -17.5 - 30.$$

Thus a leveled commitment contract with $\rho = 60$, $a = 10$ satisfies both PIR constraints. ■

So the agents only *perceive* that this leveled commitment contract satisfies their individual rationality constraints. This is due to the fact that $f_a \ne f_b$; i.e., at least one agent's estimate of the distribution of the contractor's outside offer is biased. On the other hand, if the contractee's fallback payoff is sufficiently low, both agents know (by Proposition 5.1) that the contract cannot really be IR for both. Now which agent is going to incur the loss if the agents agree to the contract that is perceived to be IR by both? The following positive result states that an agent with unbiased beliefs has (as in COP games but unlike in SEQD, SIMUDBP, and SIMUDNP games) an expected payoff of what it thinks it has independent of the other agent's beliefs. Thus the unbiased agent will not enter an unprofitable (non-IR) contract due to the other agent's biases. It also means that agents need not counterspeculate their negotiation partner's beliefs.

PROPOSITION 6.3 (Payoff Unaffected by Opponent's Beliefs in COBV Games). *Say that one agent's information is unbiased, i.e., either $f_a = f$ or $f_b = f$. Now that agent's expected payoffs for contracts are unaffected by the possible biases of the other agent's information. Thus the former agent's preference ordering over contracts and the null deal is unaffected.*

Proof. The contractor's expected payoff for not accepting either contract is $\int_{-\infty}^{\infty} f(\breve{a})[-\breve{a}]\, d\breve{a}$. The contractor's payoff for the full commitment contract is $-\rho_F$, and the expected payoff for the leveled commitment contract is $\int_{-\infty}^{\rho-a} f(\breve{a})[-\breve{a} - a]\, d\breve{a} + \int_{\rho-a}^{\infty} f(\breve{a})[-\rho]\, d\breve{a}$. None of these depend on the contractee's information.

The contractee's payoff for not accepting either contract is $\breve{b}$. His/her payoff for the full commitment contract is ρ_F, and the expected payoff for the leveled commitment contract is $[1 - (\int_{-\infty}^{\rho-a} f(\breve{a})\, d\breve{a})]\rho + (\int_{-\infty}^{\rho-a} f(\breve{a})\, d\breve{a})[\bar{b} + a]$. None of these depend on the contractor's information. ■

COROLLARY 6.1 (Perceived IR Contracts Are IR for the Unbiased Agent in COBV Games). *Say that at most one agent's information is biased, i.e., either $f_a = f$ or $f_b = f$. Say that the contract is perceived IR by the agent x for which $f_x = f$. Now, the contract really is IR for agent x.*

Proof. By definition, a contract is IR for the contractor if it is preferred over the null deal. But by Proposition 6.3 the preference ordering is unaffected by the contractee's information. Similarly, by definition, a contract is IR for the contractee if it is preferred over the null deal. But by Proposition 6.3 the preference ordering is not affected by the contractor's information. ■

It follows that if a contract is perceived IR by both agents, but really is not, the contract is really IR for the agent with unbiased beliefs but not for the agent with biased beliefs.

7. PRACTICAL PRESCRIPTIONS FOR MARKET DESIGNERS

Since we introduced leveled commitment contracts (Sandholm and Lesser, 1995, 1996), they have quickly been adopted into implementation. For example, Mitsubishi has applied them to an electronic market for construction waste recycling in Japan (Akiyoshi *et al.*, 1999). They have also been applied to automated negotiation in a manufacturing setting (Collins *et al.*, 1998) and in a digital library (Park *et al.*, 1996). In this section we present some practical prescriptions for designing negotiation systems that use leveled commitment contracts.

The results from the above canonical games suggest that it is worthwhile from a contract enabling and a contract Pareto-improving perspective to incorporate the decommitment mechanism into automated contracting protocols. The decommitment penalties are best chosen by the agents dynamically at contract time as opposed to statically in the mechanism. This allows the tuning of the penalties not only to specific negotiation situations and environmental uncertainties, but also to the specific belief structures of the agents. In this paper we presented the IR constraints which can be used to evaluate a given contract (price and decommitting penalties) to see whether it is beneficial to both parties.

