Leveled Commitment Contracts

Tuomas W. Sandholm *
sandholm@cs.wustl.edu
Washington University
Department of Computer Science
One Brookings Drive
St. Louis, MO 63130-4899

Victor R. Lesser
lesser@cs.umass.edu
University of Massachusetts
Department of Computer Science
Amherst, MA 01003

Abstract

In negotiation systems consisting of self-interested agents, contracts have traditionally been binding. Such contracts do not allow agents to capitalize on uncertain future events. Contingency contracts have been proposed to solve this problem. Contingency contracts are often impractical due to large numbers of interdependent and unanticipated future events on which to condition, and because some events are not mutually observable. We propose a leveled commitment contracting protocol that allows agents to capitalize on uncertain future events by having the possibility of unilaterally decommitting from a contract based on local reasoning. Decommitment penalties are assigned to both agents in a contract: to be freed from the obligations of the contract, an agent only pays the penalty to the other party. One concern is that a self-interested 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. It is shown through noncooperative equilibrium analysis that the leveled commitment feature in a contracting protocol increases the expected payoff to both contract parties and can make contracts individually rational when no full commitment contract can. Asymmetric beliefs are also discussed, and different decommitting protocols are compared. Practical prescriptions for market designers are provided. A leveled commitment contract optimizer is offered on the web.[1,2]

1 Introduction

The importance of automated negotiation systems with self-interested agents is increasing. One reason for this is the technology push of a growing standardized communication infrastructure—the Internet, WWW, EDI, KQML, FIPA, Java, Odyssey,

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

[1]Patent pending.

[2]A short early version of this paper appeared at the National Conference on Artificial Intelligence (AAAI) [34].

Voyager, Concordia, Telescript, 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 [27, 32, 14]. The second reason is 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 commerce on the Internet for purchasing goods, services, information, communication bandwidth, *etc.* [7]. 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, but 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 protocols 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 [23, 25, 28, 2, 11, 12, 13, 38, 6]. 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 [42], 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 [40, 41] and in cooperative coordination protocols [9].

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 protocols need to be considered from the perspective of noncooperative game theory: given a protocol (mechanism), what is the best strategy from a self-interested viewpoint that each agent can choose, and then what social outcomes 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 [22]. 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 regarding the use of contingency contracts in automated negotiation among self-interested agents. 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 agent may be a prohibitively laborious and error prone activity.

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, and need therefore 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 [33, 25, 23]. Thus there is a potential combinatorial explosion of possible future worlds, and each of them may need to be associated with a different contingency.

The third problem is that of verifying the unraveling of the events. Sometimes an event is only observable by one of the agents. That agent may have an incentive to lie to the other contract party about the event in case the event is associated with a disadvantageous contingency to the observing agent. 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 into contracting protocols.

Principles for assessing decommitment penalties have been studied in the economics of law [5, 21], but the purpose has been to assess a penalty on the agent that has breached the contract *after the breach has occurred*. Similarly, penalty clauses for partial failure—such as not meeting a deadline—are commonly used in contracts, but the purpose is usually to motivate the agents to 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.[3] Somewhat unintuitively, it turns out that the decommitting possibility in a contract can increase the expected payoff for all contract parties.

1.3 Practical motivations for leveled commitment

The goal of our leveled commitment contracting protocol 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.[4] 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.
- 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 to an agent, that agent can decommit from those old contracts.
- 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 (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 the tasks in that set further 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 protocol, 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. We have shown that in task allocation, contracts of a single task at a time usually lead to local optima, and that this problem can be addressed by cluster contracts

[3] 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 [10].

[4] 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 [26, 33]. The analysis of this paper focuses on one contract where the penalties are negotiated up front, and they do not change over time.

(where multiple tasks are negotiated over atomically), swap contracts (where tasks are swapped between agents), and multiagent contracts (where a contract has more than two parties) [28, 26, 2]. The leveled commitment protocol allows any of these composite contracts to be constructed from a sequence of basic contracts. For example, say that the only profitable contract is one where agent A gives task t_1 to agent B, and agent C gives task t_2 to agent A. Now A can make the unprofitable contract with B in anticipation of the contract with C which will make the combination profitable. Then, if C does not agree to the contract with A, agent A can decommit from the contract that it made with B.

- It makes local feasibility checks unnecessary. When bidding, an agent traditionally has to make sure that it can handle all of its obligations even if all of its pending bids—and the bid being constructed—get accepted. Such feasibility checks often use a major portion of a contracting software agent's deliberation resources [25, 26]. With leveled commitment protocols, agents need not carry out feasibility checks up front because if an agent ends up overcommitted, it can decommit from some of the contracts.

- It makes each agent's local deliberation control more flexible because marginal cost calculation of a contract [25, 26, 33] can go on even after the contract has already been agreed upon. If the contract turns out unbeneficial in the extended deliberation, the agent can decommit. This allows the agent to carry out only a rough marginal cost computation before bidding, and allows the agent to invest the computation resources for more thorough marginal cost calculation only if the agent wins the bidding. From a societal perspective, this reduces computation overhead.

- It can speed up the negotiation process. For example, an agent can make the same low-commitment offer (or offers that overlap in task sets) to multiple agents. In case more than one accepts, the agent has to pay the penalty to all but one of them, but the negotiation process speedup from being able to address multiple agents in committal mode may outweigh this risk.

- It allows the less risk averse agents to carry a greater portion of the risk. A more risk averse agent can trade off paying a higher price to its contractee (or get paid a lower price as a contractee) for being allowed to have a lower decommitting penalty. This increases the Pareto efficiency of the contracts.

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, its 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 social welfare (the sum of the payoffs of the contract parties) when viewed *ex post*.

Second, one might think that full commitment contracts never have higher social welfare than leveled commitment contracts because the latter incorporate new information (new events), and according to decision theory information never has negative value. However, the result that information never has negative value only holds in single agent decision making. In multiagent systems, information can have negative 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 1. Each agent's dominant strategy is to defect, so the social welfare will be $1 + 1 = 2$.

		column player	
		cooperate	defect
row	cooperate	3, 3	0, 5
player	defect	5, 0	1, 1

Table 1: *Prisoner's dilemma game. In each square, the row player's payoff is listed first.*

This is also the only solution 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 with an expected social welfare of $\frac{1}{4}(3+3)+\frac{1}{4}(0+5)+\frac{1}{4}(5+0)+\frac{1}{4}(1+1) = 4.5$. So, the 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 social welfare or Pareto efficiency than full commitment contracts.

Third, agents might decommit insincerely. A truthful 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 in decommitting. It can take into account the chance that the other party will decommit, in which case the former agent gets freed from the contract obligations, does not have to pay a decommitting penalty, and will collect a decommitting penalty from the other party. Due to such reluctant decommitting, contracts may end up being inefficiently kept even though breaking them would be best from the social welfare perspective.

In the rest of this paper, we formally show the nonobvious fact that leveled commitment contracts are indeed superior to full commitment contracts even though agents decommit insincerely. In Section 2 we present our main models and analysis of leveled commitment contracts. We show that leveled commitment contracts enable contracts and improve their efficiency. 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 protocols. Section 3 discusses settings where only one agent's future involves uncertain events. In Section 4 we discuss games where the fallback positions change for the worse before the decommitting game. Section 5 presents what happens when an agent thinks that the distributions of future events are common knowledge and acts according to the formulas derived under that assumption, but in reality at least one of the agents has biased beliefs. In Section 6, practical prescriptions are given for market designers. Section 7 concludes and presents future research directions.

2 Comparing leveled and full commitment contracts

We analyze contracting situations from the perspective of two risk neutral agents each of which attempts to maximize his own expected payoff: the *contractor* who pays to get a task done, and the *contractee* who gets paid for handling the task. Handling a task can mean taking on any types of obligations. The method is not specific to classical task allocation. The contractor tries to minimize the contract price, ρ, that he has to pay. The contractee tries to maximize the payoff, ρ, that she receives from the contractor.

Outside offers from third parties will be explicitly discussed. The contracting setting consists of two games. First, the *contracting game* involves the agents choosing a contract—or no contract, i.e. the null deal—before any future events have unraveled. Second, the *decommitting game* involves the agents deciding on whether to decommit or to carry out the obligations of the contract—after the future events have unraveled. The decommitment game is a subgame of the contracting game: the 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 [15]. The contracting game will be analyzed with respect to individual rationality: is the contract better for an agent than the null deal given that the decommitting game will be played in equilibrium?

Often there is either no contract that is individually rational (IR) for both agents or then 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. 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, deadlines, or bargaining costs into the bargaining process [37, 24] or by asserting desirable properties that the chosen contract should fulfill compared to other contracts [18, 19]. In this paper, we will not address the choice among IR contracts. In another paper we present algorithms for computing the welfare maximizing IR contracts (contract price and penalties) [36].

In this section we study a contracting setting where the future of both agents involves uncertainty. Specifically, the agents might receive full commitment outside offers.[5] The contractor's best outside offer $\breve{a}$ is only probabilistically known *ex ante* by both agents, and is characterized by a probability density function $f(\breve{a})$. If the contractor does not receive an outside offer, $\breve{a}$ corresponds to its best outstanding outside offer or its fallback payoff, i.e. payoff that it receives if no contract is made. The contractee's best outside offer $\breve{b}$ is also only probabilistically known *ex ante*, and is characterized by a probability density function $g(\breve{b})$. 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 statistically independent. The distributions $f(\breve{a})$ and $g(\breve{b})$ are assumed to be common knowledge between the contractor and the contractee.

