

Leveled-Commitment Contracting

A Backtracking Instrument for Multiagent Systems

Tuomas Sandholm and Victor Lesser

■ In (automated) negotiation systems for self-interested agents, contracts have traditionally been binding. They do not accommodate future events. Contingency contracts address this but are often impractical. As an alternative, we propose leveled-commitment contracts. The level of commitment is set by decommitting penalties. To be freed from the contract, an agent simply pays its penalty to the other contract party(ies). A self-interested agent will be reluctant to decommit because some other contract party might decommit, in which case the former agent gets freed from the contract, does not incur a penalty, and collects a penalty from the other party. We show that despite such strategic decommitting, leveled commitment increases the expected payoffs of all contract parties and can enable deals that are impossible under full commitment. Different decommitting mechanisms are introduced and compared. Practical prescriptions for market designers are presented. A contract optimizer, `ECOMMITTER`, is provided on the web.

The importance of automated negotiation systems consisting of self-interested agents is increasing. One reason is the technology push of a growing standardized communication infrastructure—the internet, the World Wide Web, EDI, HTML, KQML, FIPA, XML, JAVA, ODYSSEY, VOYAGER, CONCORDIA, AGLETS, and so on—over which separately designed agents belonging to different organizations can interact in an open environment in real time and safely carry out transactions (Sandholm and Ferrandon 2000). The second reason is strong application pull for computer support for contracting, especially at the operative decision-making level. For example, we are witnessing the advent of small transaction business-to-consumer and consumer-to-consumer commerce on the internet for purchasing goods,

services, information, communication bandwidth, and so on. Numerous electronic business-to business trading sites have also emerged, some of which already incorporate automated negotiation capability. There is also an industrial trend toward virtual enterprises: dynamic alliances of small, agile enterprises that together can take advantage of economies of scale when available (for example, by being able to respond to larger and more diverse orders than they could individually) but do not suffer from diseconomies of scale.

Multiagent technology facilitates the automated formation of such dynamic alliances on an order-by-order basis by automated contracting. Such automation can save labor time of human negotiators, but in addition, other savings are possible because computational agents are often more effective at finding beneficial contracts and contract combinations than humans are in strategically and combinatorially complex settings. In traditional automated negotiation mechanisms for self-interested agents, once a contract is made, it is binding (see, for example Andersson and Sandholm [1999]; Cheng and Wellman [1998]; Kraus [1993]; Monderer and Tennenholtz [1998]; Rosenschein and Zlotkin [1994]; Sandholm [1993]; Shoham and Tennenholtz [2001]). The contract parties cannot back out no matter how future events unravel. Although a contract might be profitable to an agent when viewed *ex ante*, it need not be profitable when viewed *ex post*. Similarly, a contract that is unprofitable *ex ante* can become profitable *ex post*. Full-commitment contracts are unable to capitalize on the gains that such future events can provide.

However, many multiagent systems consisting of cooperative agents incorporate some

form of decommitting possibility to allow the agents to accommodate new events. For example, in the original contract net protocol (Smith 1980), the agent that had contracted out a task could send a termination message to cancel the contract even when the contractee had already partially fulfilled the contract. This was viable because the agents were not self-interested: the contractee did not mind losing part of its effort without compensation. Similarly, decommitting in automated contracting among cooperative agents has been studied in a meeting scheduling domain (Sen and Durfee 1998, 1994).

Unlike systems with cooperative agents, multiagent systems consisting of self-interested agents require that we consider agents that do not follow externally specified strategies but choose their own strategies. Therefore, the *interaction mechanism* (that is, rules of the game) has to be designed from the perspective of noncooperative game theory. Specifically, the mechanism should be designed so that (1) it is possible to determine what each agent's best strategy is from a self-interested, expected utility-maximizing perspective (this can depend on the other agents' strategies as we demonstrate in this article) and (2) a desirable social outcome will follow even though every agent uses its self-interested strategy.

One can think of *negotiation* as search among multiple agents where the search focus is the set of commitments made by the agents. *Decommitting* can then be viewed as backtracking in this multiagent search. Although backtracking has been used in multiagent systems among cooperative agents, for example, in distributed constraint satisfaction (Yokoo 2000), the self-interest of agents imposes additional requirements on a backtracking instrument because the system designer cannot control the committing and decommitting behavior of the agents. In this article, we discuss backtracking instruments that have desirable properties even when used among self-interested agents.

Contingency Contracts

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

their fallback positions, but there is a contingency contract that both agents prefer over their fallbacks.

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

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

Second, although contingency contracts can be useful in anticipating a small number of key events, they become cumbersome as the number of relevant future events to monitor increases. In the extreme, all domain events (new tasks arriving, resources breaking down and becoming back online, and so on) and all negotiation events (offers, acceptances, rejections, and so on, from related negotiations) can affect the value of the original contract to the agent and should therefore be conditioned on. Furthermore, these future events might not affect the value of the original contract independently. The value of the original contract might depend on combinations of the future events (Rosenschein and Zlotkin 1994; Sandholm 1993; Sandholm and Lesser 1995). Thus, there is a potential combinatorial explosion of possible future worlds, and each of them might need to be associated with a different contingency, leading to a potential combinatorial explosion of the contract (for example, the size of the contingency table that represents the contract).

In addition to these two practical difficulties associated with contingency contracts, there is a third, fundamental problem. An event might be observable by only one of the contract parties. That agent might lie to the other party about the event in case the event is associated with a disadvantageous contingency to the observing party. Thus, to be viable, contingency contracts would require an event verification mechanism that is not manipulable and not prohibitively complicated or costly.

Leveled-Commitment Contracts

To avoid the drawbacks of contingency contracts, we propose another instrument for capitalizing on the possibilities provided by probabilistically known future events. Instead of conditioning the contract on future events, a mechanism is built into the contract that allows unilateral decommitting. This is achieved by specifying in the contract decommitting penal-

ties, one for each agent. If an agent wants to decommit—that is, to be freed from the obligations of the contract—it can do so simply by paying the decommitting penalty to the other party. We call such contracts *leveled-commitment contracts* because the decommitting 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. Therefore, no event-verification mechanism is required. Another potential advantage is that the agents are not forced to consider all combinations of possible future events up front, but rather, each agent can react to only those events that actually occur.