In another paper we present algorithms for computing the decommitting equilibria given a contract and algorithms for computing contracts (price and decommitting penalties) that maximize the sum of the contract parties' expected payoffs (Sandholm *et al.*, 1999b). As part of *eMediator*, our next generation electronic commerce server, we provide a leveled commitment contract optimizer, *eCommitter*, on the web at http://ecommerce.cs.wustl.edu/contracts.html. It turns out that the surplus provided by leveled commitment can be divided in any way between the contract parties without reducing it, as long as both parties get a nonnegative share of it (Sandholm and Zhou, 2000). This requires that the distributive bargaining occurs over contract price and that the penalties are tailored to that price so as to maximize surplus. On the other hand, if the penalties are set first, some divisions of the surplus reduce the magnitude of the surplus.

Leveled commitment contracts allow an agent to decommit based on local reasoning: no negotiation is necessary at decommitment time. Leveled commitment contracts are simpler than contingency contracts which require, in the worst case, the specification of the contract's alternative obligations for all alternative worlds induced by possible realizations of combinations of future events. Furthermore, the proposed decommitment method does not require an event verification mechanism as contingency contracts do. These advantages, which make leveled commitment contracts more practical for (automated) negotiation systems than contingency contracts, in many settings can come at the cost of Pareto efficiency. If all possible events are known in advance and the events are mutually observable or verifiable, it is, in theory, possible to write a contingency contract that sets the contract obligations optimally for each possible future world induced by different combinations of events. Therefore, in such settings, contingency contracts are *ex post* efficient. Leveled commitment contracts are generally not *ex post* efficient because rational agents decommit strategically and therefore some contracts end up being inefficiently kept although it would be more efficient to breach them. The sum of the contract parties' *ex post* payoffs can be arbitrarily far from optimal because each agent can be made arbitrarily reluctant to decommit by setting its decommitting penalty sufficiently high.

We showed that, for any given contract, the equilibria are different for the three decommitting mechanisms. It turns out that the optimal contracts also differ across these three mechanisms (Sandholm *et al*., 1999b). However, in a recent paper we show that, surprisingly, among risk-neutral agents each of the three mechanisms leads to the same sum of the contract parties' payoffs when the contract price and penalties are optimized for each mechanism separately (Sandholm and Zhou, 2000). Therefore, the sum of the contract parties' payoffs cannot be used as the criterion for choosing between these mechanisms. On the other hand, among agents that are not risk-neutral, the three mechanisms lead to different sums of utilities, and the ranking of the mechanisms varies based on the agents' utility functions.

Because the sum of the contract parties' payoffs cannot be used as the criterion for choosing a mechanism among risk-neutral agents, other criteria are needed. One possible goal is to minimize the number of payment transfers. In another paper we show that if the contract is optimized separately for each of the mechanisms, among risk-neutral agents, the decommitting thresholds will be the same for all of the mechanisms (Sandholm and Zhou, 2000). Therefore, the mechanisms can be fairly compared based on what happens, depending on how the outside offers occur compared to these thresholds. In an outcome where one agent gets an offer that is better than its threshold, but the other agent does not, all of the mechanisms lead to one penalty being paid. If neither agent receives an offer that is better than the agent's threshold, all of the mechanisms lead to no penalty payments. If both agents receive offers that are better than the thresholds, both will decommit in the simultaneous games while only the first agent will decommit in the sequential game. In that case the SIMUDBP mechanism leads to two payments, SEQD leads to one, and SIMUDNP leads to none. One might think that this case gives an unfair advantage to the second mover in the SEQD game. However, as in the simultaneous games, in SEQD games the surplus provided by leveled commitment can be divided arbitrarily between the agents by picking the contract price and tailoring the penalties to that price (Sandholm and Zhou, 2000). In summary, from the perspective of minimizing the number of penalty payments, the SIMUDNP mechanism is best, SEQD is second, and SIMUDBP is worst.

Another mechanism evaluation criterion is the robustness of the equilibrium of the decommitting game. For the sequential decommitting game we were able to use iterated dominance as the solution concept while for the simultaneous decommitting games Nash equilibrium was used. This suggests using sequential decommitting mechanisms because iterated dominance is a more robust solution concept than Nash equilibrium. An additional consideration that promotes sequential decommitting mechanisms is that decommitting is easy at least for one party because the last party to decommit is best off by decommitting nonstrategically if the other(s) did not decommit, and not at all if the other(s) did. Finally, in a recent paper we show that the equilibrium is always unique in sequential decommitting games while multiple equilibria can exist in simultaneous decommitting games (Sandholm *et al*., 1999b). These considerations speak for using sequential decommitting mechanisms.

