



DIGITAL ACCESS TO SCHOLARSHIP AT HARVARD

Intention Reconciliation by Collaborative Agents

The Harvard community has made this article openly available.
[Please share](#) how this access benefits you. Your story matters.

Citation	Sullivan, David G., Barbara. J. Grosz and Sarit Kraus. 2000. Intention Reconciliation by Collaborative Agents. In Proceedings of the Fourth International Conference on Multi-Agent Systems (ICMAS-2000), July 10-12, 2000, Boston, Mass., ed. ICMAS, 293-300. Los Alamitos, Calif: IEEE Computer Society.
Published Version	http://doi.ieeecomputersociety.org/10.1109/ICMAS.2000.858466
Accessed	January 16, 2018 4:42:18 AM EST
Citable Link	http://nrs.harvard.edu/urn-3:HUL.InstRepos:2641073
Terms of Use	This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA

(Article begins on next page)

Intention Reconciliation by Collaborative Agents*

David G. Sullivan, Barbara J. Grosz
Division of Engineering and Applied Sciences
Harvard University, USA

Sarit Kraus
Bar-Ilan University, Israel
and Institute for Advanced Computer Studies,
University of Maryland, USA

Abstract

Research on resource-bounded agents has established that rational agents need to be able to revise their commitments in the light of new opportunities. In the context of collaborative activities, rational agents must be able to reconcile their intentions to do team-related actions with other, conflicting intentions. The SPIRE experimental system allows the process of intention reconciliation in team contexts to be simulated and studied. Prior work with SPIRE examined the effect of team norms, environmental factors, and agent utility functions on individual and group outcomes for homogeneous groups of agents. This paper extends these results to situations involving heterogeneous groups in which agents use different utility functions. The paper provides new illustrations of the ways in which SPIRE can reveal unpredicted interactions among the variables involved, and it suggests preliminary principles for designers of collaborative agents.

This paper focuses specifically on the decision making that self-interested, collaborative agents must perform when their commitment to a group activity conflicts with opportunities to commit to different actions or plans. We present the results of experiments using the SPIRE simulation system, an experimental framework that allows us to simulate and study this type of intention reconciliation by collaborative agents. In earlier work [5, 20], we used SPIRE to examine the effect of team norms, environmental factors, and agent characteristics on the decisions and outcomes of members of homogeneous groups of agents. In this paper, we examine the outcomes of simulations involving larger, heterogeneous groups of agents who use differing utility functions to assess their options, and we suggest principles for designers of agents and agent environments based on the results.

1 Introduction

Computer systems increasingly are elements of complex, distributed environments in which human and computer agents interact. In applications that require agents to work collaboratively to satisfy a shared goal [3, 19, 21, inter alia], agents form teams to carry out actions, making commitments to their team's activity and to their individual actions in service of that activity. Agents in these environments may operate on behalf of individuals and organizations with different interests, and thus need to behave as rational, self-motivated individuals. These agents are also resource-bounded and must adapt their commitments and plans in the light of changing circumstances. However, collaborative agents also need to be able to count on each other. Therefore, team norms and other means may be needed to encourage agents to consider the good of the group when making decisions about their commitments [2, 5, 20]. As a result, decision-making in the context of collaborative activity presents a number of challenging problems to designers of collaborative multi-agent systems.

2 Intention reconciliation and collaboration

Research on collaboration in multi-agent systems [6, 7, 13, 15, 22] has established that commitment to the joint activity is a defining characteristic of collaboration. Although theories differ in the ways they encode this commitment, they agree on its centrality. At the same time, research on rationality and resource-bounded reasoning [4, 10, inter alia] has established the need for agents to dynamically adapt their plans to accommodate new opportunities and changes in the environment; sometimes, commitments need to be dropped. However, efforts in this area have mainly focused on plan management and evolution in the context of individual plans. Our work brings these two threads of research together. It addresses the need for collaborative agents to manage plans and intentions in multi-agent contexts, reasoning jointly about commitments to individual plans and commitments to group activities.

Our investigation focuses on the problem of intention reconciliation that arises because rational agents cannot adopt conflicting intentions [1, 6, inter alia]. If an agent has adopted an intention to do some action β and is given the opportunity to do another action γ that would in some way preclude its being able to do β , then the agent must decide

* This research was supported by National Science Foundation grants IIS-9978343, IRI-95-25915, CDA-94-01024, and IIS-9907482. Alyssa Glass and Michael Epstein helped to develop the SPIRE system.

between doing β and doing γ . It must *reconcile* intentions, deciding whether to maintain its intention to do β or to replace that intention with an intention to do γ .

