The Influence of Social Norms and Social Consciousness on Intention Reconciliation*

Barbara J. Grosz[1], Sarit Kraus[2], David G. Sullivan[1], Sanmay Das[3]

[1] Division of Engineering and Applied Sciences
Harvard University, USA
[2] Bar-Ilan University, Israel, and Institute for Advanced Computer Studies,
University of Maryland, USA
[3] MIT Artificial Intelligence Laboratory, USA

Abstract

Research on resource-bounded agents has established that rational agents need to be able to revise their commitments in 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. Initial work with SPIRE examined the impact of environmental factors and agent utility functions on individual and group outcomes in the context of one set of social norms governing collaboration. This paper extends those results by further studying the effect of environmental factors and the agents' level of social consciousness and by comparing the impact of two different types of social norms on agent behavior and outcomes. The results show that agents operating under both sets of norms respond similarly to varying environmental factors, but that the choice of social norms influences the effectiveness of social consciousness and other aspects of agents' utility functions. In experiments using heterogeneous groups of agents, both sets of norms were susceptible to the free-rider effect. However, the gains of the less responsible agents were minimal, suggesting that agent designers would have little incentive to design agents that deviated from the standard level of responsibility to the group.

1 Introduction

A number of applications have been proposed that require agents to work collaboratively to satisfy a shared goal [11, 17, 37, 41, *inter alia*]. In the context of such collaborative activities, agents need to be able to revise their commitments and plans as new opportunities arise, handling situations in which intentions to do team-related actions conflict with other possible actions or plans. This paper focuses on the process of decision making that agents perform when reconciling conflicting intentions in these group-activity contexts, and it examines the

* This research was supported by National Science Foundation grants IIS-9978343, IRI-95-25915, CDA-94-01024, and IIS-9907482. The work presented was completed when the fourth author was a student at Harvard University. Alyssa Glass and Michael Epstein helped to develop the SPIRE system. We thank Luke Hunsberger and Tim Rauenbusch for helpful comments on previous drafts.

use of both external social norms (e.g., imposing penalties when agents fail to honor their team commitments) and internal measures of social consciousness (e.g., an agent's sense of its reputation as a responsible collaborator) to influence the agents' decisions.

The problem of intention reconciliation arises because rational agents cannot adopt conflicting intentions [7, 15, *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 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. This paper examines instances of intention reconciliation in which at least one of the conflicting intentions is related to an agent's commitment to a team plan. While much of the prior work on agent collaboration and negotiation [25, 27, 31] has assumed that commitments to collaborative activity are binding, we are interested in situations in which agents are allowed to renege occasionally on such commitments. For example, in the domain of automated systems administration (see [39]), it might be reasonable to allow an agent committed to performing a file-system backup to break that commitment (to *default*) so that it can assist with crash recovery on another system.

Intention reconciliation in team contexts requires that agents weigh purely individual costs and benefits of their decisions with team-related concerns. Defaulting on a team-related commitment for the sake of another opportunity may at times appear beneficial from a short-term, selfish perspective, even with the imposition of immediate penalties for defaulting. However, because we assume that agents have relationships that persist over time, an agent must also consider the impact of defaulting on its ability to collaborate in the future and, more generally, on its future expected outcomes.

In a given society of agents, *social-commitment policies* [39]—domain-independent social norms that govern various aspects of collaboration—may be used to influence an agent's decision making. These policies include both rewards and penalties for individual acts in the context of group activities, and the details of a given team's policy are made known to its members so they can adjust their decision making accordingly. By stipulating ways in which current decisions affect future as well as current 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. Social-commitment policies may include immediate penalties for defaulting like the ones proposed by Sandholm et al. [32, 1, 33, 34], as well as longer-term policies based on the agent's behavior over time. Section 4 describes two examples of the latter type of policy that were studied in the work reported in this paper.

Social factors may also function in another way. If agents get part of their utility from the team, they have a stake in maximizing group utility [9]. 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" (refusing offers that improve the agent's immediate individual income) 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. Therefore, it may be beneficial for agents to incorporate this type of *social consciousness* into their decision making, considering both this internal factor and the externally imposed social-commitment policies when deciding whether to default. We have used the *brownie-points model* of Glass and Grosz [13] to study the effect of social consciousness on outcomes, as well as its susceptibility to manipulation. Our experiments show that including social consciousness in agents' decision making can lead to improved outcomes when agents face a high degree of uncertainty about the values of their future outside offers and task assignments.

To enable the simulation and study of intention reconciliation by collaborative agents, we have developed the SPIRE experimental framework. In earlier work [13, 39, 40], we used SPIRE to study the effects of a single social-commitment policy, as well as various environmental factors and agent characteristics, on the decisions and outcomes of both homogeneous and heterogeneous groups of agents. In this paper, we extend this work by further examining the effect of environmental factors and the agents' level of social consciousness, and by studying the outcomes of simulations using two different social-commitment policies.

This work with SPIRE brings two threads of research together, joining research on collaboration in multi-agent systems [15, 16, 24, 27, 42]—which has established that commitment to the joint activity is a defining characteristic of collaboration—with research on rationality and resource-bounded reasoning [8, 12, 21, 30, *inter alia*]—which has established the need for agents to dynamically adapt their individual plans to accommodate new opportunities and changes in the environment. Our work 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. In addition, SPIRE allows us to consider repeated agent interactions that change over time in non-trivial ways; these long-term interactions are more realistic than both the repeated games typically considered in game-theory research and the one-shot deals that have been the focus of most prior work on team formation and intention reconciliation. The contributions of this paper include two examples of social-commitment policies that could be employed by agent societies, and a comparison of the effectiveness of these policies in a variety of contexts. In addition, we examine the interaction of these external policies with an internal measure of social consciousness. Based on the results of the experiments presented in this work, we suggest principles for designers of agents and agent environments. The experiments suggest that designers of agents for environments employing either of our social-commitment policies would have little incentive to develop agents that deviated from a standard level of responsibility to the group.

