Advantages of a Leveled Commitment Contracting Protocol

Tuomas W. Sandholm and Victor R. Lesser
Computer Science Department
University of Massachusetts at Amherst
CMPSCI Technical Report 95-72
September 7, 1995

Advantages of a Leveled Commitment Contracting Protocol

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

Abstract

In automated negotiation systems consisting of self-interested agents, contracts have traditionally been binding. Such contracts do not allow agents to efficiently accommodate future events. Game theory has proposed contingency contracts to solve this problem. Among computational agents, contingency contracts are often impractical due to large numbers of interdependent and unanticipated future events to be conditioned on, and because some events are not mutually observable. This paper proposes a leveled commitment contracting protocol that allows self-interested agents to efficiently accommodate future events by having the possibility of unilaterally decommitting from a contract based on local reasoning. A decommitment penalty is assigned to both agents in a contract: to be freed from the obligations of the contract, an agent only pays this penalty to the other party. It is shown through formal analysis of multiple contracting settings that this leveled commitment feature in a contracting protocol increases Pareto efficiency of deals and can make contracts individually rational when no full commitment contract can. This advantage holds even if the agents decommit manipulatively.[1]

[1]Supported by ARPA contract N00014-92-J-1698. The content does not necessarily reflect the position or the policy of the Government and no official endorsement should be inferred. T. Sandholm also supported by the Finnish Culture Foundation, Honkanen Foundation, Ella and George Ehrnrooth Foundation, Finnish Science Academy, Leo and Regina Wainstein Foundation, Finnish Information Technology Research Foundation, and Jenny and Antti Wihuri Foundation.

1 Introduction

The importance of automated negotiation systems is likely to increase [14]. One reason is the growth of a fast and inexpensive standardized communication infrastructure (EDI, NII, KQML [3], Telescript [5] etc.), over which separately designed agents belonging to different organizations can interact in an open environment in real-time, and safely carry out transactions [8, 22, 10, 9]. Secondly, there is an industrial trend towards *agile enterprises*: small, organizational overhead avoiding enterprises that form short term alliances to be able to respond to larger and more diverse orders than they individually could. Such ventures can take advantage of economies of scale when they are available, but do not suffer from diseconomies of scale.

In multiagent systems consisting of self-interested agents, contracts have traditionally been binding [18, 19, 6]. Once an agent agrees to a contract, it has to follow through with it no matter how future events unravel. Although a contract may be profitable to an agent when viewed *ex ante*, it need not be profitable when viewed after some future events have occurred, i.e. *ex post*. Similarly, a contract may have too low expected payoff *ex ante*, but in some realizations of the future events, the same contract may be desirable when viewed *ex post*. Normal full commitment contracts are unable to efficiently take advantage of the possibilities that such probabilistically known future events provide.

On the other hand, many multiagent systems consisting of cooperative agents incorporate some form of decommitment possibility in order to allow the agents to accommodate new events. For example, in the original Contract Net Protocol [25], 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 [24] and in cooperative coordination protocols [2]. Again, the agents did not require a monetary compensation for their efforts: an agent agreed to cancel a contract merely based on the fact that some other agent wanted to decommit. This research was descriptive: what will happen if agents use certain externally specified strategies.

Multiagent systems consisting of self-interested agents require that we consider the case where agents choose their own strategies and do not necessarily follow externally specified ones. Thus the interaction protocols need to be considered from a normative perspective: given a protocol, what is the best strategy from a self-interested viewpoint that each agent should (and thus will) choose, and then what are the social outcomes.

Some normative research in game theory has focused on utilizing the potential provided by probabilistically known future events by *contingency contracts* among self-interested agents [17]. 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 corresponding fall-back positions, but there is a contingency contract that each agent prefers over its fall-back.

There are at least three problems regarding the use of contingency contracts in automated negotiation among self-interested agents. Though useful in anticipating a small number of key events, contingency contracts get cumbersome as the number of relevant events to monitor from the future increases. In the limit, 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 should 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 [23, 19, 18]. Thus there is a potential combinatorial explosion of events to be conditioned on. Secondly, even if it were feasible to use such cumbersome contingency contracts among the computerized agents, it is often impossible to enumerate all possible relevant future events in advance. The third problem is that of verifying the unraveling of the events. Sometimes an event is only observable by one of the agents. This agent may have an incentive to lie to the other party of the contract about the event in case the event is associated with an unadvantageous contingency to the directly observing agent. Thus, to be viable, contingency contracts would require an event verification mechanism that is not manipulable and not prohibitively complicated.

We propose another method for taking advantage of the possibilities provided by probabilistically known future events in a contracting setting with self-interested agents. 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 this decommitment feature into a contracting protocol. In [23] we presented an example protocol that uses this feature.

Principles for assessing decommitment penalties have been studied in law [1, 16], but the purpose has been to assess a penalty on the agent that has breached the contract *after the breach has occurred.* Similarly, penalty clauses for partial failure (such as not meeting a deadline) are commonly used in contracts, but the purpose is usually to motivate the agents to follow the contract. To our knowledge, the possibility of explicitly allowing decommitment from the whole contract for a predetermined price has not been studied as an active method for utilizing the potential provided by an uncertain future. Somewhat unintuitively, it turns out that the mere existence of a decommitment possibility in a contract can increase each agent's expected payoff.

To discuss these issues more specifically, some concepts from microeconomics are introduced. First, *social welfare* measures the sum of the agents' payoffs [7, 4]. It is a measure of how good the payoffs are to the society of agents: it does not measure distribution aspects. *Pareto efficiency* measures both the societal goodness of a solution and distribution aspects [7, 4]. A vector of payoffs to the agents Pareto dominates another vector if each agent's payoff in the first vector is no less than in the second, and there exists an agent whose payoff in the first vector is greater than in the second. Social welfare and Pareto efficiency can be measured either *ex ante* as expected values or *ex post* as the realizations. Third, the strategies (mappings from observed history of the game to actions) S_a of the contractor and S_b of the contractee are in *Nash equilibrium* if S_a is a best (expected payoff maximizing) response to S_b, and S_b is a best response to S_a [12, 7, 4]. Finally, an agent's strategy is a *dominant strategy* if it is the best response to *any* strategy that the other agent can choose [7, 4].

We will analyze the contracting situation from the perspective of two agents: the contractor (who pays to get a task done) and the contractee (who gets paid for handling the task). The contractor tries to minimize the contract price ρ that it has to pay in the contract. The contractee tries to maximize the payoff ρ that it receives from the contractor. Outside offers from third parties will be explicitly discussed. The contracting setting consists of two games. First, the *contracting game* involves the agents choosing a contract (or no contract, i.e. the *null deal*) before any future events have unraveled. Secondly, the *decommitting game* involves the agents deciding on whether to decommit or to carry out the obligations of the contract—after the future events have unraveled. The decommitment game is a subgame of the contracting game: the expected outcomes of the decommitting game affect the agents' preferences over contracts in the contracting game. The decommitting game will be analyzed using the Nash equilibrium and the dominant strategy concepts. The contracting game will be analyzed with respect to *individual rationality* (IR): is the contract better for an agent than the null deal? Often there is either no contract that is IR for both agents or then there are many such contracts. When there are many IR contracts to choose from, there are infinitely 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. Actually, *axiomatic bargaining theory* [11, 13, 15] studies the question of choosing among these IR contracts by asserting desirable properties that the chosen contract should fulfill compared to the other contracts.

In Sections 2 and 3 we analyze the advantage of the leveled commitment contracting protocol compared to the full commitment one using different contracting settings. These sections and their subsections are ordered so that simpler settings precede the more complex ones. Section 2 describes settings where only one agent's future involves uncertainty, while Section 3 describes settings where both agents' futures involve uncertainty. The symbols used in the upcoming sections are summarized in Table 1. In Section 4 some practical prescriptions are given to builders of automated negotiation systems. Section 5 concludes and presents future research directions.

