Autonomous Agents and Multi-Agent Systems, 6, 317–339, 2003

Socially Conscious Decision-Making

ALYSSA GLASS glass@parc.com
Palo Alto Research Center (PARC), Palo Alto, CA 94304, USA

BARBARA J. GROSZ grosz@eecs.harvard.edu
Division of Engineering and Applied Sciences, Harvard University, Cambridge, MA 02138, USA

Abstract. For individually motivated agents to work collaboratively to satisfy shared goals, they must make decisions about actions and intentions that take into account their commitments to group activities. This paper examines the role of social consciousness in the process of reconciling intentions to do group-related actions with other, conflicting intentions. We operationalize the notion of social consciousness and provide a first attempt to formally add social consciousness to a cooperative decision-making model. We define a measure of social consciousness; describe its incorporation into the SPIRE experimental system, a simulation environment that allows the process of intention reconciliation in team contexts to be studied; and present results of several experiments that investigate the interaction in decision-making of measures of group and individual good. In particular, we investigate the effect of varying levels of social consciousness on the utility of the group and the individuals it comprises. A key finding is that an intermediate level of social consciousness yields better results in certain circumstances than an extreme commitment. We suggest preliminary principles for designers of collaborative agents based on the results.

Keywords: multi-agent systems, simulation, decision-making, social consciousness, collaboration, intention reconciliation

1. Introduction

In many situations, computer agents must interact with other systems and with people to accomplish tasks, and for many applications, it is necessary to form teams that comprise people and computer agents that work collaboratively to satisfy a shared goal. As rational agents, individual team members must be able to make individually rational decisions about their commitments and plans [16]. They must also, however, be responsible to the teams in which they participate. In such settings, agent designers need to construct computer agents that include a sense of group commitment in their reasoning about actions and plans.

Various possibilities arise for externally influencing agents to act in the group's interest, including the imposition of sanctions on agents that default [19]. In this paper, we investigate a different possible influence on decision-making, one that is internally motivated and considers the consequences of multiple decisions within a context in which team members may join forces repeatedly to work on group activities. We examine the design of agents that incorporate a notion of *social consciousness* into the ways they measure benefit and maximize individual outcome. In informal terms, these socially conscious agents may make decisions that are locally, individually suboptimal, because doing so engenders a society that is globally better off, and thus optimal for them in the long term. Castelfranchi has argued for such social consciousness in agents [3]. After presenting the model of social consciousness, we discuss the ways this notion differs from the notions of benevolence in prior multi-agent work.

Our research aims to specify ways agent designers can create optimally socially conscious, individually rational agents by examining the trade-off between individual, group, and total payoffs. We first describe the SPIRE (SharedPlans Intention Reconcilation Experiments) simulation system which enables investigation of the effectiveness of different decision-making strategies under various environmental conditions [23]. We then provide a method for assigning a measure to social consciousness in collaborative agents and define a model in which nonmonetary social factors play a role in how agents make decisions about their intentions. We show how this measure of social consideration can be made compatible with monetary considerations to create agents whose level of social commitment can be altered and studied using SPIRE.

The paper presents the results of several experiments in which the level of social consciousness is varied, and the effect on different measures of group income is determined. The experiments also consider environments with different densities of tasks to be accomplished by the group. Contrary to expectations, we find under certain conditions a maximal, intermediate level of social consciousness. These experiments have provided an initial foundation for future, more complex experiments studying social consciousness in group decision-making [6]. We conclude by providing suggestions for creating agents that maximize group income subject to the constraint of remaining individually rational.

2. Intention reconciliation

2.1. The problem

The experiments we describe in this paper extend previous work [8, 9, 17, 22, 25] by addressing 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 which arises because rational agents cannot adopt conflicting intentions [2, 8]. If an agent has adopted an intention to do some action β, and is presented the opportunity to do another action γ, that would in some way preclude its being able to do β, then the agent must decide between doing β and doing γ: it must *reconcile* intentions.

We will use an example from one of our application domains, Systems Administration [10, 23] to illustrate the problem of intention reconciliation in the context of group activities and to motivate our experiments. In this domain, teams of people and computer systems work together to maintain a cluster of workstations. Their overall group activity (which we will refer to as α) comprises many subtasks (β_i) including such actions as

— reading and replying to technical questions,
— upgrading hardware,

— restoring files from backups,
— checking system security,
— maintaining printers (adding paper, toner).

The need for intention reconciliation would arise, for instance, if a human agent who intends to perform an operating system upgrade as part of his commitment to α, is subsequently offered an opportunity to attend a lecture by a Nobel Prize winner at the same time. Similarly, in the case of a computer agent, the agent could be asked to do a web search for an individual web user. Such a task is not part of the job of a Systems Administration agent, but the agent would receive an individual payment for completing the task. The agent must decide whether to remain committed to the group, or to renege on its original intention in favor of attending the lecture.

Because the agent is a member of a team, its income, and therefore its utility, depends not just on the completion of its own subtasks, but also on the completion of the full task, and thus, the successful completion of the subtasks of other team members. Similarly, the utility of the other team members depends on the actions of this agent. To reconcile intentions in a group context, an agent must have a method of reasoning about future utility and the effect of its decisions on the group, as well as about current utility. The agent must consider how other group members will view its failure to honor a commitment, and must reason about future utility, considering the costs it may incur as a result of the group's reaction to its defaulting on a group-related task. For instance, if the agent reneges on its task assignment, it may receive less valuable task assignments in the future, decreasing its overall utility. The focus of this paper is on another component of the agent's calculations: in addition to monetary utility from income earned, the agent may receive utility simply from being a "good guy" and honoring its commitment to the group. Depending on the level of social consciousness that the agent has, it may assign different weights to this "good guy" factor in its utility calculation. As will be explained in Section 3.2, this "good guy" utility is not just a "feel good" factor, but may ultimately, though indirectly, affect an agent's monetary utility.

2.2. Social-commitment policies

In interacting with one another, and particularly in working together, agents adopt, either explicitly or implicitly, what we term *social-commitment policies* (SCPs) [23]. These policies govern various aspects of group behavior, including both rewards and penalties for individual acts in the context of group activities. For instance, they may specify such things as what defines a fair distribution of tasks among agents, the distribution of benefits from a group activity, and the penalties imposed on an agent who defaults on its commitment to a group activity. While we could assume that these policies are agreed on by a team when it forms, it seems more natural and efficient to require that the community of agents embody these principles because, in computational settings, we expect agent designers will build many agents, different ones of which will, at different times, come together to form teams.

Social-commitment policies are important in light of the assumption that agents, even as team members, continue to be self-motivated; although they may consider team utility, they are not necessarily team-utility maximizers. Therefore, there may be tension between what is best for an individual in isolation, and what is best for the team. By stipulating ways in which current decisions will affect both current and future utility, social-commitment policies can influence the way agents evaluate trade-offs between individual and group-oriented tasks. Thus, they provide a mechanism for constraining individuals, so that group good plays a role in their decision-making.

Social-commitment policies can also function in an additional way. Even though agents are self-motivated, they may have a stake in maximizing the group utility, because they may get a portion of their utility from the group. In such cases, a larger group benefit means a larger share for each agent, and thus a larger individual 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 taking into account the impact on the group may appear suboptimal by itself, everyone's doing so when faced with similar choices may lead to optimal outcomes for everyone in the group, a situation referred to in game theory literature as a Prisoner's Dilemma [1]. In such situations, if all team members opt to do what is best for the group, even though it is individually suboptimal, then the team as a whole will benefit, and each individual ultimately gains.