2 The SPIRE framework

The SPIRE (SharedPlans Intention Reconciliation Experiments) simulation system was designed to enable the study of intention reconciliation in collaborative contexts. SPIRE allows us to examine the impact of environmental factors, social-commitment policies, and different 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. Below, we present an overview of the system; more details can be found in an earlier paper [39].

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 make the simplifying assumptions that each task lasts one time unit and 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 is divided into a series of *weeks*, each of which is in turn divided into some number of *time slots*. The tasks in the group activity each take one time slot to complete. Each week, the same group activity is assigned to the same team of agents, because varying either the group activity or the team members could obscure sources of variation in the outcomes. However, the individual tasks within the group activity are not necessarily done by the same agents each week. Because negotiation over which agents should perform which tasks is not the focus of our work, we simulate this process with a central scheduler. At the start of each week, this scheduler assigns tasks from the group activity to agents according to a fixed policy (see Sect. 3.1 and 3.2 for more details).

To model the type of conflicts we are interested in studying, agents are chosen at random each week and given the chance to do one of a series of *outside offers* that conflict with tasks from the group activity. (In the experiments reported in this paper, the offers are tasks that involve a single agent in isolation; in the future, we plan to explore conflicts that arise when agents are committed to multiple group activities.) 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 one of the other agents in the group is capable of doing β and is available at the time slot for which it is scheduled, the task is given to that agent (whom we will refer to in this context as a *replacement agent*); otherwise, β goes undone. If the replacement agent honors its commitment to do β, it receives the value of β as part of its income for that week; no additional rewards are obtained for serving as a replacement. The team as a whole incurs a cost whenever an agent defaults; these *group costs* are divided equally among the team's members, and the magnitude of the costs is larger when no replacement agent is available. Group costs model the real-world losses that result from searching for a replacement and, if none can be found, from the group's failure to complete its group task.

The framework described above addresses scenarios similar to ones considered in the economics and game-theory literatures (e.g., [20, 22, 28]). For example, honoring or defaulting on group commitments is comparable to cooperating or defecting in the classic prisoner's-dilemma game [2, 3, 4]. However, many of the assumptions made in prior work are not valid in the environments that SPIRE is designed to consider. In the repeated prisoner's dilemma [29, 26, 35, 5], for example, agents face the same scenario at each decision point, whereas SPIRE agents generally play a different game on each iteration. Moreover, the games differ in nontrivial ways; in particular, an agent's utilities and task assignments can depend on both its own past actions and the actions of other agents (see Section 4). SPIRE agents also have imperfect information about the actions, utility functions, and outside-offer values of other agents, and they typically face multiple decisions in a single time period, before any new information is received.

Goldman and Kraus [14] have developed formal models of the type of scenario that SPIRE is designed to address for situations from an e-commerce domain involving two sellers and three buyers. The complexity stemming from the above-mentioned features of the SPIRE game made it difficult to treat a larger number of agents. In particular, the fact that agents' actions affect their future utilities means the individual games are not independent, a fact which introduces serious complications into the analysis. In addition, SPIRE models situations that may involve heterogeneous groups of agents; such heterogeneity makes it difficult to collapse a group of agents into a single agent for the purpose of simplifying the analysis. Another possible theoretical approach involves applying classic decision-theoretic frameworks like Markov decision processes to SPIRE scenarios. However, while researchers have begun to adapt these frameworks to multi-agent systems, a number of important research challenges need to be addressed before these frameworks will be able to handle problem domains with collaborating but self-interested agents [6]. It is also worth noting that while Sandholm et al. [32, 1, 33, 34] apply a formal approach to scenarios involving two or three agents in their related work on leveled commitment contracts (see Section 7), they too employ simulations when considering larger groups of agents like the ones we have studied with SPIRE.

A SPIRE simulation requires that a number of parameters be set, some of which are central to the decision-making problem being studied, such as those involved in the social-commitment policies and the agent utility functions. Varying these parameters allows insights about the problem of intention reconciliation to emerge. Other parameters, such as the number of time slots and the number of task types, should be chosen based on the particular application domain being modeled (e.g., we have modeled a group of agents engaged in computer-systems administration; see our earlier work [39] for more detail). Still other parameters, such as the values of the tasks and outside offers, are less central, and the values chosen for them are in some sense arbitrary. Provided that certain relationships are maintained (e.g., that the outside offers are worth more on average than the group activity tasks), the particular choices do not seem to significantly affect the nature of the results (see, e.g., subsequent work by Das [10], who modified the distributions of the task and offer values).

3 Social-commitment policies in SPIRE

Our experiments compare two social-commitment policies that encourage agents to consider the good of the group when reconciling intentions. In one policy, agents are ranked on the basis of their past behavior, and agents who are higher in the rankings receive more valuable tasks. In the other policy, each agent's income is discounted by a factor that primarily reflects its own past behavior, regardless of what other agents have done. The following sections describe each of these policies in turn.