In asynchronous negotiation systems without a trusted third party, implementing a decommitting mechanism carries the risk that each agent wants to be the last to reveal its decommitting decision. This is because if the other agent decommits, the last agent can always say that it does not want to decommit. This causes the last agent not to have to pay a penalty and to collect a penalty from the other agent. This problem can be overcome using the following mechanism.[5] Each agent encrypts its decommitting decision and mails it to the other contract parties. Once an agent has received an encrypted message from every other party, it sends out its key to the others so they can observe the agent's decommitting decision. Under this scheme the others cannot read the agent's decision before they have committed to their own decisions. Similarly, the agent cannot change its decision.[6]

In a web of multiple mutual contracts among multiple agents, classical full commitment contracts induce one negotiation focus consisting of the obligations of the contracts. Under the mechanism proposed in this paper, there are multiple such foci and any agent involved in a contract can switch from one such focus to another by decommitting from a contract by paying the decommitment penalty. It may happen that such a switch makes it beneficial for another agent to decommit from another contract, and so on. To avoid loops of decommitting and recommitting in practice, recommitting could possibly be disabled. This could be implemented by choosing a mechanism that specifies that if a contract offer is accepted and later either agent decommits, the original offer becomes void as opposed to

[5] Similar mechanisms have been used for other purposes; see e.g. Zlotkin and Rosenschein (1994) and Sandholm *et al*. (1999a).

[6] The agent could manipulate this scheme if the agent were capable of tailoring its key *ex post* so that the agent's encrypted message would change meaning sensibly from "I want to decommit" to "I do not want to decommit." In cryptography such tailoring is believed to be difficult.

staying valid according to its original deadline which might not have been reached at the time of decommitment.

Even though two agents cannot explicitly recommit to a contract, it is hard to specify and monitor in a mechanism that they will not make another contract with identical content. This gives rise to the possibility of the equivalent of useless decommit–recommit loops. Such loops can lead to extended negotiation due to deep backtracking and infinite loops (Andersson and Sandholm, 1998). This can be avoided by a mechanism where the decommitment penalties increase with time (real time or a number of domain events or negotiation events). This allows a low commitment negotiation focus to be moved in the joint search space while still making the contracts meaningful by some level of commitment. The increasing level of commitment causes the agents to not backtrack very deeply in the negotiations, which can also save computation. Recent research shows that in the clock that is used as the basis of price increments, at least some element of absolute time needs to be used: to avoid infinite loops among myopic agents, it does not suffice to count time from the moment when each contract is made (Andersson and Sandholm, 1998). Similar issues for agents that carry out strategic lookahead were also recently studied (Andersson and Sandholm, 2001).

The initially low commitment to contracts can also be used as a mechanism to facilitate linking of deals. Often, there is no contract over a single item that is beneficial, but a combination of contracts among two agents would be (Sandholm, 1993; Sandholm and Lesser, 1995). Even if explicit clustering of issues into contracts (Sandholm, 1996; Sandholm, 1993; Sandholm and Lesser, 1995) is not used, an agent can agree to an initially unbeneficial low commitment contract in anticipation of synergic future contracts from the other agent that will make the first contract beneficial (Sandholm and Lesser, 1995). If no such contracts appear, the agent can decommit. In a similar way the initially low commitment to contracts can be used as a mechanism to facilitate contracts among more than two agents. Even without explicit multiagent contract mechanisms (Sandholm and Lesser, 1995), multiagent contracts can be implemented by one agent agreeing to an initially unbeneficial low commitment contract in anticipation of synergic future contracts from third parties that will make the first contract beneficial (Sandholm and Lesser, 1995). Again, if no such contracts appear, the agent can decommit.

Even if explicit clustering of issues into atomic contracts and/or multiagent contracts allowed, using leveled commitment mechanisms on top of such contracts can be beneficial. For example, identifying profitable combinatorial contracts can be difficult computationally and due to lack of a global view. In such settings an agent may be better off trying to construct profitable combinations from sequences of individual contracts, each with leveled commitment. The best contracting mechanisms will probably involve both mechanisms: explicit linking of issues and leveled commitment. We recently compared combinations of these mechanisms via a simulation study (Andersson and Sandholm, 1998). We also developed the theory of leveled commitment contracts involving more than two contract parties in a single contract (Sandholm *et al.*, 1999b).