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) full commitment outside offer.
$\breve{b}$	Price of the contractee's (best) full commitment outside offer.
$f(\breve{a})$	*Ex ante* probability density function of $\breve{a}$.
$g(\breve{b})$	*Ex ante* probability density function of $\breve{b}$.
p_a	Probability that the contractor decommits.
p_b	Probability that the contractee decommits.
$\bar{b}$	Contractee's fall-back, i.e. payoff that it gets if no contract is made.
B	Big positive number.

Table 1: *Symbols used in the analysis of the leveled commitment protocol in the example contracting settings. We restrict our analysis to contracts where $a \geq 0$ and $b \geq 0$, i.e. we rule out contracts that specify that the decommitting agent receives a payment (for decommitting) from the victim of the decommitment.*

2 Uncertainty about one agent's outside offer

This section presents games where one agent's outside offer is fixed and known to both agents at contract time but uncertainty prevails about the price of the other agent's outside offer. Let the contractee have the deterministically known outside offer $\breve{b}$. Let the contractor's (best) outside offer be only known probabilistically by the agents via a probability density function $f(\breve{a})$. In case the contractor receives no outside offer, $\breve{a}$ is the best of its outstanding outside offers and its fall-back payoff. The case where the contractor's outside offer is deterministically known but the contractee's outside offer is only probabilistically known is analogous. The contract is made when the contractee's outside offer is known but the contractor's is not. On the other hand, the decommitting game takes place when the contractor has found out the value of $\breve{a}$. We assume that at this point the contractee does not know the contractor's outside offer $\breve{a}$. This seems realistic in automated contracting systems. Now there are two cases depending on whether the contractee's outside offer stays valid up to the point when the contractor finds out about its outside offer—i.e. up to when the contractor decides between decommitting and following through with the obligations of the contract.

2.1 The deterministic offer prevails (DOP)

In this case, the contractee's outside offer stays valid up to the point when the contractor finds out about its outside offer. So if the contractor decides to decommit, the contractee can still accept the outside offer. We will call this situation the DOP game, Figure 1.

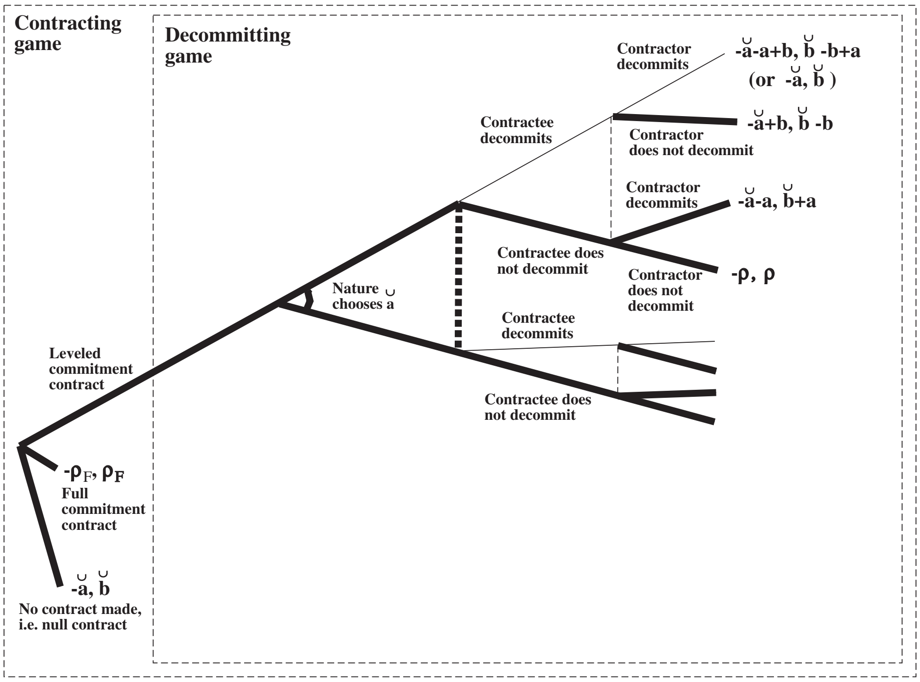

Figure 1: *The Deterministic Offer Prevails (DOP) game. If the contractor decommits, the contractee can still accept the outside offer. In the figure, the contractor's payoff is listed before the contractee's. The bold solid lines show choices that may actually occur in any subgame. The bold dashed line represents the contractee's information set: it does not know ă in the decommitting game. The thin dashed lines represent the alternative situation where both agents reveal decommitment simultaneously: when deciding on decommitting, the contractor has not observed whether the contractee decommitted.*

In a sequential decommitting DOP game where the contractee reveals decommitment first (Figure 1), 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 1). 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. If the contractor does not decommit, the contractee is motivated to abide to the contract because—by the very fact that the contractee accepted the contract—its outstanding outside offer is worth less than the contract, i.e. $\breve{b} \leq \rho$. 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 1), 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. 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 $E[-\breve{a}] \leq E[max[-\breve{a} - a, -\rho]]$:

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

Now that the game has been specified, the advantages of a leveled commitment contract can be analyzed. Obviously in this game the full commitment contracts are a subset of leveled commitment ones because a leveled commitment contract can emulate a full commitment one by choosing a high decommitment penalty $a = B$ that motivates the contractor to surely not decommit (assuming that $\breve{a}$ is bounded from below). Therefore, for any full commitment contract, there exists a leveled commitment contract that has no worse payoff to either agent. Furthermore, the following two theorems state the strict superiority of leveled commitment contracts over full commitment ones in this game. The first theorem states that in some DOP games the agents can not make a beneficial contract under the full commitment protocol, but can under the leveled commitment one.

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

Proof. Example. Let $\breve{b} = 55$ and let $f(\breve{a}) = \begin{cases} 0.01 & \text{if } 0 \leq \breve{a} \leq 100 \\ 0 & \text{otherwise} \end{cases}$. Now there exists no full commitment contract F that satisfies both IR constraints because that would require $\breve{b} \leq \rho_F \leq E[\breve{a}]$ which is impossible because $55 = \breve{b} \geq E[\breve{a}] = 50$. Let us analyze a leveled commitment contract where $\rho = 60$, $a = 10$, (and $b \geq 0$). The contractor's IR constraint is satisfied:

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

$$\Leftrightarrow \int_{0}^{100} f(\breve{a})[-\breve{a}]d\breve{a} \leq \int_{0}^{60-10} f(\breve{a})[-\breve{a} - 10]d\breve{a} + \int_{60-10}^{100} f(\breve{a})[-60]d\breve{a}$$

$$\Leftrightarrow -50 \leq -17.5 - 30$$

Now, $p_a = \int_{-\infty}^{\rho-a} f(\breve{a})d\breve{a} = \int_{0}^{60-10} f(\breve{a})d\breve{a} = 0.5$. Thus the contractee's IR constraint is also satisfied:

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

$$\Leftrightarrow 55 \leq [1-0.5]60 + 0.5[55+10]$$

$$\Leftrightarrow 55 \leq 62.5$$

□

The next theorem states that even if both protocols would allow a beneficial contract, the leveled commitment protocol may allow higher expected payoffs for both agents than the full commitment protocol. This holds as long as there is some chance that the contractor's outside offer will be lower than the contractee's.

Theorem 2.2 ***Pareto efficiency improvement in a DOP game.*** *Let F be an arbitrary full commitment contract ($a_F = B$) that satisfies both IR constraints, i.e. $\breve{b} \leq \rho_F \leq E[\breve{a}]$. If $\exists \breve{a} < \breve{b}$ s.t. $f(\breve{a}) \neq 0$ (i.e. $f(\breve{a})$ extends below $\breve{b}$), then there exists a leveled commitment contract that increases the contractor's payoff and the contractee's payoff (and thus also satisfies the IR constraints).*

Proof. Under F, the contractor's payoff is $-\rho_F$, and the contractee's ρ_F. Now we construct a leveled commitment contract L defined by ρ, a, and b. Let $a = \rho_F - \breve{b}$, (and $b \geq 0$). Choose $\rho = \rho_F + \epsilon$ for some $\epsilon > 0$. Now the contractee's expected payoff has increased because the contractee will get $\breve{b} + a = \rho_F$ at worst and there is a nonzero probability that the contractee will get $\rho > \rho_F$. The contractor's payoff has increased by

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

Case 1: $\breve{b} = \rho_F$: Now,

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

$$= \underbrace{\int_{-\infty}^{\rho_F} f(\breve{a})[\underbrace{-\breve{a}+\rho_F}_{>0}]d\breve{a}}_{>0} + \int_{\rho_F}^{\rho_F+\epsilon} f(\breve{a})[-\breve{a}+\rho_F]d\breve{a} + \int_{\rho-a}^{\infty} f(\breve{a})[-\epsilon]d\breve{a}$$

Now $\epsilon > 0$ can be chosen small enough so that this expression is positive.