In particular, this paper examines the problem of intention reconciliation in the context of collaborative activities, i.e., situations in which at least one of the conflicting intentions is related to an agent’s commitment to a team plan. Much of the prior work on agent collaboration and negotiation [14, 16] has assumed that commitments to collaborative activity are binding, and work in game theory has shown that it is possible to enforce this assumption by imposing harsh enough punishments when agents break a commitment. We are interested in situations in which some amount of defaulting is acceptable. For example, in the domain of automated systems administration (see [20]), it might be reasonable to allow an agent committed to performing a filesystem backup to default on that commitment so that it can assist with crash recovery on another system. To build collaboration-capable agents [8] that interact in environments where defaulting is possible, agent designers need to consider the problem of intention reconciliation in collaborative contexts.

Intention reconciliation in team contexts requires that agents weigh the purely individual costs and benefits of their decisions with team-related concerns. We assume that agents have relationships that persist over time, and thus the extent to which agents are trusted not to default may influence their long-term good. Defaulting on a team-related commitment for the sake of another opportunity may at times appear beneficial from a purely individualistic perspective, but an agent must consider the impact of such behavior on its ability to collaborate in the future and, more generally, on its future expected outcomes. In earlier work [20], we describe our assumptions about agents in greater detail, and we present scenarios that illustrate the issues involved in this type of decision-making.

In a given society of agents, the group-related impact of defaulting stems from what we term *social-commitment policies* [20], domain-independent policies that govern various aspects of collaboration, including both rewards and penalties for individual acts in the context of group activities. By stipulating ways in which current decisions affect both current and future utility, social-commitment policies change the way agents evaluate trade-offs. They provide a mechanism for constraining individuals so that the good of the team plays a role in their decision making. Section 3 describes one such social-commitment policy.

Social factors can also function in an additional way. If agents get part of their utility from the team, they have a stake in maximizing group utility. Therefore, when facing a choice, it may be useful for an agent to consider not only this single choice, but also the larger context of similar choices by itself and others. While being a “good guy” may

appear suboptimal by itself, everyone’s being a good guy when faced with similar choices may lead to better outcomes for everyone in the team. Our experimental framework uses the *brownie points model* developed by Glass and Grosz [5] to incorporate this type of consideration into decision making, allowing us to study its effect on outcomes, as well as its susceptibility to manipulation.

3 The SPIRE framework

The SPIRE simulation system allows us to study the impact of environmental factors, social-commitment policies, and agent utility functions on individual and group outcomes. The many variables involved and the often unexpected ways in which they interact make a system like SPIRE useful for testing hypotheses, uncovering relationships, and gaining insight into the issues involved in the intention-reconciliation problem. In the following paragraphs, we present an overview of the system. More details can be found in our earlier work [20].

SPIRE models situations in which a team of agents works together on group activities, each of which consists of doing a set of tasks. We currently assume that each task lasts one time unit and is a single-agent action, i.e., that it can be performed by an individual agent. Agents receive income for the tasks that they do; this income can be used to determine an agent’s current and future expected utility.

A SPIRE simulation run consists of a single group activity done repeatedly by the same team, because varying either the group activity or the team members would obscure sources of variation in the outcomes. However, the individual tasks within the activity will not necessarily be done by the same agents each time. To simplify the description, we assume that the group activity maps to a “weekly task schedule” (WTS). At the start of each week, a central scheduler takes the elements of this weekly task schedule and assigns them to agents.¹

After the scheduler has assigned all of the tasks, agents are chosen at random and given the chance to do one of a series of *outside offers*. Outside offers correspond to actions that an agent might choose to do apart from the group activity. Each outside offer conflicts with a task in the WTS; to accept an offer, an agent must default on one of its assigned tasks. The values of the outside offers are chosen randomly from a distribution that gives agents an incentive to default. If an agent chooses an outside offer, it defaults on its assigned task β . If there is another agent capable of doing β and available at the time for which it is scheduled, the task is given to that agent; otherwise, β goes undone.