The contractor's options are either to make a contract with the contractee or to wait for $\breve{a}$. Similarly, the contractee's options are either to make a contract with the contractor or to wait for $\breve{b}$. If the two agents decide to make a contract, they can choose a leveled commitment contract or a full commitment contract. A leveled commitment contract is specified by the contract price ρ, the contractor's decommitment penalty a, and the contractee's decommitment penalty b. A full commitment contract is specified by the contract price only: there is no possibility to decommit. Under a leveled commitment contract, the contractor has to decide on decommitting when it knows its outside offer, $\breve{a}$, but does not know the contractee's outside offer, $\breve{b}$. Similarly, the contractee has to decide on decommitting when it knows its outside offer, $\breve{b}$, but does not know the contractor's outside offer, $\breve{a}$. This seems realistic from a practical automated contracting perspective.

ρ	Contract price.
$a \geq 0$	Contractor's decommitment penalty.
$b \geq 0$	Contractee'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 that the contractor decommits.
p_b	Probability that the contractee decommits.

Table 2: *Symbols used in this paper.*

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 another paper we show that this restriction is not crucial because if there exists a social welfare maximizing IR contract, then there exists one with $a \geq 0$ and $b \geq 0$ [39].

We do not assume that the agents decommit truthfully. For example, an agent may not decommit although its outside offer is better (even after paying the decom-

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

mitment penalty) for itself than the contract because the agent believes that there is a high probability that the partner will decommit. This would save the agent its decommitment penalty and make the agent receive a decommitment penalty from the partner. 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 the same 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 (SEQD)

In our sequential decommitting (SEQD) game, one agent has to declare its decommitment decision before the other agent. We will study the case where the contractee has to decommit first. The case where the contractor has to decommit first is analogous. The game tree is presented in Figure 1. There are two alternative types of leveled commitment contracts that differ on what happens if both agents decommit. In the first, both agents have to pay the decommitment penalties to each other 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 backwards to the beginning of the game. In the subgame where the contractee has decommitted, the contractor's best move is not to decommit because $-\breve{a} - a + b \leq -\breve{a} + b$ (because $a \geq 0$). This also holds for a contract where neither agent has to pay a decommitment penalty if both decommit—because $-\breve{a} \leq -\breve{a} + b$, (Figure 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

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

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

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

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

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

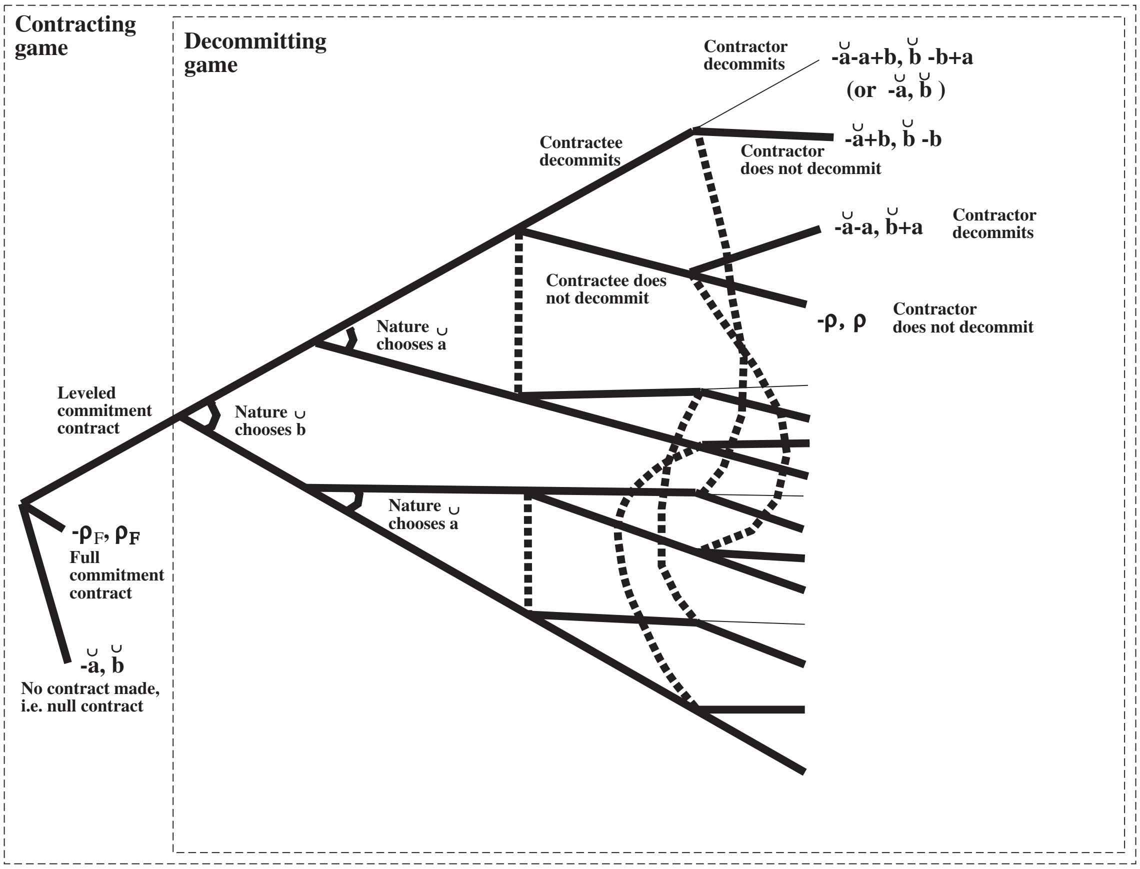

Figure 1: *The "SEQuential Decommitting" (SEQD) game. The game tree represents two alternative protocols, i.e. two different games. In the first, both agents have to pay the decommitment penalties to each other if both decommit. In the second, neither agent has to pay if both decommit. The payoffs of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

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:

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

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

$$\begin{aligned}\pi_a &= \int_{\breve{b}^*(\rho,a,b)}^{\infty} g(\breve{b}) \int_{-\infty}^{\infty} f(\breve{a})[-\breve{a}+b]d\breve{a}d\breve{b} \\ &+ \int_{-\infty}^{\breve{b}^*(\rho,a,b)} g(\breve{b})[\int_{-\infty}^{\rho-a} f(\breve{a})[-\breve{a}-a]d\breve{a} + \int_{\rho-a}^{\infty} f(\breve{a})[-\rho]d\breve{a}]d\breve{b}\end{aligned}$$

$$\geq \quad E[-\breve{a}] = \int_{-\infty}^{\infty} f(\breve{a})[-\breve{a}]d\breve{a} \tag{4}$$

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, full commitment contracts are a subset of leveled commitment ones. This reasoning holds for contracts where both agents have to pay the penalties if both decommit, and for contracts where neither agent has to pay a penalty if both decommit. Because full commitment contracts are a subset of leveled commitment contracts, the former can be no better in the sense of Pareto efficiency or social welfare than the latter. It follows that if there exists an IR full commitment contract, then there also exist IR leveled commitment contracts.

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 $f(\breve{a})$ and $g(\breve{b})$) where no full commitment contract satisfies the IR constraints but a leveled commitment contract does.*

Proof. Let $f(\breve{a}) = \begin{cases} \frac{1}{100} & \text{if } 0 \leq \breve{a} \leq 100 \\ 0 & \text{otherwise} \end{cases}$ and $g(\breve{b}) = \begin{cases} \frac{1}{110} & \text{if } 0 \leq \breve{b} \leq 110 \\ 0 & \text{otherwise.} \end{cases}$ Now a full commitment contract F does not satisfy both IR constraints because that would require $E[\breve{b}] \leq \rho_F \leq E[\breve{a}]$ which is impossible because $55 = E[\breve{b}] > E[\breve{a}] = 50$. 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

$$\begin{aligned}
& \int_{\breve{b}^*(\rho,a,b)}^{\infty} g(\breve{b}) \int_{-\infty}^{\infty} f(\breve{a})[-\breve{a}+b]d\breve{a}d\breve{b} \\
& + \int_{-\infty}^{\breve{b}^*(\rho,a,b)} g(\breve{b})[\int_{-\infty}^{\rho-a} f(\breve{a})[-\breve{a}-a]d\breve{a} + \int_{\rho-a}^{\infty} f(\breve{a})[-\rho]d\breve{a}]d\breve{b} \\
\geq \quad & \int_{-\infty}^{\infty} f(\breve{a})[-\breve{a}] \\
\Leftrightarrow \quad & \int_{\breve{b}^*(\rho,a,b)}^{110} g(\breve{b}) \int_{0}^{100} f(\breve{a})[-\breve{a}+20]d\breve{a}d\breve{b} \\
& + \int_{0}^{\breve{b}^*(\rho,a,b)} g(\breve{b})[\int_{0}^{52.5-30} f(\breve{a})[-\breve{a}-30]d\breve{a} + \int_{52.5-30}^{100} f(\breve{a})[-52.5]d\breve{a}]d\breve{b} \\
\geq \quad & \int_{0}^{100} f(\breve{a})[-\breve{a}] \\
\Leftrightarrow \quad & \int_{\breve{b}^*(\rho,a,b)}^{110} \frac{1}{110}\frac{1}{100}[\frac{-(100)^2}{2} + 100 \cdot 20]d\breve{b} +
\end{aligned}$$

$$
\begin{aligned}
& \int_0^{\breve{b}^*(\rho,a,b)} \frac{1}{110}[\frac{1}{100}[\frac{-(22.5)^2}{2} - 22.5 \cdot 30] + \frac{1}{100}[-52.5 \cdot (100-22.5)]]d\breve{b} \\
\geq \; & -50 \\
\Leftrightarrow \quad & \int_{\breve{b}^*(\rho,a,b)}^{110} \frac{15.5}{110} d\breve{b} + \int_0^{\breve{b}^*(\rho,a,b)} -\frac{49.96875}{110} d\breve{b} \geq -50
\end{aligned}
$$