Case 2: $\breve{b} < \rho_F$**:** Now let $\epsilon \leq \rho_F - \breve{b}$. Thus,

$$\begin{aligned} & \int_{-\infty}^{\epsilon+\breve{b}} f(\breve{a})[-\breve{a}+\rho_F]d\breve{a} + \int_{\rho-a}^{\infty} f(\breve{a})[-\epsilon]d\breve{a} \\ = \; & c - \int_{\rho-a}^{\infty} f(\breve{a})\epsilon d\breve{a} \end{aligned}$$

where $c > 0$. Now $\epsilon > 0$ can be chosen small enough so that this expression is positive. $\square$

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

2.1.1 Effect of biased asymmetric information in DOP games

The following theorem shows the desirable property that an agent cannot be hurt by its negotiation partner's biased beliefs in a DOP game. 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. In the theorem, let $f_a(\breve{a})$ be the contractor's belief of $f(\breve{a})$, and let $f_b(\breve{a})$ be the contractee's belief of $f(\breve{a})$.

Theorem 2.3 ***Contract payoff unaffected by other agent's beliefs in a DOP game.*** *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 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. $\square$

2.2 The certain offer becomes void (COBV)

This section discusses the 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 fall-back payoff). Otherwise the $\breve{b}$-offer becomes void. Thus, to agree to the contract, the contractee has to have a higher expected payoff when 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 fall-back payoff $\bar{b}$ plus the contractor's decommitment penalty payment a. The fall-back $\bar{b}$ 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 COBV game, Figure 2.

In a sequential decommitting COBV game where the contractee reveals decommitment first (Figure 2), 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 2). 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. If the contractor does not decommit, the contractee is motivated to abide to the contract because—by the very fact that the contractee accepted the contract—its fall-back payoff is worth less than the contract,

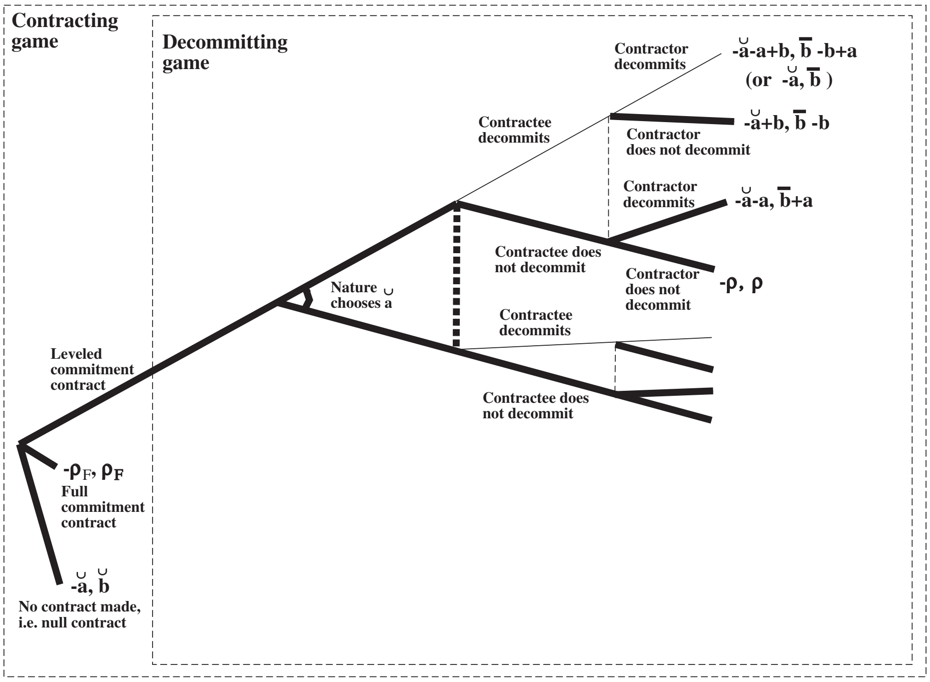

Figure 2: *The "Certain Offer Becomes Void" (COBV) game. If at all, the contractee's outside offer $\breve{b}$ has to be accepted before the contractor's outside offer $\breve{a}$ becomes known. The bold solid lines show choices that may actually occur in any subgame. The bold dashed line represents 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 decommitment simultaneously: when deciding on decommitting, the contractor has not observed whether the contractee decommitted.*

i.e. $\bar{b} \leq \rho$. 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 2), 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 fall-back position, i.e. the payoff it gets if it does not get its outside offer $\breve{b}$ or the contract with the contractor. Note that if $\bar{b} = \breve{b}$, this situation reduces to a DOP game (Section 2.1), and thus the DOP results in favor of leveled commitment hold. 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.

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

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

Full commitment contracts are a subset of leveled commitment ones because the contractor's decommitment penalty can be chosen so high ($a = B$) that the contractor will surely not decommit (assuming that $\breve{a}$ is bounded from below). As discussed earlier, the contractee will not decommit either for any $b \geq 0$. Thus, the class of leveled commitment contracts is no worse than the class of full commitment ones. Although there are games where a leveled commitment contract is Pareto superior to any full commitment contract if the contractee's fall-back is sufficiently high (a COBV game reduces to a DOP game if $\bar{b} = \breve{b}$), the following two theorems show that if the contractee's fall-back is too low, leveled commitment contracts are not helpful in COBV games.

Theorem 2.4 ***No enabling in a COBV game.*** *Let us restrict to COBV games where* $\bar{b} \leq \frac{\int_{-\infty}^{\rho - a} f(\breve{a}) \breve{a} d\breve{a}}{\int_{-\infty}^{\rho - a} f(\breve{a}) d\breve{a}}$. *In such a game (defined by* $\breve{b}$*,* $f(\breve{a})$*, and* $\bar{b}$*), if no* *full commitment contract satisfies the IR constraints, no leveled commitment contract satisfies them either.*

Proof. For a full commitment contract F, $\check{b} \leq \rho_F \leq E[\breve{a}]$. No such F exists iff $\check{b} > E[\breve{a}]$. Now say that $\check{b} > E[\breve{a}]$, and assume—for contradiction—that some leveled commitment contract L defined by ρ, a, and b satisfies both IR constraints. Thus,

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

Contradiction. Thus no such contract L satisfies both IR constraints. □

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 theorems with negative results have limited scope.

Theorem 2.5 ***No Pareto efficiency improvement in a COBV game.*** *Let us restrict to COBV games where* $\bar{b} \leq \frac{\int_{-\infty}^{\rho-a} f(\breve{a})\breve{a}d\breve{a}}{\int_{-\infty}^{\rho-a} f(\breve{a})d\breve{a}}$. *Let* F *be an arbitrary full commitment contract* ($a_F = B$) *that satisfies both IR constraints, i.e.* $\check{b} \leq \rho_F \leq E[\breve{a}]$. *Now 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

L (defined by ρ, a, and b) that increases at least one of these payoffs while not decreasing the other, i.e.

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

Contradiction. Thus no such L exists. □

2.2.1 Effect of biased asymmetric information in COBV games

This section discusses COBV games where the agents' beliefs differ. Specifically, let $f_a(\breve{a})$ be the contractor's belief of $f(\breve{a})$, and let $f_b(\breve{a})$ be the contractee's belief of $f(\breve{a})$, and 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 theorem 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.

Theorem 2.6 ***Perceived enabling in a COBV game.*** *There are 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 ($a_F = B$) 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 Theorem 2.4 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. $\square$

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 fall-back payoff is sufficiently low, both agents know (by Theorem 2.4) 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 an expected payoff of what it thinks it has independent of the other agents beliefs (stemming from a reasoning process or information sources). 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.