For example, in the Systems Administration domain, an individual may benefit from choosing to abandon the group and attend the lecture, but if everyone in the group made a similar choice, the group utility would suffer severely. Such effects could occur either within a single interaction (for instance, if the whole group defaults to attend the same lecture) or, more typically, over the long-term (different members of the group default at different times in favor of such "outside" opportunities). In either case, ultimately, the individual's utility would be lower, because the group utility from which it derives part of its individual utility is lower. Such social dilemmas have been studied among self-interested agents in computational environments over single interactions [11]. The problem changes significantly in our domain, where there are multiple decision points over time. We analyze this problem using methods from decision-theory so that agents can computationally reason about larger number of agents without needing to unnecessarily simplify complex decisions.

2.3. The SPIRE framework

The SPIRE system [23] enables manipulation of various agent properties and environmental conditions and the examination of the effect of different decision-making strategies under those conditions. In SPIRE, a team of agents ($G_1, \dots, G_n$) work together on group activities, called *GroupTasks*, each of which consists of doing a set of tasks (task instances). Each task instance is of one of the types $\beta_1, \dots, \beta_k$ and occurs at one of the times $T_1, \dots, T_m$. For example, a GroupTask in the Systems Administration domain might consist of a week's work (with the times T_i being the hours of the work week) doing various tasks β_i. Some task types may have only one instance in the week (e.g., printer maintenance); others may have multiple instances (e.g., running backups). Agents receive income for the tasks they do, and this income can be used in determining an agent's current and future expected utility.

A SPIRE simulation consists of a sequence of GroupTasks. To simplify the simulations and the analysis, the same GroupTask is done repeatedly by the same group, although the individual tasks within the GroupTask will not necessarily be done by the same agent each time. SPIRE considers a given GroupTask to consist of a set of tasks with time constraints on the tasks and capability requirements for agents doing the tasks. To simplify the description that follows, we assume that a GroupTask maps to a weekly task schedule. A simulation then consists of activity over a sequence of weeks.

A *weekly task schedule* (WTS) is a set of pairs $\langle task_i, time_i\rangle$ where $task_i$ is to be done at $time_i$, and a *weekly task schedule assignment* (WTSA) is a set of triples $\langle task_i, time_i, agent_i\rangle$ where $task_i$ is to be done at $time_i$ by $agent_i$. Each agent has a set of task capabilities and a set of available times that constrain the assignment of tasks in the WTS to produce a WTSA. An agent can only be assigned tasks for which it has the needed capabilities and the time availability.

To model the need to reconcile intentions, a sequence of *outside offers* is generated. These offers correspond to actions that an agent might choose to do apart from the GroupTask. Each outside offer γ conflicts with some task β in the WTSA; to accept an outside offer, an agent must default on one of its assigned tasks. The central question we investigate is the ways in which different levels of social consciousness influence the rates at which agents default, and their individual and collective incomes, given a particular group time horizon and configuration of environmental factors.

Each week, agents, chosen randomly, are offered opportunities to do γs that conflicts with tasks (β) in the context of doing the overall group activity α that they have been assigned in the WTSA. The income value of γ is also chosen randomly from a distribution with approximately the same shape as the distribution of task values in the WTS. Using the same distribution for the choices of tasks β, and the choices of outside offers γ, allows us to vary the distribution between simulation runs while still permitting us to maintain the comparative values of the two. To provide an incentive to default, the distribution of outside offers is shifted, so that it has a mean value that exceeds the mean value of the WTS tasks. If the agent chooses the new opportunity, it defaults on the task β with which γ conflicts. If there is an agent that is available and capable of doing β, the task is given to that agent; otherwise, β cannot be completed by the group, and therefore goes undone.

The group as a whole incurs a cost whenever an agent defaults. In our simulations, this cost is divided equally among the group's members. The cost of a particular default depends on its impact on the group. At a minimum, it equals a baseline value that represents the cost of finding a replacement agent. If no replacement is available, and so the task will not be done, the cost is increased by an amount proportional to the value of the task.

In the SPIRE system, the group's reaction to an agent defaulting on its commitment is represented by a social-commitment policy that constrains assignment of tasks to agents, as defined in the previous section. Each week, agents are assigned a portion of their tasks based on how responsible they have been thus far in the simulation. Each agent has a rank 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 rank, the more valuable the tasks it receives. Because there is a greater impact on the group when tasks go undone, an agent's rank is reduced by a larger amount if it defaults when no one can replace it.

To assign group tasks to agents, SPIRE makes use of an omniscient scheduler that has total information about all of the agents' ranks and history of defaults.[1] It is important to note, however, that while the central scheduler has complete information about all of the agents, the individual agent does *not* have access to this knowledge. Thus, it must estimate the behavior of other agents based on publicly-available information, as described in the next section.

3. Decision-making in SPIRE

3.1. *Estimating monetary utility*

In deciding whether to default on a task β and accept an outside offer γ, an agent weighs the impact of the choice on two monetary factors: current income (CI) and future expected income (FEI).

3.1.1. Current income. CI only considers the income from the task or outside offer in question, as well as the agent's share of the group cost should it default. The value of CI in the default (CI_{def}) and no-default(CI_{no_def}) cases is thus:

$$CI_{def}(\beta, \gamma) = value(\gamma) - \frac{group\_cost(\beta)}{n}$$
$$CI_{no\_def}(\beta) = value(\beta),$$

where n is the number of agents in the group, $value(x)$ is the payoff from completing task x, and $group\_cost(\beta)$ is the cost to the entire group from task β not being completed.

3.1.2. Estimating income one week ahead. For the Future Expected Income (FEI) calculation, an agent first needs to estimate its income in the following week, both if it defaults on β, and if it does not. There are many factors that affect an 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 ranks 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 future ranking using such limited information, and the fact that any attempt to derive more accurate estimates from past experience runs into game-theoretic complications, we adopted the simple approach described below. Subsequent work explores other Social Commitment Policies which are easier for agents to estimate, providing more accurate predictions of the effect of defaulting on future income [24].

An agent begins its estimation by using the difference between its previous and current weeks' rankings to approximate the number of agents who defaulted last week. 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 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 then adds itself to one of these equivalence classes. To approximate what will happen if the agent does not default, it adds itself to the second class. If it defaults with another agent replacing it, it creates a new class for itself between the second and third classes. And if it defaults with no replacement, it adds itself to the third class. Because we assume that agents know the policy used by the scheduler to assign tasks, we allow them to call the scheduling function. An agent calls this function once to compute the value of its rank-assigned tasks if it does not default, and a second time to compute the value of its rank-assigned tasks if it does default.

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 calls to the scheduler. Despite these assumptions and limitations, we have found the agent's estimations of F to be accurate enough for the *FEI* calculation described below.

3.1.3. Future expected income. Once the agent has its estimation of the following week's income, F, the agent then extrapolates beyond the following week to make a more complete estimation, discounting its estimates of subsequent weeks' income by an uncertainty factor $\delta < 1$. In this approach, if F is the original estimate of the following week's income, then the discounted estimate for the kth week after the current one is $\delta^k F$. The full *FEI* estimate in week i of a finite-horizon simulation is thus:

$$\begin{aligned} FEI_{finite}(F, i) &= \delta F + \delta^2 F + \cdots + \delta^{M-i} F \\ &= (\delta + \delta^2 + \cdots + \delta^{M-i}) F, \end{aligned}$$

where M is the number of weeks in the simulation.

Note that the expression in parentheses decreases as the simulation progresses, reflecting the fact that an agent has less to lose from defaulting when there are fewer weeks left in the GroupTask.