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

$$
\begin{aligned}
& \int_{\breve{b}^*(\rho,a,b)}^{\infty} g(\breve{b})[\breve{b} - b]d\breve{b} + \int_{-\infty}^{\breve{b}^*(\rho,a,b)} g(\breve{b})[\int_{-\infty}^{\rho-a} f(\breve{a})[\breve{b}+a]d\breve{a} + \int_{\rho-a}^{\infty} f(\breve{a})\rho d\breve{a}]d\breve{b} \\
\geq \; & \int_{-\infty}^{\infty} g(\breve{b})\breve{b}d\breve{b} \\
\Leftrightarrow \quad & \int_{\breve{b}^*(\rho,a,b)}^{110} \frac{1}{110}[\breve{b} - 20]d\breve{b} \\
& + \int_0^{\breve{b}^*(\rho,a,b)} \frac{1}{110}[\int_0^{52.5-30} \frac{1}{100}[\breve{b}+30]d\breve{a} + \int_{52.5-30}^{100} \frac{1}{100} 52.5 d\breve{a}]d\breve{b} \geq 55 \\
\Leftrightarrow \quad & \int_{\breve{b}^*(\rho,a,b)}^{110} \frac{1}{110}[\breve{b} - 20]d\breve{b} + \int_0^{\breve{b}^*(\rho,a,b)} \frac{1}{110}[\frac{22.5}{100}[\breve{b}+30] + \frac{77.5}{100} 52.5]d\breve{b} \geq 55 \\
\Leftrightarrow \quad & \int_{\breve{b}^*(\rho,a,b)}^{110} \frac{1}{110}[\breve{b} - 20]d\breve{b} + \int_0^{\breve{b}^*(\rho,a,b)} 0.43125 + \frac{0.225}{110}\breve{b}d\breve{b} \geq 55
\end{aligned}
$$

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. □

Actually, in the example game of the proof, both IR constraints are satisfied by a wide range of leveled commitment contracts—and for no full commitment contract. Which leveled commitment contracts defined by ρ, a, and b satisfy the 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 (Equations 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 Figure 2 left. Furthermore, the boundaries of the programmed model need to be checked. The boundaries $a = 0$, $a = \rho$, and $\breve{b}^* = 110$ are plotted in Figure 2 left. The constraint

$\breve{b}^* > 0$ is always satisfied in this case and is thus not plotted. To summarize, in the gray area of Figure 2 left, the contracts are IR for both agents, given that the agents decommit according to the iterated dominance equilibrium.

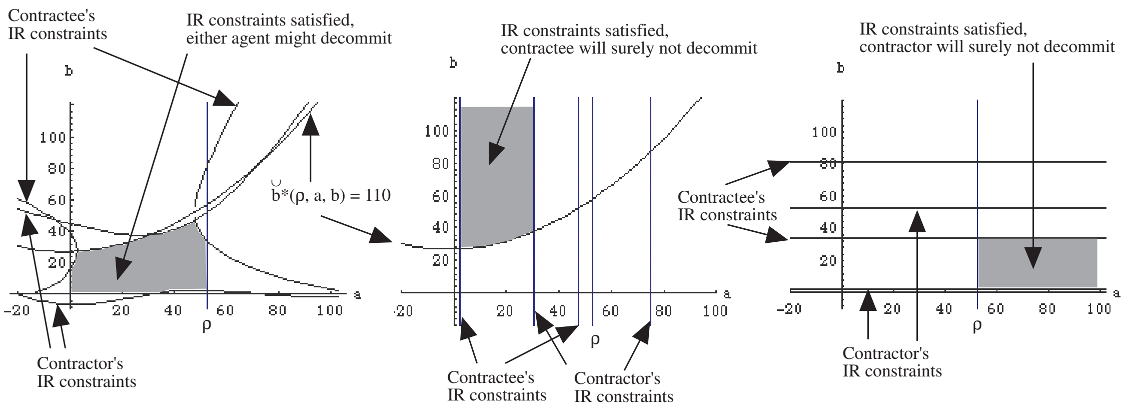

Figure 2: *The decommitment penalties a and b that satisfy both agents' IR constraints in the example SEQD game. Right: case where either agent might decommit ($a < \rho$, and $\breve{b}^*(\rho, a, b) < 110$). Middle: case where the contractor might decommit but the contractee will not ($a < \rho$, and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

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 $\breve{b}^*(\rho, a, b) = \rho + \frac{b + \int_{-\infty}^{\rho - a} f(\breve{a}) d\breve{a}[a]}{\int_{\rho - a}^{\infty} f(\breve{a}) d\breve{a}} = \rho + b$, i.e. the contractee decommits truthfully. 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

$$\begin{aligned}
& \frac{1}{110} \frac{1}{100} [(110 - (\rho + b)) \cdot (\frac{-(100)^2}{2} + 100b) + (\rho + b) \cdot (-100\rho)] \geq E[-\breve{a}] \\
\Leftrightarrow \quad & \frac{1}{110} \frac{1}{100} [(57.5 - b) \cdot (-5000 + 100b) + (52.5 + b) \cdot (-5250)] \geq -50 \\
\Leftrightarrow \quad & 2.5 \leq b \leq 52.5
\end{aligned}$$

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

$$\begin{aligned}
& \int_{\rho+b}^{110} g(\breve{b})[\breve{b} - b] \int_0^{100} f(\breve{a}) d\breve{a} d\breve{b} + \int_0^{\rho+b} g(\breve{b})[\rho] \int_0^{100} f(\breve{a}) d\breve{a} d\breve{b} \geq E[\breve{b}] \\
\Leftrightarrow \quad & \frac{1}{110}\frac{1}{100}[(\frac{110^2}{2} - 110b - (\frac{(\rho+b)^2}{2} - (\rho+b)b)) \cdot 100 + (\rho+b)\rho \cdot 100] \geq 55 \\
\Leftrightarrow \quad & b \leq \text{ approximately } 34.05 \text{ or } b \geq \text{ approximately } 80.95
\end{aligned}$$

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 contracts are IR for both agents, given that the agents decommit according to the iterated dominance equilibrium. This region is colored gray in Figure 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

$$\begin{aligned}
& \int_{-\infty}^{\breve{b}^*(\rho,a,b)} g(\breve{b})[\int_{-\infty}^{\rho-a} f(\breve{a})[-\breve{a} - a] d\breve{a} + \int_{\rho-a}^{\infty} f(\breve{a})[-\rho] d\breve{a}] d\breve{b} \geq E[-\breve{a}] \qquad (7) \\
\Leftrightarrow \quad & \int_0^{110} \frac{1}{110}[\int_0^{\rho-a} f(\breve{a})[-\breve{a} - a] d\breve{a} + \int_{\rho-a}^{100} f(\breve{a})[-\rho] d\breve{a}] d\breve{b} \geq -50 \\
\Leftrightarrow \quad & \int_0^{\rho-a} f(\breve{a})[-\breve{a} - a] d\breve{a} + \int_{\rho-a}^{100} f(\breve{a})[-\rho] d\breve{a} \geq -50
\end{aligned}$$

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

$$\begin{aligned}
& \frac{1}{100}[\int_0^{\rho-a} [-\breve{a} - a] d\breve{a} + \int_{\rho-a}^{100} [-\rho] d\breve{a}] \geq -50 \\
\Leftrightarrow \quad & \frac{1}{100}[(\frac{-(\rho-a)^2}{2} + (\rho-a)(-a)) + ((100 - (\rho-a)) \cdot (-\rho)) \geq -50 \\
\Leftrightarrow \quad & a \leq \text{ approximately } 30.14 \text{ or } a \geq \text{ approximately } 74.86
\end{aligned}$$

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

Similarly, the contractee's IR constraint becomes

$$\begin{aligned}
& \int_{-\infty}^{\breve{b}^*(\rho,a,b)} g(\breve{b})[\int_{-\infty}^{\rho-a} f(\breve{a})[\breve{b} + a] d\breve{a} + \int_{\rho-a}^{\infty} f(\breve{a})[\rho] d\breve{a}] d\breve{b} \geq E[\breve{b}] \qquad (8) \\
\Leftrightarrow \quad & \int_0^{110} \frac{1}{110}[[\breve{b} + a] \int_0^{\rho-a} f(\breve{a}) d\breve{a} + [\rho] \int_{\rho-a}^{100} f(\breve{a}) d\breve{a}] d\breve{b} \geq 55 \\
\Leftrightarrow \quad & \frac{1}{110}[[\frac{110^2}{2} + 110a] \int_0^{\rho-a} f(\breve{a}) d\breve{a} + 110\rho \int_{\rho-a}^{100} f(\breve{a}) d\breve{a}] \geq 55 \\
\Leftrightarrow \quad & [55 + a] \int_0^{\rho-a} f(\breve{a}) d\breve{a} + \rho \int_{\rho-a}^{100} f(\breve{a}) d\breve{a} \geq 55
\end{aligned}$$

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

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

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 Figure 2 middle.

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 subsume full commitment ones). This occurs always 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 (Efficiency improvement in SEQD games) *If a SEQD game has at least one IR full commitment contract and*

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

2. *$\breve{a}$ is bounded from below, and $\int_{E[\breve{a}]+\epsilon}^{\infty} g(\breve{b})d\breve{b} > 0$ for some $\epsilon > 0$,*

then that game has a leveled commitment contract that increases the contractor's expected payoff as well as the contractee's expected payoff over any full commitment contract. Therefore, the leveled commitment contract is Pareto superior and IR.

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

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.

