

Bargaining with Deadlines

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

Nir Vulkan†
n.vulkan@bristol.ac.uk
Department of Economics
University of Bristol
8 Woodland Road
Bristol BS8 1TN, England

April 12, 1999

Abstract

This paper analyzes automated distributive negotiation where agents have firm deadlines that are private information. The agents are allowed to make and accept offers in any order in continuous time. We show that the only sequential equilibrium outcome is one where the agents wait until the first deadline, at which point that agent concedes everything to the other. This holds for pure and mixed strategies. So, interestingly, rational agents can never agree to a nontrivial split because offers signal enough weakness of bargaining power (early deadline) so that the recipient should never accept. Similarly, the offerer knows that it offered too much if the offer gets accepted: the offerer could have done better by out-waiting the opponent. In most cases, the deadline effect completely overrides time discounting and risk aversion: an agent's payoff does not change with its discount factor or risk attitude. Several implications for the design of negotiating agents are discussed. We also present an effective protocol that implements the equilibrium outcome in dominant strategies.

*Supported by the National Science Foundation under CAREER Award IRI-9703122, Grant IRI-9610122, and Grant IIS-9800994.

†Supported by the EPSRC Award GR M07052.

1 Introduction

Multiagent systems for automated negotiation between self-interested agents are becoming increasingly important. One reason for this is the technology push of a growing standardized communication infrastructure—Internet, WWW, NII, EDI, KQML, FIPA, Concordia, Voyager, Odyssey, Telescript, Java, *etc*—over which separately designed agents belonging to different organizations can interact in an open environment in real-time and safely carry out transactions. Another form of technology push comes from recent advances in automated negotiation technology itself [Sandholm, 1993, Rosenschein & Zlotkin, 1994, Sandholm & Lesser, 1996, Sandholm, Sikka, & Norden, 1999, Sandholm & Zhou, 1999, Sen, 1993, Sandholm & Lesser, 1995, Sandholm, 1997, Andersson & Sandholm, 1998a, Andersson & Sandholm, 1998b, Sandholm, 1998, Andersson & Sandholm, 1999]. The second reason is strong application pull for computer support for negotiation at the operative decision making level. For example, we are witnessing the advent of small transaction electronic commerce on the Internet for purchasing goods, information, and communication bandwidth. There is also an industrial trend toward virtual enterprises: dynamic alliances of small, agile enterprises which together can take advantage of economies of scale when available (e.g., respond to more diverse orders than individual agents can), but do not suffer from diseconomies of scale. Multiagent technology facilitates such negotiation both in terms of processes and outcomes.

For many-to-many negotiation settings, market mechanisms are often used, and for one-to-many negotiation, auctions are often appropriate. The competitive pressure on the side with many agents often reduces undesirable strategic effects. On the other hand, market mechanisms often have difficulty in “scaling down” to small numbers of agents [Osborne & Rubinstein, 1990]. In the limit of one-to-one negotiation, strategic considerations become prevalent. At the same time, one-to-one negotiation settings that crave software agents are ubiquitous. Consider, for example, two scheduling agents negotiating meeting times on behalf of their users, or any e-commerce application where agents negotiate the final price of a good, or a scenario where agents representing different departments bargain over the details of a service which they provide jointly (for example, the ADEPT multi-agent system, currently used by British Telecom, see [Jennings *et al.*, 1996]).

One-to-one negotiation generally involves both integrative and distribu-

tive bargaining. In integrative bargaining the agents search for *Pareto efficient* agreements, i.e. deals such that no other deal exists for making one agent better off without making the other worse off. Intuitively, integrative bargaining is the process of making the joint cake as large as possible. Enumerating and evaluating the Pareto efficient deals can be difficult especially in combinatorially complex settings. Automated negotiating agents hold significant promise in this arena due to their computational speed [Sandholm, 1993].

In distributive bargaining, the focus of this paper, the agents negotiate on how to split the surplus provided by the deal, i.e. how to divide the cake. A continuum of splits is possible at least if the agents can exchange sidepayments. We call any split where each agent gets a nonnegative benefit from the deal *individually rational*, i.e. each agent would rather accept the deal than no deal. Splitting the gains of an optimal contract in an individually rational way can be modeled generically as follows. Without loss of generality, the surplus provided by the contract is normalized to 1, and each agent's fallback payoff that would occur if no contract is made is normalized to 0. Then, distributive bargaining can be studied as the process of “splitting-a-dollar”. This paper focuses on designing software agents that optimally negotiate on the user's behalf in distributive bargaining.

The designer of a multiagent system can construct the interaction *protocol* (aka. mechanism) which determines the legal actions that agents can take at any point in time. Violating the protocol can sometimes be made technically impossible—e.g. disallowing a bidder from submitting multiple bids in an auction—or illegal actions can be penalized e.g. via the regular legal system. To maximize global good, the protocol needs to be designed carefully taking into account that each self-interested agent will take actions so as to maximize its own utility regardless of the global good. In other words, the protocol has to provide the right incentives for the agents. In the extreme, the protocol could specify everything, i.e. give every agent at most one action to choose from at any point. However, in most negotiation settings, the agents can choose whether to participate or not. So, to have the protocol used, the designer has to provide incentives for participation as well. We will return to this question in Section 8.

Perhaps the most well-known economic model of strategic bargaining is the infinite horizon alternating offers game [Rubinstein, 1982]. Since it has a unique solution where agents agree on a split immediately, it seems attractive

for automated negotiation, see e.g. [Kraus, Wilkenfeld, & Zlotkin, 1995]. However, the infinite horizon assumption is often not realistic in the same context. Most ecommerce applications simply cannot allow for negotiations to continue forever. Moreover, the results change considerably if there is a known last period, or even if the distribution from which the number of negotiation rounds is drawn is known and has a finite support. A second disadvantage of the model is that the predictions of the model are specific to exponential time discounting and to the assumption that the dollar is infinitely divisible. For example, under linear time discounting—i.e. a fixed bargaining cost per round of offers—the results change dramatically. The first mover either gets the whole surplus or most of it depending on the ratio of the fixed costs of the two agents. The assumption of perfect information is also of limited use to designers of agents and negotiation protocols. In practice, agents have private information. In such settings, the alternating offers model leads to multiple equilibria, including some where the true types are revealed after long delays, or never. The length of the delay depends on the number of types. See e.g. [Fudenberg & Tirole, 1983], [Fudenberg, Levine, & Tirole, 1985], and [Rubinstein, 1985]. The usefulness of the model as a blueprint for designing agents and protocols is questionable when it allows for such qualitatively different outcomes.