This *FEI* calculation, as noted in Sullivan et al. [23], assumes that an agent knows the number of weeks the group will continue to work together. In some situations, however, agents will only have indefinite information about how long they will form teams. These situations resemble group activity in an infinite time horizon, because without knowledge of a final week, agents are forced to treat every week as an intermediate week, equally far from the start as the finish [7]. To simulate this infinite horizon, the *FEI* calculation can be modified so that it is an infinite sum. If F is the original estimate of the following week's income, and δ is the uncertainty factor, *FEI* in an infinite time horizon is:

$$\begin{aligned} FEI_{infinite}(F) &= \delta F + \delta^2 F + \delta^3 F + \cdots \\ &= (\delta + \delta^2 + \delta^3 + \cdots)F \\ &= \left(\frac{\delta}{1-\delta}\right)F. \end{aligned}$$

The relationship between FEI_{finite} and $FEI_{infinite}$ is examined in Section 4.1.

3.1.4. Total estimated income. Once *CI* and *FEI* have been calculated, they are combined into a total estimated income (*TEI*) to provide a way of measuring total utility from these monetary measures. *TEI* in week i of the simulation in the default and no-default cases is:

$$\begin{aligned} TEI_{def}(\beta, \gamma, i) &= CI_{def}(\beta, \gamma) + FEI(F_{def}, i) \\ TEI_{no\_def}(\beta, i) &= CI_{no\_def}(\beta) + FEI(F_{no\_def}, i), \end{aligned}$$

where F_{def} and F_{no_def} are the agent's estimates of its income for the following week, if it does and does not default, respectively.

3.2. Utility from brownie points

In addition to direct monetary utility, agents' decision making processes may take into account a group commitment or social consciousness factor. The major challenges addressed in this paper are (1) how to compute a measure of social consciousness, and (2) how to combine this measure with the monetary factors discussed in the previous section. First, to compute social consciousness, we introduce a *brownie point* model, in which socially conscious "good guys"—that is, agents willing to sacrifice short-term personal gain for the good of the group and, ultimately, their own individual good—measure this social consciousness in addition to strictly monetary factors. In this model, both monetary factors and social nonmonetary utility are weighed during decision-making.

In the brownie point model, good guy agents maintain their own internal measure of social consciousness, giving themselves brownie points (*BP*) each time they choose not to default. In addition, an agent loses brownie points when it does default. This loss models an agent's disappointment in itself for failing to be the kind of agent it wants on its teams, *i.e.*, failing to live up to its commitment to the group. The number of brownie points an agent gains or loses at each decision point depends on the value of the task that it is considering defaulting on (β), and the value of the outside offer that it is considering (γ).

To represent an agent's concern with its historical reputation, its utility from brownie points is also dependent on the total number of brownie points that it has stored up over time. If an agent has remained faithfully committed to the group for a long time, storing up a lot of brownie points, the relative percentage loss in brownie points that it will experience from defaulting will be lower.

Brownie points represent an agent's own private 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 or its direct monetary rewards. Rather, brownie points allow agents to incorporate a measure of *social consciousness* into their decisions. Experimenting with brownie points allows us to investigate the conditions under which such an internal constraint on defaulting could be advantageous for agent design. While it might be possible to express this element of an agent's utility in monetary terms, using a nonmonetary measure is simpler and more intuitive.

Keeping these desirable properties in mind, we define one such function to compute the *BP* value in the default and no-default cases. Other, similar functions which also meet these properties would likely yield similar results, and are an area for future research.

For these experiments, we define *BP* in the default and no-default cases as follows:

$$BP_{def}(\beta, \gamma, currentBP) = currentBP - \frac{value(\beta)^2}{value(\gamma)},$$

$$BP_{no\_def}(\beta, \gamma, currentBP) = currentBP + \frac{value(\gamma)}{value(\beta)},$$

where *currentBP* is the total number of brownie points that the agent has accumulated thus far in the simulation, based on some initial amount allocated at the beginning of the simulation.

These definitions have several important properties. First, if an agent does not default, its *BP* value increases as $value(\beta)$ goes down or $value(\gamma)$ rises, reflecting an agent's greater pride in turning down highly valued outside offers, and in committing itself to less-valued tasks for the good of the group. Second, if an agent defaults on some task β in favor of γ, its *BP* value decreases as $value(\beta)$ goes up or $value(\gamma)$ goes down. This factor thus reflects an agent's greater willingness to forgive itself for defaulting in favor of a highly valued outside offer, or a less-valued β task, and vice versa. For example, an agent will punish itself more for neglecting to check for security breaks than it would for not putting toner in the printer. Alternatively, an agent will more readily forgive itself for defaulting on its group commitment in order to do a highly-valued activity, like attend to a sick family member, than it would for defaulting in order to attend a baseball game.

In addition, $BP_{def}(\beta, \gamma, currentBP)$ changes quadratically in $value(\beta)$, rather than linearly as in the no-default case. This change guarantees an adequate base deduction in BP even for small $value(\beta)$, since the average $value(\beta)$ is smaller than the average $value(\gamma)$. In addition, a quadratic change in $value(\beta)$, as opposed to any type of linear change, provides desirable percent changes in BP over time, punishing agents for defaulting even at the beginning of a simulation when their initial BP level is still quite high.

BP is not just a constant measure over time, but rather is a variable metric based on the situation faced by both the group and the specific agent. As a result, it is not just a measure of individual persistence. Furthermore, by taking into account both group benefit and individual gain in the long-term, BP differs from notions of benevolence that are concerned with helping other agents and increasing their utility [4, 14, inter alia].

3.3. Social consciousness in decision-making

In our studies, an agent's intention-reconciliation strategy takes into account the three factors just described: current income (CI); future expected income (FEI); and good guy stature in the community (BP). CI and FEI have been aggregated to form TEI. The previous section showed how our brownie point model addressed the first challenge, that of computing a measure of social consciousness, BP. Combining this social nonmonetary factor with TEI, which is entirely monetary, is the second challenge addressed in this paper.

To combine these factors into a single utility function, we use an approach from multi-attribute decision-making [26]. This approach is used so that factors with different units can be considered together within a single decision, without unit differences unexpectedly giving more weight to one than the other. The first step in this approach is to normalize both factors using linear normalization, which effectively removes unit differences in the two factors. To normalize TEI, which is in monetary units, and BP, which is not, each factor is divided by the maximum of the two possible values for that factor.

For instance, since $BP_{no\_def}(\beta, \gamma, currentBP) > BP_{def}(\beta, \gamma, currentBP)$ in all cases, normalized BP in each case is:

$$normBP_{no\_def}() = 1,$$

$$normBP_{def}(\beta, \gamma, currentBP) = \frac{BP_{def}(\beta, \gamma, currentBP)}{BP_{no\_def}(\beta, \gamma, currentBP)}.$$

An analogous calculation is used to normalize TEI, although no similar assumptions can be made concerning whether $TEI_{def}(\beta, \gamma, i)$ or $TEI_{no\_def}(\beta, i)$ is larger.

Once these factors have been thus normalized, the second step in this multi-attribute decision-making approach is to weight the factors relative to each other, allowing varying emphasis to be placed on each factor. The impact of these relative weights is empirically analyzed and discussed in Section 4.