1. This central scheduler is used only for convenience. In many domains, agents would likely need to negotiate each week’s schedule. Because negotiation is not our focus, we simplified this aspect of the problem.

The team as a whole incurs a cost whenever an agent defaults; this cost is divided equally among the team’s members. In addition, SPIRE currently applies a social-commitment policy in which a portion of each agent’s weekly tasks is assigned based on its behavior over the course of the simulation. Each agent has a *score*² that reflects the total number of times it has defaulted, with the impact of past weeks’ defaults diminishing over time. The higher an agent’s relative score, the more valuable the tasks it receives. Both the group costs and score reductions are larger when no replacement agent is available.

The scheduler assigns N tasks per agent on the basis of the agents’ scores; we refer to these tasks as *score-assigned tasks*. The scheduler gives each agent the most valuable tasks remaining in the pool of unscheduled tasks. The agent with the highest score receives the N highest-valued tasks that it can perform (given its capabilities and the times when it is available), the agent with the second-highest score receives the next N tasks, and so on. If there is more than one agent with the same score, the scheduler randomly orders the agents in question and cycles through them, giving them tasks one at a time. After all agents receive N tasks, the remaining tasks are assigned to agents picked at random. The strength of the social commitment policy can be varied by modifying the value of N .

4 Decision making in SPIRE

In deciding whether to default on a task β so as to accept an outside offer γ , an agent determines the utility of each option. SPIRE currently provides for up to three factors to be considered in utility calculations: current income (CI), future expected income (FEI), and brownie points (BP). Below we review each of them in turn.

4.1 Current and future expected income

Current income only considers the income from the task or outside offer in question, as well as the agent’s share of the group cost if it defaults. Future expected income represents the agent’s estimate of its income in future weeks, based on its position in the rankings when the agents are ordered according to their scores. The agent first approximates the impact that defaulting will have on one week of its income. Section 4.2 describes this estimate in more detail. The agent then extrapolates beyond that week to compute a more complete estimation, using a discount factor $\delta < 1$. Discounting reflects agents’ increasing uncertainty about their predictions [20]. For the experiments in this paper, we assume that agents are also uncer-

2. In prior work [5, 20], we refer to this score as the agent’s *rank*. Using *score* avoids confusion with the agent’s *ranking*, which is its position relative to the other agents when they are ordered according to their scores.

tain about the duration of their collaboration, and therefore use the infinite-horizon version of the FEI formula described by Glass and Grosz [5]. If F is the estimate of next week’s income and δ is the discount factor, then:

$$FEI(F) = \delta F + \delta^2 F + \delta^3 F + \dots = \left(\frac{\delta}{1-\delta}\right)F \quad (1)$$

We refer to the factor in parentheses as the *FEI weight*.

4.2 Estimating the loss in next week’s income

An agent estimates its potential loss in income during the following week by approximating its new position in the rankings both if it defaults and if it does not default, and determining the score-assigned tasks it would receive in each case. There are many factors that affect the agent’s actual position in the rankings, including the behavior of other agents and the offers that the agent receives later in the same week. To model situations in which agents have only limited information about each other, we assume that agents do not know the scores of other agents nor the total number of defaults in a given week, but only their own ranking in both the current and the previous week. Given the difficulty of estimating an agent’s ranking using such limited information, and the fact that it is unclear whether more sophisticated methods would lead to better estimations, we adopted the simple approach described below.

An agent begins its estimation by using its previous and current weeks’ rankings to approximate the number of agents who defaulted last week. For example, if an agent’s position in the rankings improved and it did not default last week, it assumes that some of the agents who were previously above it in the rankings must have defaulted. It carries this estimate over to the current week, assuming that the same number of agents will again default. Using this estimate, the agent creates four agent equivalence classes: (1) the agents currently above it who will not default, (2) the agents above it who will default, (3) the agents below it who will not default, and (4) the agents below it who will default. The agent adds itself to the equivalence classes using the following rules:³

- (a) To approximate what will happen if it does not default, it adds itself to the second class.
- (b) To approximate what will happen if it defaults when there is an agent available to replace it, it adds itself to a new class between the second and third classes.
- (c) To approximate what will happen if it defaults with no replacement, it adds itself to the third class.

It then calls the scheduler once to compute the value of its score-assigned tasks if it does not default ($F_{\text{no-def}}$,

3. These rules may underestimate the impact of defaulting, since agents can drop even further in the rankings when they default.