Still, the tools of game theory and mechanism design can be used to study new types of bargaining models, inspired by the various applications of automated negotiation. In fact, game theory and mechanism design theory are more suitable for software agents than for humans because agents can be designed off-line to act rationally even if determining rational strategies is complex [Rosenschein & Zlotkin, 1994]. Also, computational agents do not suffer from emotional irrationality. Finally, the bounded rationality of computational agents can be more systematically characterized than that of humans [Sandholm & Lesser, 1997]. Whether a given communication protocol for self-interested agents will lead to efficient outcomes typically depends on the underlying strategic structure of these interactions. The traditional economic approach of mechanism design which studies first the set of equilibria is therefore of particular relevance to the engineering of such protocols.

In many of the applications we have come across, agents must reach an agreement by a given deadline, which is private information. For example in ADEPT, agents negotiate for services which must be supplied and used

by a certain time [Jennings *et al.*, 1996]. In ADEPT, as in many similar applications, deadlines - and not time discounting - are the main determinant of the agent's bargaining power. In our model therefore agents face firm deadlines in their bargaining. This is also an appealing assumption from a practical perspective since users easily understand deadlines and can trivially communicate them to software agents. Since each agent's deadline is private information, there is a disadvantage in making offers. Any offer—with the exception of demanding everything for oneself—reveals some information about the proposer's deadline, namely that it cannot be very long. If it were, the proposer would stand a good chance of being able to out-wait the opponent, and therefore would ask for a bigger portion of the surplus than it did. Similarly, the offerer knows that it offered too much if the offer gets accepted: the offerer could have done better by out-waiting the opponent. To model real world automated negotiation, we replace the temporal monopoly assumption of alternating offers at fixed intervals, which underlies Rubinstein's model, by a protocol where each agent can make and accept offers in continuous time. However, our results go through even if time is discrete.

Our model resembles war of attrition games where two agents compete for an object and winner-takes-all when one concedes. Those games exhibit multiplicity of equilibria. If the value of the object is common knowledge, there is a symmetric equilibrium where one agent concedes immediately. There is also a symmetric equilibrium where agents concede at a rate which depends on the value of the object. [Hendricks, Weiss, & Wilson, 1988] study the general class of war of attrition games with perfect information. [Hendricks & Wilson, 1989] study the case of incomplete information, normally about the object's value, see also [Riley, 1980]. The game was introduced by [Smith, 1974] in a biological context and has been applied in industrial economics [Fudenberg *et al.*, 1983], [Kreps & Wilson, 1982]).

Our model differs from the war of attrition. Agents are allowed to split the dollar instead of winner-takes-all. One would expect that to enlarge the set of equilibria and equilibrium outcomes. This intuition turns out to be false. We show that the only equilibrium outcome is one where agreement is delayed until one of the deadlines is reached, and then one agent gets the entire surplus.

We show that there exists a sequential equilibrium where agents do not agree to a split until the first deadline, at which time the agent with the later deadline receives the whole surplus. Conversely, we show that there do not

exist any other Bayes-Nash equilibria where agents agree to any other split at any other time. Therefore both our positive and negative results are strong with respect to the degree of sequential rationality agents are assumed to have. Intuitively speaking, in these equilibria the agents update their beliefs rationally, and neither agent is motivated to change its strategy at any point of the game given that the other agent does not change its strategy.

Our results are robust in other ways as well. First, there does not exist even a mixed strategy equilibrium where an agent concedes at any rate before its deadline. This, again, is in contrast with the usual equilibrium analysis of war of attrition games. Second, the results hold even if the agents discount time in addition to having deadlines. Third, even if agents have different risk attitudes, they will not agree to any split before their deadline. That is, even risk averse agents will refuse safe and generous offers and will instead prefer to continue to the risky “waiting game”.

The rest of the paper is organized as follows. Section 2 describes our formal model of bargaining with deadlines. Section 3 presents our main results for pure-strategy equilibria. Section 4 extends them to mixed strategies. Sections 5 and 6 present the results with time discounting and with agents with risk attitudes. Section 7 describes the entailed insights for designing automated negotiating agents. Section 8 discusses implications for the design of interaction protocols. Finally, Section 9 concludes.

2 A model of bargaining under deadlines

Our bargaining game, (a, b) , has two agents, 1 and 2. The type of agent 1 is its deadline d_1 . The type of agent 2 is its deadline d_2 . The types are private information: each agent only knows its own type. The type d_1 is drawn from a distribution $a(d_1)$, and d_2 is independently drawn from a distribution $b(d_2)$, where a and b are common knowledge. Without loss of generality we scale the time axis so that

$$\min(\min(d_1), \min(d_2)) = 0, \text{ and} \tag{1}$$

$$\max(\max(d_1), \max(d_2)) = 1 \tag{2}$$

That is, $d_1, d_2 \in (0, 1)$.

A history $\langle H(t) \rangle$ is a (possibly empty) list of all unaccepted offers up to time t . An *action* at time t describes the threshold of what offers are acceptable to the agent from then on, and what offer the agent has outstanding from then on. We do not assume that the agents have to improve their offers over time. They can even offer less than they offered earlier. Each agent also has a belief updating rule. Finally, each agent chooses a strategy for the game. A *pure strategy* is an agent's deterministic mapping from its history to action, i.e. it defines what offers the agent would accept and make as a function of what offers and rejections/acceptances the other agent has made so far. A *mixed strategy* is an agent's probability mixture over pure strategies.

The agents can agree to any split $(x, 1 - x)$ where $x \in [0, 1]$. The payoffs from an agreement $(x, 1 - x)$ at time $t < d_1, d_2$ are x and $1 - x$ for agents 1 and 2 respectively.¹ The payoff for an agent from any agreement which takes place after its deadline is 0. We assume that an agent strictly prefers to hand over the whole surplus than to miss its deadline. In other words, if the agent will get zero payoff anyway, it will rather give the surplus to the other agent than not. This discontinuity of preferences at the deadline is not crucial, although an additional set of equilibria appears where agents miss their deadline with probability 1 if this assumption is removed. Still, this assumption can be defended on the grounds that there is usually a small (possibly infinitesimally small), but positive utility associated with reaching an agreement.