3.1 Ranking-based social-commitment policy (RSCP)

In the *ranking-based social-commitment policy (RSCP)*, agents are ranked on the basis of their past behavior, and the scheduler assigns a portion of each week's tasks on the basis of these rankings. More specifically, each agent has a score[1] that reflects the number of times it has defaulted. Each default in the current week results in an absolute decrement in the agent's score. The impact of past weeks' defaults diminishes over time. In addition, the score reduction that a defaulting agent incurs is larger if no agent available is to replace it. At the end of each week, an agent's score is updated as follows:

$$s(w+1) = \alpha s(w) - \rho_1 d - \rho_2 D \tag{1}$$

where w is the week number, $s(0) = 0$, α is a constant decay factor that reduces the impact of past defaults, d represents the number of defaults with replacement in the current week, D represents the number of defaults without replacement in the current week, and ρ_1 and ρ_2 are constant score reductions for defaulting with and without replacement respectively. This formula results in scores that are rational values less than or equal to zero. For the experiments in this paper, we used $\alpha = 0.5$ to represent a moderate level of decay, and we chose $\rho_1 = 1$ and $\rho_2 = 5$ to reflect the fact that defaulting without a replacement is much more costly to the group. These values can be modified to create other versions of this social-commitment policy.

The scheduler assigns N tasks per agent on the basis of the agents' scores; we refer to these tasks as *score-assigned tasks*. The agent with the highest score receives the N highest-valued tasks that it can perform (given its capabilities and availability), 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 each agent receives N tasks, the remaining tasks are assigned to agents picked at random. The strength of the RSCP can be varied by modifying the value of N.

This policy was used in our prior work involving SPIRE [13, 39, 40]. It reflects the intuitive notion that teams of agents would be more likely to entrust their most valuable group-related tasks to collaborators who had been

1. In some of our prior work [13, 39], 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.

most responsible in their past interactions with the group. One difficulty of this policy is that an agent's future income is based on its ranking, which in turn depends on the behavior of other agents. As a result, obtaining an accurate estimate of future income requires a game-theoretic analysis, and such analyses are known to be difficult in situations involving large numbers of agents. Section 4.2 outlines the approach to income estimation that we have adopted in light of these difficulties.

3.2 Discount-based social-commitment policy (DSCP)

In the *discount-based social-commitment policy (DSCP)*, the income that an agent receives from a portion of its tasks is discounted using a factor that depends on the agent's individual reputation as a collaborator and the reputation of the group as a whole, but not on the individual reputations of other agents. Because an agent's income does not depend on its ranking, this social-commitment policy allows agents to estimate their future income more accurately. In addition, discounting allows irresponsible agents to be punished in cases when they are the only ones capable of performing a valuable task; under the RSCP, such agents would receive the full value of the task, regardless of their position in the rankings.

The reputation of an agent is represented using a numeric score that *increases* as the agent's reputation decreases. In addition, a defaulting agent incurs a larger score increase if no agent is available to replace it. At the end of each week, an agent's score is again updated using equation 1, where ρ_1 and ρ_2 have negative values. The group as a whole also has a reputation score that is maintained in a similar way. For the experiments in this paper, we used $\alpha = 0.5$ as we did under the RSCP, and we chose $\rho_1 = -1$ and $\rho_2 = -1.1$. The ratio of ρ_2 to ρ_1 cannot be as large here as it is under the RSCP, because an agent's score leads directly to the magnitude of its task discounts. As a result, larger values of ρ_2 would drastically skew the income that agents receive.

In the DSCP, the scheduler assigns all tasks randomly, but *N* tasks per agent are selected at random and discounted by a factor that depends on both the agent's individual reputation score and the reputation score of the group. The actual discount factor is computed by scaling these scores to be in the range [0, 1] and taking their product. The *individual scaling function* (*ISF*) is the following sigmoid function:

$$ISF(x) = \frac{1.04}{1 + e^{(x-5)/1.5}} \tag{2}$$

where *x* is the individual's score; it is shown in the left half of Figure 1. The parameters of the function were chosen to accommodate typical individual-score values, which range from 0 to 10. Using a sigmoid function allows agents with good reputations to default a moderate number of times without being punished too severely, while still providing a disincentive against defaulting too often. In addition, agents whose reputations are already very poor do not suffer much additional harm if they continue to default.

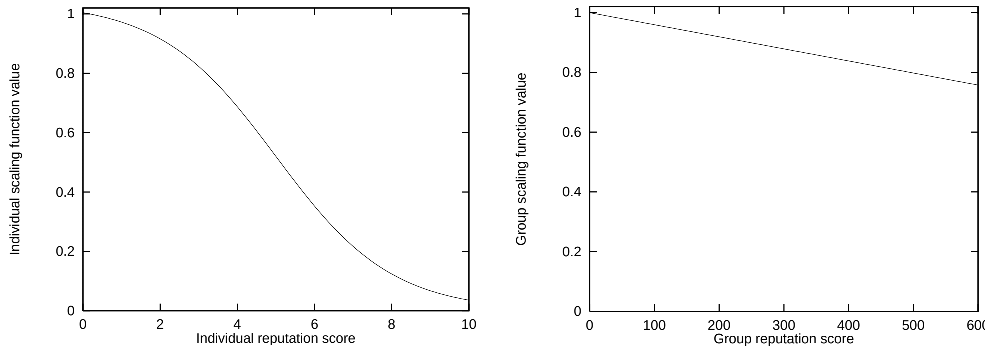

Figure 1. The individual scaling function or ISF (*left*) and the group scaling function or GSF (*right*) used in experiments involving the DSCP.

The *group scaling function* (*GSF*) is the following linear function:

$$GSF(x) = 1 - \frac{x}{41.25n} \tag{3}$$

where x is the group's score and n is the number of agents in the group; it is shown in the right half of Figure 1. Dividing by the number of agents effectively scales the cost of defaulting based on the number of opportunities to default, because the number of outside offers in a SPIRE simulation is directly proportional to the number of agents. The parameter values were chosen to give *GSF* values that are typically between 0.8 and 1. This range of values ensures that group reputation matters less than individual reputation when determining the income from discounted tasks. As a result, agents need not consider the behavior of individual agents or subgroups of agents when estimating their future income, and they can thus avoid the type of game-theoretic analysis that is required for accurate future-income estimates in the RSCP (Sect. 3.1).