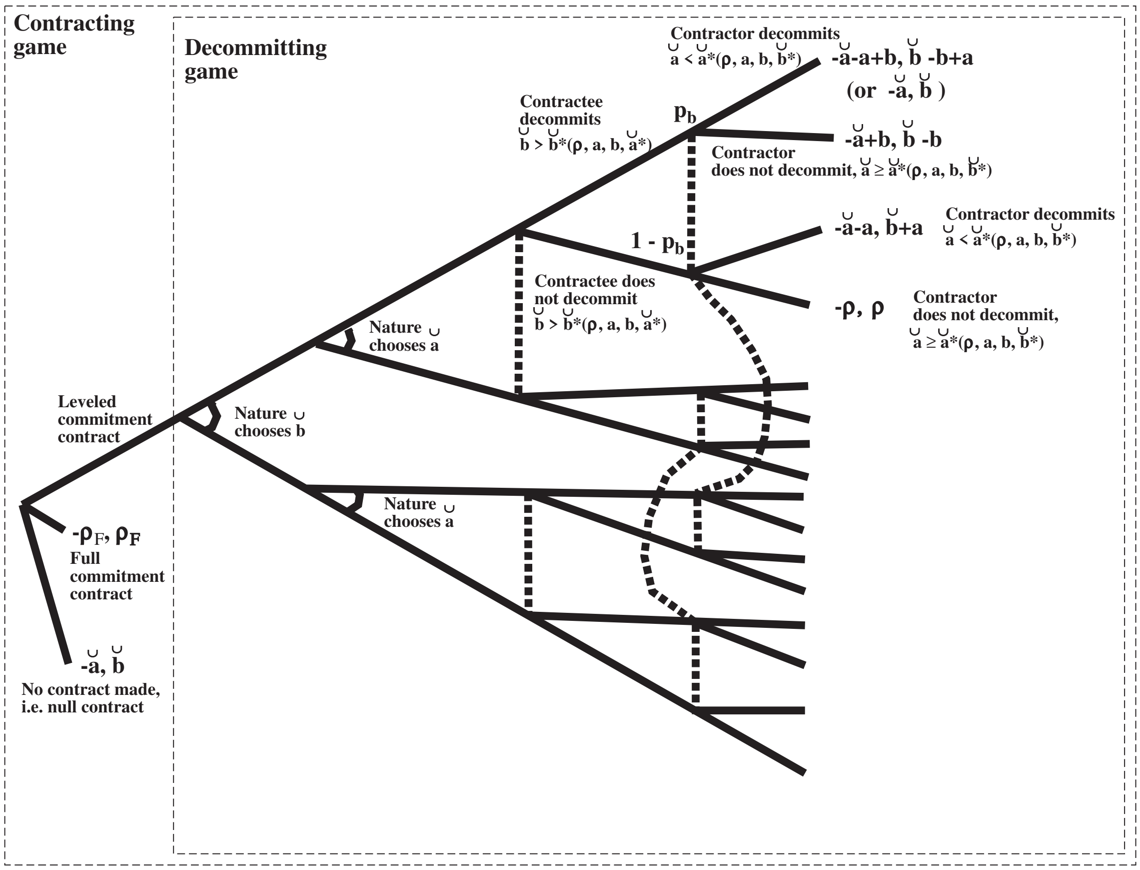

Figure 3: *The "SIMUltaneous Decommit - Both Pay if both decommit" (SIMUDBP) game. The parenthesized payoffs represent the "SIMUltaneous Decommit - Neither Pays if both decommit" (SIMUDNP) game. The dashed lines represent the agents' information sets. When decommitting, the contractor does not know the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

2.2.1 Both pay if both decommit (SIMUDBP)

This section discusses simultaneous decommitting games where a protocol is used where both agents have to pay the decommitting penalties if both decommit. Such settings will be called SIMUDBP games, Figure 3. Let p_b be the probability that the contractee decommits. The value of this variable depends on $f(\breve{a})$, $g(\breve{b})$, 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 contracts where either one of the agents gets paid for decommitting. On the other hand, the above inequality is equivalent to

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

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

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

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 contracts where either one of the agents gets paid for decommitting. On the other hand, the above inequality is equivalent to

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

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 Equation 12. Condition 11 uses the variable p_a which is defined by Equation 10. So, together, Equations 9, 10, 11, and 12 define the Nash equilibria of the decommitting game.

Now the contractor's IR constraint becomes

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

$\Leftrightarrow$

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

The first row corresponds to the contractee decommitting, while the second corresponds to the contractee not decommitting. The second integral in each row corresponds to the contractor decommitting, while the third integral corresponds to the contractor not decommitting.

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

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

If $\breve{a}$ is bounded from below, the contractor's decommitment penalty, a, can be chosen so high that the contractor's decommitment threshold, $\breve{a}^*(\rho, a, b, \breve{b}^*)$, becomes lower than any 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, full commitment contracts are a subset of leveled commitment ones. Therefore, the former can be no better in the sense of Pareto efficiency or social welfare (i.e. sum of the agents' payoffs) than the latter.

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 $f(\breve{a})$ and $g(\breve{b})$) where no full commitment contract is individually rational to both agents but a leveled commitment contract is.*

Proof. Let $f(\breve{a}) = \begin{cases} \frac{1}{100} & \text{if } 0 \leq \breve{a} \leq 100 \\ 0 & \text{otherwise} \end{cases}$ and $g(\breve{b}) = \begin{cases} \frac{1}{110} & \text{if } 0 \leq \breve{b} \leq 110 \\ 0 & \text{otherwise} \end{cases}$. No full commitment contract F satisfies both IR constraints because that would require $E[\breve{b}] \leq \rho_F \leq E[\breve{a}]$ which is impossible because $55 = E[\breve{b}] > E[\breve{a}] = 50$. 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 Figure 4. The Nash equilibrium decommitment thresholds $\breve{a}^*$ and $\breve{b}^*$ differ from the truthful 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 Equations 9, 10, 11, and 12 and the IR constraints. To make the algebra tractable (constant $f(\breve{a})$ and $g(\breve{b})$), versions of

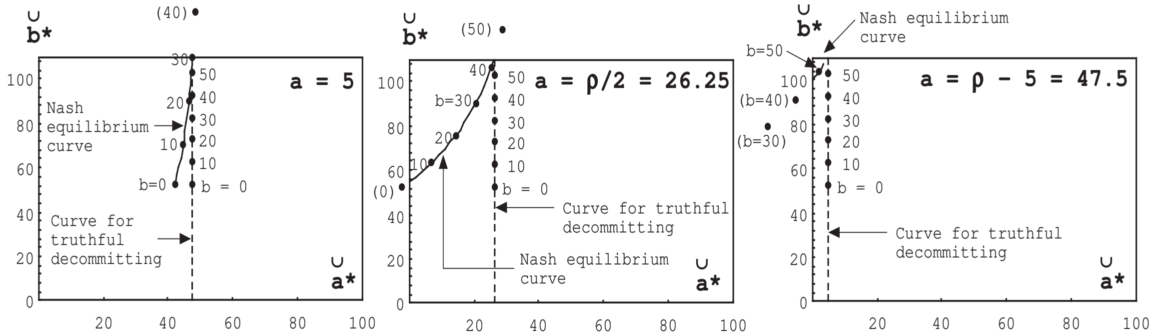

Figure 4: *The Nash equilibrium decommitment thresholds* $\breve{a}^*$ *and* $\breve{b}^*$ *of our example SIMUDBP game for different values of the decommitment penalties* a *and* b*. The Nash equilibrium deviates from truthful decommitting. If* $0 < \breve{a}^* < 100$*, and* $0 < \breve{b}^* < 110$*,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

these equations were used that assumed $0 < \breve{a}^* < 100$, and $0 < \breve{b}^* < 110$, without loss of generality. Therefore the first task was to check the boundaries of the validity of the model. The boundaries $\breve{a}^* = 0$ and $\breve{b}^* = 110$ are plotted in Figure 5. The boundary $\breve{a}^* = 100$ turns out to be the line $b = 0$. There exists no boundary $\breve{b}^* = 0$ because $\breve{b}^*$ was always greater than zero. After plotting the validity boundaries of the model,

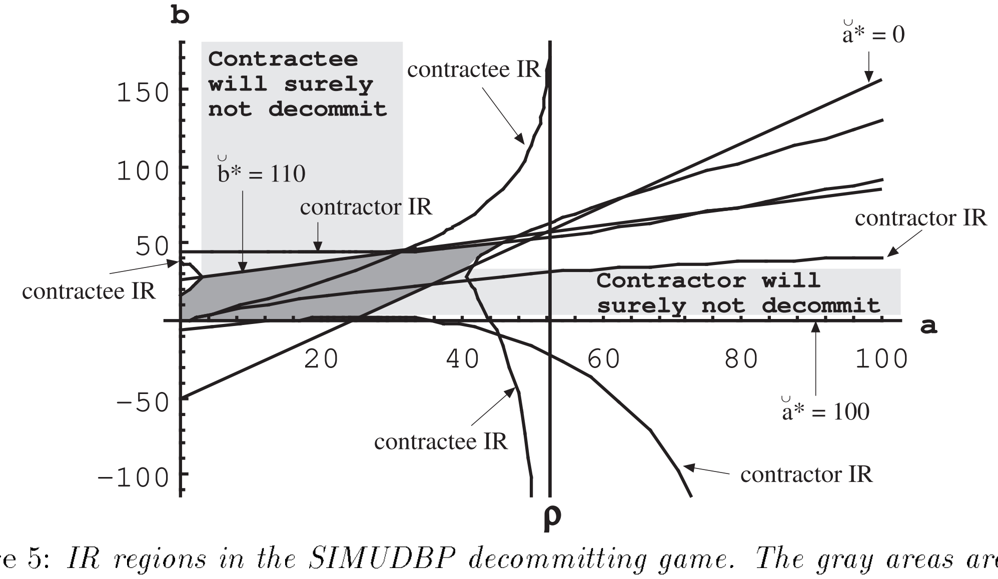

Figure 5: *IR regions in the SIMUDBP decommitting game. The gray areas are three qualitatively different regions of contracts that are IR for both agents and allow an equilibrium in the SIMUDBP decommitting game. In the dark gray area either agent might decommit while in the light gray areas only one 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.*

the curves at which the IR constraints held with equality were plotted, see Figure 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 Figure 5 represents the values of the decommitment penalties a and b for which the validity constraints of the programmed model and the IR constraints are satisfied. In other words, for any such a and b, there exists decommitment thresholds $\breve{a}^*$ and $\breve{b}^*$ such that these form a Nash equilibrium, and there is a nonzero probability for either agent to decommit, and each agent has higher expected payoff with the contract than without it.