Principles for assessing breach penalties have been studied in the economics of law (Calamari and Perillo 1977; Posner 1977), but the purpose has usually been to assess a penalty on the agent that has breached the contract after the breach has occurred. Similarly, penalty clauses for partial failure—such as not meeting a deadline—are commonly used in contracts, but the purpose is usually to motivate the agents to abide by the contract. Instead, in leveled-commitment contracts, explicitly allowing decommitting from the contract for a predetermined price is used as an active method for using the potential provided by an uncertain future.¹ As we discuss later, it turns out that, somewhat unintuitively, the decommitting possibility can increase the expected payoff for both contract parties.

Practical Motivations for Leveled Commitment

The goal of the leveled-commitment contracting mechanism is to allow some flexibility, as in the case with no commitment while agents are guaranteed some level of security, as in the case of full commitment. Full-commitment contracts can be viewed as one end of a spectrum with commitment-free contracts at the other. Leveled-commitment contracts span this entire spectrum based on how the decommitting penalties are chosen.²

There are several practical reasons why leveled commitment is desirable:

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

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

It provides a backtracking instrument for distributed search (in the AI sense) that works among self-interested agents, unlike traditional backtracking techniques for distributed search (see Yokoo [2000] for a review of the traditional techniques). It allows more controlled profitable risk taking; thus, in terms of search, a low-commitment search focus is moved around in the global search space of commitments (because decommitting is not unreasonably expensive) so that more of the space can be explored among self-interested agents that would otherwise avoid risky commitments. For example, in multiagent task allocation, an agent can accept a task set and later try to subcontract out the tasks in this set separately. With full commitment, an agent would need to have standing offers from the agents it will subcontract the tasks to, or it has to be able to handle them (profitably) itself. With leveled commitment, the agent can accept the task set even if it is not sure about its chances of getting the tasks handled. If it does not get them handled, it can decommit.

It allows profitable construction of combinatorial contracts from basic contracts. Often, the value of a contract to an agent depends on which other contracts the agent will get (Rosenschein and Zlotkin 1994; Sandholm 1993). Using leveled commitment, an agent can take on unbeneficial contracts in anticipation of later synergic contracts that will make the sequence beneficial overall. If the later contracts in the sequence do not occur, the agent can backtrack out of the first part of the sequence.

This method of constructing a combinatorial contract through a sequence of basic contracts also applies to cases where the combinatorial contracts involve more than two parties. Contracts involving more than two parties are important for avoiding local optima in automated negotiation: Even if no contract among k agents is beneficial, a contract among $k + 1$ agents can be (Andersson and Sandholm 2000, 1999; Sandholm 1998).

It saves computation and time. In many automated negotiation applications, computing the value of taking on a contract is intractable, so the agents have to resort to approximation (Sandholm 1996, 1993; Sandholm and Lesser 1995). Leveled commitment allows an agent to bid based on a rough-value calculation. If the agent wins the bid, the agent can invest a more thorough value calculation. If the contract no longer looks beneficial in light of this more refined calculation, the agent can decommit. The fact that only the winning bidders carry out a refined calculation can save

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

Table 1. Prisoners' Dilemma Game.

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

computation systemwide. Also, the negotiations can be carried out faster because agents can bid based on less computation.

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

It speeds up the negotiation process by increasing parallelism. An agent can make (low-commitment) offers to multiple recipients, although these offers are mutually exclusive from the agent's perspective. In case more than one recipient accepts, the agent can backtrack from all but one, allowing the agent to address the recipients in parallel instead of addressing them one at a time and blocking to wait for an answer before addressing the next. (Note that the alternative of sending an offer with no commitment would be strategically meaningless).

It can be used to increase the contract parties' welfare by reallocating risk. By choosing the contract price and decommitting penalties appropriately, every agent's expected utility can increase by making the less risk-averse agents carry more of the risk. The more risk-averse agents would be willing to compensate the former agents with a higher expected payoff. As we show, leveled commitment can increase the contract parties' expected payoffs even if the agents are risk neutral. The use of this instrument as a risk reallocation tool provides an **additional** utility increase to the contract parties. For simplicity, we do not address risk reallocation in this article.

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

Why Are the Advantages of Leveled Commitment Not Obvious?

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

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

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

Third, agents might decommit strategically. Consider a contractor agent that can award one contract and a contractee agent that can take on one contract. A nonstrategic agent would decommit whenever its best outside offer, plus the decommitting penalty, is better than the current contract. However, a rational self-interested agent would be more reluctant

in decommitting because there is a chance that the other party will decommit, in which case, the former agent gets freed from the contract obligations, does not have to pay a decommitting penalty, and will collect a decommitting penalty from the other party. Similarly, the other contract party will be reluctant to decommit for the same reason. Because of such reluctant decommitting, a contract might end up being inefficiently kept even though each party would be better off decommitting and paying the penalty (and therefore, the sum of the contract parties' payoffs would be higher if either agent alone, or both agents, would decommit).

Next, we show that leveled-commitment contracts are superior to full-commitment contracts despite these concerns.

The Game

The benefits of leveled commitment can be demonstrated already in a relatively simple contracting setting with two risk-neutral agents, each of which attempts to maximize its own expected payoff. We call one of the agents the *contractor* and the other agent the *contractee*.

A full-commitment contract between the contractor and the contractee is defined by the contract obligations, which include (1) a description of what each of the two agents has to perform (handling tasks, contributing goods, lending resources, and so on) and (2) a contract price that the contractor has to pay to the contractee.

Say that the value (or cost) of executing the contract description can change for each agent separately. At contract time, each agent only has probabilistic information about this value. The change in value can stem from changes in the agent's own characteristics, such as resources failing or becoming available. The change can also be the result of computation: Based on a rough computation, a contract's value might have looked different than it looks after a more refined computation. Furthermore, the value of a contract can change because of changes in outside options such as offers from third parties. Our framework and results are not specific to any particular source of change. However, we present our results in the setting where the change stems from outside options. Specifically, the agents might receive outside offers. For simplicity, say that all offers have the same description, so price is the only concern, and each of the two agents only want to be involved in one contract (the contractor gets

its task handled and does not need to have it handled more than once, and the contractee has limited resources and can only take on one contract).