In many practical automated contracting settings, agents are bounded rational—for example, because limited computation resources bound their capability to solve combinatorially complex problems (Sandholm, 1993; Sandholm and Lesser, 1995, 1997; Larson and Sandholm, 2000). The very fact that an agent's computation is bounded induces uncertainty. For example, the value of a contract may only be probabilistically known to the agent at contract time. The leveled commitment contracting mechanism allows the agent to continue deliberation regarding the value of the contract after the contract is made. If the value of the contract turns out to be lower than expected, the agent can decommit. On the other hand, a leveled commitment contracting mechanism where the decommitment penalties increase quickly in time may be appropriate with bounded rational agents so that the agents do not need to consider the combinatorial number of possible future worlds where alternative combinations of decommitments have occurred (Sandholm and Lesser, 1995).

Multi-item auctions where the bidders have preferences over combinations of items are one important potential application of leveled commitment contracts. It is known that by certain combinatorial auction designs one can obtain efficient outcomes in dominant strategy equilibrium; see, e.g., Sandholm (2000) and Monderer and Tennenholtz (1999). However, this can require a bidder to compute its valuation for each combination of items and to bid for each combination. Also, the auctioneer's task of determining the winners is computationally complex (Sandholm, 1999; Sandholm and Suri, 2000). A potentially more practical alternative is to use sequential or ascending auctions with bidding on individual items or restricted combinations only. Leveled commitment contracts could be used as a mechanism for bidders to put back items if they do not (or if they project that they will not) obtain the combinations that they want. Similarly, the auctioneer may want to exercise a take-back—for example if it receives a better bid later. Decommitting could be allowed during the auction process or even after (some of) the items have been allocated.[7]

[7] In the Federal Communications Commission's bandwidth auction the bidders were allowed to retract their bids (McAfee and McMillan, 1996). In case of a retraction, the item was opened for reauction. If the new winning price was lower than the old one, the bidder that retracted the bid had to pay the difference. This guarantees that retractions do not decrease the auctioneer's payoff. However, it exposes the retracting bidder to considerable risk because he does not know the penalty when decommitting. The mechanism also does not allow the auctioneer to take back items.

8. CONCLUSIONS AND FUTURE RESEARCH

We presented a decommitting mechanism for contracting protocols that allows the agents to accommodate future events profitably, unlike full commitment contracts.

The contract specifies decommitment penalties for both parties. To decommit, an agent just pays that penalty to the other agent. This mechanism is more practical than a contingency contract, especially for complex computerized contracting settings: potentially combinatorial and hard to anticipate contingencies need not be considered and no event verification mechanism is necessary. Decommitting can be decided based on local *ex post* deliberation.

One concern is that a rational agent would be reluctant to decommit because there is a chance that the other party will decommit. In this case the former agent gets freed from the contract, does not have to pay a penalty, and collects a penalty from the breacher. Since both agents will be reluctant in this way, some contracts will be inefficiently kept. Via a noncooperative equilibrium analysis of the decommitting games, we showed that despite this, leveled commitment contracts can enable deals in settings where no full commitment contract is beneficial for both parties. We also showed conditions under which a leveled commitment contract can be constructed that yields higher expected payoffs to both contract parties than any full commitment contract. Specifically, such a contract can be constructed if there is some chance that the contractor's outside offer is lower than the contractee's expected outside offer, or some chance that the contractee's outside offer is higher than the contractor's expected outside offer. Furthermore, for any full commitment contract a leveled commitment contract can always be constructed that is no worse than the full commitment contract (assuming that the contractor's outside offer is bounded from below and the contractee's from above). This is achieved trivially by setting the decommitting penalties so high that neither party will decommit.

We showed these results for six decommitting mechanisms which differ based on whether the decommitting decisions have to be revealed sequentially (contractor first or contractee first) or simultaneously, and whether or not the agents have to pay the penalties if both decommit. In the case of sequential decommitting, the latter distinction does not matter because the second mover will never decommit if the first mover does. So, four distinct mechanisms are left. Of the sequential mechanisms, we studied the one where the contractee has to reveal its decommitting decision first (the one where the contractor goes first is analogous). Therefore, there were in essence three mechanisms to study. For each one, the decommitting equilibrium is different, but the advantages hold under the same conditions.