3 Pure strategy equilibria

It turns out that a “sit-and-wait” way of playing the game is one rational way for agents to behave:

Proposition 3.1 *There exists a sequential equilibrium of (a, b) , where the agent with the latest deadline receives the whole surplus exactly at the earlier deadline.*

Proof. Consider the following (symmetric) strategies: $s_i(d_i)$: demand $x = 1$ (everything) at any time $t < d_i$. At $t = d_i$ accept any offer. At any

¹Payoffs could be defined in several ways. One approach, by [Smith, 1997], is to define payoffs in terms of aspiration levels, or acceptance thresholds. However, it turns out to be sufficient for the results of this paper to express payoffs not in terms of strategies, but in the traditional way of stating them in terms of preferences over final agreements.

time $t < d_i$, reject offers $x_i \geq d_i$. Update beliefs according to Bayes rule, putting all the weight of the posterior distribution over the values of the opponent's deadline, d_{-i} , over the interval $(1 - x_i, 1)$. At any time $t < d_i$, reject offers $x_i < d_i$. Update beliefs in the following way. If $1 - x_i$ is already in the prior of d_{-i} then the posterior is simply equal to the prior. If $1 - x_i$ is not in the prior of d_{-i} then add the point $1 - x_i$ and update according to Bayes rule.

We first show that the beliefs specified above are consistent: Let s_i^ϵ (for $i = 1, 2$) assign probability $1 - \epsilon$ to the above specified posterior beliefs, and probability ϵ to the rest of the support of d_{-i} . As $\epsilon \rightarrow 0$ the fully mixed strategy pair converges to (s_1, s_2) , while the beliefs generated by the fully mixed strategy pair converge to the beliefs described above. It is now easy to see that, given these beliefs, actions are sequentially rational. Along the equilibrium path agents will always demand the full surplus and therefore no agreement will be achieved before the first deadline. At the first deadline, the agent with the later deadline will receive the whole surplus. \square

The equilibrium described above is clearly also an equilibrium of any subgame beginning at time t onwards regardless of the history, $\langle H(t) \rangle$. That is, at any stage of the game each agent can move to this waiting game, and the other agent's best response is to do the same. We use this property in the proof of our main result which states that the "sit-and-wait" equilibrium is the only sequential equilibrium of (a, b) . In fact, we prove a stronger claim that there is no equilibrium even using a weaker definition of equilibrium, the Bayes-Nash equilibrium. It follows that no sequential equilibrium exists either. We first present the result for pure strategies. We generalize it to mixed strategies in Section 4.

Theorem 3.1 *If $d_1 > 0$ or $d_2 > 0$, there does not exist a pure strategy Bayes-Nash equilibrium of (a, b) , where agents agree to a split other than $(1, 0)$ or $(0, 1)$.*

Proof. Assume, for contradiction, that there exist types $d_1 > 0$ and $d_2 > 0$ and a pure strategy Bayes-Nash equilibrium (s_1, s_2) where the agents agree to a split $(\pi_1, \pi_2) = (x, 1 - x)$, where $x \in (0, 1)$ at time $t \geq 0$. We assume, without loss of generality, that agent 1 receives at least one half, i.e. $x \geq \frac{1}{2}$. We can therefore write $x = \frac{1}{2} + \epsilon$, where $\frac{1}{2} > \epsilon \geq 0$.

Let $g_0(d_2)$ denote agent 1's beliefs about d_2 at time t . Similarly, let $f_0(d_1)$ denote agent 2's beliefs about d_1 at time t . Denote by $G(d_2) \equiv \int_t^1 g_0(d_2) dd_2$

the cumulative distribution of g_0 , and by $F(d_1) \equiv \int_t^1 f_0(d_1) dd_1$ the cumulative distribution of f_0 .

In equilibrium, agent 2 will accept $1-x$ only if it does not expect to receive more by unilaterally moving to the waiting game. The expected payoff from the waiting game is simply agent 2's subjective probability that $d_2 < d_1$. Hence agent 2 would accept only if

$$\frac{1}{2} - \epsilon \geq \int_t^{d_2} f_0(d_1) dd_1 = F(d_2) - F(t) = F(d_2) \quad (3)$$

In other words, agent 2's type, d_2 , must not be too high. Let $\alpha(y) \equiv \inf[d_2 \mid y \geq F(d_2)]$. With this notation, (3) can be rewritten as $d_2 \leq \alpha(\frac{1}{2} - \epsilon)$.

Now, agent 1 will only accept this offer if it will give it an expected payoff at least as large as that of the waiting game, which equals its subjective probability of winning the waiting game. There are two cases. First, if $d_1 > \alpha(\frac{1}{2} - \epsilon)$, agent 1 knows that it will win the waiting game with probability one, so the split $(x, 1-x)$ could not occur in equilibrium. The second case occurs when $d_1 \leq \alpha(\frac{1}{2} - \epsilon)$. Agent 1 can use the fact that $d_2 \leq \alpha(\frac{1}{2} - \epsilon)$ to update its beliefs about agent 2's deadline as follows:

$$\begin{aligned} g_1(d_2) &= 0 \text{ if } d_2 > \alpha(\frac{1}{2} - \epsilon). \text{ Otherwise,} \\ g_1(d_2) &= g_0(d_2) \frac{\int_t^1 g_0(d_2)}{\int_t^{\alpha(\frac{1}{2} - \epsilon)} g_0(d_2)} \\ &= g_0(d_2) \underbrace{\left[1 + \frac{\int_{\alpha(\frac{1}{2} - \epsilon)}^1 g_0(d_2)}{\int_t^{\alpha(\frac{1}{2} - \epsilon)} g_0(d_2)} \right]}_{\geq 2 \text{ because } \alpha(\frac{1}{2} - \epsilon) \leq \text{median}(g_0)} \end{aligned} \quad (4)$$

Based on these updated beliefs, agent 1 would accept only if

$$\frac{1}{2} + \epsilon \geq \int_t^{d_1} g_1(d_2) dd_2 \geq 2[G(d_1) - G(t)] = 2G(d_1) \quad (5)$$

In other words, agent 1's type, d_1 , must also not be too high. Let $\beta(y) \equiv \inf[d_1 \mid y \geq G(d_1)]$. With this notation, (5) can be rewritten as $d_1 \leq \beta(\frac{1}{2} + \epsilon)$.