Table 1. Summary of functions used in decision-making

No Default Computations:	$CI_{no\_def}(\beta) = value(\beta)$
	$FEI_{finite}(F, i) = (\delta + \delta^2 + \cdots + \delta^{M-i})F$
	$FEI_{infinite}(F) = \left(\frac{\delta}{1-\delta}\right)F$
	$TEI_{no\_def}(\beta, i) = CI_{no\_def}(\beta) + FEI(F_{no\_def}, i)$
	$BP_{no\_def}(\beta, \gamma, currentBP) = currentBP + \frac{value(\gamma)}{value(\beta)}$
	$U_{no\_def}(\beta, i) = (TEIweight \times normTEI_{no\_def}(\beta, i)) + (BPweight \times 1)$
Default Computations:	$CI_{def}(\beta, \gamma) = value(\gamma) - \frac{group\_cost(\beta)}{n}$
	$TEI_{def}(\beta, \gamma, i) = CI_{def}(\beta, \gamma) + FEI(F_{def}, i)$
	$BP_{def}(\beta, \gamma, currentBP) = currentBP - \frac{value(\beta)^2}{value(\gamma)}$
	$U_{def}(\beta, \gamma, i, currentBP) = (TEIweight \times normTEI_{def}(\beta, \gamma, i))$ $+ (BPweight \times normBP_{def}(\beta, \gamma, currentBP))$
Default Condition:	$U_{def}(\beta, \gamma, i, currentBP) > U_{no\_def}(\beta, i).$

The SPIRE system thus uses the following formulas for the utility an agent receives from defaulting, and from not defaulting, in week i of the simulation:

$$\begin{aligned} U_{def}(\beta, \gamma, i, currentBP) &= TEIweight \times normTEI_{def}(\beta, \gamma, i) \\ &\quad + BPweight \times normBP_{def}(\beta, \gamma, currentBP), \\ U_{no\_def}(\beta, i) &= TEIweight \times normTEI_{no\_def}(\beta, i) + BPweight \times 1, \end{aligned}$$

where *TEIweight* and *BPweight* can be adjusted to create agents with varying levels of social consciousness. Agents default when:

$$U_{def}(\beta, \gamma, i, currentBP) > U_{no\_def}(\beta, i).$$

For clarity, Table 1 contains a summary of the functions used by SPIRE during decision-making.

4. Socially conscious agents

In earlier work [23], we were able to show that agents default less often and increase their individual and group income when more tasks are assigned based on rank or more weight is given by the agents to future income. We also showed a complex relationship between the number of tasks scheduled concurrently (*task density*), and agent default behavior and income. In this section, we present several experiments that examine various aspects of socially conscious decision-making. Base values for important SPIRE parameters in the majority of experiments are given in Table 2; any departures from these values are noted in individual experiment descriptions. These values were chosen after speaking with people who work in Systems Administration, in an effort to faithfully model actual conditions in this domain. These experiments all have the simplifying assumption that the agents are

Table 2. Default SPIRE settings. Departures from these values are noted in experiment descriptions

52	weeks per simulation
12	agents
20	task types (values = 5, 10, 15, ..., 100)
40	time slots per week
10	tasks per time slot = 400 tasks per week
10	tasks per agent per week assigned based on the agent's rank; the rest assigned randomly
250–350 outside offers per week	
δ weighting factor for $FEI = 0.8$	
$TEIweight = 0.5$	
$BPweight = 0.5$	

homogeneous. Some early results in which we have begun to relax this assumption can be found in another paper [24]. To maximize the contrast between socially conscious and socially unconcerned agents, SPIRE parameters for these experiments were set to make a relatively large number of outside offers, and to impose relatively large rank deductions and group costs when agents default.

The results presented below are averages of 30 runs that used the same parameter settings, but had different, randomly-chosen, starting configurations (the values of the tasks in the WTS, and the possible values of the outside offers). For each of these runs, the WTS was formed by randomly selecting tasks from the stated task types. The values of the outside offers for that run were then chosen to have the same distribution as the associated random task types for that week. In each run, the first ten weeks serve to put the system into a state in which the agents have different ranks; these weeks are not included in the statistics that SPIRE gathers.

4.1. *Varying time horizons*

To provide a basis for further experiments, we investigated the effect of different time horizons on the average number of defaults over the course of a multi-week simulation (MWS). We compared two sets of runs: one with a known finite horizon of 52 weeks, and one with an infinite horizon. The two differed in their *FEI* calculations: the finite horizon runs used FEI_{finite} with $M = 52$; the infinite horizon runs used $FEI_{infinite}$.

A comparison of these two time horizons is shown in Figure 1. For the early weeks of an MWS, when there are many weeks left, FEI_{finite} and $FEI_{infinite}$ are very close in value. The values diverge as the finite horizon approaches, because FEI_{finite} becomes much smaller than $FEI_{infinite}$. This change reflects the fact that agents have less to lose by defaulting in later weeks, when there is little time left during which they can be punished for their defaulting on their commitment to the group. As a result, in our experiments, during the first several weeks of the MWS, the two default rates coincide, but later they diverge.

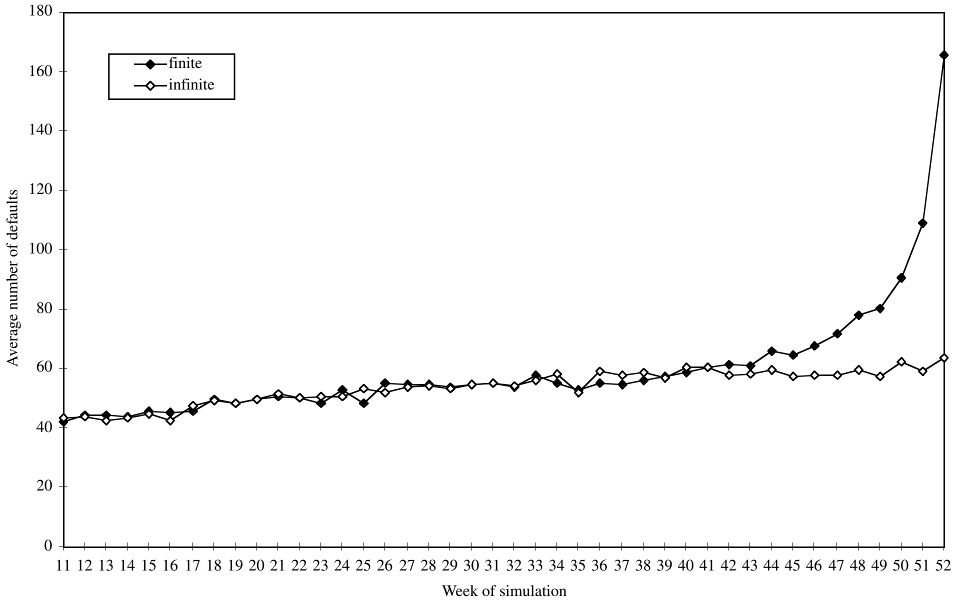

Figure 1. Effect of different time horizons on the average number of defaults each week.

In particular, for our experiments, the number of defaults in both the finite and infinite cases appears to generally increase until about week 25, when it reaches a plateau. After this initial upward trend, the infinite case appears to reach a steady number of defaults, while the finite case continues with an increasingly upward trend. The major divergence in number of defaults occurs at week 41 (*i.e.*, about 10 weeks from the end of the simulation). By the last week, the number of defaults in the finite case is nearly three times higher than in the infinite horizon.

The plateau phenomenon may be explained in terms of an adjustment of brownie point level. At the beginning of the simulation, agents are initialized with a starting *BP* level. Different values for *BPweight*, however, dictate a different *BP* equilibrium level for the agent. The initial rise in defaults indicates a period of time in which the agents are gradually moving toward this stable *BP* level. The plateau occurs when this stable level is reached.