Theorem 2.7 ***Contract payoff unaffected by other agent's beliefs in a COBV game.*** *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 2.1 ***Perceived IR contracts are IR for the agent with unbiased information in a COBV game.*** *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 that agent.*

Proof. By definition, a contract is IR for the contractor if it is preferred over the null deal. But by Theorem 2.7 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 Theorem 2.7 the preference ordering is not affected by the contractor's information. $\square$

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 about $f(\breve{a})$.

3 Uncertainty about both agents' outside offers

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

The DOP game that was discussed in Section 2.1 is a special case of this type of games. In the DOP game, all of the probability mass of $g(\breve{b})$ is on one value $\breve{b}$. The DOP game is also a special case of the COBV game described in Section 2.2—when $\bar{b} = \breve{b}$. On the other hand, COBV games are not a subset of this type of games: in COBV games some opportunity (outside offer) may be missed due to waiting for the unraveling of the new outside offers.

We do not assume that the agents decommit truthfully. For example, an agent may not decommit although its outside offer is better for itself than the contract, because the agent believes that there is a high probability that the opponent is going to decommit. This would save the agent its decommitment penalty and in fact make the agent receive a decommitment penalty from the opponent. Games of this type differ significantly based on whether the agents decommit sequentially or simultaneously. The next two sections analyze these cases in detail.

3.1 Sequential decommitting (SEQD)

In our sequential decommitting (SEQD) game, one agent has to declare decommitment 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 3. 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 dominance as the solution concept: reasoning about the agents' actions starts from the leaves of the game tree and proceeds backwards. 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 $a \geq 0$). The same holds even for a contract that specifies that neither agent has to pay a decommitment penalty if both decommit—because $-\breve{a} \leq -\breve{a} + b$, (Fig. 3 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 $\int_{-\infty}^{\rho-a} f(\breve{a}) d\breve{a}$. Put together, the contractee gets payoff $\breve{b} - b$ if it decommits, $\breve{b} + a$ if it does not decommit but the contractor does, and ρ if neither decommits. Thus the contractee decommits if

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

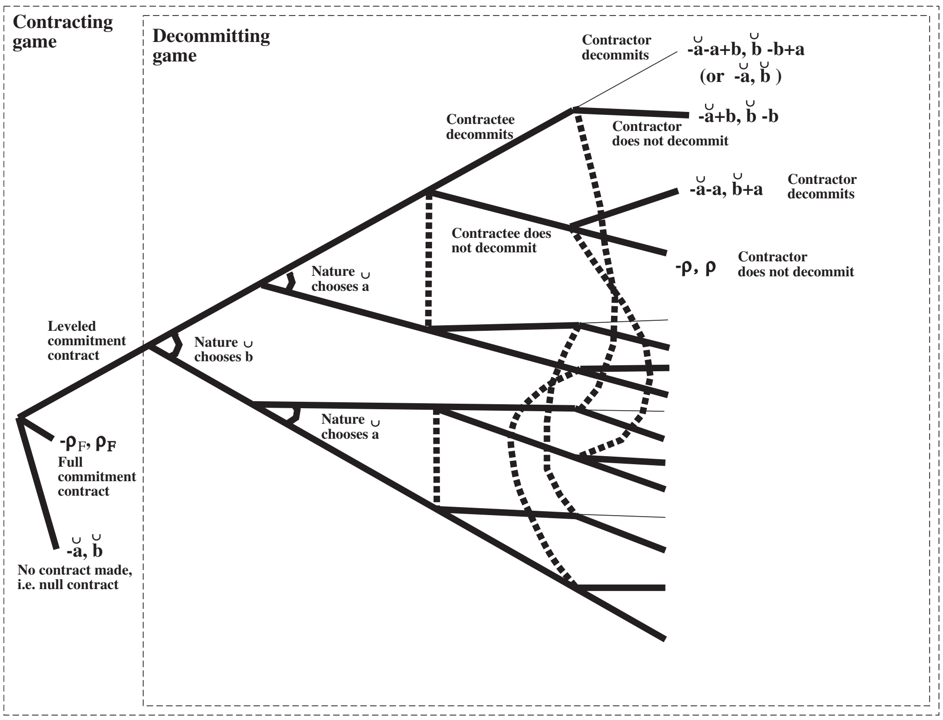

Figure 3: *The "Sequential Decommitting" (SEQD) game. The game tree of the figure represents two alternative protocols (and therefore two different games). In the first, both agents have to pay the decommitment penalties to each other if both decommit. In the second, neither agent has to pay if both decommit. The payoffs of the latter protocol (when different from the other protocol's) are in parentheses. The dotted lines represent information sets: the contractor does not know the contractee's outside offer and vice versa.*

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

$$\breve{b} > \rho + \frac{b + \int_{-\infty}^{\rho-a} f(\breve{a})d\breve{a}[a]}{\int_{\rho-a}^{\infty} f(\breve{a})d\breve{a}} \stackrel{def}{=} \breve{b}^*(\rho, a, b) \text{ when } \int_{\rho-a}^{\infty} f(\breve{a})d\breve{a} > 0 \quad (1)$$

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:

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

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

$$\geq E[\breve{b}] = \int_{-\infty}^{\infty} g(\breve{b})\breve{b}d\breve{b} \tag{2}$$

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

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

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

$$\geq E[-\breve{a}] = \int_{-\infty}^{\infty} f(\breve{a})[-\breve{a}]d\breve{a} \tag{3}$$

Because the contractor can want to decommit only if $-\breve{a} - a > -\rho$, its decommitment penalty can be chosen so high ($a = B$) 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 ($b = B$) that it will surely not decommit. Thus, full commitment contracts are a subset of leveled commitment ones. Note that 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 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 SEQD games, leveled commitment contracts can enable a deal that is not possible via full commitment contracts.

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

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 analyze 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

$$
\begin{aligned}
& \int_{\breve{b}^*(\rho,a,b)}^{\infty} g(\breve{b}) \int_{-\infty}^{\infty} f(\breve{a})[-\breve{a} + b] d\breve{a} d\breve{b} \\
& + \int_{-\infty}^{\breve{b}^*(\rho,a,b)} g(\breve{b}) [\int_{-\infty}^{\rho-a} f(\breve{a})[-\breve{a} - a] d\breve{a} + \int_{\rho-a}^{\infty} f(\breve{a})[-\rho] d\breve{a}] d\breve{b} \\
\geq \; & \int_{-\infty}^{\infty} f(\breve{a})[-\breve{a}] \\
\Leftrightarrow \; & \int_{\breve{b}^*(\rho,a,b)}^{110} g(\breve{b}) \int_{0}^{100} f(\breve{a})[-\breve{a} + 20] d\breve{a} d\breve{b} \\
& + \int_{0}^{\breve{b}^*(\rho,a,b)} g(\breve{b}) [\int_{0}^{52.5-30} f(\breve{a})[-\breve{a} - 30] d\breve{a} + \int_{52.5-30}^{100} f(\breve{a})[-52.5] d\breve{a}] d\breve{b} \\
\geq \; & \int_{0}^{100} f(\breve{a})[-\breve{a}] \\
\Leftrightarrow \; & \int_{\breve{b}^*(\rho,a,b)}^{110} \frac{1}{110} \frac{1}{100} [\frac{-(100)^2}{2} + 100 \cdot 20] d\breve{b} + \\
& \int_{0}^{\breve{b}^*(\rho,a,b)} \frac{1}{110} [\frac{1}{100} [\frac{-(22.5)^2}{2} - 22.5 \cdot 30] + \frac{1}{100} [-52.5 \cdot (100 - 22.5)]] d\breve{b} \\
\geq \; & -50 \\
\Leftrightarrow \; & \int_{\breve{b}^*(\rho,a,b)}^{110} \frac{15.5}{110} d\breve{b} + \int_{0}^{\breve{b}^*(\rho,a,b)} -\frac{49.96875}{110} d\breve{b} \geq -50
\end{aligned}
$$