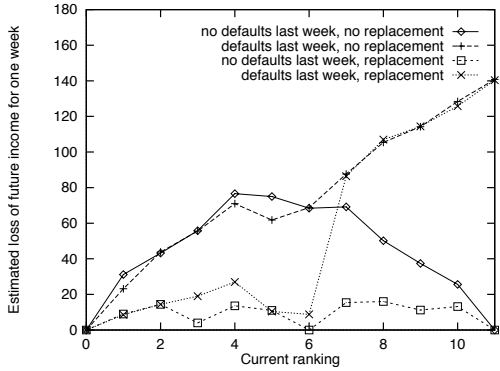


Figure 1. Possible estimates of FEI loss (using a discount factor δ of 0.8) for an agent that was ranked seventh out of twelve agents in the previous week.

obtained using the classes formed from rule a), and a second time to determine the value of its score-assigned tasks if it does default (F_{def} , using the classes from (b) or (c)).

The agent’s estimate of its one-week income loss from defaulting ($F_{\text{no-def}} - F_{\text{def}}$) thus depends on five factors: its previous and current rankings, whether it defaulted last week, whether there is an agent available to replace it, and the number of agents with which it is collaborating (since this affects the sizes of the equivalence classes). The estimated loss of income can vary greatly, as Figure 1 shows. Agents occasionally estimate an income loss of 0, which means that factors such as the strength of the social-commitment policy and the value of the FEI discount factor will not affect their decisions. We experimentally determined that the average actual income loss in these situations is 20, and thus we modified the decision-making code so that income-loss estimates of 0 are increased to 20. As a result, we obtained somewhat fewer defaults in this paper’s experiments than we did in our earlier work.

In the current system, an agent’s estimation does not consider the number of times that it has already defaulted in the current week. Although this simplification ignores the fact that agents should expect to drop more in the rankings the more they default, it saves considerable computation by allowing agents to reuse their estimations, avoiding repeated, expensive calls to the scheduler. We plan to explore estimation methods that consider prior defaults without requiring repeated scheduler invocations. However, given the inherent difficulty of making these estimates and the fact that our approach has been shown to respond reasonably to changes in the social-commitment policy [20], the estimates seem adequate for our current purposes.

4.3 Non-monetary factor: brownie points

In addition to being concerned about its income, an agent may also derive utility from being a “good guy.” Glass and Grosz’s *brownie point* model captures this

aspect of agents’ utilities [5], providing a measure of an agent’s sense of its reputation as a responsible collaborator. Agents begin a simulation with an identical, non-zero number of brownie points. When they default, agents lose brownie points. In addition, agents gain brownie points when they choose not to default, reflecting the fact that they are doing what is good for the group. Because an agent’s reputation is affected not only by whether or not it defaults, but also by the context of the decision, the changes in brownie points take into account the values of the task and offer involved in the decision. If an agent defaults on a low-valued task, its brownie points are reduced less than if it defaults on a high-valued task; if it defaults for the sake of a high-valued offer, its brownie points are affected less than if it defaults for a low-valued offer. Similarly, the increases in brownie points when agents choose not to default are greater for low-valued tasks and for high-valued offers.

Note that brownie points represent an agent’s *own* evaluation of its reputation as a responsible collaborator, not the perception of other agents. This factor is not a social-commitment policy: it does not directly affect the value of the tasks that an agent receives in the current collaboration. Rather, brownie points allow agents to incorporate a measure of *social consciousness* in their decisions. In informal terms, socially conscious agents may make decisions that are locally, individually suboptimal, because doing so enables the group as a whole—and perhaps, indirectly, the agent itself—to be better off. While it might be possible to express this element of an agent’s utility in monetary terms, using the non-monetary measure described above is simpler and more intuitive.

4.4 Combining the factors

To compare the overall utility of an agent’s options, the CI and FEI values for each option are combined to give the total estimated income (TEI). Next, the TEI and brownie point (BP) values are normalized: the default and no-default TEI values, TEI_{def} and $TEI_{\text{no-def}}$ respectively, are each divided by $\max(TEI_{\text{def}}, TEI_{\text{no-def}})$, and the default and no-default BP values (BP_{def} and $BP_{\text{no-def}}$) are similarly adjusted. Finally, the normalized values are weighted based on the agent’s social consciousness:

$$U_{\text{def}} = TEIweight \times normTEI_{\text{def}} + BPweight \times normBP_{\text{def}}$$

$$U_{\text{no-def}} = TEIweight \times normTEI_{\text{no-def}} + BPweight \times normBP_{\text{no-def}}$$

where $TEIweight$ and $BPweight$ sum to 1.