As a numerical example, pick a contract where $a = \frac{\rho}{2} = 26.25$, and $b = 30$. Now in Nash equilibrium, the decommitment thresholds are $\breve{a}^* \approx 20.50$, and $\breve{b}^* \approx 90.24$, Figure 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 $\breve{b}^*(\rho, a, b, \breve{a}^*) = \rho + \frac{b}{\int_{\breve{a}^*}^{\infty} f(\breve{a}) d\breve{a}} = \rho + b$, i.e. the contractee decommits truthfully. The contractor's IR constraint becomes 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 $b \leq$ approximately 34.05. Thus the open region $2.5 \leq b \leq 34.05$, $\breve{a}^* \leq 0$ is where this type of contracts are IR for both agents and in equilibrium. This region is colored light gray in Figure 5.

Case 3, Contractee will surely not decommit. If $\breve{b}^* \geq 110$, the contractee will surely not decommit ($p_b = 0$). Now $\breve{a}^*(\rho, a, b, \breve{b}^*) = \rho - \frac{a}{1-p_b} = \rho - a$, i.e. the contractor decommits truthfully. 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 $a \leq$ approximately 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 contracts are IR for both agents, and in equilibrium. This region is colored light gray in Figure 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. Though 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 ones). This occurs always 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 (Efficiency improvement in SIMUDBP games) *If a SIMUDBP game (defined by $f(\breve{a})$ and $g(\breve{b})$) has at least one IR full commitment contract and*

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

2. *$\breve{a}$ is bounded from below, and $\int_{E[\breve{a}]+\epsilon}^{\infty} g(\breve{b})d\breve{b} > 0$ for some $\epsilon > 0$,*

then that game has a leveled commitment contract that increases the contractor's expected payoff as well as the contractee's expected payoff over any full commitment contract. Therefore, the leveled commitment contract is Pareto superior and IR.

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 protocol is used where neither agent has to pay a decommitting penalty if both agents decommit. Such settings will be called SIMUDNP games, Figure 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 contracts where either one of the agents gets paid for decommitting. On the other hand, the above inequality is equivalent to

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

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 contracts where either one of the agents gets paid for decommitting. On the other hand, the above inequality is equivalent to

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

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 Equation 16. Condition 15 uses the variable p_a which is defined by Equation 14. So together, Equations 13, 14, 15, and 16 define the Nash equilibria of the decommitting game.

Now the contractor's IR constraint becomes

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

The first row corresponds to the contractee decommitting, while the second corresponds to the contractee not decommitting. The second integral in each row corresponds to the contractor decommitting, while the third integral corresponds to the contractor not decommitting.

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

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

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, full commitment contracts are a subset of leveled commitment ones. Therefore, the former can be no better in the sense of Pareto efficiency or in the social welfare sense than the latter. 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 exist SIMUDNP games (defined by $f(\breve{a})$ and $g(\breve{b})$) where no full commitment contract is individually rational to both agents but a leveled commitment contract is.*

Proof. Let $f(\breve{a}) = \begin{cases} \frac{1}{100} & \text{if } 0 \leq \breve{a} \leq 100 \\ 0 & \text{otherwise} \end{cases}$ and $g(\breve{b}) = \begin{cases} \frac{1}{110} & \text{if } 0 \leq \breve{b} \leq 110 \\ 0 & \text{otherwise} \end{cases}$. No full commitment contract, F, satisfies both IR constraints because that would require $E[\breve{b}] \leq \rho_F \leq E[\breve{a}]$ which is impossible because $55 = E[\breve{b}] > E[\breve{a}] = 50$. 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 Figure 6. Note that

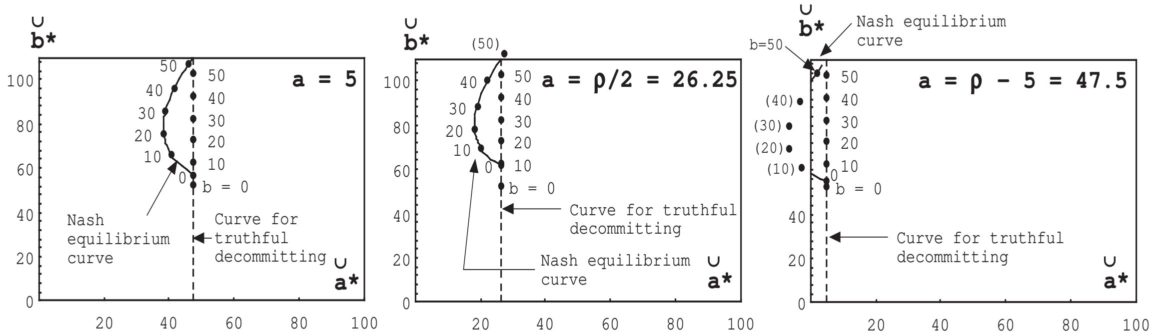

Figure 6: *The Nash equilibrium decommitment thresholds $\breve{a}^*$ and $\breve{b}^*$ of our example SIMUDNP game for different values of the decommitment penalties a and b. The Nash equilibrium deviates from truthful decommitting. If $0 < \breve{a}^* < 100$, and $0 < \breve{b}^* < 110$, there is some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

the Nash equilibrium decommitment thresholds $\breve{a}^*$ and $\breve{b}^*$ indeed do differ from the truthful ones. It is not guaranteed that all of these Nash equilibria satisfy the agents' IR constraints. Therefore, we programmed a model of Equations 13, 14, 15, and 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 Figure 7. After plotting the validity boundaries of the model, the curves at which the IR constraints held with equality were plotted in Figure 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 Figure 7 represents the values of the

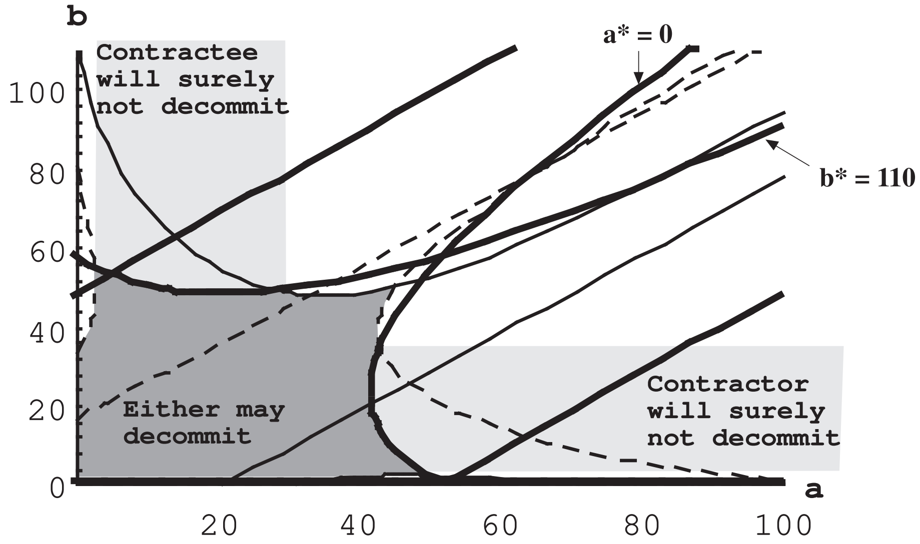

Figure 7: *IR regions in the SIMUDNP decommitting game. The gray areas are three qualitatively different regions of contracts that are IR for both agents and allow an equilibrium in the SIMUDNP decommitting game. The bold lines are the validity constraints for the programmed model that requires* $0 < \breve{a}^* < 100$*, 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.*

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 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 $a = \frac{\rho}{2} = 26.25$, and $b = 30$. Now in Nash equilibrium, the decommitment thresholds are $\breve{a}^* \approx 19.03$, and $\breve{b}^* \approx 88.67$, Figure 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 $\breve{b}^*(\rho, a, b, \breve{a}^*) = \rho + b - \frac{a \int_{-\infty}^{\breve{a}^*(\rho,a,b,\breve{b}^*)} f(\breve{a}) d\breve{a}}{\int_{\breve{a}^*(\rho,a,b,\breve{b}^*)}^{\infty} f(\breve{a}) d\breve{a}} = \rho + b$, i.e. the contractee decommits truthfully. The contractor's IR constraint becomes exactly the same as in case 2 of the example SEQD game (Eq. 5). (It is also 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 $b \leq$ approximately 34.05. Thus the open region $2.5 \leq b \leq 34.05$, $\breve{a}^* \leq 0$ is where this type of contracts are IR for both agents and in equilibrium. This region is colored light gray in Figure 7.

Case 3, Contractee will surely not decommit. If $\breve{b}^* \geq 110$, the contractee will surely not decommit ($p_b = 0$). Now $\breve{a}^*(\rho, a, b, \breve{b}^*) = \rho - a - \frac{bp_b}{1-p_b} = \rho - a$, i.e. the contractor decommits truthfully. The contractor's IR constraint becomes exactly the same as in case 3 of the example SEQD game (Eq. 8). (It is also same as in case 3 of the example SIMUDBP game). This constraint was proven equivalent to $a \leq$ approximately 30.14. Similarly, the contractee's IR constraint becomes the same as in the SEQD game (Eq. 9). (It is also 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 contracts are IR for both agents, and in equilibrium. This region is colored light gray in Figure 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. Though 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. □

Figures 4 and 6 show that the Nash equilibrium decommitting strategies differ significantly in the game where both agents have to pay if both decommit (SIMUDBP) and the game where neither agent has to pay a penalty if both decommit (SIMUDNP). The qualitative shapes of the Nash equilibrium curves using these two protocols also differ. Most notably, as the contractee's penalty approaches zero, the contractor's decommitting strategy approaches truthful decommitting in the SIMUDNP game unlike in the SIMUDBP game. This can also be seen from Equation 13. Similarly, as the contractor's penalty approaches zero, the contractee's decommitting strategy approaches truthful decommitting in the SIMUDNP game unlike in the SIMUDBP game. This can also be seen from Equation 15.