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

$$
\begin{aligned}
& \int_{\breve{b}^*(\rho,a,b)}^{\infty} g(\breve{b})[\breve{b} - b] d\breve{b} + \int_{-\infty}^{\breve{b}^*(\rho,a,b)} g(\breve{b}) [\int_{-\infty}^{\rho-a} f(\breve{a})[\breve{b} + a] d\breve{a} + \int_{\rho-a}^{\infty} f(\breve{a}) \rho d\breve{a}] d\breve{b} \\
\geq \; & \int_{-\infty}^{\infty} g(\breve{b}) \breve{b} d\breve{b} \\
\Leftrightarrow \; & \int_{\breve{b}^*(\rho,a,b)}^{110} \frac{1}{110} [\breve{b} - 20] d\breve{b} \\
& + \int_{0}^{\breve{b}^*(\rho,a,b)} \frac{1}{110} [\int_{0}^{52.5-30} \frac{1}{100} [\breve{b} + 30] d\breve{a} + \int_{52.5-30}^{100} \frac{1}{100} 52.5 d\breve{a}] d\breve{b} \geq 55
\end{aligned}
$$

$$\Leftrightarrow \quad \int_{\breve{b}^*(\rho,a,b)}^{110} \frac{1}{110}[\breve{b} - 20]d\breve{b} + \int_0^{\breve{b}^*(\rho,a,b)} \frac{1}{110}[\frac{22.5}{100}[\breve{b} + 30] + \frac{77.5}{100}52.5]d\breve{b} \geq 55$$

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

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 game of the above proof, both IR constraints are satisfied by a wide range of leveled commitment contracts—and for no full commitment contract. Which leveled commitment contracts defined by ρ, a, and b satisfy the IR constraints? There are many values of ρ for which some a and b exist such that the constraints are satisfied. As in the above 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 3 and 2) for this case. To make the algebra tractable (constant $f(\breve{a})$ and $g(\breve{b})$), versions of these IR constraint equations were used that assumed $0 \leq a < \rho$, and $0 < \breve{b}^* < 110$. The corresponding decommitment penalties a and b that satisfy the IR constraints are plotted in Figure 4 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 4 left. Obviously the constraint $\breve{b}^* > 0$ is always satisfied in this case and is thus not plotted. To summarize, in the gray area of Figure 4 left, the contracts are IR for both agents, given that the agents decommit optimally—not necessarily truthfully.

Case 2, Contractor will surely not decommit. In the case where $a \geq \rho$, 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

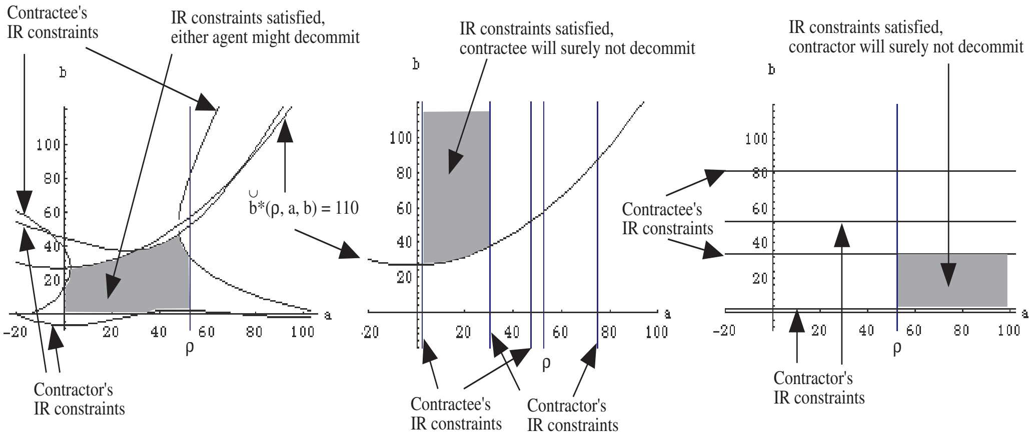

Figure 4: *The decommitment penalties a and b that satisfy both agents IR constraints in the example. Right: case where either agent might decommit ($a < \rho$, and $\breve{b}^*(\rho, a, b) < 110$). Middle: case where the contractor might decommit but the contractee will not ($a < \rho$, and $\breve{b}^*(\rho, a, b) \geq 110$). Left: case where $a \geq \rho$, i.e. the contractor will surely not decommit but the contractee might.*

truthfully. Now the contractor's IR constraint becomes

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

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

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

by the quadratic equation solution formula.

Similarly, the contractee's IR constraint becomes

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

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 optimally (not necessarily truthfully). This region is colored gray in Figure 4 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 $-\breve{a} - a > -\rho \Leftrightarrow \breve{a} < \rho - a$. Let us define $\breve{a}^* = \rho - a$. The contractor's IR constraint becomes

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

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

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

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 optimally (not necessarily truthfully). This region is colored gray in Figure 4 middle.

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 3.2 ***Pareto efficiency improvement in a SEQD game.*** *There exist SEQD games with IR full commitment contracts where the best full commitment contract has lower payoff to each agent than the best leveled commitment contract (which is also thus IR).*

Proof. A DOP game is equivalent to a SEQD game where all the probability mass of $g(\breve{b})$ is on one $\breve{b}$. Thus the result follows from Theorem 2.2. $\square$

3.1.1 Effect of biased asymmetric information in SEQD games

Unlike in the DOP and COBV games of Sections 2.1 and 2.2, in SEQD games, one agent's expected payoff for a given contract may depend on the other agent's—possibly biased—beliefs. For example, the contractee's decision of whether to decommit depends on its belief $f_b(\breve{a})$ of the contractor's 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):

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

3.2 Simultaneous decommitting

In our simultaneous decommitting games, both agents have to declare decommitment 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 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 5 presents the games induced by both of these contract types. The next two sections analyze these alternative games in detail.

3.2.1 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 5. Let us define the variable p_b as 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

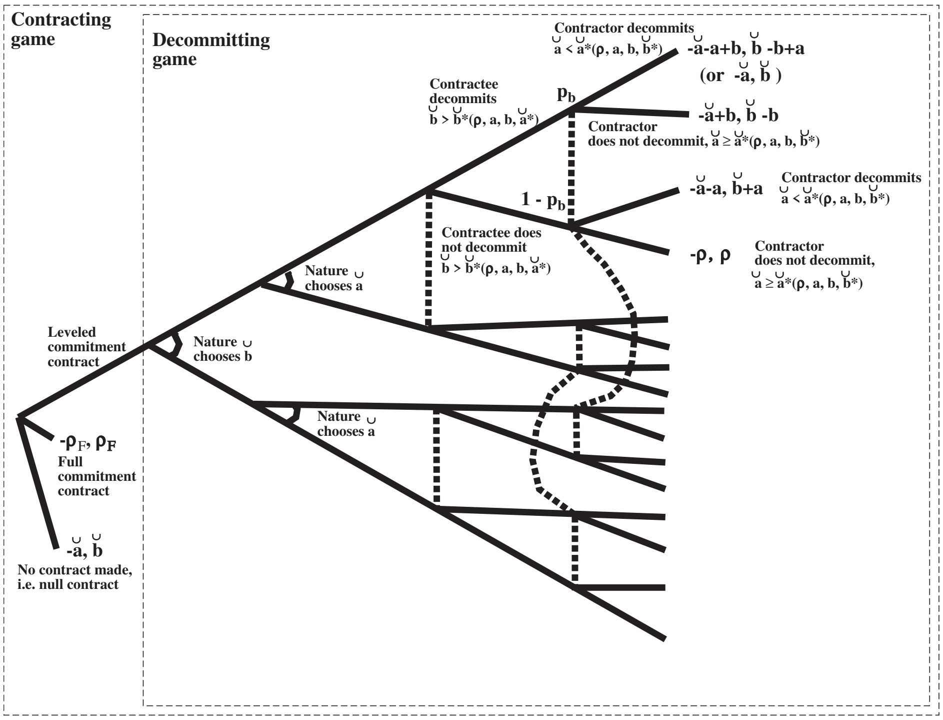

Figure 5: *The "SIMUltaneous Decommit - Both Pay if both decommit" (SIMUDBP) game. The parenthesized payoffs represent the "SIMUltaneous Decommit - Neither Pays if both decommit" (SIMUDNP) game. The dashed lines represent the agents' information sets. When decommitting, the contractor does not know the 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).*

decommit if

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

If $p_b = 1$, this 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

$$\breve{a} < \rho - \frac{a}{1 - p_b} \stackrel{def}{=} \breve{a}^*(\rho, a, b, \breve{b}^*) \text{ when } p_b < 1 \tag{4}$$