Agents default when $U_{\text{def}} > U_{\text{no-def}}$. Agents who do not use brownie points (corresponding to a $BPweight$ of 0) may compare their unnormalized, unweighted TEI values.

52 weeks per simulation run
20 task types (values=5, 10, ..., 100)
40 time slots per week
5n/6 tasks per time slot (n = # of agents),
of randomly chosen types

10 score-assigned tasks per agent
per week, the rest assigned randomly
5t/8-7t/8 offers per week (t = # tasks):
• number & values chosen randomly
• possible values = task values + 95

Figure 2. SPIRE settings used for most of the experiments in this paper.

5 Experimental results

In the following experiments, we make the simplifying assumptions that all agents are capable of doing all tasks and that all agents are initially available at all times. To maximize the contrast between different settings, we configured most of the experiments to make a large number of outside offers (see Sect. 5.3 for a discussion of how this affects the results). Figure 2 summarizes the settings used for most of the experiments; departures from these values are noted in each experiment’s description. Several of the settings, including the number of task types and the number of time slots, were chosen to model the work week of a systems administration team. Other settings were chosen based on prior experimentation. For example, the number of tasks per time slot was chosen so that there would be a significant number of outside offers for which no replacement agent is available.

The results presented are averages of 30 runs that used the same parameter settings but had different, randomly-chosen starting configurations (the values of tasks in the weekly task schedule, and the number and possible values of the outside offers). Error bars on the graphs indicate the end points of 95% confidence intervals. In each run, the first ten weeks serve to put the system into a state in which agents have different scores; these weeks are not included in the statistics SPIRE gathers.

5.1 Heterogeneity in the weight given to FEI

In this set of experiments, we considered heterogeneous groups of agents who use different δ values to weight their estimates of income in the following week (see eq. 1). In prior work [20], we showed that increasing δ in homogeneous groups of agents leads to fewer defaults and to increased individual and group incomes. In the current experiments, some agents use a δ value of 0.8, others a δ of 0.95. All of them also use brownie points, with a *BPweight* of 0.4 (a setting that falls in the optimal range determined by Glass and Grosz [5]). We used a total of 60 agents, and we varied the percentage of agents in each subgroup, considering cases in which none, 1/12, 4/12, 6/12, 8/12, 11/12, and all of the agents use $\delta = 0.95$, and the rest use $\delta = 0.80$. Figure 3 displays the results.

In general, agents who put a higher value on future expected income default less than agents who do not value

future income as much. For agents who value FEI highly, expected gains from outside offers are often outweighed by estimates of future losses from defaulting. In fact, if such agents give enough weight to FEI, they may never default, and this phenomenon can be accentuated by inaccurate estimates on the part of the agents.

Because the agents who use $\delta = 0.80$ (for an FEI weight of 4) default more often than agents who use $\delta = 0.95$ (for an FEI weight of 19), they tend to occupy the bottom rankings. Therefore, they receive lower valued score-assigned tasks, and outside offers become more attractive. In addition, brownie-point computations have less influence on agents who are at the bottom of the rankings. The BP loss from defaulting on a task is quadratic in the task’s value [5], so agents who receive lower-valued tasks compute smaller BP losses. The 0.8-delta agents thus enter a cycle of more and more defaulting as a simulation progresses: defaulting moves them down in the rankings, and having a lower ranking makes them more likely to default. As the number of 0.80-delta agents decreases, a larger percentage of them are found at the bottom of the rankings, and thus both of the effects mentioned above are increased. Therefore, the 0.8-delta agents default more frequently as they become a smaller percentage of the group.

The impact of the changing subgroup sizes on both mean individual income (from both tasks and outside offers) and subgroup income (from tasks only) is shown in the right half of Figure 3. The normalization factor used for these results differs from that used in our earlier work. The prior work divided an agent’s income by the income that the agent would have earned if all of its originally assigned tasks had been done (its *base income*). This normalization factor works well for homogeneous groups of agents, but it is problematic when dealing with heterogeneous groups. In particular, in heterogeneous groups, agents (and subgroups) that are assigned predominately low-valued tasks will have low base incomes and thus artificially high normalized incomes. Hence, for the normalizations reported in this paper, we compute the total of the tasks assigned to the top-ranked agent in each week. This total (which no single agent may have as its actual base income) is used to normalize the income of every agent.