4.2. Temporal changes in TEI and BP inclinations

When designing intelligent agents to cooperate in group activity over time, a sudden increase in defaults at the end of the group commitment is usually not desirable. Instead, as in the Systems Administration domain, agents may be expected to behave uniformly over all of the weeks of the simulation, as a sudden increase in defaults would deteriorate the overall accomplishments of the group. For instance, a portion of the agent's commitment decision is based on what it expects other agents to do in the following week, as shown in its calculation of *FEI*. If the total number of defaults rises dramatically from week to week during the end of the simulation, the agent's estimations of the default behavior of the other agents will be extremely inaccurate.

In particular, an agent's estimation of income for the following week will be imprecise as an indicator of future weeks. At the end of a simulation, when agents are able to default with few remaining weeks in which to be punished, each step closer to the end of the collaboration means the penalty becomes a smaller and smaller percentage of the possible payoff from outside offers. This can be seen clearly in the dramatic increase in defaults during the last 10 weeks. During this rapidly-changing period of the simulation, the discounting factor δ, which is meant to account for uncertainty, becomes dwarfed by the large week-to-week differences that do not occur earlier in the collaboration. This inability to predict future weeks further weakens the agent's ability to make rational decisions. Thus, it will be unlikely to correctly predict maximal outcomes, even for its own individual utility.

To investigate how this deterioration in group commitment can be overcome in a finite time horizon, and to provide more background data on the plateau effect discussed in the previous section, we investigated the effect of increasing *BPweight* in a finite horizon.

It is expected that as *BPweight* increases in the agent's utility function, decreasing the agent's utility more for each default, an agent's group commitment also increases, and thus the agent's overall number of defaults will decrease. Figure 2 validates this expectation, showing a drop in default levels as *BPweight* is increased in an infinite time horizon.

The effect of increasing *BPweight* in the finite horizon is less clear. As explained above and illustrated in Figure 1, the number of defaults per week in a finite horizon will coincide with the analogous situation in an infinite horizon, for some time at

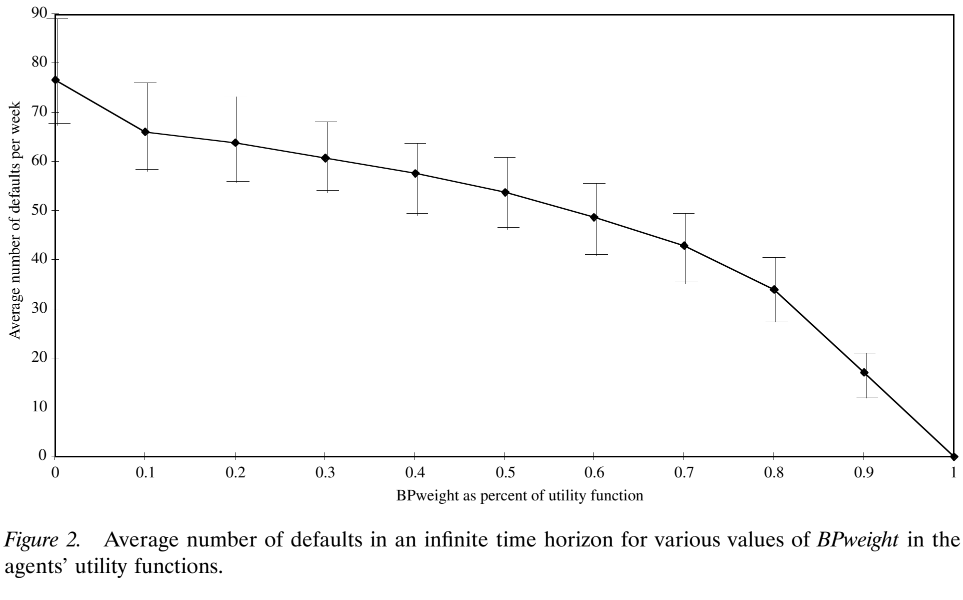

Figure 2. Average number of defaults in an infinite time horizon for various values of *BPweight* in the agents' utility functions.

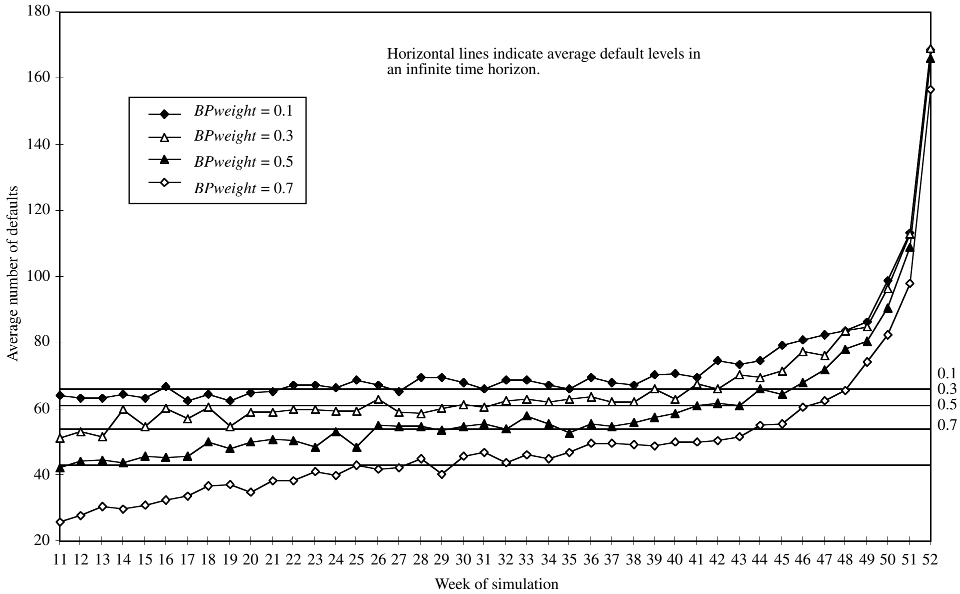

Figure 3. Average number of defaults per week in a finite time horizon for various values of *BPweight* in the agents' utility functions.

the beginning of the MWS before diverging. To analyze factors contributing to this divergence, our next experiments compared where this divergence happened for communities of agents with various levels of *BPweight* in their utility functions. The results can be found in Figure 3.

We found that, for high values of *BPweight*, the number of defaults does not "plateau" as it does for lower values, and as was noted above. As *BPweight* decreases, this plateau slowly becomes more and more pronounced, starting earlier in the simulation and continuing for more weeks.

This result reflects the relative weight of different factors in the utility function over time, and the resulting effect on default equilibrium. When *BPweight* is very low, the agent is making its decision based almost entirely on monetary factors (*TEI*). In a finite horizon, the agent will experience pressure to increase its number of defaults as the simulation proceeds, and its drop in rank will affect fewer weeks.

However, *BP* is also indirectly affected by the week number in the simulation, since the agent's total store of *BP* changes with time. The more an agent defaults, the lower its *BP* becomes. However, the absolute change in *BP* at each decision point is *not* affected by the timing of that decision, relative to other decision points or its occurence early or late in the simulation. By extension, then, late in the simulation, if *BP* is low, each default will induce a larger *percent* change in total *BP*.

The result of this effect is that, although the *TEI* portion of the agent's utility function encourages it to default more often over time, the *BP* factor also changes over time, encouraging the agent to default less and less. The plateau region indicates where these two dueling factors, based on the relative percent change in each at a given decision point, have reached equilibrium.