Thus we have characterized a decommitting threshold $\breve{a}^*$ for the contractor. If the contractor's outside offer $\breve{a} < \breve{a}^*$, the contractor is best off by decommitting. The contractee decommits if

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

If $\int_{\breve{a}^*(\rho,a,b,\breve{b}^*)}^{\infty} f(\breve{a})d\breve{a} = 0$, this 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

$$\breve{b} > \rho + \frac{b}{\int_{\breve{a}^*}^{\infty} f(\breve{a})d\breve{a}} \stackrel{def}{=} \breve{b}^*(\rho, a, b, \breve{a}^*) \text{ when } \int_{\breve{a}^*(\rho,a,b,\breve{b}^*)}^{\infty} f(\breve{a})d\breve{a} > 0 \qquad (5)$$

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

$$p_b = \int_{\breve{b}^*(\rho,a,b,\breve{a}^*)}^{\infty} g(\breve{b})d\breve{b} \qquad (6)$$

Condition 4 states the contractor's best response (defined by $\breve{a}^*$) to the contractee's strategy that is defined by $\breve{b}^*$. Condition 5 states the contractee's best response $\breve{b}^*$ to the contractor's strategy that is defined by $\breve{a}^*$. Condition 4 uses the variable p_b which is defined by Equation 6. So together, Equations 4, 5, and 6 define the Nash equilibria of the decommitting game.

Now the contractor's IR constraint becomes

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

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

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

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

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

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 the Nash equilibrium decommitment thresholds $\breve{a}^*$ and $\breve{b}^*$ really do differ from the truthful ones. Yet there exist Nash equilibria

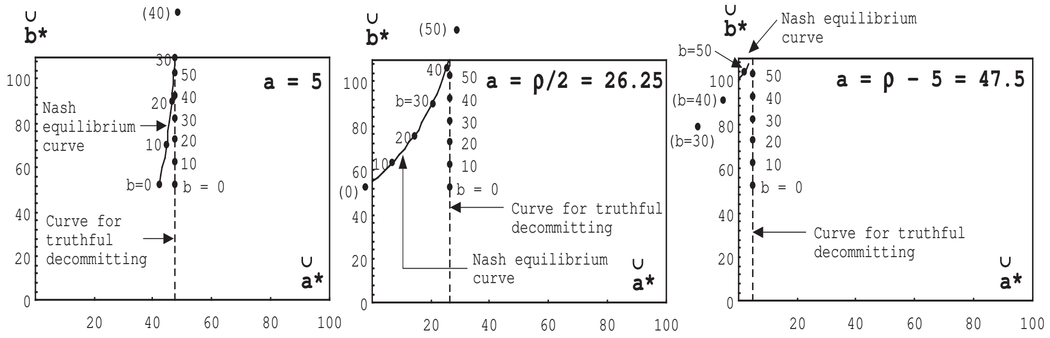

Figure 6: *The Nash equilibrium decommitment thresholds* $\breve{a}^*$ *and* $\breve{b}^*$ *of our example SIMUDBP game for different values of the decommitment penalties* a *and* b*. The Nash equilibrium deviates from truthful decommitting. If* $0 < \breve{a}^* < 100$*, and* $0 < \breve{b}^* < 110$*,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

within the proper range of $\breve{a}^*$ and $\breve{b}^*$. These Nash equilibria need not satisfy the agents' IR constraints however.

We programmed a model of Equations 4, 5, and 6 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 boundaries of the validity of the model. The boundaries $\breve{a}^* = 0$ and $\breve{b}^* = 110$ are plotted in Figure 7. The boundary $\breve{a}^* = 100$ turns out to be the line $b = 0$. There exists no boundary $\breve{b}^* = 0$ because $\breve{b}^*$ was always greater than zero.

After plotting the validity boundaries of the model, the curves at which the IR constraints held with equality were plotted, see 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 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.

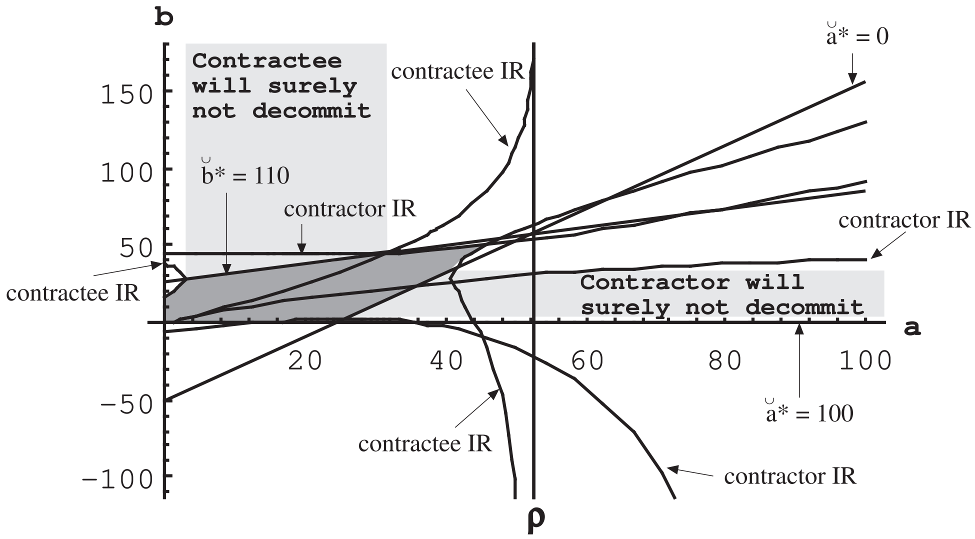

Figure 7: *The gray areas are three qualitatively different regions of contracts that are IR for both agents and allow an equilibrium in the decommitting game. In the dark gray area either agent might decommit while in the light gray areas only one agent might decommit. The curves represent the IR constraints 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.*

As a numeric 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 6. 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, we proceed to present the other types of equilibria that can occur in the example.

Case 2, 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 payoff is transferred.

Case 3, 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

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

This is the same situation as in the case of sequential decommitting which was discussed earlier: once it is known that the contractor will not decommit, the contractee is best off by decommitting truthfully. In the discussion of sequential decommitting, the above inequality was shown to be equivalent to $2.5 \leq b \leq 52.5$. Similarly, the contractee's IR constraint becomes

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

This is also same as in the case of sequential decommitting, where the above was shown to be 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 with light gray in Figure 7.

Case 4, 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

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

This is the same situation as in the case of sequential decommitting which was discussed earlier: once it is known that the contractee will not decommit, the contractor is best off by decommitting truthfully. In the discussion of sequential decommitting, the above inequality was shown to be equivalent to $a \leq$ approximately 30.14. Similarly, the contractee's IR constraint becomes

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

This is also same as in the case of sequential decommitting, where the above was shown to be 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 with light gray in Figure 7. □

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 3.4 ***Pareto efficiency improvement in a SIMUDBP game.*** *There exist SIMUDBP games with IR full commitment contracts where the best full commitment contract has lower payoff to each agent than the best leveled commitment contract (which is also thus IR).*

Proof. A DOP game is equivalent to a SIMUDBP game where all the probability mass of $g(\breve{b})$ is on one $\breve{b}$. Specifically, in such a game, if the contractee has found the contract IR, it will surely not decommit. The result of the theorem follows from Theorem 2.2. □

3.2.2 Effect of biased asymmetric information in SIMUDBP games

In SIMUDBP games—like SEQD games but unlike DOP and COBV games—one agent's expected payoff for a given contract may depend on the other agent's—possibly biased—beliefs. For example, the contractee's decision of whether to decommit depends on its belief $f_b(\breve{a})$ of the contractor's upcoming outside offer. For example, 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.