In equilibrium, agent 2 only accepts if it gives it an expected payoff at least as large as that of the waiting game, which equals its subjective probability

of winning the waiting game. There are two cases. First, if $d_2 > \beta(\frac{1}{2} + \epsilon)$, agent 2 knows that it will win the waiting game with probability one, so the split $(x, 1 - x)$ could not occur. The second case occurs when $d_2 \leq \beta(\frac{1}{2} + \epsilon)$. Agent 2 can use the fact that $d_1 \leq \beta(\frac{1}{2} + \epsilon)$ to update its beliefs about agent 1's deadline as follows:

$$\begin{aligned}
f_1(d_1) &= 0 \text{ if } d_1 > \beta(\frac{1}{2} + \epsilon). \text{ Otherwise,} \\
f_1(d_1) &= f_0(d_1) \frac{\int_t^1 f_0(d_1)}{\int_t^{\beta(\frac{1}{2} + \epsilon)} f_0(d_1)} \\
&= f_0(d_1) \left[1 + \frac{\int_{\beta(\frac{1}{2} + \epsilon)}^1 f_0(d_1)}{\int_t^{\beta(\frac{1}{2} + \epsilon)} f_0(d_1)} \right] \\
&\geq 2 \text{ because } \beta(\frac{1}{2} + \epsilon) \leq \text{median}(f_0)
\end{aligned} \tag{6}$$

Based on these updated beliefs, agent 2 would accept only if

$$\frac{1}{2} - \epsilon \geq \int_t^{d_2} f_1(d_1) dd_1 \geq 2[F(d_2) - F(t)] = 2F(d_2) \tag{7}$$

i.e. $d_2 \leq \alpha(\frac{1}{2} - \epsilon)$.

This process of belief update and acceptance threshold resetting continues to alternate between agents. After r rounds of this alternation, all types have been eliminated except those that satisfy $d_1 \leq \beta(\frac{1}{2} - \frac{\epsilon}{2^r})$ and $d_2 \leq \alpha(\frac{1}{2} + \frac{\epsilon}{2^r})$. This process can continue for an unlimited number of steps, r , so the upper bounds approach zero. Therefore the equilibrium cannot exist if $d_1 > 0$ or $d_2 > 0$. Contradiction. \square

4 Mixed strategy equilibria

We now strengthen our impossibility result by showing that it holds for mixed strategies as well, i.e. that there is no other rational way of playing the game than “sit-and-wait” even if randomization is possible. This is yet another difference between our setting and war of attrition games. In the latter,

mixed strategies play an important role: typically the unique symmetric equilibrium has concession rates that are mixed strategies.

Theorem 4.1 *If $d_1 > 0$ or $d_2 > 0$, there does not exist a mixed strategy Bayes-Nash equilibrium of , (a, b) , where the agents agree to a split other than $(1, 0)$ or $(0, 1)$ with positive probability.*

Proof. Assume, for contradiction, that there exist types $d_1 > 0$ and $d_2 > 0$ and a mixed strategy Bayes-Nash equilibrium where there is positive probability of an agreement other than $(1, 0)$ or $(0, 1)$. Now there has to exist at least one point in time, t , where there is positive probability of an agreement other than $(1, 0)$ or $(0, 1)$. We analyze the equilibrium at such a time t . Recall f, g, F, G, α , and β from the proof of Theorem 3.1.

Agent 1 will accept an agreement if it gets a share $x \geq a_1$, where a_1 is its acceptance threshold. That threshold depends on its type. Since we are analyzing a mixed strategy equilibrium, the threshold can also depend on randomization. We therefore say that a_1 is randomly chosen for time t from a probability density function $m(a_1)$. Similarly, agent 2 will accept an agreement if it has to offer 1 a share of $x \leq a_2$ where a_2 is agent 2's offering threshold. We say that a_2 is chosen for time t from a probability density function $n(a_2)$.

Without loss of generality, we assume that there is positive probability that the agreement is made in the range $x \geq \frac{1}{2}$. This implies that there is positive probability that $a_2 \geq \frac{1}{2}$.

Let \underline{a}_1 be the smallest a_1 in the support of m (alternatively let \underline{a}_1 be the infimum of m). The assumption that there is positive probability of an agreement other than $(1, 0)$ or $(0, 1)$ means that $\underline{a}_1 = 1 - \epsilon$ for some $\epsilon > 0$.

Because the strategies are in equilibrium, m and n must be best responses to each other. For n to be a best response, each threshold, a_2 , in the support of n has to give agent 2 at least the same payoff as it would get by going to the waiting game. Focusing on those a_2 for which $a_2 \geq \frac{1}{2}$ this means

$$\frac{1}{2} \geq E[\pi_2^{wait}] = \int_t^{d_2} f_0(d_1) dd_1 = F(d_2) - F(t) = F(d_2) \quad (8)$$

So, $d_2 \leq \alpha(\frac{1}{2})$.

Now, in equilibrium, every strategy in the support of m has to give agent 1 at least the same payoff that it would get by going to the waiting game.

There are two cases. First, if $d_1 > \alpha(\frac{1}{2})$, agent 1 knows that it will win the waiting game with probability one, so the split $(x, 1 - x)$ could not occur in equilibrium. The second case occurs when $d_1 \leq \alpha(\frac{1}{2})$. Agent 1 can use the fact that $d_2 \leq \alpha(\frac{1}{2})$ to update its beliefs about agent 2's deadline:

$$\begin{aligned}
g_1(d_2) &= 0 \text{ if } d_2 > \alpha(\frac{1}{2}). \text{ Otherwise,} \\
g_1(d_2) &= g_0(d_2) \frac{\int_t^1 g_0(d_2)}{\int_t^{\alpha(\frac{1}{2})} g_0(d_2)} \\
&= g_0(d_2) \left[1 + \frac{\int_{\alpha(\frac{1}{2})}^1 g_0(d_2)}{\int_t^{\alpha(\frac{1}{2})} g_0(d_2)} \right] \\
&\quad \underbrace{\hspace{10em}}_{\geq 2 \text{ because } \alpha(\frac{1}{2}) \leq \text{median}(g_0)} \tag{9}
\end{aligned}$$

Based on these updated beliefs, the support of m can include \underline{a}_1 only if

$$\underline{a}_1 = 1 - \epsilon \geq \int_t^{d_1} g_1(d_2) dd_2 \geq 2[G(d_1) - G(t)] = 2G(d_1) \tag{10}$$

In other words, agent 1's type, d_1 , cannot be too high. Specifically, this can be written as $d_1 \leq \beta(\frac{1-\epsilon}{2})$.