In addition to the quantitative results presented in this paper, several qualitative observations were made about the decommitting equilibrium. First, in sequential mechanisms, the last agent to decommit is best off by decommitting nonstrategically if the other(s) did not decommit and not at all if the other(s) did. Second, in the simultaneous decommitting mechanism where both pay the penalties to each other if both decommit, as an agent's decommitting penalty approaches zero the agent becomes nonstrategic. On the other hand, in the simultaneous decommitting mechanism where neither pays a penalty if both decommit, as an agent's decommitting penalty approaches zero the other agent becomes nonstrategic. Third, in both simultaneous mechanisms, as an agent's penalty increases so that the agent's decommitting probability approaches zero, the other agent becomes nonstrategic. Fourth, for a given contract and probability that the other contract party will decommit, a rational agent is most reluctant to decommit (i.e., most strategic) in the sequential mechanism (the first mover), next reluctant in the simultaneous mechanism where both pay the penalties to each other if both decommit, and least reluctant in the simultaneous mechanism where neither pays a penalty if both decommit.

Clearly it is also possible to construct game instances where the agents' fallback positions are so good that no deal is possible between them, not even a leveled commitment contract. Also, if an agent's fallback position decreases too much between the time of contracting and the time of decommitting, the agent would not agree to a leveled commitment contract.

We also discussed the decommitting mechanisms comparatively. From the perspective of minimizing the number of penalty payments, the simultaneous decommitting mechanism where neither pays a penalty if both decommit is best, the sequential mechanism is second, and the simultaneous mechanism where both pay if both decommit is worst. From the perspective of robustness and uniqueness of the equilibrium and ease of rational play, the sequential mechanism is best.

In addition to studying games where both agents' future involves uncertainty, we analyzed games with one-sided uncertainty. In those games the agent with a certain future prefers not to decommit if the contract was originally individually rational to that agent. Thus only one agent may want to decommit. In those games an agent's payoff from a contract is unaffected by the other agent's beliefs. Therefore the preference order over contracts is unaffected by the other agent's possibly biased beliefs. It follows that an agent need not counterspeculate its negotiation partner's beliefs and that an agent cannot incur a loss due to the other agent's erroneous beliefs. On the other hand, in the games where both agents' future involves uncertainty, an agent's payoff from a contract may depend on the negotiation partner's possibly biased beliefs. Therefore an agent may need to counterspeculate the other agent's beliefs.

While we presented the uncertainty as stemming from potential future outside offers, the model can be applied to any setting where the uncertainty stems from fluctuations in the worth of the contract to the contract parties.

Options are another technique for capitalizing on the gains that probabilistically known future events provide. In an option, the option price is paid whether or not the option is exercised. In leveled commitment contracts it is paid only if exercised. Also, leveled commitment contacts provide both (all) contract parties the possibility of decommitting while a basic option allows for flexibility only for one contract party. While it is possible to construct packages of options, we do not currently know of a way of emulating a leveled commitment contract with any package of options. The fact that leveled commitment contracts lead to a strategic decommitting game while options do not suggests that leveled commitment contracts are a fundamentally different new financial instrument.

Extensions of this research include studying more closely the best pace for increasing the decommitment penalties with time or with occurring events. A normative theory relating deliberation control to the issues of this paper is also desirable. We have already taken steps toward devising normative theories of deliberation in other games (Sandholm and Lesser, 1997; Larson and Sandholm, 2000). Finally, the relationship between leveled commitment contracting and explicit multi-issue and multiagent contracts should be studied in more detail.

REFERENCES

Akiyoshi, M., Abe, K., Ono, T., and Tabinoki, Y. (1999). "Agent-Based Decision Support Framework for Asynchronous Negotiation of Resource Transfer," in *Proceedings of the IEEE Systems, Man, and Cybernetics Conference*, Tokyo, Japan.

Andersson, M. R., and Sandholm, T. W. (1998). "Leveled Commitment Contracting among Myopic Individually Rational Agents," in *Proceedings of the Third International Conference on Multi-agent Systems (ICMAS), Paris*, pp. 26–33.

Andersson, M. R., and Sandholm, T. W. (1999). "Time–Quality Tradeoffs in Reallocative Negotiation with Combinatorial Contract Types," in *Proceedings of the National Conference on Artificial Intelligence (AAAI), Orlando, FL*, pp. 3–10.