3.2.3 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 5. 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

$$\breve{a} < \rho - a - \frac{bp_b}{1 - p_b} \stackrel{def}{=} \breve{a}^*(\rho, a, b, \breve{b}^*) \text{ when } p_b < 1 \tag{7}$$

The contractee decommits if

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

If $\int_{\breve{a}^*(\rho,a,b,\breve{b}^*)}^{\infty} f(\breve{a})d\breve{a} = 0$, this 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

$$\breve{b} > \rho + b - \frac{a\int_{-\infty}^{\breve{a}^*(\rho,a,b,\breve{b}^*)} f(\breve{a})d\breve{a}}{\int_{\breve{a}^*(\rho,a,b,\breve{b}^*)}^{\infty} f(\breve{a})d\breve{a}} \stackrel{def}{=} \breve{b}^*(\rho, a, b, \breve{a}^*) \text{ when } \int_{\breve{a}^*(\rho,a,b,\breve{b}^*)}^{\infty} f(\breve{a})d\breve{a} > 0 \tag{8}$$

The probability that the contractee will decommit is

$$p_b \;=\; \int_{\breve{b}^*(\rho,a,b,\breve{a}^*)}^{\infty} g(\breve{b})d\breve{b} \tag{9}$$

Condition 7 states the contractor's best response (defined by $\breve{a}^*$) to the contractee's strategy that is defined by $\breve{b}^*$. Condition 8 states the contractee's best response $\breve{b}^*$ to the contractor's strategy that is defined by $\breve{a}^*$. Condition 7 uses the variable p_b which is defined by Equation 9. So together, Equations 7, 8, and 9 define the Nash equilibria of the decommitting game.

Now the contractor's IR constraint becomes

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

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

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

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 3.5 ***Enabling in a SIMUDNP game.*** *There exist SIMUDNP games (defined by $f(\breve{a})$ and $g(\breve{b})$) where no full commitment contract satisfies* *the IR constraints but where a leveled commitment contract where neither pays a penalty if both decommit satisfies both IR constraints.*

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 8. Note that the Nash equilibrium decommitment

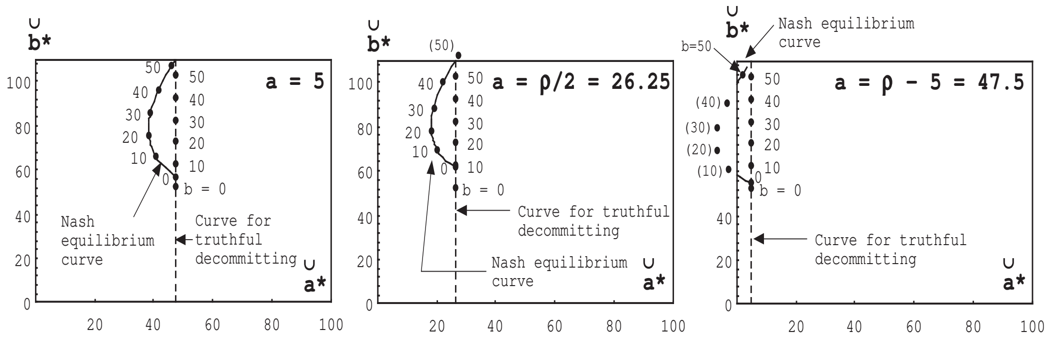

Figure 8: *The Nash equilibrium decommitment thresholds $\breve{a}^*$ and $\breve{b}^*$ of our example SIMUDNP game for different values of the decommitment penalties a and b. The Nash equilibrium deviates from truthful decommitting. If $0 < \breve{a}^* < 100$, and $0 < \breve{b}^* < 110$, there is some chance that either agent will decommit.*

thresholds $\breve{a}^*$ and $\breve{b}^*$ really do differ from the truthful ones. They also differ from (are closer to the truthful ones than) what they were when a protocol where both agents pay if both decommit was used, Figure 6. The shapes of the curves using these two protocols also differ significantly. Yet there exist Nash equilibria within the proper range of $\breve{a}^*$ and $\breve{b}^*$. These Nash equilibria need not satisfy the agents' IR constraints however.

We programmed a model of Equations 7, 8, and 9 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 9.

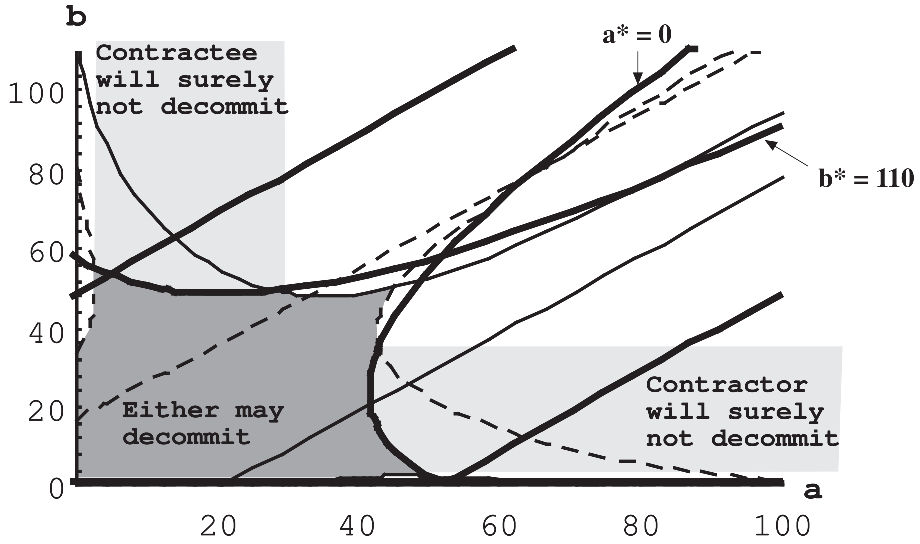

Figure 9: *Three qualitatively different regions of contracts that are IR for both agents and allow an equilibrium in the decommitting game. The bold lines are the validity constraints for the programmed model that requires* $0 < \breve{a}^* < 100$, *and* $0 < \breve{b}^* < 110$. *One of the constraints 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.*

After plotting the validity boundaries of the model, the curves at which the IR constraints held with equality were plotted, see Figure 9. 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 9 represents the values of the decommitment penalties a and b for which the validity constraints of the programmed model and the IR constraints are satisfied. In other words, for any such a and b, there exist decommitment thresholds $\breve{a}^*$ and $\breve{b}^*$ such that these form a Nash 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 numeric 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 8. 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, we proceed to present the other types of equilibria that can occur in the example.

Case 2, 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.

Case 3, 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. Now the contractor's IR constraint becomes

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

This is the same situation as in the case of sequential decommitting which was discussed earlier: once it is known that the contractor will not decommit, the contractee is best off by decommitting truthfully. It is also same as the situation where the simultaneous decommitment leveled commitment protocol was used where both pay the decommitment penalties if both decommit. In the discussion of sequential decommitting, the above inequality was shown to be equivalent to $2.5 \leq b \leq 52.5$. Similarly, the contractee's IR constraint becomes

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

This is also the same as in sequential decommitting and as with the simultaneous decommitment leveled commitment protocol where both pay the decommitment penalties if both decommit. In the discussion of sequential decommitting, the above inequality was shown to be equivalent to $b \leq$ approximately 34.05. Put together, 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 with light gray in Figure 9.