For the analysis, say that the contractor's best (lowest) outside offer is not known *ex ante*, but the probability distribution from which it will be drawn is known. Similarly, the contractee's best (highest) outside offer is unknown, but it will be drawn from another known distribution. If an agent does not receive an outside offer, its "best outside offer" corresponds to its best outstanding outside offer or its fallback payoff. The advantages of leveled commitment prevail even if the outside offers are statistically independent, and the figures in this article depict such settings. For the analysis, the prior distributions of outside offers are assumed to be common knowledge between the contractor and the contractee. This assumption of commonly known priors is arguably close to accurate in applications where the agents have good statistical information on how much it costs to have the contract description filled. For example, if one were to sell an airline ticket using a leveled-commitment contract, there are good statistics about how much an airline can charge for a seat on a certain leg as well as statistics on how much it costs a passenger to travel this leg using a choice of any airline.

The contractor and the contractee can either make a contract or wait for their outside offers. If the contractor and the contractee do make a contract, they can choose to use some full-commitment contract or some leveled-commitment contract. A leveled-commitment contract between a contractor and a contractee is defined by four components: (1) a description of what each of the two agents has to perform (handling tasks, contributing goods, lending resources, and so on), (2) a contract price that the contractor has to pay to the contractee, (3) the contractor's decommitting penalty, and (4) the contractee's decommitting penalty.

If an agent decommits and pays its penalty, then the contract description does not need to be fulfilled, and the contract price does not need to be paid. We say that each of the two agents observes its own outside offers before having to decide whether to decommit, but the agents do not observe each other's outside offers. This seems realistic from a practical (automated) contracting perspective.

An easy way to think about the outside offers is to consider them to be full-commitment contracts. Alternatively, they can be

expected utilities that the agents get under outside leveled-commitment contracts.

Our game consists of three stages. In the first stage, which we call the *contracting game*, the agents choose a contract (or no contract). In the second stage, each agent receives its outside offers. In the third stage, which we call the *decommitting game*, the agents decide on whether to decommit or not. Clearly, how the agents will play in the decommitting game affects their preferences over contracts in the contracting game. We say that a contract is *individually rational* for an agent if the agent's expected payoff under the contract is higher than this agent's expected payoff would be if the agent made no contract and waited for its best outside offer. A contract is called *individually rational* if it is individually rational for both agents.

A nonstrategic (truthful) agent would decommit whenever its outside offer is better than the existing contract plus its decommitting penalty. However, we do not assume that the agents decommit nonstrategically. Instead, a self-interested agent can decommit strategically. It might be more reluctant to decommit than a nonstrategic agent would be because there is a chance that the other contract party will decommit; in which case, the agent gets freed from the contract obligations, does not have to pay a penalty, and will collect a penalty from the other agent. At the same time, the other agent will be reluctant to decommit as well.

Each agent's decommitting strategy is characterized by the agent's decommitting threshold. If the contractor's outside offer is lower than the contractor's threshold, then the contractor will decommit. Similarly, if the contractee's outside offer is greater than the contractee's threshold, then the contractee will decommit.

To analyze such situations, one needs to determine which strategy would be best for a self-interested agent, and only then can one expect a self-interested agent to use this strategy. The interesting aspect of the decommitting game is that neither agent has a dominant strategy. How an agent should set its decommitting threshold depends on how the other agent sets its threshold. Therefore, we analyzed the (Nash) equilibrium of the decommitting game where the contractor's threshold is a best response to the contractee's, and the contractee's threshold is a best response to the contractor's. If others design their software agents to act in this manner, then no agent designer can do better than designing his/her agent to act this way.

Decommitting Strategies under Different Leveled-Commitment Mechanisms

The equilibrium of the decommitting game differs significantly in different variants of leveled-commitment contracting mechanisms. We identified two natural dimensions along which these mechanisms can differ.

Along one dimension, there are three different orders in which the agents have to reveal their decommitting decisions: contractor first, contractee first, or simultaneously. The key distinction is not whether decommitting happens at the same point in time but whether an agent knows the other agent's decommitting decision by the time it has to reveal its own decommitting decision. If neither agent knows the other's decision at this stage, the decommitting is in essence simultaneous.⁵

Along the second dimension, there are two choices of what happens if both agents decommit. Either both pay the penalties to each other, or neither pays a penalty.

These dimensions set up a space of $3 \times 2 = 6$ leveled-commitment contracting mechanisms.

In the sequential games, if the first agent decommits, the second agent would never decommit because it could only reduce its payoff by doing so. Therefore, in the sequential games (contractor first or contractee first), the variant where both pay if both decommit is equivalent to the variant where neither pays if both decommit. Thus, four distinct mechanisms are left. Of these, the sequential mechanism where the contractor moves first and the sequential mechanism where the contractee moves first are analogous, which leaves three fundamentally different mechanisms: (1) *sequential decommitting* (say, contractee first); (2) *simultaneous decommitting*, both pay if both decommit; and (3) *simultaneous decommitting*, neither pays if both decommit.

The equilibrium differs across these mechanisms. It is easy to observe that in the sequential game, if the first agent does not decommit, then the second agent will decommit nonstrategically (truthfully). However, the first agent's behavior is strategic and requires quantitative analysis. In the simultaneous games, both agents' strategies require quantitative equilibrium analysis. We conducted these analyses (Sandholm 1996; Sandholm and Lesser 2001), and the results that we discuss in this article are based on this work. These results are general and apply to any distributions of uncertainty. However, to keep the presentation nonmathematical in this article, we only demonstrate the equilibria of the games graph-