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 subsume full commitment ones):

Theorem 2.6 (Efficiency improvement in SIMUDNP games) *If a SIMUDNP game (defined by $f(\breve{a})$ and $g(\breve{b})$) has at least one IR full commitment contract and*

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

2. *$\breve{a}$ is bounded from below, and $\int_{E[\breve{a}]+\epsilon}^{\infty} g(\breve{b})d\breve{b} > 0$ for some $\epsilon > 0$,*

then that game has a leveled commitment contract that increases the contractor's expected payoff as well as the contractee's expected payoff over any full commitment contract. Therefore, the leveled commitment contract is Pareto superior and IR.

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 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). Games where at least one agent's future is certain are a subset of these games. In such games all of the probability mass of $f(\breve{a})$ and/or $g(\breve{b})$ is on one point. Our results above that showed that leveled commitment protocols can make a deal individually rational in setting where no full commitment contract can apply directly to this setting of one-sided uncertainty. So do the results that showed sufficient conditions under which leveled commitment contracts improve the expected payoff of both contract parties compared to any full commitment contract. According to the results this occurs if there is some chance that the contractor's outside offer is lower than the contractee's.

Because the result for two-sided uncertainty apply directly to the case of one-sided uncertainty, they do not have to be reproven for the latter case. However, 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, Figure 8.

In a sequential decommitting COP game where the contractee reveals decommitment first (Figure 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.

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 Figure 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

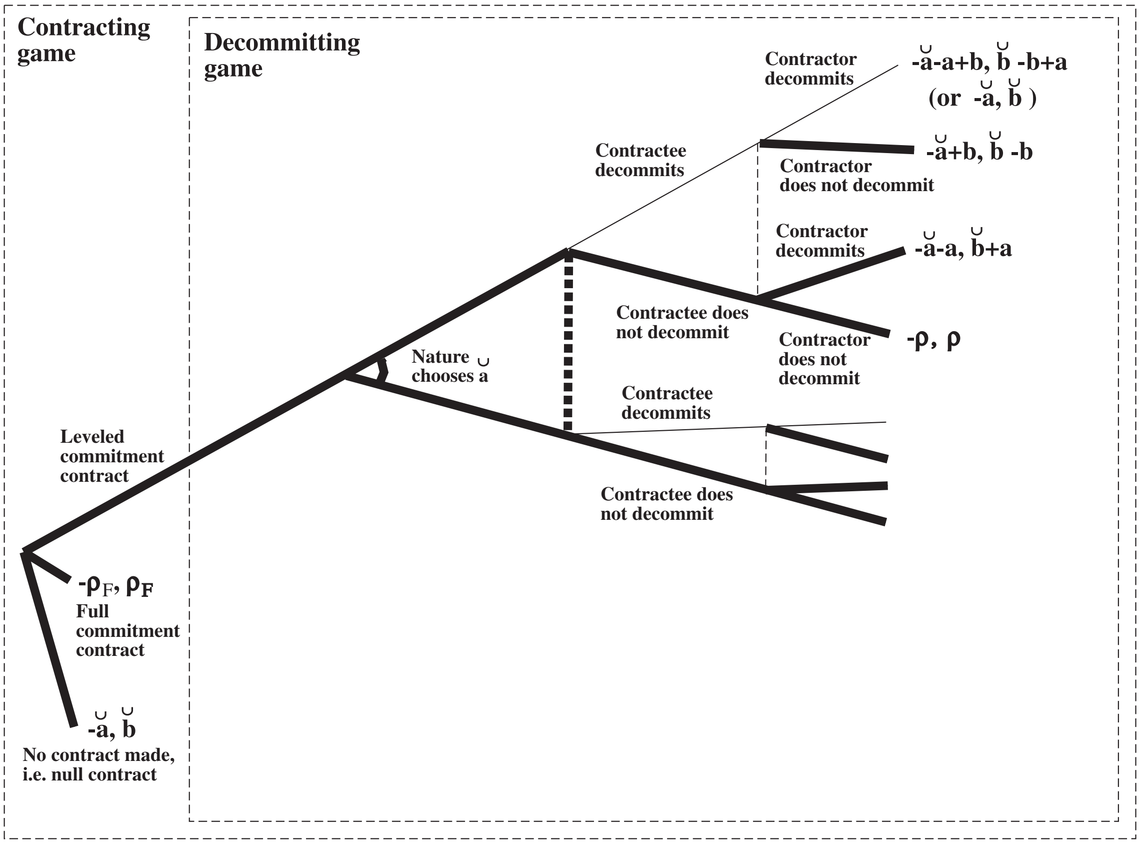

Figure 8: *The Certain Offer Prevails (COP) game. If the contractor decommits, the contractee can still accept the outside offer. In the figure, the contractor's payoff is listed before the contractee's. The bold solid lines show choices that may actually occur in any subgame. The bold dashed line represents the contractee's information set: it does not know $\breve{a}$ in the decommitting game. The thin dashed lines represent the alternative protocol where both agents have to reveal their decommitment decisions simultaneously: when deciding on decommitting, the contractor has not observed whether the contractee decommitted.*

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 is not IR for the contractee. Therefore the contractee never decommits.

Even if a protocol is used that specifies that neither agent has to pay the decommitment penalty if both decommit (payoffs in parentheses in Figure 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 protocols, the contractor can reason that the contractee will not decommit. Therefore the three cases become equivalent. This holds for the protocol that specifies that both have to pay if both decommit and for the protocol that specifies that neither has to pay if both decommit.

The contractor's cost is ρ if it does not decommit, and $\breve{a} + a$ if it does. In other words the contractor's payoff is $-\rho$ or $-\breve{a} - a$. 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 $E[-\breve{a}] \leq E[max[-\breve{a} - a, -\rho]]$:

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

4 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 higher expected payoff by passing on the $\breve{b}$-offer and agreeing to the risky contract than by accepting the $\breve{b}$-offer. If the contract is made, decommit 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 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 Figure 9.

In a sequential decommitting COBV game where the contractee reveals decommitment first (Figure 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 Figure 9.

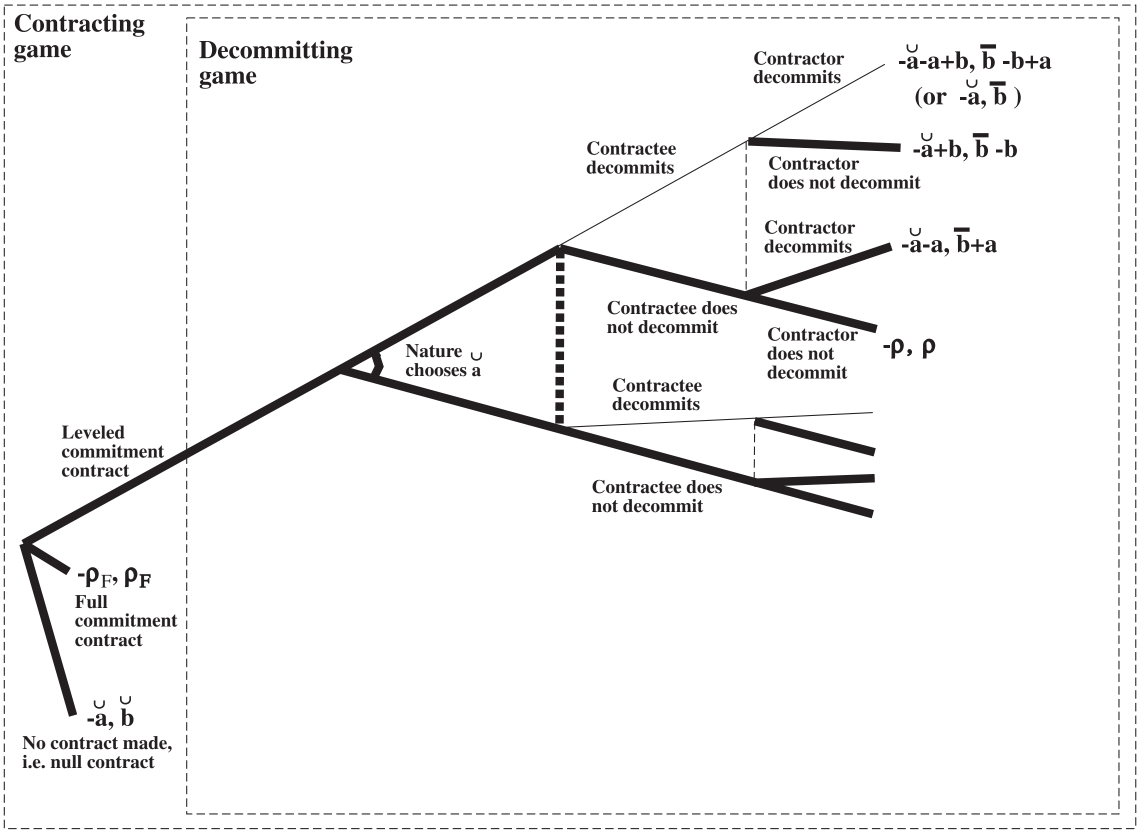

Figure 9: *The "Certain Offer Becomes Void" (COBV) game. If at all, the contractee's outside offer $\breve{b}$ has to be accepted before the contractor's outside offer $\breve{a}$ becomes known. The bold solid lines show choices that may actually occur in any subgame. The bold dashed line represents the contractee's information set: it does not know $\breve{a}$ in the decommitting game. The thin dashed lines represent the alternative situation where both agents reveal their decommitting decisions simultaneously: when deciding on decommitting, the contractor has not observed whether the contractee decommitted.*

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

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 protocol that specifies that both have to pay if both decommit and for the protocol that specifies that neither has to pay if both decommit.