In equilibrium, every strategy in the support of n has to give agent 2 at least the same payoff that it would get by going to the waiting game, which equals its subjective probability of winning the waiting game. There are two cases. First, if $d_2 > \beta(\frac{1-\epsilon}{2})$, agent 2 knows that it will win the waiting game with probability one, so the split $(x, 1 - x)$ could not occur. The second case occurs when $d_2 \leq \beta(\frac{1-\epsilon}{2})$. Agent 2 can use the fact that $d_1 \leq \beta(\frac{1-\epsilon}{2})$ to update its beliefs about agent 1's deadline:

$$\begin{aligned}
f_1(d_1) &= 0 \text{ if } d_1 > \beta(\frac{1-\epsilon}{2}). \text{ Otherwise,} \\
f_1(d_1) &= f_0(d_1) \frac{\int_t^1 f_0(d_1)}{\int_t^{\beta(\frac{1-\epsilon}{2})} f_0(d_1)} \\
&= f_0(d_1) \left[1 + \frac{\int_{\beta(\frac{1-\epsilon}{2})}^1 f_0(d_1)}{\int_t^{\beta(\frac{1-\epsilon}{2})} f_0(d_1)} \right] \\
&\quad \underbrace{\hspace{10em}}_{\geq 2 \text{ because } \beta(\frac{1-\epsilon}{2}) \leq \text{median}(f_0)} \tag{11}
\end{aligned}$$

Based on these updated beliefs, and focusing on those a_2 for which $a_2 \geq \frac{1}{2}$ we can rule out high values of d_2 (otherwise agent 2 would be better off by waiting):

$$\frac{1}{2} \geq E[\pi_2^{wait}] = \int_t^{d_2} f_1(d_1) dd_1 \geq 2F(d_2) \quad (12)$$

i.e. $d_2 \leq \alpha(\frac{1}{4})$.

This process of belief update and acceptance threshold resetting continues to alternate between agents. After r rounds of this alternation, all types have been eliminated except those that satisfy $d_1 \leq \beta(\frac{1-\epsilon}{2^{r+1}})$ and $d_2 \leq \alpha(\frac{1}{2^{r+1}})$. This process can continue for an unlimited number of steps, r , so the upper bounds approach zero. Therefore the equilibrium cannot exist if $d_1 > 0$ or $d_2 > 0$. Contradiction. \square

5 Incorporating discounting

Time discounting is a standard way of modeling settings where the value of the good decreases over time, e.g. due to inflation or due to perishing. In the previous sections we assumed that agents do not discount time. However, we now show that our results are robust to the case where agents do discount time in addition to having firm deadlines. Let δ_1 be the discount factor of agent 1, and δ_2 be the discount factor of agent 2. The utility of agent i from an agreement where he receives a share x at time $t < d_i$ is then $\delta_i^t x$. We denote by $(a, b, \delta_1, \delta_2)$ the bargaining game where a , b , δ_1 , and δ_2 are common knowledge. We now prove that our previous result for (a, b) holds also for a large range of parameters in $(a, b, \delta_1, \delta_2)$. So, interestingly, the bargaining power of an agent does not change with its discount factor, in contrast to the results of most other bargaining games. In other words, the deadline effect completely suppresses the discounting effect. This result is important in its own right for the design of automated negotiating agents, and it also motivates the study of deadline-based models as opposed to focusing only on discounting-based ones.

Proposition 5.1 *For any $\delta_1, \delta_2, 0 < \delta_1 \leq 1, 0 < \delta_2 \leq 1$, there exists a sequential equilibrium of $(a, b, \delta_1, \delta_2)$ where the agent with the latest deadline receives the whole surplus exactly at the earlier deadline.*

Proof. The equilibrium strategies and proof of sequential equilibrium are identical to those in the proof of Proposition 3.1 with the difference that the threshold is no longer d_i but $\delta_i^t d_i$. Also the posteriors are now defined only until δ_i^t and not 1. \square

Theorem 5.1 *If $\delta_1 \delta_2 > \frac{1}{2}$, there does not exist a Bayes-Nash equilibrium of $(a, b, \delta_1, \delta_2)$, (in pure or mixed strategies) where agents agree to a split other than $(1, 0)$ or $(0, 1)$.*

Proof. We prove the case for pure strategy equilibrium. The extension for mixed strategy equilibrium is identical to that in Theorem 4.1. The proof is a variant of the proof in Theorem 3.1, and we keep the notation from there.

Assume, for contradiction, that there exist types $d_1 > 0$ and $d_2 > 0$ and a pure strategy Bayes-Nash equilibrium (s_1, s_2) where the agents agree to a split of the total surplus available at time t , according to proportions $(\pi_1, \pi_2) = (x, 1 - x)$, where $x \in (0, 1)$ at time $t \geq 0$. We assume, without loss of generality, that agent 1 receives at least one half, i.e. $x \geq \frac{1}{2}$. We can therefore write $x = \frac{1}{2} + \epsilon$, where $\frac{1}{2} > \epsilon \geq 0$.

In equilibrium, agent 2 will accept $1 - x$ only if it does not expect to receive more by unilaterally moving to the waiting game. The expected payoff from the waiting game is now agent 2's subjective probability that $d_2 < d_1$, multiplied by the discounted value of winning. Hence agent 2 would accept only if

$$\delta_2^t \left(\frac{1}{2} - \epsilon \right) \geq \int_t^{d_2} \delta_2^{d_1} f_0(d_1) dd_1 \geq \delta_2 \int_t^{d_2} f_0(d_1) dd_1 \quad (13)$$

Dividing both sides by δ_2^t , we get:

$$\frac{1}{2} - \epsilon \geq \delta_2^{1-t} \int_t^{d_2} f_0(d_1) dd_1 = \delta_2^{1-t} F(d_2) \geq \delta_2 F(d_2) \quad (14)$$

In other words: $d_2 \leq \alpha \frac{(\frac{1}{2} - \epsilon)}{\delta_2}$.

Now, agent 1 can use this to update its beliefs about agent 2's deadline in the same way as in equation (4), with the difference that now $g_1 \geq 2\delta_2 g_0$. Since $\delta_2 > 0.5$ (by the assumption that $\delta_1 \delta_2 > 0.5$), we know that $g_1 > g_0$ (when g_1 is not zero).