The motivation for discounting task values is economically grounded. If workers fail to meet their obligations, they will tend to receive lower pay raises or larger pay cuts. Similarly, if a group (e.g., a company) fails to meet its obligations, it will be less likely to win bids for new jobs and will be forced to bid lower, resulting in lower group income. Strictly speaking, the portion of the discount that stems from the group's reputation is actually part of the environment, not the social-commitment policy, because it comes from outside the group. One problem that we have yet to address is what should be done with the income that agents lose from discounting. In our current implementation, this income is simply discarded.

4 Decision making in SPIRE

In deciding whether to default on a task β so as to accept an outside offer γ, an agent must determine 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 should it default. Future expected income represents the agent's estimate of its income in future weeks, based on the social-commitment policy and the agent's score. The agent first approximates the impact that defaulting will have on one week of its income. Sections 4.2 and 4.3 describe this estimate in more detail. The agent then extrapolates beyond that week to obtain a more complete estimate, using a discount factor $\delta < 1$. This use of discounting can be viewed in at least two different lights: as a reflection of an agent's internal uncertainty about its predictions [39], or as an external factor—much like an interest rate—that allows an agent to assess the present value of income that will be earned in the future.

For the experiments in this paper, we assume that agents are uncertain about the duration of their collaboration, and therefore we use the infinite-horizon version of the FEI formula described by Glass and Grosz [13]. If F is the estimate of next week's income and δ is the discount factor, then:

$$FEI(F) = \delta F + \delta^2 F + \delta^3 F + \ldots = \left(\frac{\delta}{1-\delta}\right)F \tag{4}$$

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

4.2 Estimating next week's income under the RSCP

When operating under the RSCP, an agent estimates its two possible incomes 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 value of 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. It is difficult for an agent to accurately estimate its ranking using such limited information. To compute a future ranking, the agent needs to reason not only about its own behavior, but about the behavior of other agents, and about how its own behavior will influence the other agents, and so on. To avoid these game-theoretic complications, 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:[2]

(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.

Because members of a collaborative group are aware of their group's social-commitment policy—and thus will know the policy used to assign tasks under the RSCP—we allow the agents to call the scheduling function. An agent calls this function once to compute the value of its score-assigned tasks if it does not default ($F_{\text{no-def}}$, 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)).

An 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 (because this affects the sizes of the equivalence classes). The estimated loss of income can vary greatly [40]. 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 $F_{\text{no-def}}$ is increased by 20 in such cases.

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.

4.3 Estimating next week's income under the DSCP

When operating under the DSCP, an agent estimates its two possible incomes during the following week by computing its new reputation score both if it defaults and if it does not default, and determining the expected value of its discounted tasks in each case. We assume that an agent knows both its individual score and the ISF function, but not the GSF function or the behavior of other agents. Because an agent could infer its group's prior GSF *values* from the values of its discounted tasks, we provide these prior GSF values to the agents, who average them to compute an estimate of the current week's GSF value. Each individual score is first converted to an ISF

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

value and then multiplied by both the expected value of the tasks that will be discounted and the estimated GSF values. Because all tasks are assigned randomly, agents can compute the expected value of the tasks that will be discounted without the expense of calling the scheduling function. The lower computation costs also allow agents to compute estimated incomes that reflect their prior defaults during the current week (cf. Sect. 4.2).

4.4 Social consciousness using brownie points

In addition to being concerned about its income, an agent may also derive utility from being a “good guy” and considering the good of the group. Glass and Grosz’s *brownie-points model* [13] captures this aspect of agents’ utilities, providing a measure of an agent’s sense of its reputation as a responsible collaborator. Agents begin a simulation run 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, each change in brownie points takes 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 increase in brownie points when an agent chooses not to default is greater for low-valued tasks and for high-valued offers.

Note that 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. 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. 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 the non-monetary measure described above is simpler and more intuitive. And although the monetary and non-monetary elements of an agent’s utility have different units, the various factors can be combined using a technique from multi-attribute decision-making [45] described below.

4.5 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 a 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

all experiments:
60 agents
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
3t/10 offers per week (t = # tasks):
- values chosen randomly
- possible values = task values + 95

experiments using the RSCP:
10 score-assigned tasks per agent per week
δ (factor used to weight FEI) = 0.4
BPweight = 0.1

experiments using the DSCP:
10 discounted tasks per agent per week
δ (factor used to weight FEI) = 0.85
BPweight = 0

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

default and no-default BP values (BP_{def} and $BP_{\text{no-def}}$) are similarly normalized. This normalization allows us to compare TEI with BP and to combine each (TEI, BP) pair into a single utility value. Finally, the normalized values are weighted based on the agent's social consciousness:

$$\begin{aligned} U_{\text{def}} &= TEIweight \times normTEI_{\text{def}} + BPweight \times normBP_{\text{def}} \\ U_{\text{no-def}} &= TEIweight \times normTEI_{\text{no-def}} + BPweight \times normBP_{\text{no-def}} \end{aligned} \qquad (5)$$

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.

5 Experimental results

In this section, we present the results of experiments designed to compare the two social-commitment policies and to further examine the impact of environmental factors and agent characteristics—including social consciousness—on individual and group outcomes. We have grouped the experiments according to whether they use the RSCP (Sect 5.1) or the DSCP (Sect. 5.2). In Section 6, we discuss the results in their entirety.

The experiments make the simplifying assumptions that all agents are capable of doing all tasks and that all agents are initially available at all times. (SPIRE can also handle the more general situation in which agents have different capabilities and availabilities, but we have yet to investigate this type of scenario.) 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.