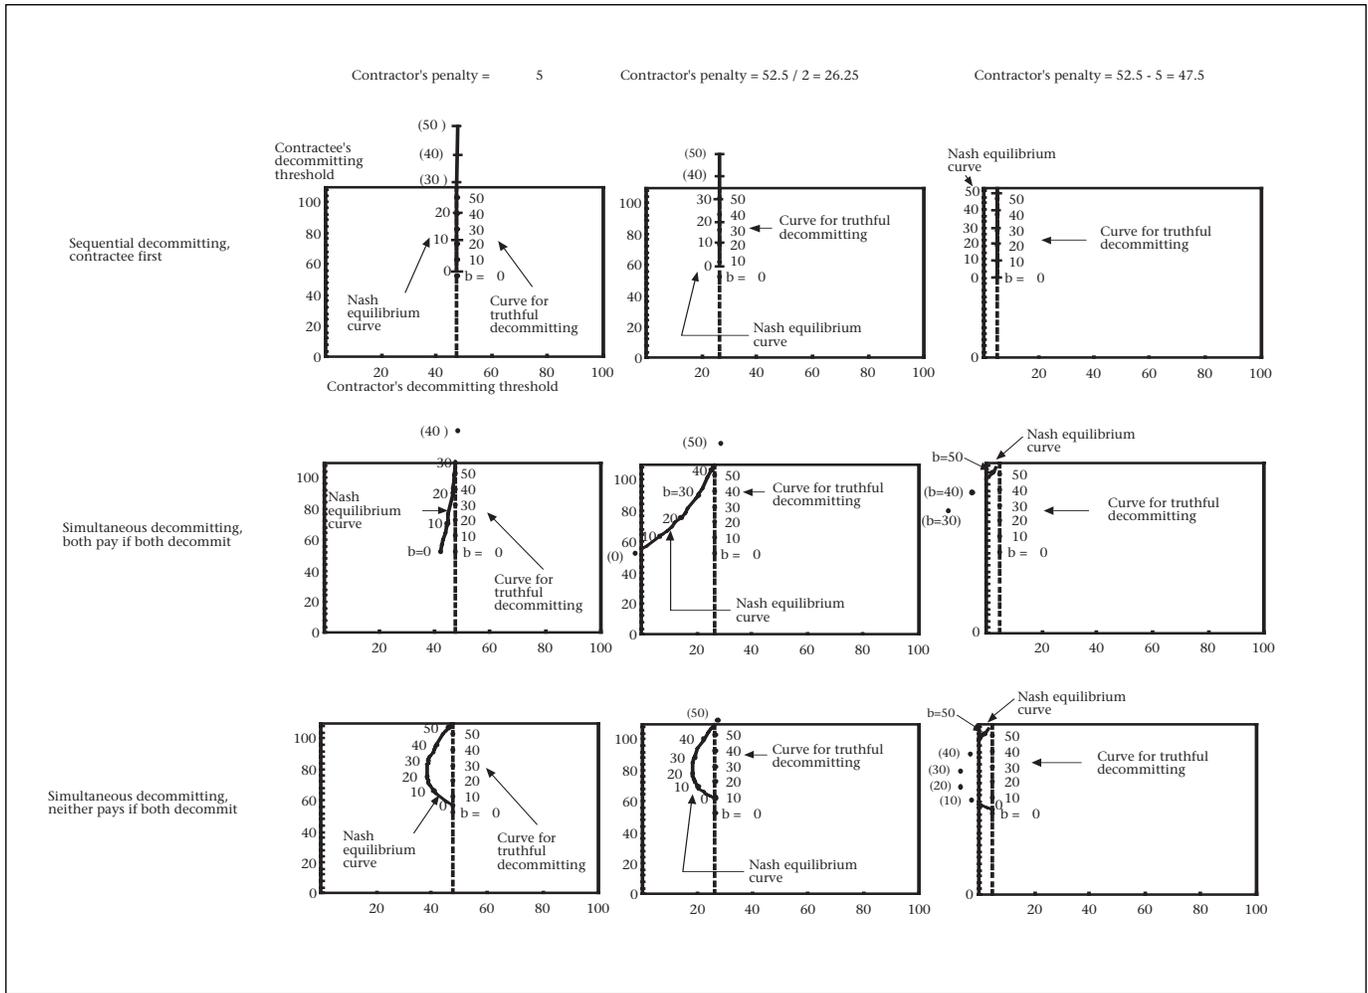


Figure 1. The Nash Equilibrium Decommitting Thresholds of Our Example.

ically in the context of a particular example.

Consider a contracting setting where the contractor’s best (lowest) outside offer is drawn from a uniform distribution on $[0, 100]$, and the contractee’s best (highest) outside offer is drawn from a uniform distribution on $[0, 110]$. Figure 1 illustrates the equilibria for the three different leveled-commitment mechanisms. To avoid drawing in three dimensions, we fixed the value of the contract price (to 52.5, which is halfway between the contractor’s expected best outside offer [50] and the contractee’s expected best outside offer [55]).

Figure 1 shows the Nash equilibrium decommitting thresholds of our example. As we move from the left graphs to the right graphs, we increase the contractor’s penalty. Within each graph, we vary the contractor’s penalty b , and the equilibrium moves in the plane accordingly. In each graph, a point (x, y) corresponds to an equilibrium; that is, the contractor decommits if its outside offer is lower than x , and the

contractee decommits if its outside offer exceeds y . (However, in the sequential game, the contractor never decommits if the contractee decommits.) In each graph, the Nash equilibrium deviates from nonstrategic (“truthful”) decommitting. In cases where the equilibrium point is outside the frame (that is, outside the support of the distribution of the agent’s best outside offer), it is sure that the corresponding agent will not decommit.

Figure 1 shows that in the simultaneous game where both pay the penalties if both decommit, as an agent’s penalty approaches zero, the agent becomes truthful, but the opponent does not. On the contrary, in the simultaneous game where neither pays if both decommit, as an agent’s penalty approaches zero, the agent does not become truthful, but the opponent does! We showed that these phenomena prevail, in fact, under any distributions of outside offers, not only in this example (Sandholm and Lesser 2001).