Based on these updated beliefs, agent 1 would accept only if

$$\delta_1^t \left(\frac{1}{2} + \epsilon \right) \geq \int_t^{d_1} \delta_1^{d_2} g_1(d_2) dd_2 \geq \delta_1 \int_t^{d_1} g_1(d_2) dd_2 \quad (15)$$

Dividing both sides by δ_1^t and using the updated beliefs, g_1 , we can now rule out "high" types of agent 1. Formally, $d_1 \leq \beta(\frac{\frac{1}{2}+\epsilon}{2\delta_2\delta_1})$.

Once more, belief updating by agent 2 (in the same way as in (6)) yields $f_1 \geq 2\delta_2\delta_1 f_0$. Since $\delta_1\delta_2 > 0.5$ we get that $f_1 > f_0$ (when it is not zero). Based on these updated beliefs, agent 2 would accept only if

$$\delta_2^t(\frac{1}{2} - \epsilon) \geq \int_t^{d_2} \delta_2^{d_1} f_1(d_1) dd_1 \geq \delta_2 \int_t^{d_2} f_0(d_1) dd_1 \quad (16)$$

Dividing by δ_2^t and using the updated beliefs, f_1 we can rule out the following types: $d_2 \leq \alpha(\frac{\frac{1}{2}-\epsilon}{2\delta_2^2\delta_1}) < \alpha(\frac{\frac{1}{2}-\epsilon}{2\delta_2})$. This process of belief update and acceptance threshold resetting continues to alternate between agents. After r rounds of this alternation, all types have been eliminated except those that satisfy $d_1 \leq \beta(\frac{\frac{1}{2}-\epsilon}{(2\delta_1)^r})$ and $d_2 \leq \alpha(\frac{\frac{1}{2}+\epsilon}{(2\delta_2)^r})$. This process can continue for an unlimited number of steps, r , so the upper bounds approach zero. Therefore the equilibrium cannot exist if $d_1 > 0$ or $d_2 > 0$. Contradiction. \square

To satisfy the condition $\delta_1\delta_2 > \frac{1}{2}$, time discounting must not be too significant compared to the deadlines, i.e. the cake cannot shrink too much before the deadline. Since most deadline bargaining situations will certainly have shorter deadlines than the time it will take for the cake to shrink to, say, one half of its original size, Theorem 5.1 shows that "sit-and-wait" remains the only equilibrium. So, in practice, the effect of deadlines suppresses that of discount factors. This is even more commonly true in automated negotiation because that is most likely going to be used mainly for fast negotiation at the operative decision making level instead of strategic long-term negotiation.

6 Robustness to risk attitudes

We now generalize our results to agents that are not necessarily risk neutral. Usually in bargaining games the equilibrium split of the surplus depends on the agents' risk attitudes. However, we show that this does not happen in our setting. This is surprising at first since a risk averse agent generally prefers a smaller but safe share to the risky option of the waiting game even if she expects to win it with high probability. However, we show that the type-elimination effect described in the theorems so far is still present and dominates any concessions which may be consistent with risk aversion.

Let the agents' risk attitudes be captured by utility functions, u_i where $i = 1, 2$. Without loss of generality we let $u_i(0) = 0$ and $u_i(1) = 1$ for both agents.

Proposition 6.1 *There exists a sequential equilibrium of, (a, b, u_1, u_2) , where the agent with the latest deadline receives the whole surplus exactly at the earlier deadline.*

Proof. The equilibrium strategies and proof of sequential equilibrium are identical to those in the proof of Proposition 3.1 with the difference that the threshold is no longer d_i but $u_i(d_i)$. \square

The following definition is used to state our main result for the case of different risk attitudes.

Definition. 6.1 *The maximum risk aversion of agent i is*

$$\rho_i \equiv \max_x \frac{u_i(x)}{x} \quad (17)$$

We can now show that our impossibility result applies to a large range of risk attitudes of the agents:

Theorem 6.1 *If $\rho_1 \rho_2 < 2$, there does not exist a Bayes-Nash equilibrium (pure or mixed) of, (a, b, u_1, u_2) , where agents agree to a split other than $(1, 0)$ or $(0, 1)$.*

Proof. We prove the case for pure strategy equilibrium. The extension to mixed strategies is identical to that in Theorem 4.1. We keep the notations from Theorem 3.1. Assume, for contradiction, that there exist types $d_1 > 0$ and $d_2 > 0$ and a pure strategy Bayes-Nash equilibrium (s_1, s_2) where the agents agree to a split $(\pi_1, \pi_2) = (x, 1 - x)$, where $x \in (0, 1)$ at time $t \geq 0$. Assume, without loss of generality, that agent 1 receives at least half, i.e. $x \geq \frac{1}{2}$. Thus we can write $x = \frac{1}{2} + \epsilon$, where $\frac{1}{2} > \epsilon \geq 0$. In equilibrium, agent 2 accepts $1 - x$ only if it does not expect to receive more by unilaterally moving to the waiting game. The expected payoff from the waiting game is agent 2's subjective probability that $d_2 < d_1$. So, agent 2 would accept only if

$$\rho_2 \left(\frac{1}{2} - \epsilon \right) \geq u_2 \left(\frac{1}{2} - \epsilon \right) \geq \int_t^{d_2} f_0(d_1) dd_1 = F(d_2) \quad (18)$$

In other words, $d_2 \leq \alpha(\rho_2(\frac{1}{2} - \epsilon))$. Now, agent 1 can use this to update its beliefs about agent 2's deadline in the same way as in equation (4), with the difference that now $g_1 \geq \frac{2}{\rho_2}g_0$. Since (by assumption) $\rho_2 < 2$, then $g_1 > g_0$ (when g_1 is not zero). Based on these updated beliefs, agent 1 would accept only if:

$$u_1(\frac{1}{2} + \epsilon) \geq \int_t^{d_1} g_1(d_2)dd_2 \quad (19)$$

Using ρ_1 and the updated beliefs, g_1 , we can rule out "high" types of agent 1. Formally, $d_1 \leq \beta(\rho_2\rho_1(\frac{1}{2} + \epsilon))$.