Andersson, M. R., and Sandholm, T. W. (2001). "Leveled Commitment Contracts with Myopic and Strategic Agents," *J. Econ. Dynamics and Control* (Special Issue on Agent-Based Computational Economics), **25**, 615–640. [(1998) An early version was presented at the *National Conference on Artificial Intelligence (AAAI), Madison, WI*, pp. 38–45].

Calamari, J. D., and Perillo, J. M. (1977). *The Law of Contracts*, 2nd ed. West Publishing.

Cheng, J. Q., and Wellman, M. P. (1998). "The WALRAS Algorithm: A Convergent Distributed Implementation of General Equilibrium Outcomes," *Comput. Econ.* **12**, 1–24.

Choi, S.-Y., Stahl, D. O., and Whinston, A. B. (1997). *The Economics of Electronic Commerce*. MacMillan Technical.

Collins, J., Youngdahl, B., Jamison, S., Mobasher, B., and Gini, M. (1998). "A Market Architecture for Multi-Agent Contracting," in *Proceedings of the Second International Conference on Autonomous Agents (AGENTS), Minneapolis/St. Paul, MN*, pp. 285–292.

Decker, K., and Lesser, V. R. (1995). "Designing a Family of Coordination Algorithms," in *Proceedings of the First International Conference on Multi-Agent Systems (ICMAS), San Francisco, CA*, pp. 73–80.

Diamond, P. A., and Maskin, E. (1979). "An Equilibrium Analysis of Search and Breach of Contract. I. Steady States," *Bell J. Econ.* **10**, 282–316.

Ephrati, E., and Rosenschein, J. S. (1991). "The Clarke Tax as a Consensus Mechanism among Automated Agents," in *Proceedings of the National Conference on Artificial Intelligence (AAAI), Anaheim, CA*, pp. 173–178.

Kraus, S. (1993). "Agents Contracting Tasks in Noncollaborative Environments," in *Proceedings, Eleventh National Conference on Artificial Intelligence (AAAI-93)*, pp. 243–248.

Kraus, S., Wilkenfeld, J., and Zlotkin, G. (1995). "Multiagent Negotiation under Time Constraints," *Artificial Intelligence* **75**, 297–345.

Larson, K., and Sandholm, T. W. (2000). "Deliberation in Equilibrium: Bargaining in Computationally Complex Problems," in *Proceedings of the National Conference on Artificial Intelligence (AAAI), Austin, TX*, pp. 48–55.

Low, S. H., Maxemchuk, N. F., and Paul, S. (1996). "Anonymous Credit Cards and Their Collusion Analysis," *IEEE/ACM Trans. Networking* **4**(6).

Mas-Colell, A., Whinston, M., and Green, J. R. (1995). *Microeconomic Theory*. Oxford, UK: Oxford Univ. Press.

McAfee, R. P., and McMillan, J. (1996). "Analyzing the Airwaves Auction," *J. Econ. Perspectives* **10**(1), 159–175.

Monderer, D., and Tennenholtz, M. (1999). "Asymptotically Optimal Multi-object Auctions for Risk-Averse Agents," Technical report. Faculty of Industrial Engineering and Management, Technion, Haifa, Israel.

Nash, J. (1950). "The Bargaining Problem," *Econometrica* **18**, 155–162.

Osborne, M. J., and Rubinstein, A. (1990). *Bargaining and Markets*. New York: Academic Press.

Park, S., Durfee, E., and Birmingham, W. (1996). "Advantages of Strategic Thinking in Multiagent Contracts (A Mechanism and Analysis)," in *Proceedings of the Second International Conference on Multi-agent Systems (ICMAS), Keihanna Plaza, Kyoto, Japan*, pp. 259–266.

Posner, R. A. (1977). *Economic Analysis of Law*, 2nd ed. New York: Little, Brown, and Co.

Raiffa, H. (1982). *The Art and Science of Negotiation*. Cambridge, MA: Harvard Univ. Press.

Rosenschein, J. S., and Zlotkin, G. (1994). *Rules of Encounter: Designing Conventions for Automated Negotiation among Computers*. Cambridge, MA: MIT Press.

Rubinstein, A. (1982). "Perfect Equilibrium in a Bargaining Model," *Econometrica* **50**, 97–109.

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), Washington, D.C.*, pp. 256–262.

Sandholm, T. W. (1996). *Negotiation among Self-Interested Computationally Limited Agents*, Ph.D. thesis. University of Massachusetts, Amherst [Available at http://www.cs.wustl.edu/~sandholm/dissertation.ps].