Since the effect of each factor is governed by its relative weight, different values of *BPweight* will cause this plateau to occur at different times, and for varying durations. For all values of *BPweight*, as time progresses, there will be downward pressure on *BP* from defaults, since defaulting causes a bigger change in *BP* than not defaulting. In cases with a high value for *BPweight*, however, the number of defaults typically starts out rather low. These cases also have fewer defaults overall. Since high values of *BPweight* are thus accompanied by lower default rates, an agent's total store of *BP* will change much more slowly than for lower values of *BPweight*. It also follows, then, that these agents will take longer to reach an equilibrium plateau, where the weighted change in *BP* equalizes the weighted change in future income (*FEI*). As *BPweight* decreases, and an agent's overall level of default increases, the agent reaches this equilibrium much more quickly, causing a more prolonged plateau in number of defaults. Thus, in addition to altering average number of defaults overall, manipulating *BPweight* allows us to alter agent behavior over time, in a finite horizon environment.

4.3. Optimal group commitments

While altering behavior with respect to number of defaults is useful to agent designers in some domains, in many other domains the number of defaults will be of less interest than the actual income earned by the agents. Our next set of experiments examined this income effect in the infinite time horizon case, where *BPweight* has a more steady influence. The results in Figure 4 show how average *group task income*, a phrase we use to refer to income just from group-assigned tasks, varies as *BPweight* is increased. In this figure, and subsequently in the paper, incomes are normalized with respect to the income that would have been earned if the originally assigned tasks had all been completed. Group task income is the income earned from β-tasks assigned by the group, minus the penalties incurred by the group from defaults.

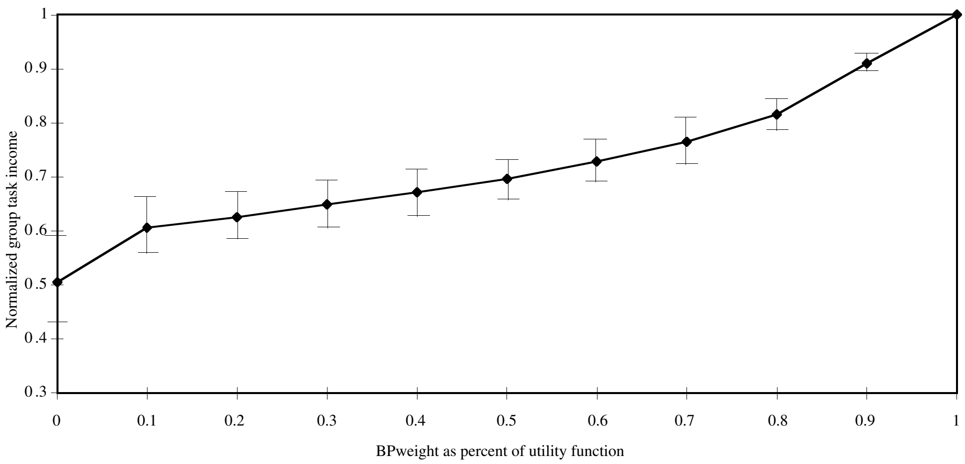

Figure 4. Effect of *BPweight* on mean group task income.

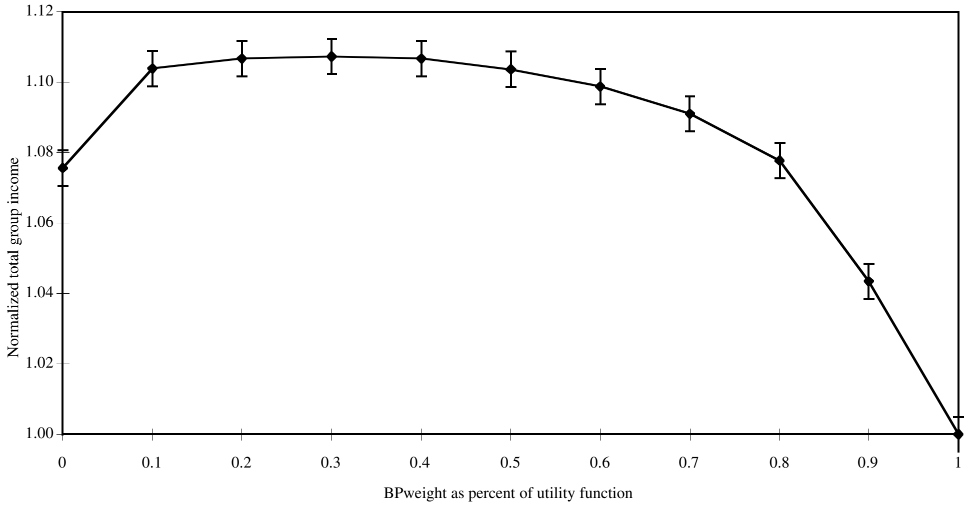

Figure 5. Effect of *BPweight* on mean total group income.

It does not include income earned by agents for any outside offers (γ) that they complete. As expected, the results show that group task income strictly increases as *BPweight* increases, and agents thus default less often.

In contrast to group task income, *total income* is an agent's income from group-assigned tasks and outside offers accepted. When we examine total income, we find rather surprising results (Figure 5). Note that, since total group income is normalized with respect to the income that would have been earned if the originally assigned tasks had all been completed, agents may sometimes earn a normalized total group income greater than 1, in cases where outside tasks completed more than made up for group tasks that were defaulted on.

Instead of strictly increasing with increased *BPweight*, as with group task income, total income reaches a maximum at *BPweight* = 0.3 (although differences between this *BPweight* value and a few surrounding values are within error), and then begins to drop. This result contradicts our original hypothesis, which was that the optimal number of defaults for maximizing total income would correspond with a high *BPweight*, near 1, where the impact on the group is the biggest factor in the agent's decision. Instead, this result shows that the optimal number of defaults actually occurs when *BPweight* is less than 1.

On further analysis, this result can be explained based on task density and the range of values of outside offers available in a given WTS. As explained in Section 2.3, the mean value of the outside offers exceeds the mean value of the group-assigned tasks. Thus, when deciding between β and γ, it is possible that an agent may be deciding between jobs that vary widely in value. When *BPweight* is low, an agent's decisions are driven mainly by *TEI*. For comparatively high outside offers, then, the current income (*CI*) gained by defaulting will be much bigger than the potential loss in future income (*FEI*). Thus, the agent will default.

When *BPweight* is very high and social factors are being weighted very heavily in the agent's utility function, the brownie point factor in the agent's utility function may overcome this potential gain in income. As a result, the agent may give up some very lucrative outside offers in order to stay committed to the group. In these cases, *value*(γ) is actually high enough to offset any group penalties suffered from the default, but the agent is, in a way, blinded to this fact by its unusually strong conscience.

In addition, with a very high density of group tasks assigned each week, the effect on future income of defaulting is decreased, because even irresponsible agents are guaranteed to be assigned a large number of group tasks, and each individual task therefore represents a smaller percentage of the agent's overall possible income [23]. In these situations, defaulting has a very small effect and brownie point considerations may tend to over-emphasize this effect. When combined with a high *BPweight*, then, the agent may not default even when it is truly beneficial to do so. Thus, it seems that "good guys" do not finish first. Rather, communities made up of *less* socially conscious, more balanced agents actually do better.

4.4. Optimality across different environments

Our final set of experiments extends this finding by varying an environmental factor, and observing the effect on the income-maximizing number of defaults. In these experiments, the number of tasks scheduled in each time slot (*task density*) was varied. As we varied the task density in these experiments, we also varied the number of tasks assigned based on rank so that the *percent* of rank-based tasks was constant across different task densities. In these experiments, the number of rank-based tasks was kept to approximately 30 percent of the total group tasks, as was also the case in the other experiments. The results are shown in Figure 6.