We fixed the number of outside offers for these experiments at 30 percent of the number of tasks, matching the percentage of tasks affected by the social-commitment policies. When there are a large number of outside offers, as was the case in our earlier experiments [13, 39, 40], the impact of the social-commitment policies on the agents' incomes is diminished, because agents will frequently be able to offset their lower incomes from group tasks with income earned from outside offers.

For each social-commitment policy, we first present experiments involving homogeneous groups of agents. The initial experiments determine the optimal utility-function parameters for homogeneous groups. The resulting values for δ—the parameter that agents use to weight future income (Sect. 4.1)—and for *BPweight*—the parameter that determines an agent's level of social consciousness (Sect. 4.5)—are used as the standard settings in the remaining experiments. Earlier work with SPIRE [13, 40] showed that homogeneous groups of agents do better as individuals when they have an intermediate amount social consciousness (i.e., when they consider both brownie points and income by using a *BPweight* that is greater than 0 and less than 1), and when they give a moderate weight to future income, avoiding δ values close to 1 that prevent most defaults. Therefore, we expected to see comparable results for both the RSCP and the DSCP under the new outside-offer rate.

Other homogeneous-group experiments assess the impact of two environmental factors: the number of tasks scheduled for each time slot (task density), and the rate at which outside offers are made. These experiments test the following hypotheses: (1) as task density increases, the rate of defaulting will decrease, provided that the percentage of tasks affected by the social-commitment policy is held constant; (2) as task density increases, mean individual income (from tasks and offers) and group income (from tasks only) will both decrease; (3) as the outside-offer rate increases, the *rate* of defaulting will stay the same but the absolute *number* of defaults will increase; and (4) as the outside-offer rate increases, group income will decrease, but mean individual income will increase. The motivations for these hypotheses are discussed in the sections describing the actual experiments (Sect. 5.1.2, 5.1.3, 5.2.2, and 5.2.3).

In addition to these homogeneous-group experiments, we also present the results of experiments that assess the robustness of the optimal homogeneous-group utility parameters by testing them in heterogeneous contexts in which some agents attempt to take advantage of the group. Initial SPIRE experiments involving heterogeneous groups [40] showed that the *free-rider effect*—in which less responsible agents take advantage of their more responsible collaborators—is a potential problem in such settings. However, preliminary tests indicated that this effect might be avoided in environments with lower outside-offer rates. Therefore, we used the heterogeneous-group experiments in this work to test the hypothesis that agents who use the optimal homogeneous-group parameters for social consciousness and the weight given to future income will do better than their less responsible collaborators in environments with a moderate number of outside offers.

All of the experiments also serve to compare the two social-commitment policies. We hypothesized that agents operating under the DSCP would achieve higher individual and group outcomes than agents operating under the RSCP because of the better estimates of future income—and thus the better decisions—that the DSCP makes possible (see Sect. 3.2). Similarly, we expected that the ability to make better predictions would lead agents operating under the DSCP to respond better to environmental changes than agents operating under the RSCP.

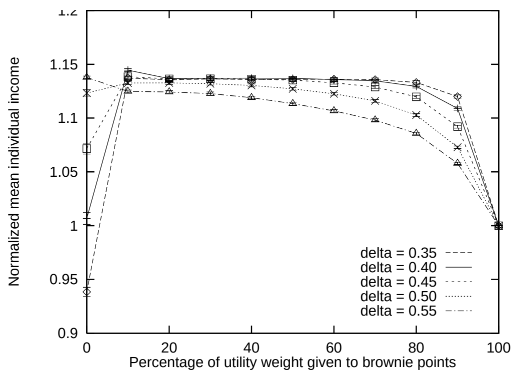

Figure 3. The impact of δ and *BPweight* on the normalized mean individual income of homogeneous groups of agents under the RSCP. To make the distinctions between the curves clearer, we have focused on a portion of the y-axis from [0.9:1.2].

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-percent confidence intervals; in many cases, the intervals are small enough that the errorbars are difficult to see. 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 RSCP experiments

We first present the results of experiments using the ranking-based social-commitment policy or RSCP (Sect. 3.1). The incomes that we present are normalized by dividing each agent's income by the *average* total value of the tasks received by an individual agent. In computing mean individual income and group task income, this approach is equivalent to dividing each agent's income by the value of the tasks assigned to the agent itself, but it avoids this latter method's tendency to artificially inflate the incomes of less responsible subgroups in experiments involving heterogeneous groups. In addition, because the average total value of the tasks assigned to an agent does not vary greatly across the task schedules used in a simulation, it adds less variance to the results than an alternative normalization method that divides by the total value of the tasks assigned to a particular agent.

5.1.1 Determining the baseline utility-function parameters for the RSCP

In selecting the standard settings for δ—the parameter that agents use to compute FEI (Sect. 4.1)—and for *BPweight*—the parameter that agents use to weight brownie points (Sect. 4.5)—we chose values that allow homogeneous groups of agents to receive optimal mean individual incomes under a 30 percent outside-offer rate. We tested many different pairs of values for δ and *BPweight*; Figure 3 displays the results of some of these experiments.