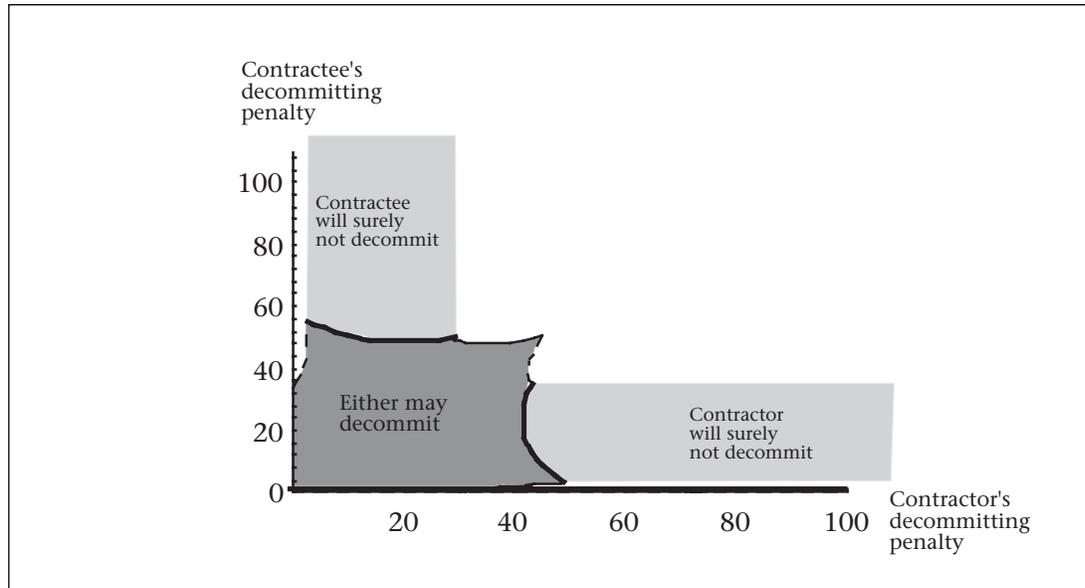


Figure 2. The Gray Regions Correspond to Three Qualitatively Different Classes of Leveled-Commitment Contracts That Both Parties Prefer over Waiting for Their Outside Offers.

To avoid drawing in three dimensions, we fix the contract price to a particular value (52.5). Other choices work as well: The set of individually rational contracts forms a three-dimensional volume.

Another general fact is the relative amount of strategy that the mechanisms induce. Let p in the simultaneous mechanisms be the probability that the opponent decommits or, in the sequential mechanism, the probability that the second mover decommits given that the first mover did not decommit. Now, for given p , given contract price, and given decommitting penalties, the (first) agent is most strategic (reluctant to decommit) in the sequential mechanism, less strategic in the simultaneous mechanism where both pay if both decommit, and least strategic in the simultaneous mechanism where neither pays if both decommit (Sandholm and Lesser 2001).

Beyond these general principles, Sandholm et al. designed and built a system, *E-COMMITTER*,³ which—given the contract price, decommitting penalties, and distributions of outside offers—finds the equilibrium decommitting thresholds for each of the six leveled-commitment mechanisms (Sandholm, Sikka, and Norden 1999).

Leveled Commitment Increases the Contract Parties' Payoffs and Enables Deals

Because a leveled-commitment contract can emulate full commitment by setting the penalties high enough, there always exists a leveled-commitment contract that is at least as good as

any full-commitment contract. Furthermore, it turns out that each one of the leveled-commitment mechanisms can be used to strictly increase the expected payoffs of both contract parties. It follows that there are games (defined by the distributions of outside offers) where no full-commitment contract is possible (because at least one agent would rather just wait for its best outside offer), but a leveled-commitment contract is possible because each contract party prefers it over waiting for the best outside offer. In other words, leveled commitment enables deals, which can be shown even in the simple example earlier where the contractor's best outside offer is uniformly distributed on $[0,100]$ and the contractee's on $[0,110]$. The contractor would accept a full-commitment contract if the price were below 50. The contractee would accept a full-commitment contract if the price were above 55; so, there is no mutually acceptable full-commitment contract in this example. However, we can construct a leveled-commitment contract that works. We proved this analytically (Sandholm 1996; Sandholm and Lesser 2001]; here, we simply illustrate the final results. Figure 2 shows the leveled-commitment contracts that are individually rational in this game where no full-commitment contract is. The figure shows them for the simultaneous decommitting game where neither pays a penalty if both decommit. The figures for the other decommitting mechanisms are similar but not identical.

In addition to enabling deals that are impossible using full-commitment contracts, leveled-commitment contracts can increase the economic efficiency of a deal even if a full-commitment contract were possible. The reverse cannot occur because leveled-commitment contracts can emulate full-commitment contracts by setting the penalties high enough. We quantitatively characterized the uncertainty that leveled-commitment contracts can profitably capitalize on and uncovered the following simple rule: Leveled commitment increases the contract parties' payoffs if there is some chance that the contractor's outside offer is lower than the contractee's expected outside offer or some chance that the contractee's outside offer is higher than the contractor's expected outside offer (Sandholm 1996; Sandholm and Lesser 2001). Note that this condition can be satisfied even if only one agent's future involves uncertainty.

Optimizing the Contract Parameters and Dividing the Surplus

Usually there is either no contract that is individually rational for the agents (better for both than waiting for outside offers), or there are many such contracts. Bargaining theory addresses the choice among individually rational deals, usually by modeling time discounting (Kraus, Wilkenfeld, and Zlotkin 1995; Rubinstein 1982), deadlines (Sandholm and Vulkan 1999), or bargaining costs (Rubinstein 1982) in the bargaining process or by asserting desirable properties that the chosen contract should have [Nash 1950; Osborne and Rubinstein 1990]. In this article, we do not address the choice among individually rational contracts. Algorithms for determining the individually rational contracts (contract price and decommitting penalties) that maximize the sum of the contract parties' payoffs are presented in Sandholm et al. (1999).⁴

One conceivable concern is that this surplus might only be maximized under certain splits of the surplus between the contract parties. However, it was recently shown that for any of the six leveled-commitment mecha-

nisms, the surplus can be maximized under any split of the surplus by appropriately setting the contract price and decommitting penalties (Sandholm and Zhou 2000). This makes the leveled-commitment contract technology a modular component for (automated) negotiation. It can be used as a tool for increasing surplus in settings with uncertainty and using any (bargaining) mechanism for deciding how to divide the surplus between the contract parties.

Another concern is what happens if a contract party misperceives the distributions of outside offers. We showed that in settings where only one agent's outside offers are uncertain, a misperceiving agent can only hurt itself, but in games where more than one agent's outside offers are uncertain, a misperceiving agent can also hurt the other contract party by bad decommitting decisions (Sandholm and Lesser 2001).

Which Leveled-Commitment Mechanism Is Best?

As shown in figure 1, for any given contract, the equilibria are different for the different decommitting mechanisms. It turns out that the surplus maximizing contract parameters also differ across the mechanisms (Sandholm, Sikka, and Norden 1999). However, a recent paper shows that surprisingly, among risk-neutral agents, each of the six mechanisms leads to the same sum of the contract parties' payoffs when the contract price and penalties are optimized for each mechanism separately (Sandholm and Zhou 2000). (Among agents that are not risk neutral, the three mechanisms lead to different sums of utilities, and the ranking of the mechanisms varies based on the agents' utility functions.)