The contractor's cost is ρ 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 3), 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 4.1 No enabling in a COBV game. *Let* $\bar{b} \leq \frac{\int_{-\infty}^{\rho-a} f(\breve{a})\breve{a}d\breve{a}}{\int_{-\infty}^{\rho-a} f(\breve{a})d\breve{a}}$ *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,

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

Contradiction. Thus no leveled commitment contract satisfies both IR constraints. □

Proposition 4.2 No Pareto efficiency improvement in a COBV game. *Let* $\bar{b} \leq \frac{\int_{-\infty}^{\rho-a} f(\breve{a})\breve{a}d\breve{a}}{\int_{-\infty}^{\rho-a} f(\breve{a})d\breve{a}}$ *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 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.

$$\begin{aligned}
& \int_{-\infty}^{\rho-a} f(\breve{a})[-\breve{a} - a]d\breve{a} + \int_{\rho-a}^{\infty} f(\breve{a})[-\rho]d\breve{a} \geq -\rho_F \\
\text{and} \quad & [1 - (\int_{-\infty}^{\rho-a} f(\breve{a})d\breve{a})]\rho + (\int_{-\infty}^{\rho-a} f(\breve{a})d\breve{a})[\bar{b} + a] \geq \rho_F \\
& \text{and at least one of the above inequalities is strict.} \\
\Rightarrow \quad & \int_{-\infty}^{\rho-a} f(\breve{a})[\breve{a} + a]d\breve{a} + \int_{\rho-a}^{\infty} f(\breve{a})\rho d\breve{a} \\
& < [1 - (\int_{-\infty}^{\rho-a} f(\breve{a})d\breve{a})]\rho + (\int_{-\infty}^{\rho-a} f(\breve{a})d\breve{a})[\bar{b} + a] \\
\Leftrightarrow \quad & \int_{-\infty}^{\rho-a} f(\breve{a})\breve{a}d\breve{a} < (\int_{-\infty}^{\rho-a} f(\breve{a})d\breve{a})\bar{b}
\end{aligned}$$

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

Contradiction. Thus no such leveled commitment contract exists. □

The constraint $\bar{b} \leq \frac{\int_{-\infty}^{\rho-a} f(\breve{a})\breve{a}d\breve{a}}{\int_{-\infty}^{\rho-a} f(\breve{a})d\breve{a}}$ 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 fall-back. The former requirement does not seem very restrictive, but the latter does. Thus these two propositions with negative results have relatively limited scope.

5 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 formulas 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 two-sided uncertainty an agent's biased beliefs can hurt the other contract party as well, while in the games with one-sided uncertainty they cannot.

5.1 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):

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

5.2 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 truthfully. 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 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.

5.3 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 SIMUDBP 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 truthfully. 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 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 counterspeculate the other agent's beliefs in order to determine a preference order over contracts.

5.4 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 agents reasoning process or information sources. Thus an agent need not counterspeculate 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 5.1 Payoff unaffected by opponent's beliefs in COP games. *Say that one agent's information is unbiased, i.e. either $f_a(\breve{a}) = f(\breve{a})$ or $f_b(\breve{a}) = f(\breve{a})$. 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. $f_a(\breve{a}) = f(\breve{a})$. 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 $\int_{-\infty}^{\rho-a} f_a(\breve{a})[-\breve{a}-a]d\breve{a} + \int_{\rho-a}^{\infty} f_a(\breve{a})[-\rho]d\breve{a}$. None of these depend on the contractee's information.

Now say the contractee's information is unbiased, i.e. $f_b(\breve{a}) = f(\breve{a})$. 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. □

5.5 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 drops 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

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

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 5.2 Perceived enabling in COBV games. *There are COBV games (defined by $f(\breve{a})$, $f_a(\breve{a})$, $f_b(\breve{a})$, $\bar{b}$, and $\breve{b}$) where no full commitment contract satisfies both IR constraints, no leveled commitment contract satisfies both IR constraints, 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 $f(\breve{a}) = f_a(\breve{a}) = \begin{cases} 0.01 & \text{if } 0 \leq \breve{a} \leq 100 \\ 0 & \text{otherwise} \end{cases}$ and $\breve{b} = 55$, i.e. no full commitment contract satisfies both IR constraints. Now let $\bar{b} = 0$. It follows by Proposition 4.1 that no leveled commitment contract satisfies both IR constraints. No full commitment contract satisfies both PIR constraint because it would require $55 = \breve{b} \leq \rho_F \leq E_a[\breve{a}] = 50$. Now we show a leveled commitment contract that satisfies both PIR constraints. Let $f_b(\breve{a}) = \begin{cases} 0.01 & \text{if } 50 \leq \breve{a} \leq 150 \\ 0 & \text{otherwise} \end{cases}$. Now the contractee's PIR constraint is

$$55 \leq [1 - (\int_{-\infty}^{\rho - a} f_b(\breve{a})d\breve{a})]\rho + (\int_{-\infty}^{\rho - a} f_b(\breve{a})d\breve{a})[0 + a]$$

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

$$55 \leq [1 - (\int_{-\infty}^{60-10} f_b(\breve{a})d\breve{a})]60 + (\int_{-\infty}^{60-10} f_b(\breve{a})d\breve{a})10 \Leftrightarrow 55 \leq 60 + 0$$

The contractor's PIR constraint is

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

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

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

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 $f_a(\breve{a}) \neq f_b(\breve{a})$, 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 4.1) 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 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 5.3 Payoff unaffected by opponent's beliefs in COBV games. *Say that one agent's information is unbiased, i.e. either $f_a(\breve{a}) = f(\breve{a})$ or $f_b(\breve{a}) = f(\breve{a})$. 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 its expected payoff for the leveled commitment contract is $\int_{-\infty}^{\rho-a} f(\breve{a})[-\breve{a}-a]d\breve{a} + \int_{\rho-a}^{\infty} f(\breve{a})[-\rho]d\breve{a}$. None of these depend on the contractee's information.

The contractee's payoff for not accepting either contract is $\breve{b}$. Its payoff for the full commitment contract is ρ_F, and its 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 5.1 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 $f_a(\breve{a}) = f(\breve{a})$ or $f_b(\breve{a}) = f(\breve{a})$. Say that the contract is perceived IR by the agent x for which $f_x(\breve{a}) = f(\breve{a})$. 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 5.3 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 5.3 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.

6 Practical prescriptions for market designers

Since we introduced leveled commitment contracts [34, 33], they have quickly been adopted into implementations. For example, Mitsubishi has applied them to a construction waste recycling problem in Japan [1]. They have also been applied to automated negotiation in a manufacturing setting [8] and in a digital library [20]. 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 protocol. This allows the tuning of the penalties not only to specific negotiation situations and environmental uncertainties, but also to specific belief structures of the agents. In this paper we presented the IR constraints which can be used to evaluate a given contract (price and 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 that maximize the sum of the contract parties' expected payoffs [36]. 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 [39]. This requires that the distributive bargaining occurs over contract price, and that the penalties are tailored to that price. 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 that 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.

In the presented instance of the two-sided uncertainty game, the Nash equilibrium decommitting strategies under a simultaneous decommitting protocol were closer to truthful ones when a protocol was used where neither pays if both decommit (SIMUDNP) than when a protocol was used where both pay if both decommit (SIMUDBP). Also, in the former protocol, an agent's equilibrium decommitting strategy approaches truthful decommitting when the other agent's decommitting penalty approaches zero. In the sequential decommitting game (SEQD), the agent that decommits last should decommit truthfully. For any given contract, the equilibria are different for these three protocols. The optimal contracts are also different. However, in a recent paper we show that, surprisingly, among risk neutral agents each of the three protocols leads to the same expected social welfare (sum of the contract parties' payoffs) when the contract price and penalties are optimized for each protocol separately [39]. Therefore, social welfare cannot be used as the criterion for choosing between these protocols. Among agents that are not risk neutral, the three protocols lead to different sums of utilities, and the ranking of the protocols varies based on the shapes of the utility functions.

Because social welfare cannot be used as the criterion for choosing a protocol among risk neutral agents, other criteria have to be invoked. 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 protocols, among risk neutral agents, the decommitting thresholds will be the same for all of the protocols [39]. Therefore, the protocols 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 protocols lead to one penalty being paid. If neither agent receives an offer that is better than the agent's threshold, all of the protocols 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 protocol 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 [39]. In summary, from the perspective of minimizing the number of penalty payments, the SIMUDNP protocol is best, SEQD is second, and SIMUDBP is worst.

Another protocol 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 protocols because iterated dominance is a more robust solution concept than Nash equilibrium. An additional consideration that promotes sequential decommitting protocols 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 [36]. These considerations speak for using sequential decommitting protocols.

In asynchronous negotiation systems without a trusted third party, implementing a decommitting protocol 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 protocol.[6] 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

[6] Similar protocols have been used for other purposes, see e.g. [43, 31].

decisions. Similarly, the agent cannot change its decision.[7]

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 protocol 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 protocol that specifies that if a contract offer is accepted and later either agent decommits, the original offer becomes void as opposed to staying valid according to its original deadline that may 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 protocol 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 [3]. This can be avoided by a mechanism where the decommitment penalties increase with time (real-time or 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 [3]. Similar issues for agents that carry out strategic lookahead were also recently studied [4].

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 [25, 33]. Even if explicit clustering of issues into contracts [26, 25, 33] 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 [33]. 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 protocols [33], 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 [33]. Again, if no such contracts appear, the agent can decommit.