The tests show that mean individual income is maximized for a δ value of 0.4 and a *BPweight* of 0.1. Agents

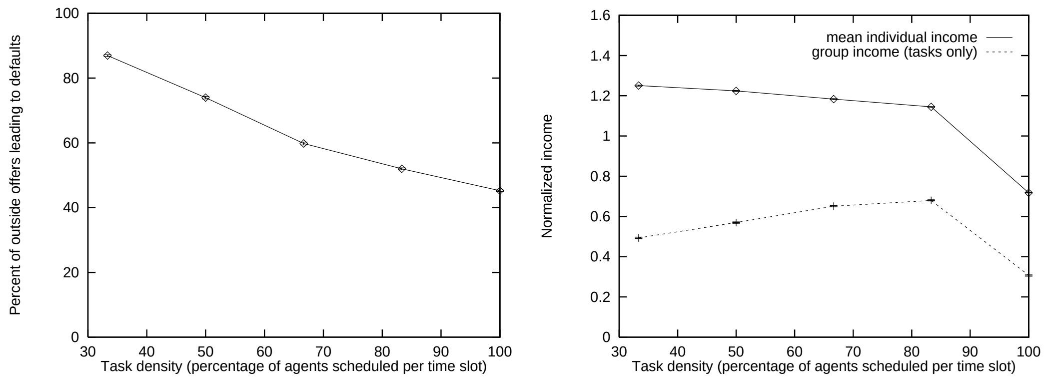

Figure 4. The impact of task density on the rate of defaulting (*left*) and on normalized individual and group incomes (*right*) when using the RSCP.

who use lower δ or *BPweight* values tend to default more often than is optimal, which results in group costs that more than counterbalance the extra income from outside offers, while agents who use higher δ or *BPweight* values do not default enough to take full advantage of the potential utility gains from outside offers. This result confirms our hypothesis that agents operating in homogeneous groups do better when they have an intermediate amount of social consciousness and give a moderate weight to future income. Unless otherwise stated, the settings of $\delta = 0.4$ and *BPweight* = 0.1 were used throughout the experiments in this section.

5.1.2 Behavior of the RSCP under different task densities

In this set of experiments using the RSCP, we varied the number of tasks scheduled for each time slot (task density). This factor can also be expressed as the percentage of agents who are scheduled to perform a task in each time slot. Earlier [39], we showed that task density can effect at least two factors that directly influence an agent's decision making: (1) the degree to which other agents are available to replace an agent when it defaults on a group-related commitment; and (2) the percentage of tasks affected by the RSCP, and thus the range of possible values for the score-assigned tasks. In these experiments, we eliminate the second of these factors by maintaining a constant *percentage* of score-assigned tasks as we increase the task density. This allows us to test the following hypotheses: (1) agents will default less often as task density is increased, provided that the percentage of tasks affected by the social-commitment policy is held constant; and (2) increasing task density will lead to lower individual and group incomes. Both of these hypotheses are based on the fact that increasing task density should lead to a decrease in the availability of replacement agents (as an agent is scheduled to perform more tasks, it is less available to serve as a replacement) and thus to an increase in the group costs from defaulting (see Sect. 2 and Sect. 3.1). Figure 4 displays the results.

With a fixed percentage of score-assigned tasks, agents default less often as task density is increased (Fig. 4, *left*), confirming our first hypothesis. However, this decrease in defaulting cannot be fully explained by a decrease in the availability of replacement agents, the rationale used in forming the hypothesis. As Table 1

Table 1. Effect of task density on the average percentage of outside offers for which no replacement agent is available under the RSCP.

Task density	Offers with no replacement
33.3%	0.00%
50.0%	0.00%
66.7%	0.00%
83.3%	5.15%
100.0%	100.00%

shows, this factor does not play a role at the lower densities, for which replacement agents are always available. Instead, the consistent decrease in defaulting is primarily caused by the influence of the score-assigned tasks. While maintaining a constant percentage of score-assigned tasks does ensure that the values of these tasks are drawn from a consistent range of values as task density is varied, it also causes the *number* of these tasks to increase as task density increases. As a result, the FEI losses from defaulting also increase, and agents are less likely to default.

The effect of task density on both mean individual income (from both tasks and outside offers) and group income (from tasks only) is shown in the graph on the right side of Figure 4. As task density increases, agents do worse as individuals, despite the fact that their rates of defaulting decrease and that the availability of replacement agents at lower densities means that the group costs of defaulting do not increase for these densities. To understand this result, it is important to realize that agents can afford to default at very high rates when replacement agents are always available. In such cases, the group costs of defaulting are lower, no income from tasks is lost because defaulted tasks are still completed, and agents receive added income from serving as replacements. While reducing the rate of defaulting lowers group costs, it also reduces the amount of extra income from outside offers, and that loss ends up being greater than the savings in group costs. When replacements are not always available, group costs are higher and income is lost from tasks that go undone. As a result, agents do even worse. The individual-income results thus confirm our hypothesis, but for reasons that are more complex than we anticipated.

As task density increases from 33.3 to 83.3 percent, group task income (i.e., income from tasks only, minus any group costs) also increases, disproving our hypothesis. Here again, the unexpected availability of replacement agents for these task densities means that group costs do not increase significantly as task density is varied within this range of values. Therefore, a reduction in defaults leads to lower group costs and higher group task income. At a task density of 100 percent, replacements are never available, and this leads to the anticipated increase in group costs, a loss in income from uncompleted tasks, and lower group task incomes.

5.1.3 Behavior of the RSCP under different outside-offer rates

Our earlier work [40] gave some indication of the effect that the outside-offer rate can have on outcomes. In an