Because the sum of the contract parties' payoffs cannot be used as the criterion for choosing a mechanism among risk-neutral agents, other criteria are needed.

One practical goal is to minimize the number of payment transfers. A recent paper shows that if the contract is optimized separately for each of the mechanisms, among risk-neutral agents, the optimal decommitting thresholds will be the same for all six mechanisms

(Sandholm and Zhou 2000). Therefore, the mechanisms can be compared based on what happens depending on where the outside offers fall with respect to these thresholds. If one agent gets an offer that is better than its threshold but the other agent does not, all the mechanisms lead to one penalty being paid. If neither agent receives an offer that is better than its threshold, all the mechanisms lead to no penalty payments. If both agents receive offers that are better than the thresholds, both will decommit in the simultaneous games, but only the first agent will decommit in the sequential game. In this case, the simultaneous mechanism where both pay leads to two payments, the sequential mechanism leads to one, and the simultaneous mechanism where neither pays leads to none. In summary, from the perspective of minimizing the number of penalty payments, the simultaneous mechanism where neither pays if both decommit is best, the sequential mechanism is in the middle, and the simultaneous mechanism where both pay if both decommit is worst.

Another evaluation criterion is robustness of the equilibrium in the decommitting game. For the sequential decommitting game, we were able to use iterated dominance as the solution concept, but for the simultaneous decommitting games, the Nash equilibrium was used. This suggests using sequential decommitting mechanisms because iterated dominance is a more robust solution concept than the Nash equilibrium (Mas-Colell, Whinston, and Green 1995).

An additional consideration that favors sequential decommitting mechanisms is that playing optimally is easy for one of the agents (the second mover). The second mover is best off by decommitting nonstrategically if the first mover did not decommit and not at all if the first mover did.

Finally, a recent article shows that the equilibrium is always unique in the sequential mechanisms, but multiple equilibria can exist in the simultaneous mechanisms (Sandholm, Sikka, and Norden 1999). This also speaks in favor of using the sequential mechanisms.

Decommitting Cascades, Infinite Decommit- Recommit Loops, and Increasing Penalties

In a web of multiple contracts (potentially involving several agents each), full-commitment contracts induce one negotiation search focus consisting of the obligations of the contracts. With leveled-commitment contracts, there are multiple such foci, and any agent involved in a contract can switch from one such focus to another by decommitting from some contract. This decommitment can make it beneficial for another agent to decommit from another contract, and so on, leading to cascades of backtracks. An appropriate amount of backtracking is desirable in the search for good allocations of obligations among the agents.

However, with myopic agents that use leveled-commitment contracts, the multiagent system can get stuck in an infinite loop of states (where each state is defined by the commitments that the agents hold) (Andersson and Sandholm 1998). The agents make a sequence of deals and then backtrack out of them in nonchronological order. Then the process repeats. To avoid such useless loops of decommitting and recommitting, recommitting could be disabled. The mechanism could specify that if a contract offer is accepted and later either agent decommits from the contract, the original offer becomes void as opposed to staying valid according to its original deadline that might not have been reached at the time of decommitment. However, even if agents cannot explicitly recommit to a contract, it is hard to monitor that they will not make another identical deal, again giving rise to the possibility of the equivalent of useless decommit-recommit loops. Useless decommit-recommit loops can be tackled using decommitting penalties that increase with time (real time or number of domain events or negotiation events).⁶ This allows a low-commitment negotiation focus to be moved in the joint search space and still make the contracts meaningful by some level of commitment. The increasing level of commitment causes the agents to not backtrack too deeply

in the negotiations, which can also save computation.

Interestingly, some element of global clock is required as the basis for increasing the penalties: To avoid infinite loops among myopic agents, it does not suffice to count time from the moment when each contract is made (Andersson and Sandholm 1998). There are also interesting issues in how to increase the penalties to get enough backtracking to reach good solutions but not too much to waste time. The results differ depending on whether the agents are myopic (Andersson and Sandholm 1998) or carry out strategic look ahead (Andersson and Sandholm 2001).

Conclusions

Leveled-commitment contracts are new backtracking instruments for multiagent systems that work even among self-interested agents that decommit strategically. Leveled commitment can increase the payoffs of all contract parties when at least one of the agents faces uncertainty. These contracts are more practical than contingency contracts. However, they generally achieve a lesser sum of payoffs than the optimal complete contingency contract (that uses event verification) because sometimes the contract will be kept although it should be dissolved. The results and algorithms that we discussed in this article also extend to leveled-commitment contracts that involve more than two contract parties in a single contract (Sandholm, Sikka, and Norden 1999).

Since we introduced leveled-commitment contracts (Sandholm and Lesser 1995), they have been used in several applications. For example, Mitsubishi has applied them to an electronic market for construction waste recycling in Japan (Akiyoshi et al. 1999). They have also been applied to automated negotiation in a manufacturing setting (Collins et al. 1998) and in a digital library (Park, Durfee, and Birmingham 1996).

Multiitem auctions where bidders have preferences over combinations of items are another important potential application. With certain auction de-

signs that allow bidding on bundles, one can obtain efficient outcomes in dominant strategy equilibrium (see, for example, Sandholm [2000]). However, this can require a bidder to compute its valuation for each combination of items and to bid for each combination. Also, the auctioneer's task of determining the winners is computationally complex (Sandholm 2002). A potentially more practical alternative is to use sequential or ascending auctions with bidding on individual items or restricted combinations only. Leveled-commitment contracts could be used as a mechanism for bidders to put back items if they do not (or if they project that they will not) obtain the combinations that they want. Similarly, the auctioneer might want to exercise a take-back, for example, if it receives a better bid later. Different one-sided, put-back mechanisms have already been tried in auction contexts (McAfee and McMillan 1996; Walsh and Wellman 1999), but leveled-commitment contracts allow both put-backs and take-backs. This gives rise to additional phenomena—such as strategic decommitting. Leveled-commitment contracts can also prove useful in online buying scenarios where a shopping agent can buy an item from multiple alternative sites (auctions, catalogs, and so on), and the suppliers can sell their goods to multiple alternative buyers. If an agent (buyer or seller) then gets offered a better deal than the one it has obtained to this point, it can decommit.

Future Research

There are several interesting directions for future research on leveled-commitment contracts.