The results show that agents who put less weight on their future income do better as individuals than agents who weight it more heavily. The extra income from the outside offers that they accept more than compensates for




Figure 3. Outcomes for agent subgroups as the percentage of agents placing a higher value on future income is increased. The percentage of defaults (*left*) and the normalized individual and group incomes (*right*) are shown for each subgroup. Section 5.1 details the income normalization.

their loss of future income and their share of the group costs. This result holds even when the homogeneous-group results (the 0% and 100% cases) are compared. While our earlier results showed that increasing δ from 0.4 to 0.8 leads to increasing individual incomes in homogeneous groups, these new experiments indicate that increasing δ to values close to 1 leads to less individual income, even when agents are homogeneous. This result is similar to one from Glass and Grosz’s earlier work on brownie points [5], which showed that a moderate level of social consciousness (i.e., a moderate *BPweight*) is better from the individual perspective, since otherwise agents turn away even the most highly-valued outside offers.

The 0.8-delta agents also see an increase in individual income as the percentage of 0.95-delta agents increases. The 0.8-delta agents benefit from the fact that their more responsible colleagues keep the group costs low, and as the number of 0.95-delta agents increases, this effect only increases. The resulting disparity in individual income is an example of the “free-rider” problem, in which agents can afford to be less responsible in the presence of more responsible collaborators. Designers of agents and agent environments need to take the possibility of this type of free-rider behavior into account.

Although the 0.8-delta agents do better as individuals, their income from group-activity tasks alone is much lower than that of the 0.95-delta agents because they have defaulted on more tasks. Moreover, the overall group income, which is effectively a weighted average of the two subgroup incomes, increases as the number of 0.95-delta agents increases (Fig. 4), showing that group outcomes are better when more agents weight future income heavily.

5.2 Heterogeneity in social consciousness

We next considered heterogeneous groups of agents in which some of the agents use brownie points and some do

not. We again used an overall group size of 60 agents, and we varied the percentage of the agents in each subgroup, considering cases in which none, 1/12, 4/12, 6/12, 8/12, 11/12, and all of the agents use brownie points, and the rest do not. All of the agents use an FEI discount factor, δ , of 0.8. The results are shown in Figure 5.

As expected, the agents who use brownie points default less often than those who do not (Fig. 5, *left*). The difference between the two subgroups remains more or less the same as the number of BP agents increases. This approximately constant difference in defaulting contrasts with the mixed-delta experiments in the previous section, in which the 0.80-delta agents defaulted more frequently as they became a smaller percentage of the overall group. In the mixed-delta experiments, 0.95-delta agents tend always to be ranked above 0.80-delta agents, since they default so infrequently. In the mixed-BP experiments, on the other hand, the two subgroups are less segregated in the rankings. The no-BP agents can thus avoid the cycle of increasing defaults that affects the 0.80-delta agents as they become confined to the lower rankings (cf. Sect. 5.1).

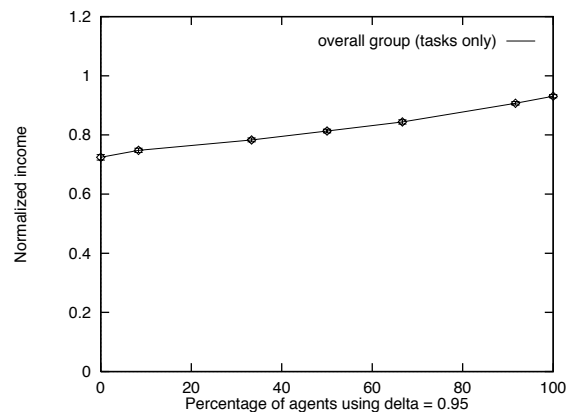


Figure 4. Task-only income for an entire group of heterogeneous agents as the percentage of agents placing a higher value on future income is increased. Section 5.1 describes the income normalization.




Figure 5. Outcomes for agent subgroups as the percentage of socially conscious agents (see Sect. 4.3) is increased. The percentage of defaults (*left*) and the normalized individual and group incomes (*right*) are shown for each subgroup. Section 5.1 describes the income normalization.