Even if explicit clustering of issues into atomic contracts and/or multiagent contracts are allowed, using leveled commitment protocols on top of such contracts can

[7]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 virtually impossible in practice.

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 protocols will probably involve both mechanisms: explicit linking of issues and leveled commitment. We recently compared combinations of these mechanisms via a simulation study [3]. We also developed the theory of leveled commitment contracts involving more than two contract parties in a single contract [36].

In many practical automated contracting settings, agents are bounded rational—for example because limited computation resources bound their capability to solve combinatorially complex problems [33, 35, 25]. 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 protocol 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 protocol 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 [33].

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. [30, 17]. 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 [29]. 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.[8]

7 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

[8] In the Federal Communications Commission's bandwidth auction the bidders were allowed to retract their bids [16]. 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.

agent just pays that penalty to the other agent. This mechanism is better suited for complex computerized contracting settings than contingency contracts: 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. This is achieved trivially by setting the decommitting penalties so high that neither party will decommit.

We showed these results for multiple decommitting protocols which differ based on whether the decommitting decisions have to be revealed sequentially 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. Therefore, there were in essence three protocols to study. For each one, the decommitting equilibrium is very different, but the advantages hold under the same conditions. Two general qualitative observations were made about the decommitting equilibrium. First, in sequential protocols, 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 protocol where neither agent pays if both decommit, an agent's best decommitting strategy approaches nonstrategic decommitting as the opponent's decommitting penalty approaches zero.

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 protocols comparatively. From the perspective of minimizing the number of penalty payments, the simultaneous decommitting protocol where neither pays a penalty if both decommit is best, the sequential protocol is second, and the simultaneous protocol where both pay if both decommit is worst. From the perspective of robustness and uniqueness of the equilibrium, and ease of optimal play, the sequential protocol 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 contracts provide both (all) contract parties the possibility to decommit 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 the performance profiles of the algorithms of bounded rational agents to the issues of this paper is also desirable. We have already taken initial steps toward relating the performance profiles and optimal negotiation actions in a coalition formation problem [35]. Finally, the relationship between leveled commitment contracting and explicit multi-issue and multiagent contracts should be studied in more detail.

References

[1] M. Akiyoshi, K. Abe, T. Ono, and Y. Tabinoki. Agent-based decision support framework for asynchronous negotiation of resource transfer. In *IEEE Systems, Man, and Cybernetics Conference*, Tokyo, Japan, Oct. 1999.

[2] M. R. Andersson and T. W. Sandholm. Contract types for satisficing task allocation: II experimental results. In *AAAI Spring Symposium Series: Satisficing Models*, pages 1–7, Stanford University, CA, Mar. 1998.

[3] M. R. Andersson and T. W. Sandholm. Leveled commitment contracting among myopic individually rational agents. In *Proceedings of the Third International*

Conference on Multi-Agent Systems (ICMAS), pages 26–33, Paris, France, July 1998.

[4] M. R. Andersson and T. W. Sandholm. Leveled commitment contracts with myopic and strategic agents. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 38–45, Madison, WI, July 1998.

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

[6] J. Q. Cheng and M. P. Wellman. The WALRAS algorithm: A convergent distributed implementation of general equilibrium outcomes. *Computational Economics*, 12:1–24, 1998.

[7] S.-Y. Choi, D. O. Stahl, and A. B. Whinston. *The Economics of Electronic Commerce*. MacMillan Technical Publishing, 1997.

[8] J. Collins, B. Youngdahl, S. Jamison, B. Mobasher, and M. Gini. A market architecture for multi-agent contracting. In *Proceedings of the Second International Conference on Autonomous Agents (AGENTS)*, pages 285–292, Minneapolis/St. Paul, MN, May 1998.

[9] K. Decker and V. R. Lesser. Designing a family of coordination algorithms. In *Proceedings of the First International Conference on Multi-Agent Systems (ICMAS)*, pages 73–80, San Francisco, CA, June 1995.

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

[11] E. Ephrati and J. S. Rosenschein. The Clarke tax as a consensus mechanism among automated agents. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 173–178, Anaheim, CA, 1991.

[12] S. Kraus. Agents contracting tasks in non-collaborative environments. In *Proc. Eleventh National Conference on Artificial Intelligence (AAAI-93)*, pages 243–248, July 1993.

[13] S. Kraus, J. Wilkenfeld, and G. Zlotkin. Multiagent negotiation under time constraints. *Artificial Intelligence*, 75:297–345, 1995.

[14] S. H. Low, N. F. Maxemchuk, and S. Paul. Anonymous credit cards and their collusion analysis. *IEEE/ACM Transactions on Networking*, 4(6), Dec. 1996.

[15] A. Mas-Colell, M. Whinston, and J. R. Green. *Microeconomic Theory*. Oxford University Press, 1995.

[16] R. P. McAfee and J. McMillan. Analyzing the airwaves auction. *Journal of Economic Perspectives*, 10(1):159–175, 1996.

[17] D. Monderer and M. Tennenholtz. Asymptotically optimal multi-object auctions for risk-averse agents. Technical report, Faculty of Industrial Engineering and Management, Technion, Haifa, Israel, Feb. 1999.

[18] J. Nash. The bargaining problem. *Econometrica*, 18:155–162, 1950.

[19] M. J. Osborne and A. Rubinstein. *Bargaining and Markets.* Academic Press, Inc., 1990.

[20] S. Park, E. Durfee, and W. Birmingham. Advantages of strategic thinking in multiagent contracts (a mechanism and analysis). In *Proceedings of the Second International Conference on Multi-Agent Systems (ICMAS)*, pages 259–266, Keihanna Plaza, Kyoto, Japan, Dec. 1996.

[21] R. A. Posner. *Economic Analysis of Law.* Little, Brown and Company, 2nd edition, 1977.

[22] H. Raiffa. *The Art and Science of Negotiation.* Harvard Univ. Press, Cambridge, Mass., 1982.

[23] J. S. Rosenschein and G. Zlotkin. *Rules of Encounter: Designing Conventions for Automated Negotiation among Computers.* MIT Press, 1994.

[24] A. Rubinstein. Perfect equilibrium in a bargaining model. *Econometrica*, 50:97–109, 1982.

[25] T. W. Sandholm. An implementation of the contract net protocol based on marginal cost calculations. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 256–262, Washington, D.C., July 1993.

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

[27] T. W. Sandholm. Unenforced E-commerce transactions. *IEEE Internet Computing*, 1(6):47–54, Nov–Dec 1997. Special issue on Electronic Commerce.

[28] T. W. Sandholm. Contract types for satisficing task allocation: I theoretical results. In *AAAI Spring Symposium Series: Satisficing Models*, pages 68–75, Stanford University, CA, Mar. 1998.

[29] T. W. Sandholm. An algorithm for optimal winner determination in combinatorial auctions. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence (IJCAI)*, Stockholm, Sweden, 1999. Extended version: Washington University, Department of Computer Science technical report WUCS-99-01.

[30] T. W. Sandholm. *eMediator*: A next generation electronic commerce server. In *AAAI-99 Workshop on AI in Electronic Commerce*, Orlando, FL, July, 1999.

[31] T. W. Sandholm, K. S. Larson, M. R. Andersson, O. Shehory, and F. Tohmé. Coalition structure generation with worst case guarantees. *Artificial Intelligence*, 1999. To appear.

[32] T. W. Sandholm and V. R. Lesser. Equilibrium analysis of the possibilities of unenforced exchange in multiagent systems. In *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence (IJCAI)*, pages 694–701, Montreal, Canada, Aug. 1995.

[33] T. W. Sandholm and V. R. Lesser. Issues in automated negotiation and electronic commerce: Extending the contract net framework. In *Proceedings of the First International Conference on Multi-Agent Systems (ICMAS)*, pages 328–335, San Francisco, CA, June 1995. Reprinted in *Readings in Agents*, Huhns and Singh, eds., pp. 66–73, 1997.

[34] T. W. Sandholm and V. R. Lesser. Advantages of a leveled commitment contracting protocol. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 126–133, Portland, OR, Aug. 1996. Extended version: University of Massachusetts at Amherst, Computer Science Department technical report 95-72.

[35] T. W. Sandholm and V. R. Lesser. Coalitions among computationally bounded agents. *Artificial Intelligence*, 94(1):99–137, 1997. Special issue on Economic Principles of Multiagent Systems. Early version appeared at the International Joint Conference on Artificial Intelligence, pages 662–669, 1995.

[36] T. W. Sandholm, S. Sikka, and S. Norden. Algorithms for optimizing leveled commitment contracts. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence (IJCAI)*, Stockholm, Sweden, 1999. Extended version: Washington University, Department of Computer Science technical report WUCS-99-02.

[37] T. W. Sandholm and N. Vulkan. Bargaining with deadlines. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, Orlando, FL, 1999.

[38] T. W. Sandholm and F. Ygge. On the gains and losses of speculation in equilibrium markets. In *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence (IJCAI)*, pages 632–638, Nagoya, Japan, Aug. 1997.

[39] T. W. Sandholm and Y. Zhou. Revenue equivalence of leveled commitment contracts. In *AAAI-99 Workshop on Negotiation: Settling conflicts and identifying opportunities*, Orlando, FL, July, 1999. Extended version: Washington University, Department of Computer Science technical report WUCS-99-03.

[40] S. Sen and E. Durfee. The role of commitment in cooperative negotiation. *International Journal on Intelligent Cooperative Information Systems*, 3(1):67–81, 1994.

[41] S. Sen and E. Durfee. A formal study of distributed meeting scheduling. *Group Decision and Negotiation*, 7:265–289, 1998.

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

[43] G. Zlotkin and J. S. Rosenschein. Coalition, cryptography and stability: Mechanisms for coalition formation in task oriented domains. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 432–437, Seattle, WA, July 1994.