Even if explicit linking of issues using combinatorial contracts (with potentially more than two agents per contract) is allowed, leveled-commitment contracts can be beneficial. Identifying profitable combinatorial contracts can be complex computationally and difficult without a global view, so an agent might be better off trying to construct profitable combinations from sequences of individual contracts, each with leveled commit-

ment. The best contracting mechanisms will most likely combine leveled commitment and explicit linking of issues. As a first step in this direction, hybrid negotiation mechanisms were developed where leveled commitment was used on top of combinatorial contracts (Andersson and Sandholm 1998). These contracts were studied in simulation.

Another open question is how an agent should allocate its scarce computational resources when evaluating different leveled-commitment contracts. Which combinations of deals should it evaluate? How much of the evaluation should it conduct before bidding and how much after winning? Steps toward devising normative theories of deliberation in other games have already been taken (Larson and Sandholm 2001a, 2001b, 2001c; Sandholm and Lesser 1997). These models are likely to help in addressing the deliberation question in leveled-commitment contracting, but this work also uncovered some complications. For example, it turned out that how an agent is best off deliberating depends on how others deliberate. Therefore, a game-theoretic equilibrium analysis was conducted in the space of the agents' deliberation strategies.

The deliberation capabilities of the agents should also be taken into account when designing leveled-commitment mechanisms. For example, they are likely to affect the best pace of increasing the decommitting penalties, which opens up a whole new avenue in game-theoretic mechanism design.

Acknowledgments

This material is based on work supported by the National Science Foundation under CAREER award IRI-9703122 and grants IIS-9800994, ITR IIS-0081246, and ITR IIS-0121678.

Notes

1. Deccommitting has been studied in other settings, for example, where there is a constant inflow of agents, and they have a time cost for searching partners of two types: good or bad (Diamond and Maskin 1979).
2. Different parts of the contract could be associated with different decommitting penalties (Sandholm 1996; Sandholm and Lesser 1995), which could be handled by

considering the parts as separate contracts.

3. www.cs.cmu.edu/~amem/eMediator.

4. The contract optimization service, ECOMMITTER, is available at www.cs.cmu.edu/~amem/eMediator.

5. An agent would have an incentive to try to outwait the other and thus turn a simultaneous mechanism into a sequential mechanism where the opponent declares its decommitting decision first. Such postponing can be avoided by using a trusted third party that does not announce the decommitting decisions until it has received the decisions from both parties. This can also be accomplished without a third party if each agent encrypts its decision that it sends to the other and sends out the key only after receiving the opponent's encrypted decision.

6. The decommitting penalty could also decrease as a function of acceptance time of the offer or be conditioned on events in other negotiations or the environment.

References

- Akiyoshi, M.; Abe, K.; Ono, T.; and Tabinoki, Y. 1999. Agent-Based Decision Support Framework for Asynchronous Negotiation of Resource Transfer. Paper presented at the 1999 IEEE International Conference on Systems, Man, and Cybernetics, October, Tokyo, Japan.
- Andersson, M. R., and Sandholm, T. 2001. Leveled-Commitment Contracts with Myopic and Strategic Agents. *Journal of Economic Dynamics and Control* (Special Issue on Agent-Based Computational Economics) 25(3-4): 615-640.
- Andersson, M., and Sandholm, T. 2000. Contract-Type Sequencing for Reallocation Negotiation. Paper presented at the Twentieth International Conference on Distributed Computing Systems, 10-13 April, Taipei, Taiwan.
- Andersson, M. R., and Sandholm, T. 1999. Time-Quality Trade-Offs in Reallocation Negotiation with Combinatorial Contract Types. In Proceedings of the Sixteenth National Conference on Artificial Intelligence (AAAI), 3-10. Menlo Park, Calif.: American Association for Artificial Intelligence.
- Andersson, M. R., and Sandholm, T. 1998. Leveled-Commitment Contracting among Myopic Individually Rational Agents. Paper presented at the Third International Conference on Multi-Agent Systems (ICMAS), 4-7 July, Paris, France.
- Calamari, J. D., and Perillo, J. M. 1977. *The Law of Contracts*. 2d ed. Cleveland, Ohio: West.
- Cheng, J. Q., and Wellman, M. P. 1998. The WALRAS Algorithm: A Convergent Distributed Implementation of General Equilibrium Outcomes. *Computational Economics* 12:1-24.
- Collins, J.; Youngdahl, B.; Jamison, S.; Mobasher, B.; and Gini, M. 1998. A Market Architecture for Multi-Agent Contracting. Paper presented at the Second International Conference on Autonomous Agents (AGENTS), 9-13 May, Minneapolis/St. Paul, Minnesota.
- Diamond, P. A., and Maskin, E. 1979. An Equilibrium Analysis of Search and Breach of Contract, I: Steady States. *Bell Journal of Economics* 10:282-316.
- Kraus, S. 1993. Agents Contracting Tasks in Noncollaborative Environments. In Proceedings of the Eleventh National Conference on Artificial Intelligence (AAAI-93), 243-248. Menlo Park, Calif.: American Association for Artificial Intelligence.
- Kraus, S.; Wilkenfeld, J.; and Zlotkin, G. 1995. Multiagent Negotiation under Time Constraints. *Artificial Intelligence* 75: 297-345.
- Larson, K., and Sandholm, T. 2001a. Bargaining with Limited Computation: Deliberation Equilibrium. *Artificial Intelligence* 132(2): 183-217.
- Larson, K., and Sandholm, T. 2001b. Computationally Limited Agents in Auctions. Paper presented at AGENTS-01 Workshop of Agents for B2B, 28 May-1 June, Montreal, Canada.
- Larson, K., and Sandholm, T. 2001c. Costly Valuation Computation in Auctions: Deliberation Equilibrium. Paper presented at the Theoretical Aspects of Reasoning about Knowledge Conference (TARK), 8-10 July, Siena, Italy.
- McAfee, R. P., and McMillan, J. 1996. Analyzing the Airwaves Auction. *Journal of Economic Perspectives* 10(1): 159-175.
- Mas-Colell, A.; Whinston, M.; and Green, J. R. 1995. *Microeconomic Theory*. Oxford, U.K.: Oxford University Press.
- Monderer, D., and Tennenholtz, M. 1998. Optimal Auctions Revisited. In Proceedings of the Fifteenth National Conference on Artificial Intelligence (AAAI), 32-37. Menlo Park, Calif.: American Association for Artificial Intelligence.
- Nash, J. 1950. The Bargaining Problem. *Econometrica* 18:155-162.
- Osborne, M. J., and Rubinstein, A. 1990. *Bargaining and Markets*. San Diego, Calif.: Academic.
- Park, S.; Durfee, E.; and Birmingham, W. 1996. Advantages of Strategic Thinking in Multiagent Contracts (a Mechanism and Analysis). Paper presented at the Second International Conference on Multi-Agent Systems (ICMAS), Keihanna Plaza, 10-13 December, Kyoto, Japan.