Given the results from the previous section, these results are unexpected. While at the highest task density studied in this paper[2] (10) there is a local maximum for total group income for a *BPweight* value greater than 0, this result does not hold for lower task densities. Instead, a pattern emerges in which a large range of values for *BPweight* greater than 0 provides a total group income which is approximately equal for a given task density. For example, for a task density of 2, the total group incomes associated with the range of *BPweight* values from 0.1 to 0.7 are all roughly the same, within error ranges.

This result points to an interesting conclusion. While we have shown that *BPweight* has a large and consistent effect on the average number of defaults and on group task income, Figure 6 shows that it does not have this same effect on total income. Thus, when creating cooperative agents, designers do not have to be as concerned about total group income as they are about these other factors. Instead, within a fairly large range of *BPweight* values, self-interested agents can be designed with utility functions that optimize for number of defaults or group task income, with the assurance that total group income will remain near its maximum.

This conclusion does not hold, however, if one considers total income when *BPweight* is 0. Figure 6 appears to show that income is maximized at this lowest

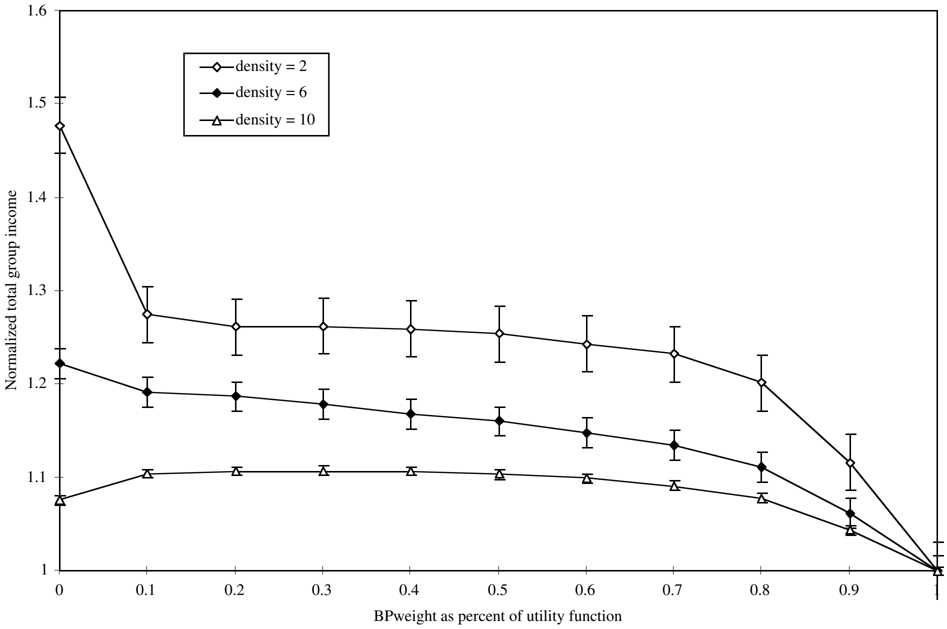

Figure 6. Effect of *BPweight* on mean total group income across various task densities.

value for *BPweight*. It is unlikely, however, that this is the behavior that would actually be desired in collaborative agents. Agents with low *BPweight* default much more frequently than more socially conscious agents. Furthermore, Figure 4 shows that group task income is at its lowest for this value of *BPweight*. Since most agent designers would attempt to optimize all of these factors, keeping defaults low and all measures of income high, this extreme behavior would be undesirable despite the apparent payoff as measured by total group income.

The seemingly inconsistent change in total income as *BPweight* initially rises from 0 to 0.1 remains an open question. A similar "jump" over this same interval in an otherwise smooth curve can also be observed in Figures 2 and 4, although to a lesser extent. While we suspect that this behavior reveals the large effect that adding brownie points at *any* weight has on agent behavior, the exact cause has not been isolated. This result seems to indicate that agents with no social consciousness (*i.e.*, brownie points do not factor at all into their utility functions) have radically different behavior from even mildly socially conscious agents. Analysis into this behavior is an area for future investigation.

5. Related work

Kalenka and Jennings [15] propose several "socially responsible" decision-making principles and examine their effects in the context of a warehouse loading scenario. Our work, like theirs, compares agents with and without social commitments and is concerned with "socially conscious" decision-making functions. However, our work differs from theirs in three ways: (1) their policies are domain-dependent and not decision-theoretic; (2) they do not vary environmental factors; and (3) they do not look at conflicting intentions or agents reneging on their tasks, but rather at whether or not agents choose to help each other.

Sen [20] also considers decision-making strategies that encourage cooperation among self-interested agents, but his work focuses on interactions between pairs of individual agents, rather than between an individual and a group. Sen's "philanthropic" agents take the good of the group into account, but do not always necessarily do what is best for the group.

There is a significant body of economics literature on rational choice and intention reconciliation. Iannaccone [13] examines social policies that alter individual utility functions to encourage group commitment. While these policies are similar in spirit to the social-commitment policies that SPIRE incorporates, they are aimed at group formation, not at conflicting intentions. Additionally, that approach is not applicable to agents that face multiple decision points over time. Höllander [12] studies incentives for encouraging group commitment and cooperation under a more limited definition of cooperation, in which an agent is required to incur a personal cost in order to cooperate. His model considers "emotional" cooperation within this limited definition, but assumes a rigid standard shared by all players, a requirement that we relax.

The social-commitment policies in SPIRE also differ from Shoham and Tennenholtz's [21] social laws. Social laws constrain the ways agents perform actions whereas social-commitment policies constrain decision-making. Social laws, are by their nature, domain specific; they constrain domain actions. In contrast, social-commitment policies affect decision-making across domains and tasks. The conventions Rosenschein and Zlotkin [18] present play a role in negotiation, similar to the role social-commitment policies play in SPIRE.

Cooper et al. [5] examine social consciousness in an approach similar to the brownie point model, referred to as the "warm glow" model, in which agents receive a constant amount of added utility when they do the "right thing." Our approach differs from this warm glow model, in that an agent's utility from brownie points is task-value dependent, and depends not just on the number of points gained for a given task, but also on the total number of brownie points that it has stored up over time.

6. Conclusions

In this paper, we show how to take into account factors other than direct monetary concerns in agent decision-making. The brownie point model provides a framework in which nonmonetary social factors can be effectively considered alongside monetary ones in collaborative, individually rational agents. By considering not just the default itself, but also task value and agent history in the calculation of brownie points, this operationalization of social consciousness realistically models the behavior in which we are interested. Additionally, because this system uses social consciousness over time, and allows it to play a variable role in utility functions, the effect of social consciousness on the behavior of the group as a whole can be studied.

Our results show that agents should be created so that they occasionally abandon their commitments in order to accept outside tasks. Their specific behavior regarding how often they do so can be controlled by setting *BPweight* appropriately. The exact weighting factor that should be used, however, depends on the goals of the agent designer. Our results empower agent creators to make these decisions on their own, based on the data that we have gathered about how agents behave in different situations. We also conclude that infinite or unknown time horizons are best, across all environments, indicating that agents should always be designed so that they do not know the exact length of their group commitment.

Keeping these results in mind, agent designers can create optimal, self-interested agents for any given domain by trading off desired default behavior with total group income. Although in many domains designers will be primarily interested in maximizing total group income, in other domains where agent designers care about defaults (whether because they care about how much of the overall group task gets done, or because they really want agents to severely limit defaults) this behavior can be controlled by adjusting social consciousness, without needing to completely redesign the agents or the group structure to get different behavior. Since total group income remains relatively constant over a large range of brownie point weights greater than 0, designers can focus more on desired default behavior, and be assured that total income will remain near its maximum.