The two subgroups exhibit small but statistically significant differences in normalized individual and group income (Fig. 5, *right*), where the normalization factor described in Section 5.1 was used. The no-BP agents again do better as individuals but worse as a group, and the group differences are more pronounced. This suggests that the free-rider effect is less of an issue when the difference in defaulting between the two types of agents is less pronounced, and that agent designers can improve group outcomes without sacrificing individual gains by building agents with a moderate amount of social consciousness.

5.3 Effect of the number of outside offers

In both sets of heterogeneous-agent experiments, agents who default more do better as individuals. One reason for this involves the large number of outside offers that were made in these experiments, between 5/8 and 7/8 of the number of group-activity tasks. Agents who default more frequently are thus able to gain a significant amount of extra income from the outside offers that they accept. When fewer outside offers are made, defaulting more often can actually lead to a smaller individual income. For example, when the number of outside offers is fixed at 10% of the number of tasks, the normalized mean individual income of the no-BP agents is only 0.935 when 11/12 of the agents use brownie points, while that of the BP agents is 0.950. We plan to investigate the effect of the number of outside offers more fully in future work.

6 Related Work

Kalenka and Jennings [12] propose several “socially responsible” decision-making principles and examine their effects in the context of a warehouse loading scenario. Our work differs from theirs in two main ways: (1) their policies are domain-dependent and not decision-theoretic; (2) they consider agents choosing whether to help each other,

not agents defaulting on their team commitments. Sen [18] also considers decision-making strategies that encourage cooperation among self-interested agents, but his work focuses on interactions between pairs of individuals, rather than those between an individual and a team.

Sandholm et al. [17] study a mechanism built into contracts between self-motivated agents that allows for decommitment through the payment of a penalty. They present an algorithm for constructing optimal contracts and for determining what an agent should do when it receives an outside offer, given the contract. They analyze single interactions between two or three agents. We study the problem of intention reconciliation in an ongoing collaboration of a large set of agents, in which no explicit arrangement for defaulting can be made. Furthermore, we use group norms to discourage agents from defaulting.

Xuan and Lesser [23] present a framework for agents to negotiate their commitments. Like Sen, they focus on interactions between pairs of individual agents and assume that the details of a commitment can be modified in a way that is satisfactory for both agents.

There is also a significant body of economics literature on rational choice and intention reconciliation [9, 11, inter alia] that space limitations preclude our reviewing here.

7 Conclusions

The SPIRE empirical framework enables us to simulate and study the process of intention reconciliation in collaborative contexts, examining the impact of environmental factors, team norms and agent utility functions on individual and group outcomes. The results of our experiments demonstrate that, in heterogeneous settings, less responsible agents do better as individuals when there are a large number of outside offers. However, their share of the group income from tasks alone is much lower than that of more responsible agents. In addition, the overall group income is better when more of the agents are responsible. Therefore,

agent designers would do well to ensure that agents have nearly equal levels of responsibility to the group, both to obtain better group outcomes and to prevent less responsible agents from taking advantage of more responsible ones. It is unclear what the ideal level of responsibility to the group should be: as our results involving 0.95-delta agents show, having a utility function that effectively prevents you from defaulting leads to extremely low individual outcomes. Agent designers could potentially use SPIRE to determine utility-function settings that lead to acceptable outcomes in a given domain.

There are many possible mechanisms for achieving greater uniformity in agent responsibility, including increasing the strength of the social-commitment policy. We plan to investigate these mechanisms in future work. One challenge is to achieve uniformity without completely sacrificing the ability of agents to default or the autonomy of agent designers.

Other classes of problems that we hope to investigate using SPIRE include: (1) the influence of information about other team members on agents' behavior; (2) heterogeneous communities where agents have different capabilities and availabilities; (3) alternative social-commitment policies; (4) the effect of decreasing the number of outside offers on the outcomes of less responsible agents in heterogeneous communities; and (5) the possibility of agents modeling and adapting to the behavior of other agents.

Because intention reconciliation in realistic multi-agent contexts is an extremely complex problem, we believe a system like SPIRE is essential for obtaining the insights needed to design collaboration-capable agents [8]. Such agents will function not merely as tools but as problem-solving partners, working as members of heterogeneous teams of people and computer-based agents in our increasingly interconnected computing environments.