Sandholm, T. W. (1997). "Unenforced E-Commerce Transactions," *IEEE Internet Computing* **1**(6, Special Issue on Electronic Commerce), 47–54.

Sandholm, T. W. (1999). "An Algorithm for Optimal Winner Determination in Combinatorial Auctions," in *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence (IJCAI), Stockholm, Sweden*, pp. 542–547 [An extended version first appeared as Tech. Rep. WUCS-99-01, Washington Univ., Dept. of Computer Science, Jan. 28th, 1999].

Sandholm, T. W. (2000). "eMediator: A Next Generation Electronic Commerce Server," in *Proceedings of the Fourth International Conference on Autonomous Agents (AGENTS), Barcelona, Spain*, pp. 73–96 [An early version appeared in the *AAAI-99 Workshop on AI in Electronic Commerce, Orlando, FL*, pp. 46–55, July 1999, and as Tech. Rep. WU-CS-99-02, Dept. of Computer Science, Washington University, St. Louis, MO, Jan. 1999].

Sandholm, T. W., Larson, K. S., Andersson, M. R., Shehory, O., and Tohmé, F. (1999a). "Coalition Structure Generation with Worst Case Guarantees," *Artificial Intelligence* **111**(1/2), 209–238 [An early version appeared at the *National Conference on Artificial Intelligence (AAAI), 1998*, pp. 46–53].

Sandholm, T. W., and Lesser, V. R. (1995). "Issues in Automated Negotiation and Electronic Commerce: Extending the Contract Net Framework," in *Proceedings of the First International Conference on Multi-agent Systems (ICMAS), San Francisco, CA*, pp. 328–335 [Reprinted in *Readings in Agents* (Huhns and Singh, Eds.), pp. 66–73. 1997].

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), Portland, OR*, pp. 126–133 [An extended version appeared as Tech. Rep. 95-72, University of Massachusetts at Amherst, Computer Science Department].

Sandholm, T. W., and Lesser, V. R. (1997). "Coalitions among Computationally Bounded Agents," *Artificial Intelligence* **94** (1, Special issue on Economic Principles of Multiagent Systems), 99–137 [An early version appeared at the *International Joint Conference on Artificial Intelligence, 1995*, pp. 662–669].

Sandholm, T. W., Sikka, S., and Norden, S. (1999b). "Algorithms for Optimizing Leveled Commitment Contracts," in *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence (IJCAI), Stockholm, Sweden*, pp. 535–540 [An extended version appeared as Tech. Rep. WUCS-99-04, Washington University, Department of Computer Science].

Sandholm, T. W., and Suri, S. (2000). "Improved Algorithms for Optimal Winner Determination in Combinatorial Auctions and Generalizations," in *Proceedings of the National Conference on Artificial Intelligence (AAAI), Austin, TX*, pp. 90–97.

Sandholm, T. W., and Vulkan, N. (1999). "Bargaining with Deadlines," in *Proceedings of the National Conference on Artificial Intelligence (AAAI), Orlando, FL*, pp. 44–51.

Sandholm, T. W., and Zhou, Y. (2000). "Surplus Equivalence of Leveled Commitment Contracts," in *Proceedings of the Fourth International Conference on Multi-agent Systems (ICMAS), Boston, MA*, pp. 247–254.

Sen, S., and Durfee, E. (1994). "The Role of Commitment in Cooperative Negotiation," *Internat. J. Intelligent Cooperative Information Systems* **3**(1), 67–81.

Sen, S., and Durfee, E. (1998). "A Formal Study of Distributed Meeting Scheduling," *Group Decision and Negotiation* **7**, 265–289.

Smith, R. G. (1980). "The Contract Net Protocol: High-Level Communication and Control in a Distributed Problem Solver," *IEEE Trans. Computers* **29**(12), 1104–1113.

Yokoo, M., Durfee, E. H., and Ishida, T. (1992). "Distributed Constraint Satisfaction for Formalizing Distributed Problem Solving," in *Proceedings of the 12th IEEE International Conference on Distributed Computing Systems*, pp. 614–621.

Zlotkin, G., and Rosenschein, J. S. (1994). "Coalition, Cryptography and Stability: Mechanisms for Coalition Formation in Task Oriented Domains," in *Proceedings of the National Conference on Artificial Intelligence (AAAI), Seattle, WA*, pp. 432–437.