Finally, experiments with varied task densities show that these results hold over multiple environments. Thus, when designing collaborative agents, designers can concentrate on creating desired default behavior and maximizing income, knowing that the basic principles outlined above will not be greatly affected by changes in the agents' task environment.

Current work in extending these findings to heterogeneous societies in which agents have different levels of social consciousness can be seen in [24]. Heterogenous societies that encourage group-conscious behavior are, however, susceptible to manipulation. In particular, the *free-rider* problem can arise: because agents are self-motivated, each would prefer that other agents do the work required to increase the group utility. A free rider wants its teammates to act in the team's interest, providing overall gains to the entire group, while opting out itself. If everyone but the free rider cooperates, its utility is maximized. However, if *all* the agents attempt to be free riders, they all will suffer. Initial investigations of the free-rider problem in the SPIRE system show that in heterogeneous groups of agents, less responsible agents are able to take advantage of their more group-conscious counterparts, reaping the full benefit of the outside offers they accept, while shifting a portion of the resulting group costs on the more responsible agents [24]. These results also show, however, that less responsible agents do only slightly better as individuals than their more responsible collaborators.

Furthermore, if all the agents choose to be less responsible, then each individual does worse than if they all chose to be responsible, *i.e.*, homogenous groups of responsible agents do better than homogenous groups of irresponsible agents. Additional work in experimenting with agents that learn and adapt to the environment, adjusting decision-making to fit the communities in which they are collaborating, can be viewed in [6].

Acknowledgments

Sarit Kraus and David Sullivan participated in the development of the SPIRE framework. This paper describes research that formed part of the first author's undergraduate honors thesis at Harvard University. Luke Hunsberger and David Sullivan provided helpful comments on earlier drafts.

This research has been supported by National Science Foundation grants IRI-95-25915, IRI-96-18848, and CDA-94-01024.

Notes

1. This central scheduler is used only for convenience. Many domains requiring cooperative agents would most likely not rely on a central scheduler in this way, but would instead negotiate each week's schedule based on (possibly incomplete) information about each agent. Since this negotiation is beyond the intended scope of the current SPIRE system, and we wish to study aspects of group-commitment scenarios which come after the initial schedule is made, we simplified this aspect of the problem for these experiments. We do not expect that this simplification has had any impact on our results.
2. In previous experiments [23], the highest possible task density in which all agents are busy all of the time, proved to be a special case, and thus was not experimented with here.

References

1. R. Axelrod, *The Evolution of Cooperation*, 1984.
2. M. E. Bratman, D. J. Israel, and M. E. Pollack, "Plans and Resource-Bounded Practical Reasoning," *Computational Intelligence*, vol. 4, no. 4, pp. 349–355, 1988.
3. C. Castelfranchi, "Modelling Social Action for AI Agents," *Artificial Intelligence*, vol. 103, pp. 157–182, 1998.
4. R. Conte and C. Castelfranchi, *Cognitive and Social Action*, UCL Press: London, 1995.
5. R. Cooper, D. V. DeJong, R. Forsythe, T. W. Ross, "Cooperation Without Reputation: Experimental Evidence from Prisoner's Dilemma Games," *Games and Economic Behavior*, vol. 12, no. 1, pp. 187–218, 1996.
6. S. Das, B. J. Grosz, and A. J. Pfeffer, "Learning and Decision-Making for Intention Reconciliation," in *Proceedings of the First Joint Conference on Autonomous Agents and Multi-agent Systems*, (to appear). Bologna, Italy, July 2002.
7. R. Gibbons, *Game Theory for Applied Economists*, Princeton University Press, 1992.
8. B. J. Grosz and S. Kraus, "Collaborative Plans for Complex Group Action," *Artificial Intelligence*, vol. 86, no. 2, pp. 269–357, 1996.
9. B. J. Grosz and S. Kraus, "The Evolution of SharedPlans," in A. Rao and M. Wooldridge (eds.), *Foundations of Rational Agency*, Kluwer Academic Press, pp. 227–262, 1999.
10. B. J. Grosz, L. Hunsberger, and S. Kraus, "Planning and Acting Together," *AI Magazine*, vol. 20, no. 4, pp. 23–34, 1999.
11. T. Hogg, "Social Dilemmas in Computational Ecosystems," *Proceedings of IJCAI-95*, pp. 711–716, 1995.
12. H. Höllander, "A Social Exchange Approach to Voluntary Cooperation," *American Economic Review*, vol. 80, no. 5, pp. 1157–1167, 1990.

13. L. R. Iannaccone, "Sacrifice and Stigma: Reducing Free-Riding in Cults, Communes, and Other Collectives," *Journal of Political Economy*, vol. 100, no. 2, pp. 271–291, 1992.
14. N. Jennings and M. Wooldridge, "Agent Theories, Architecture, and Languages: A Survey," in N. Jennings and M. Wooldridge (eds.), *Intelligent Agent: ECAI-94 Workshop on Agent Theories, Architectures, and Languages*, Springer Verlag, 1995.
15. S. Kalenka and N. R. Jennings, "Socially Responsible Decision Making in Autonomous Agents," in K. Korta et al. (eds.), *Cognition, Agency and Rationality*, Dordrecht: Kluwer, vol. 79, pp. 135–149, 1999.
16. R. Keeney and H. Raiffa, *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*, John Wiley and Sons: New York, 1976.
17. H. Levesque, P. Cohen, and J. Nunes, "On Acting Together," *Proceedings of AAAI-90*, pp. 94–99, 1990.
18. J. S. Rosenschein and G. Zlotkin, *Rules of Encounter: Designing Conventions for Automated Negotiation among Computers*, MIT Press, 1994.
19. T. Sandholm and V. Lesser, "Leveled Commitment Contracts and Strategic Breach," *Games and Economic Behavior*, vol. 35, pp. 212–270, 2001.
20. S. Sen, "Reciprocity: A Foundational Principle for Promoting Cooperative Behavior Among Self-Interested Agents," in *Proceedings of ICMAS-96*, pp. 322–329, 1996.
21. Y. Shoham and M. Tennenholtz, "On the Synthesis of Useful Social Laws for Artificial Agent Societies," in *Proceedings of AAAI-92*, pp. 276–281, 1992.
22. E. Sonenberg, G. Tidhar, E. Werner, D. Kinny, M. Ljungberg, and A. S. Rao, "Planned Team Activity," in C. Castelfranchi and E. Werner (eds.), *Artificial Social Systems, Lecture Notes in Artificial Intelligence (LNAI-830)*, Springer Verlag, pp. 227–256, 1994.
23. D. G. Sullivan, A. Glass, B. J. Grosz, and S. Kraus, "Intention Reconciliation in the Context of Teamwork: An Initial Empirical Investigation," in M. Klusch, O. Shehory, and G. Weiss (eds.), *Cooperative Information Agents III, Lecture Notes on Artificial Intelligence (LNAI-1652)*, Springer Verlag, pp. 149–162, 1999.
24. D. G. Sullivan, B. J. Grosz, and S. Kraus, "Intention Reconciliation by Collaborative Agents," in *Proceedings of ICMAS-2000*, IEEE Computer Society Press, pp. 293–300, 2000.
25. M. Tambe, "Towards Flexible Teamwork," *Journal of Artificial Intelligence Research*, vol 7, pp. 83–124, 1997.
26. K. P. Yoon and C.-L. Hwang, *Multiple Attribute Decision-Making: An Introduction*, Sage University papers, Series on Quantitative Applications in the Social Science, vol. 104, 1995.