References

- [1] Bratman, M.E. 1987. *Intention, Plans, and Practical Reason*. Harvard University Press, Cambridge, MA.
- [2] Castelfranchi, C. 1998. Modeling social action for AI agents. *Artificial Intelligence*, 103:157-182.
- [3] Decker, K. and Li, J. 1998. Coordinated hospital patient scheduling. In *Proc. of ICMAS-98*, pp. 104-111.
- [4] Doyle, J. 1991. Rational belief revision. In: *Proc. of the Second Intl. Conference on Knowledge Representation and Reasoning*, pp. 163-174.
- [5] Glass, A. and Grosz, B.J. 1999. Socially conscious decision-making. In *Proc. of the Fourth Intl. Conference on Autonomous Agents* (to appear).
- [6] Grosz, B.J. and Kraus, S. 1996. Collaborative plans for complex group action. *Artificial Intelligence*, 86(2):269-357.
- [7] Grosz, B.J. and Kraus, S. 1999. The evolution of Shared-Plans. In Wooldridge, M. and Rao, A., editors, *Foundations and Theories of Rational Agency*. Kluwer Academic Publishers, The Netherlands, pp. 227-262.
- [8] Grosz, B.J., Hunsberger, L., and Kraus, S. 1999. Planning and acting together. *AI Magazine*, 20(4):23-34.
- [9] Holländer, H. 1990. A social exchange approach to voluntary cooperation. *American Economic Review*, 80(5):1157-1167.
- [10] Horty, J. and Pollack, M.E. 1998. Option evaluation in context. In *Proc. of the 7th Conference on Theoretical Aspects of Rationality and Knowledge*, pp. 249-262.
- [11] Iannaccone, L.R. 1992. Sacrifice and stigma: reducing free-riding in cults, communes, and other collectives. *Journal of Political Economy*, 100(2):271-291.
- [12] Kalenka, S. and Jennings, N.R. 1999. Socially responsible decision making by autonomous agents. In *Proc. of the Fifth Intl. Colloquium on Cognitive Science*, pp. 153-169.
- [13] Kinny, D., Ljungberg, M., Rao, A.S., Sonenberg, E., Tidhar, G., and Werner, E. 1994. Planned team activity. In Castelfranchi, C. and Werner, E., eds., *Artificial Social Systems (LNAI-830)*, pp. 227-256. Springer Verlag.
- [14] Kraus, S., Wilkenfeld, J., and Zlotkin, G. 1995. Multiagent negotiation under time constraints. *Artificial Intelligence*, 75(2):297-345.
- [15] Levesque, H., Cohen, P. and Nunes, J. 1990. On acting together. In *Proc. of AAAI-90*, pp. 94-99.
- [16] Rosenschein, J.S. and Zlotkin, G. 1994. *Rules of Encounter: Designing Conventions for Automated Negotiation among Computers*. MIT Press, Cambridge, MA.
- [17] Sandholm, T., Sikka, S., and Norden, S. 1999. Algorithms for optimizing leveled commitment contracts. In *Proc. of IJCAI-99*, pp. 535-540.
- [18] Sen, S. 1996. Reciprocity: a foundational principle for promoting cooperative behavior among self-interested agents. In: *Proc. of ICMAS-96*, pp. 322-329.
- [19] Sen, S., Haynes, T., and Arora, N. 1997. Satisfying user preferences while negotiating meetings. *Intl. Journal on Human-Computer Studies*, 47(3):407-427.
- [20] Sullivan, D., Glass, A., Grosz, B., and Kraus, S. 1999. Intention reconciliation in the context of teamwork: an initial empirical investigation. In Klusch, M., Shehory, O., Weiss, G., eds., *Cooperative Information Agents III (LNAI-1652)*, pp. 149-162. Springer Verlag.
- [21] Sycara, J. and Zeng, D. 1996. Coordination of multiple intelligent software agents. *Intl. Journal of Intelligent and Cooperative Information Systems*, 5:181-211.
- [22] Tambe, M. 1997. Towards flexible teamwork. *Journal of Artificial Intelligence Research*, 7: 83-124.
- [23] Xuan, P. and Lesser, V. 1999. Incorporating uncertainty in agent commitments. In *Proc. of ATAL-99*, pp. 221-234.