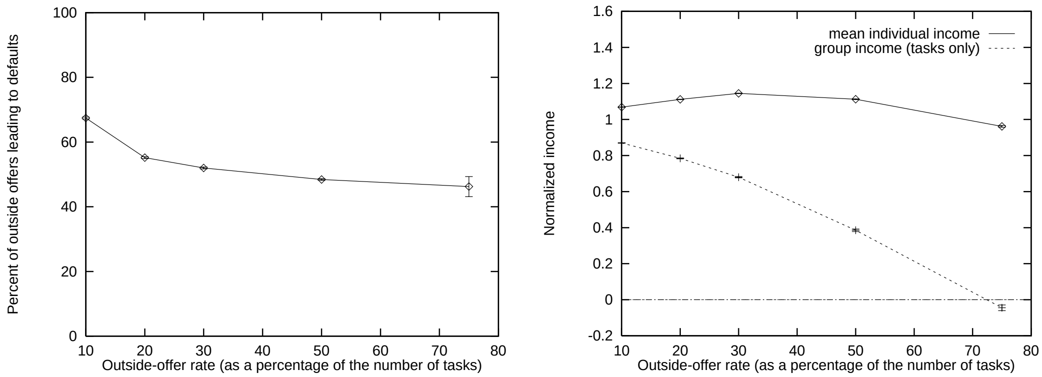

Figure 5. The impact of the outside-offer rate on the rate of defaulting (*left*) and on normalized individual and group incomes (*right*) when using the RSCP.

attempt to explore this effect more fully, we conducted a set of experiments using the RSCP in which we varied the outside-offer rate, considering cases in which the number of outside offers were 10, 20, 30, 50, and 75 percent of the total number of tasks[3]. We tested two hypotheses: (1) as the number of outside offers increases, the *rate* of defaulting will stay the same because the offers will be drawn from the same range of values, but the absolute *number* of defaults will increase; and (2) as the number of offers increases, mean individual income (from tasks and offers) will increase because of additional income from a larger number of accepted offers, but group income (from tasks only) will decrease as more tasks go undone. The results are shown in Figure 5.

Both the default-rate and income results for these experiments are heavily influenced by the availability of replacement agents, a factor whose influence we did not anticipate in this context. As the outside-offer rate increases, the same rate of defaulting leads to a larger number of defaults. Thus, the available replacements are used up more frequently, and there are more cases in which no replacement can be found (Table 2). Agents are more reluctant to default without a replacement because the resulting group costs and score reductions are larger. Thus, as the offer rate increases, agents default at a lower rate (Fig. 5, left), contrary to our hypothesis, although the increasing number of offers means that the absolute number of defaults does still increase.

The income that individual agents receive (Fig. 5, *right*) is affected by group costs, which increase as replacements become more scarce and agents accept a larger number of outside offers, and by the income that agents

Table 2. Effect of outside-offer rate on the average percentage of outside offers for which no replacement agent is available under the RSCP.

Outside-offer rate	Offers with no replacement
10%	0.03%
20%	0.91%
30%	5.15%
50%	25.45%
75%	47.81%

3. In our earlier work [13, 39, 40], the number of outside offers was chosen randomly using a uniform distribution from 5/8 to 7/8 of the number of tasks. A 75-percent outside-offer rate is the mean of this distribution.

receive from outside offers, which also increases as agents accept more of them. For the three lowest offer rates, the offer income outweighs the group costs, and the mean individual income increases as predicted. Under offer rates of 50 or 75 percent, however, the large number of offers for which no replacement is available leads group costs to dominate, and thus the mean individual income decreases. For group task income, the extra income from offers is not included, and thus increasing group costs leads to the expected lower group incomes over the entire range of offer rates.

5.1.4 Heterogeneity in social consciousness under the RSCP

We next conducted two sets of experiments in which the environmental conditions were fixed and the composition of heterogeneous groups of agents was varied. We first considered scenarios in which some of the agents are socially conscious and some are not. 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 with the optimal homogeneous-group *BPweight* setting of 0.1 (the *BP agents*), and the rest of the agents do not use brownie points (the *no-BP agents*). We hypothesized that the BP agents would do better than the no-BP agents, avoiding the free-rider effect seen in our earlier work (see the introduction to Section 5). Figure 6 displays the results.

As expected, agents who use brownie points default less often than those who do not (Fig. 6, *left*). As the more responsible, BP agents become a larger percentage of the group, the no-BP agents demonstrate small decreases in defaulting while the BP agents maintain more or less steady rates of defaulting. Increasing the percentage of more responsible agents has two conflicting effects on the agents' estimated FEI losses, and thus on their likelihood to default. On the one hand, because an agent's estimate of its ranking in the following week is based on the behavior of the other agents (see Sect. 4.2), an increased percentage of more responsible collaborators will tend to lead agents to estimate larger drops in the rankings from defaulting and *larger* resulting losses in FEI. On the other hand, a larger percentage of more responsible collaborators results in an increased availability of replace-

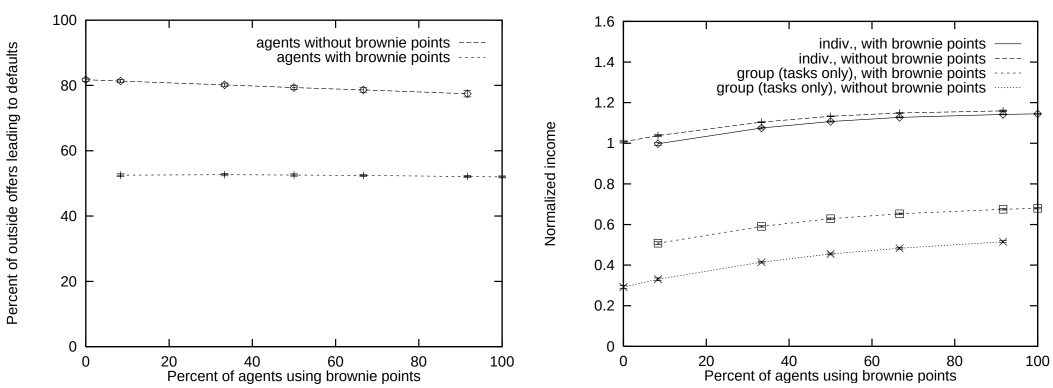

Figure 6. Outcomes of heterogeneous-group experiments in which some agents use brownie points and some do not under the RSCP. The rate of defaulting (*left*) and the normalized individual and group incomes (*right*) are shown for each subgroup as the percentage of agents that use brownie points is increased.

ment agents and thus to *smaller* estimated FEI losses. These conflicting forces have different strengths in the two subgroups, leading one to default less often and the other to default at a constant rate.