- Posner, R. A. 1977. *Economic Analysis of Law*. 2d ed. Boston, Mass.: Little, Brown.
- Raiffa, H. 1982. *The Art and Science of Negotiation*. Cambridge, Mass.: Harvard University Press.
- Rosenschein, J. S., and Zlotkin, G. 1994. *Rules of Encounter: Designing Conventions for Automated Negotiation among Computers*. Cambridge, Mass.: MIT Press.
- Rubinstein, A. 1982. Perfect Equilibrium in a Bargaining Model. *Econometrica* 50:97–109.
- Sandholm, T. 1993. An Implementation of the Contract Net Protocol Based on Marginal Cost Calculations. In Proceedings of the Eleventh National Conference on Artificial Intelligence (AAAI), 256–262. Menlo Park, Calif.: American Association for Artificial Intelligence.
- Sandholm, T. 2002. Algorithm for Optimal Winner Determination in Combinatorial Auctions. *Artificial Intelligence* 135:1–54
- Sandholm, T. 2000. eMEDIATOR: A Next-Generation Electronic Commerce Server. Paper presented at the Fourth International Conference on Autonomous Agents (AGENTS), 3–7 June, Barcelona, Spain.
- Sandholm, T. 1998. Contract Types for Satisficing Task Allocation: I. Theoretical Results. Paper presented at the AAAI Spring Symposium Series: Satisficing Models, 23–25 March, Stanford, California.
- Sandholm, T. 1996. Negotiation among Self-Interested Computationally Limited Agents. Ph.D. dissertation, Department of Computer Science, University of Massachusetts at Amherst. Available at www.cs.cmu.edu/~sandholm/dissertation.ps.
- Sandholm, T., and Ferrandon, V. 2000. Safe Exchange Planner. Paper presented at the Fourth International Conference on Multi-Agent Systems (ICMAS), 10–12 July, Boston, Massachusetts.
- Sandholm, T., and Lesser, V. R. 2001. Leveled-Commitment Contracts and Strategic Breach. *Games and Economic Behavior* (Special Issue on AI and Economics) 35:212–270.
- Sandholm, T., and Lesser, V. R. 1997. Coalitions among Computationally Bounded Agents. *Artificial Intelligence* (Special Issue on Economic Principles of Multiagent Systems) 94(1): 99–137.
- Sandholm, T., and Lesser, V. R. 1995. Issues in Automated Negotiation and Electronic Commerce: Extending the Contract Net Framework. Paper presented at the First International Conference on Multi-Agent Systems (ICMAS), 12–14 June, San Francisco, California.
- Sandholm, T., and Vulkan, N. 1999. Bargaining with Deadlines. In Proceedings of the Sixteenth National Conference on Artificial Intelligence (AAAI), 44–51. Menlo Park, Calif.: American Association for Artificial Intelligence.
- Sandholm, T., and Zhou, Y. 2000. Surplus Equivalence of Leveled-Commitment Contracts. Paper presented at the Fourth International Conference on Multi-Agent Systems (ICMAS), 10–12 July, Boston, Massachusetts.
- Sandholm, T.; Sikka, S.; and Norden, S. 1999. Algorithms for Optimizing Leveled-Commitment Contracts. In Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence (IJCAI), 535–540. Menlo Park, Calif.: International Joint Conferences on Artificial Intelligence.
- Sen, S., and Durfee, E. 1998. A Formal Study of Distributed Meeting Scheduling. *Group Decision and Negotiation* 7:265–289.
- Sen, S., and Durfee, E. 1994. The Role of Commitment in Cooperative Negotiation. *International Journal on Intelligent Cooperative Information Systems* 3(1): 67–81.
- Shoham, Y., and Tennenholtz, M. 2001. Fair Imposition. In Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence (IJCAI), 1083–1088. Menlo Park, Calif.: International Joint Conferences on Artificial Intelligence.
- Smith, R. G. 1980. The Contract Net Protocol: High-Level Communication and Control in a Distributed Problem Solver. *IEEE Transactions on Computers* C-29(12): 1104–1113.
- Walsh, W., and Wellman, M. 1999. Efficiency and Equilibrium in Task Allocation Economies with Hierarchical Dependencies. In Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence (IJCAI), 520–526. Menlo Park, Calif.: International Joint Conferences on Artificial Intelligence.
- Yokoo, M. 2000. Algorithms for Distributed Constraint Satisfaction: A Review. *Autonomous Agents and Multi-Agent Systems* 3(2): 189–212.



Tuomas Sandholm is an associate professor in the Computer Science Department at Carnegie Mellon University. He received a Ph.D. and an M.S. in computer science from the University of Massachusetts at Am-

herst in 1996 and 1994, respectively. He earned an M.S. (B.S. included) with distinction in industrial engineering and management science from the Helsinki University of Technology in 1991. While a research scientist at the Technical Research Center of Finland (1990–1992), he built the TRACONET system, one of the first distributed automated negotiation systems for self-interested agents. From then on, his research has focused on designing and building multiagent systems, mainly for electronic commerce, that handle both self-interest and computational complexity. Several of his systems have been commercially fielded. He has published over 125 technical papers and received numerous academic awards, including the NSF Career Award and the inaugural ACM SIGART Autonomous Agents Research Award. His e-mail address is sandholm@cs.cmu.edu.



Victor Lesser has been a professor of computer science at the University of Massachusetts at Amherst since 1977 and serves as the director of the Multi-Agent Systems Laboratory. He received his Ph.D. in computer

science from Stanford University in 1972 and then was a research scientist at Carnegie-Mellon University, where he was the systems architect for the HEARSAY-II speech understanding system, which was the first blackboard system developed. His major research focus is on the control and organization of complex AI systems. Lesser is a founding fellow of the American Association for Artificial Intelligence. His e-mail address is lesser@cs.umass.edu.