Once more, belief updating by agent 2 (in the same way as in (6)) yields $f_1 \geq \frac{2}{\rho_2\rho_1}f_0$. Since $\rho_2\rho_1 < 2$ we get $f_1 > f_0$ (when f_1 is not zero). Based on these updated beliefs, agent 2 would accept only if

$$u_2(\frac{1}{2} - \epsilon) \geq \int_t^{d_2} f_1(d_1)dd_1 \quad (20)$$

Using ρ_2 and the updated beliefs, f_1 , we can rule out the following types: $d_2 \leq \alpha(\rho_2^2\rho_1(\frac{1}{2} - \epsilon)) < \alpha(\rho_2(\frac{1}{2} - \epsilon))$. This process of belief update and acceptance threshold resetting continues to alternate between agents. After r rounds of this alternation, all types have been eliminated except those that satisfy $d_1 \leq \beta(\rho_1^r(\frac{1}{2} - \epsilon))$ and $d_2 \leq \alpha(\rho_2^r(\frac{1}{2} - \epsilon))$. This process can continue for an unlimited number of steps, r , so the upper bounds approach zero. Therefore the equilibrium cannot exist if $d_1 > 0$ or $d_2 > 0$. Contradiction. \square

7 Designing bargaining agents

Our motivation for studying bargaining with deadlines stems from our long-term research agenda of constructing software agents that will optimally negotiate on behalf of the real world parties that they represent. That will put experienced and poor human negotiators on an equal footing, and save human negotiation effort.

Deadlines are widely advocated and used in automated electronic commerce to capture time preference. For example when a user delegates priceline.com to find an inexpensive airline flight on the web, the user gives it one hour to complete (while priceline.com uses an agent with a deadline, the setting is a form of auction, not bargaining). Users easily understand deadlines, and it is simple to specify a deadline to an agent.

Our results show that in distributive bargaining settings with two agents with deadlines, it is not rational for either agent to make or accept offers. But what if a rational software agent receives an offer from the other party? This means that the other party is irrational, and could perhaps be exploited. However, the type-elimination argument from the proofs above applies here too, and it is not rational for the software agent to accept the offer, no matter how good it is. To exploit the other party, the agent would have to have an opponent model to model the other party’s irrationality. While game theory allows us to give precise prescriptions for rational play, it is mostly silent about irrationality, and how to exploit it.

Another classic motivation for automated negotiation is that computerized agents can negotiate faster. However, in distributive bargaining settings where the agents have deadlines, this argument does not hold because in such settings, rational software agents would sit-and-wait until one of the deadlines is reached. From an implementation perspective this suggests the use of daemons that trigger right before the deadline instead of agents that use computation before the deadline.

Finally, our results suggest that a user will be in a much stronger bargaining position by inputting time preferences to her agent in terms of a time discount function instead of a deadline, even if the discounting is significant. To facilitate this, software agent vendors should provide user interfaces to their agents that allow easy human-understandable specification of time discounting functions instead of inputting a deadline.

8 Designing bargaining protocols

The following mechanism implements, in dominant strategies, the equilibrium of the deadline bargaining game described above. First, agents report their deadlines, \tilde{d}_i , to the protocol—possibly insincerely ($\tilde{d}_i \neq d_i$). The protocol then assigns the whole dollar to the agent with the highest \tilde{d}_i , but this only takes place at time $t = \tilde{d}_{-i}$, i.e. at the earlier reported deadline. Truth-telling is a dominant strategy in this mechanism. By reporting $\tilde{d}_i < d_i$ agent i ’s probability of winning is reduced. By reporting $\tilde{d}_i > d_i$, agent i increase its probability of winning, but only in cases where $\tilde{d}_{-i} > d_i$, i.e. when i misses its deadline. Therefore, reporting $\tilde{d}_i = d_i$ is a dominant strategy.

This mechanism is efficient in several ways. First, it minimizes coun-

terspeculation. In equilibria that are based on refinements of the Nash equilibrium—such as Bayes-Nash or sequential equilibrium—an agent’s best strategy depends on what others will do. This requires speculation about the others’ strategies, which can be speculated by considering the others’ types, their rationality, what they think of the former agent, what they think the former agent thinks about them, *etc. ad infinitum*. On the other hand, in a dominant strategy mechanism an agent’s strategy is optimal no matter what others do. Therefore, counterspeculation is not useful. The agent need not waste time in counterspeculation which can be intractable or even noncomputable. In addition, it is easier to program an agent that executes a given dominant strategy than an agent that counterspeculates (see e.g. [Vulkan, 1999] for a discussion). Second, dominant strategy mechanisms are robust against irrational agents since their actions do not affect how others should behave. Finally, the mechanism minimizes communication: each agent only sends one message.

However, the mechanism is not Pareto efficient if time is being discounted, because the agreement is delayed as it was in the original free-form bargaining game. In such settings, any mechanism that results in an immediate agreement is Pareto efficient, e.g. a protocol that forces a 50:50 split up front. This protocol is efficient in all respects discussed above. It might seem like a good solution to the problem raised by our impossibility results. However, agents in e-commerce applications usually can choose whether to use a protocol or not. If agents know their types before they choose the protocol they want to use, an adverse selection problem arises. To see why, assume that types are normally distributed. This is assumed for simplicity of presentation and is not crucial. Agents with deadlines above $1/2$ will not participate in such a protocol because they can expect to do better in a free-form bargaining setting. But if only agents with deadlines below $1/2$ participate, agents with deadlines between $1/4$ and $1/2$ should not participate. Next, agents with deadlines between $1/8$ and $1/4$ would not participate, and so on. In equilibrium, no agent would participate. This argument does not rely on a 50:50 split. The adverse selection problem will affect any protocol that does not implement Nash (or stronger) equilibrium outcomes.

9 Conclusions

Automated agents have been suggested as a way to facilitate increasingly efficient negotiation. In settings where the bargaining set, i.e. set of individually rational Pareto efficient deals, is difficult to construct for example due to a combinatorial number of possible deals [Sandholm, 1993] or the computational complexity of evaluating any given deal [Sandholm & Lesser, 1997], the computational speed of automated agents can significantly enhance negotiation. Additional efficiency can stem from the fact that computational agents can negotiate with large numbers of other agents quickly and virtually with no negotiation overhead. However, this paper showed that in one-to-one negotiation where the optimal deal in the bargaining set has been identified and evaluated, and distributing the profits is the issue, an agent's power does not stem from speed, but on the contrary, from the ability to wait.

We showed that in one-to-one bargaining with deadlines, the only sequential equilibrium is one where the agents wait until the first deadline is reached. This is in line with some human experiments where adding deadlines introduced significant delays in reaching agreement [Roth, Murnighan, & Schoumaker, 1988]. We also showed that deadline effects almost always completely suppress time discounting effects. Impossibility of an interim agreement also applies to most types of risk attitudes of the agents. The results show that for deadline bargaining settings it is trivial to design the optimal agent: it should simply wait until it reaches its deadline or the other party concedes. On the other hand, a user is better off by giving her agent a time discount function instead of a deadline since a deadline puts her agent in a weak bargaining position. Finally, we discussed mechanism design, and presented an effective protocol that implements the outcome of the free-form bargaining game in dominant strategy equilibrium.