Case 4, 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

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

This is the same situation as in the case of sequential decommitting which was discussed earlier: once it is known that the contractee will not decommit, the contractor is best off by decommitting truthfully. It is also same as the situation where the simultaneous decommitment leveled commitment protocol was used where both pay the decommitment penalties if both decommit. In the discussion of sequential decommitting, the above inequality was shown to be equivalent to $a \leq$ approximately 30.14. Similarly, the contractee's IR constraint becomes

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

This is also the same as in sequential decommitting and as with the simultaneous decommitment leveled commitment protocol where both pay the decommitment penalties if both decommit. In the discussion of sequential decommitting, the above inequality was shown to be equivalent to $2.5 \leq a \leq 47.5$. Put together, 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 with light gray in Figure 9. $\square$

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 3.6 ***Pareto efficiency improvement in a SIMUDNP game.*** *There exist SIMUDNP games with IR full commitment contracts where the best full commitment contract has lower payoff to each agent than the best leveled commitment contract (which is also thus IR).*

Proof. A DOP game is equivalent to a SIMUDNP game where all the probability mass of $g(\breve{b})$ is on one $\breve{b}$. Specifically, in such a game, if the contractee has found the contract IR, it will surely not decommit. The result of the theorem follows from Theorem 2.2. $\square$

3.2.4 Effect of biased asymmetric information in SIMUDNP games

In SIMUDNP games—like SIMUDBP and SEQD games but unlike DOP and COBV games—one agent's expected payoff for a given contract may depend on the other agent's—possibly biased—beliefs. For example, the contractee's decision of whether to decommit depends on its belief $f_b(\breve{a})$ of the contractor's upcoming outside offer. For example, 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.

4 Practical prescriptions for system builders

The results from the above canonical games suggest that it is worthwhile from a contract enabling and a contract Pareto improving perspective to incorporate the decommitment mechanism into automated contracting protocols. The decommitment penalties are best chosen by the agents dynamically at contract time as opposed to statically in the protocol. This allows the tuning of the penalties not only to specific negotiation situations and environmental uncertainties, but also to specific belief structures of the agents.

The proposed decommitment mechanism allows an agent to decommit based on local reasoning: no negotiation is necessary at decommitment time. The contracts in this mechanism are simpler than traditional contingency contracts that require—in the worst case—the specification of the contract's alternative obligations for all alternative worlds induced by alternative realizations of combinations of future events. Furthermore, the proposed decommitment method does not require an event verification mechanism like contingency contracts do.

In the presented instance of the simultaneous decommitting game, the Nash equilibrium decommitting strategies 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). This suggests using the former protocol in practical systems. It also minimizes the number of payment transfers because it does not require any such transfer if both decommit.

In asynchronous negotiation systems, the judging of decommitment decisions can be implemented as follows. To decommit, an agent just sends a decommit message. To not decommit, an agent sends a no-decommit message. When both agents have sent and received one of these messages, payment transfer can take place. This method is non-manipulative. An agent cannot send a no-decommit message (which would be desirable after having received a decommit message from another agent) after having sent a decommit message. This is because the other agent would receive two messages and know that the former agent is manipulating the system.

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 swap from one such focus to another by decommitting from a contract—by paying the decommitment penalty. It may happen that one such swap makes it beneficial for another agent to decommit from another contract and so on. To avoid loops of decommitting and recommitting in practise, recommitting can be disabled. This can be implemented by 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 an identical content. This gives rise to the possibility of the equivalent of useless decommit-recommit loops. Such loops can be avoided by a mechanism where the decommitment penalties increase with 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.

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 [19, 23]. Even if explicit clustering of issues into contracts [19, 23] 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 [23]. 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 [23], 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 [23]. Again, if no such contracts appear, the agent can decommit.

In many practical automated contracting settings, agents are bounded rational—for example because limited computation resources bound their capability to solve combinatorially complex problems [23, 20, 21, 19]. 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 [23].

5 Conclusions and future research

A decommitment mechanism was presented for automated contracting protocols that—somewhat surprisingly—allows the agents to accommodate future events more profitably than traditional full commitment contracts. Each contract specifies a decommitment penalty for both agents involved. To decommit, an agent just pays that penalty to the other agent. This mechanism is better suited for complex computerized contracting settings than contingency contracts: potentially combinatorial and hard to anticipate contingencies need not be considered, no event verification mechanism is necessary, and decommitting can be decided based on local (*ex post*) deliberation. The method was analyzed using a normative approach: given the protocol, what is the strategy that each self-interested payoff maximizing agent is best off choosing, and then what are the social outcomes using those strategies. The game-theoretic analysis of the decommitting games handled the possibility that agents decommit manipulatively: an agent tries to avoid the decommitment penalty in case it believes that there is a high probability that it will be freed from the contract's obligations due to the other agent decommitting.

In all of the games studied, full commitment contracts turn out to be a subset of leveled commitment ones. A full commitment contract can be emulated by setting the decommitment penalties sufficiently high. Therefore, full commitment protocols cannot be better than leveled commitment ones in the sense of Pareto efficiency or social welfare maximization. Neither can they enable a deal that is impossible—based on individual rationality—using a leveled commitment contract. In game types where no opportunities (outstanding outside offers) become void between the contracting and the decommitting time (game types DOP, SEQD, SIMUDBP, and SIMUDNP), there are instances where the new protocol enables contracts that are impossible (not individually rational to the agents) using full commitment contracts, and improves Pareto efficiency. Obviously one can also construct game instances where the null deal is so profitable to both agents that no contract—even a leveled commitment one—is individually rational to the agents. In the COBV game where one agent loses an opportunity (outstanding outside offer) by agreeing to a contract, a leveled commitment contract can enable a deal or Pareto improve a deal over a full commitment contract only if that agent's fall-back payoff is sufficiently high.

In the DOP and COBV games where only one agent's future outside offer involves uncertainty, the agent with a certain outside offer prefers to not decommit if the contract is originally individually rational to it. Thus only one agent may want to decommit. In these games, an agent's payoff to a contract is unaffected by the other agent's beliefs. Thus also 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 SEQD, SIMUDBP, and SIMUDNP games where both agents' future outside offers involve uncertainty, an agent's payoff to a contract may depend on the negotiation partner's possibly biased beliefs.

Extensions of this research include studying more closely the best pace to increase 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 towards relating the performance profiles and optimal negotiation actions in a coalition formation problem [20, 21]. Finally, the relationship between leveled commitment contracting and explicit contracts among more than two agents should be studied in more detail.

References

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

[2] K. Decker and V. R. Lesser. Designing a family of coordination algorithms. In *Proc. First International Conference on Multiagent Systems*

(ICMAS-95), San Francisco, June 1995.

[3] T. Finin, R. Fritzson, and D. McKay. A language and protocol to support intelligent agent interoperability. In *Proc. of the CE & CALS Washington '92 Conference*, June 1992.

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

[5] General Magic, Inc. Telescript technology: The foundation for the electronic marketplace, 1994. White paper.

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

[7] D. M. Kreps. *A course in microeconomic theory.* Princeton University Press, 1990.

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

[9] S. H. Low, N. F. Maxemchuk, and S. Paul. Anonymous credit cards. 1994. Submitted.

[10] S. H. Low, N. F. Maxemchuk, and S. Paul. Collusion in a multi-party communication protocol for anonymous credit cards. *IEEE/ACM Transactions on Networking*, 1994. Submitted.

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

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

[13] J. Nash. Two-person cooperative games. *Econometrica*, 21:128–140, 1950.

[14] Office of Technology Assesment (OTA). Electronic enterprises: Looking to the future, 1994.

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

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

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

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

[19] T. W. Sandholm. An implementation of the contract net protocol based on marginal cost calculations. In *Proc. 11th National Conference on Artificial Intelligence (AAAI-93)*, pages 256–262, July 1993.

[20] T. W. Sandholm and V. R. Lesser. Coalition formation among bounded rational agents. Technical Report 95-71, University of Massachusetts at Amherst Computer Science Department, 1995. Extended version.

[21] T. W. Sandholm and V. R. Lesser. Coalition formation among bounded rational agents. In *Proc. 14th International Joint Conference on Artificial Intelligence (IJCAI-95)*, pages 662–669, Montreal, Canada, Aug. 1995.

[22] T. W. Sandholm and V. R. Lesser. Equilibrium analysis of the possibilities of unenforced exchange in multiagent systems. In *Proc. 14th International Joint Conference on Artificial Intelligence (IJCAI-95)*, pages 694–701, Montreal, Canada, Aug. 1995.

[23] T. W. Sandholm and V. R. Lesser. Issues in automated negotiation and electronic commerce: Extending the contract net framework. In *Proc. First International Conference on Multiagent Systems (ICMAS-95)*, pages 328–335, San Francisco, June 1995.

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

[25] 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.