References

- [Andersson & Sandholm, 1998a] Andersson, M. R., and Sandholm, T. W. 1998a. Leveled commitment contracting among myopic individually rational agents. In *Proceedings of the Third International Conference on Multi-Agent Systems (ICMAS)*, 26–33.

- [Andersson & Sandholm, 1998b] Andersson, M. R., and Sandholm, T. W. 1998b. Leveled commitment contracts with myopic and strategic agents. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, 38–45.
- [Andersson & Sandholm, 1999] Andersson, M. R., and Sandholm, T. W. 1999. Time-quality tradeoffs in reallocative negotiation with combinatorial contract types. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*.
- [Fudenberg & Tirole, 1983] Fudenberg, D., and Tirole, J. 1983. Sequential bargaining with incomplete information. *Review of Economic Studies* 50:221–247.
- [Fudenberg *et al.*, 1983] Fudenberg, D.; Gilbert, R.; Stiglitz, J.; and Tirole, J. 1983. Preemption, leapfrogging, and competition in patent races. *European Economic Review* 22:3–31.
- [Fudenberg, Levine, & Tirole, 1985] Fudenberg, D.; Levine, D.; and Tirole, J. 1985. Infinite-horizon models of bargaining with one-sided incomplete information. In Roth, A. E., ed., *Game-Theoretic Models of Bargaining*. Cambridge University Press. 73–98.
- [Hendricks & Wilson, 1989] Hendricks, K., and Wilson, C. 1989. The war of attrition in discrete time. Technical report, State University of New York, Stony Brook. Mimeo.
- [Hendricks, Weiss, & Wilson, 1988] Hendricks, K.; Weiss, A.; and Wilson, C. 1988. The war of attrition in continuous time with complete information. *International Economic Review* 29:663–680.
- [Jennings *et al.*, 1996] Jennings, N.; Farain, P.; Johnson, M.; Norman, T.; O’Brein, P.; and Wiegand, M. 1996. Agent-based business process management. *International Journal on Intelligent Cooperative Information Systems* 5(2–3):105–130.
- [Kraus, Wilkenfeld, & Zlotkin, 1995] Kraus, S.; Wilkenfeld, J.; and Zlotkin, G. 1995. Multiagent negotiation under time constraints. *Artificial Intelligence* 75:297–345.

- [Kreps & Wilson, 1982] Kreps, E., and Wilson, R. 1982. Reputation and imperfect information. *Journal of Economic Theory* 27:253–279.
- [Osborne & Rubinstein, 1990] Osborne, M. J., and Rubinstein, A. 1990. *Bargaining and Markets*. Academic Press, Inc.
- [Riley, 1980] Riley, J. 1980. Strong evolutionary equilibrium and the war of attrition. *Journal of Theoretical Biology* 82:383–400.
- [Rosenschein & Zlotkin, 1994] Rosenschein, J. S., and Zlotkin, G. 1994. *Rules of Encounter: Designing Conventions for Automated Negotiation among Computers*. MIT Press.
- [Roth, Murnighan, & Schoumaker, 1988] Roth, A. E.; Murnighan, J. K.; and Schoumaker, F. 1988. The deadline effect in bargaining: some experimental evidence. *American Economic Review* 78:806–823.
- [Rubinstein, 1982] Rubinstein, A. 1982. Perfect equilibrium in a bargaining model. *Econometrica* 50:97–109.
- [Rubinstein, 1985] Rubinstein, A. 1985. A bargaining model with incomplete information about time preferences. *Econometrica* 53:1151–1172.
- [Sandholm & Lesser, 1995] Sandholm, T. W., and Lesser, V. R. 1995. Issues in automated negotiation and electronic commerce: Extending the contract net framework. In *Proceedings of the First International Conference on Multi-Agent Systems (ICMAS)*, 328–335. Reprinted in *Readings in Agents*, Huhns and Singh, eds., pp. 66–73, 1997.
- [Sandholm & Lesser, 1996] Sandholm, T. W., and Lesser, V. R. 1996. Advantages of a leveled commitment contracting protocol. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, 126–133.
- [Sandholm & Lesser, 1997] Sandholm, T. W., and Lesser, V. R. 1997. Coalitions among computationally bounded agents. *Artificial Intelligence* 94(1):99–137. Special issue on Economic Principles of Multiagent Systems. Early version appeared at the International Joint Conference on Artificial Intelligence, pages 662–669, 1995.

- [Sandholm & Zhou, 1999] Sandholm, T. W., and Zhou, Y. 1999. Revenue equivalence of leveled commitment contracts. In *AAAI-99 Workshop on Negotiation: Settling conflicts and identifying opportunities*. Extended version: Washington University, Department of Computer Science technical report WUCS-99-03.
- [Sandholm, Sikka, & Norden, 1999] Sandholm, T. W.; Sikka, S.; and Norden, S. 1999. Algorithms for optimizing leveled commitment contracts. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence (IJCAI)*. Extended version: Washington University, Department of Computer Science technical report WUCS-99-02.
- [Sandholm, 1993] Sandholm, T. W. 1993. An implementation of the contract net protocol based on marginal cost calculations. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, 256–262.
- [Sandholm, 1997] Sandholm, T. W. 1997. Unenforced E-commerce transactions. *IEEE Internet Computing* 1(6):47–54. Special issue on Electronic Commerce.
- [Sandholm, 1998] Sandholm, T. W. 1998. Contract types for satisficing task allocation: I theoretical results. In *AAAI Spring Symposium Series: Satisficing Models*, 68–75.
- [Sen, 1993] Sen, S. 1993. *Tradeoffs in Contract-Based Distributed Scheduling*. Ph.D. Dissertation, Univ. of Michigan.
- [Smith, 1974] Smith, J. M. 1974. The theory of games and evolution in animal conflict. *Journal of Theoretical Biology* 47:209–221.
- [Smith, 1997] Smith, L. 1997. Re-thinking bargaining theory. Technical report, MIT Department of Economics. Mimeo.
- [Vulkan, 1999] Vulkan, N. 1999. Economic implications of agent technology and e-commerce. *The Economic Journal* 109(453):F67–F90.