The two subgroups exhibit small but statistically significant differences in normalized individual income, and even more significant differences in group task income (Fig. 6, *right*). The agents who are not socially conscious do worse as individuals in homogeneous contexts (comparing the individual-income results at the 0 and 100 percent points), but they do better than the more responsible, socially conscious agents in heterogeneous groups. These less responsible, no-BP agents take advantage of their more responsible collaborators, who reduce the overall group costs by defaulting less often. The less responsible agents are thus able to reap the full benefit of the outside offers that they accept while shifting a portion of the resulting group costs on the more responsible agents. This disproves our hypothesis that the free-rider effect would be avoided under a lower outside-offer rate; see Section 6 for a discussion of why these experiments yield different results than the preliminary experiments that formed the basis of our hypothesis.

Although the no-BP agents do better as individuals in heterogeneous groups, their income from group-related tasks alone is much lower than that of the BP agents because they default on more tasks. Moreover, the overall group task income—which is effectively a weighted average of the two subgroup incomes—as well as the individual and group task incomes of both subgroups, all increase as the number of BP agents increases, showing that everyone benefits when more agents are socially conscious. Given that the less responsible agents do only slightly better as individuals than their more responsible collaborators, these results suggest that agent designers can improve group outcomes without sacrificing much in the way of individual gains by building agents with a moderate amount of social consciousness [13].

5.1.5 Heterogeneity in the weight given to FEI under the RSCP

In this set of experiments using the RSCP, we considered heterogeneous groups of agents who use different δ values to weight their estimates of F, their income in the following week (see eq. 4). 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 the optimal homogeneous-group δ value of 0.4 (the *higher-delta agents*), and the rest use a lower δ value of 0.3 (the *lower-delta agents*). Both subgroups also used the standard *BPweight* of 0.1. We hypothesized that the higher-delta agents would do better than the lower-delta agents, avoiding the free-rider effect seen in our earlier work (see the introduction to Section 5). Figure 7 displays the results.

As expected, agents who put a higher value on future expected income default less often than agents who do not value future income as highly (Fig. 7, *left*). As the size of the more responsible subgroup increases, members of both subgroups default at fairly constant rates, for reasons discussed in the previous section. These default-rate results differ markedly from the results of our earlier mixed-delta experiments [40]. In those experiments, the

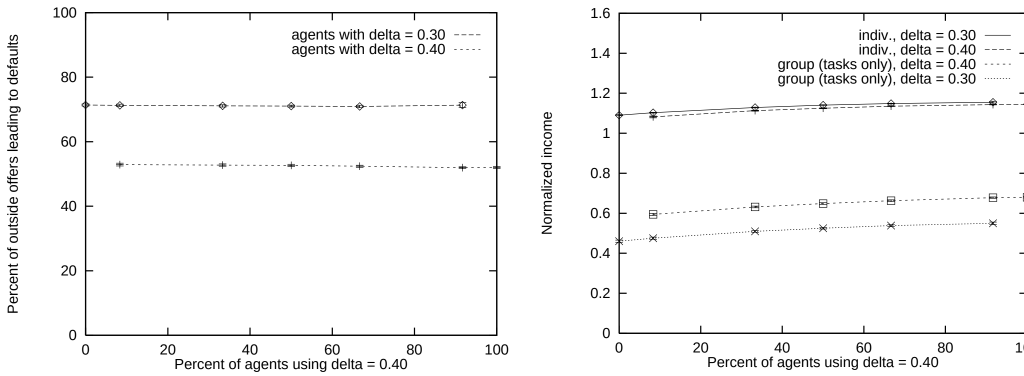

Figure 7. Outcomes of heterogeneous-group experiments in which agents use two different values of δ to weight future income under the RSCP. The rate of defaulting (*left*) and the normalized individual and group incomes (*right*) are shown for each subgroup as the percentage of agents placing a higher value on future income is increased.

more responsible agents almost never defaulted. As a result, the two agent subgroups were segregated in the rankings, with the less responsible agents always occupying the lowest rankings and thus having less incentive to behave responsibly. This effect was accentuated as the more responsible agents became a larger percentage of the group, and thus the lower-delta agents defaulted *more* often as the number of higher-delta agents increased. In the current mixed-delta experiments, the more responsible agents default often enough to prevent a strict segregation of the two subgroups in the rankings, and the less responsible agents thus avoid the cycle of increasing defaults that we saw in the earlier experiments.

The impact of changing subgroup sizes on both mean individual income and subgroup task income is shown in the right half of Figure 7. Once again, the less responsible agents do worse as individuals in homogeneous contexts (comparing the individual-income results at the 0 and 100 percent points) and better as individuals in heterogeneous groups. Like the no-BP agents, the lower-delta agents are able take advantage of the reduced group costs that their more responsible collaborators bring about. This further disproves our hypothesis that the free-rider effect would be avoided under a lower outside-offer rate; see Section 6 for a discussion of why these experiments yield different results than the preliminary experiments that formed the basis of our hypothesis. Both subgroups also see an increase in both individual and group task income as the percentage of more responsible agents increases and overall group costs decline.

5.2 DSCP experiments

We next present the results of experiments using the discount-based social-commitment policy or DSCP (Sect. 3.2). The incomes that we present are normalized values obtained by dividing each agent's income by the undiscounted value of the tasks that it was assigned. Because tasks are assigned randomly under the DSCP, this normalization factor is independent of the agent's behavior and thus will not artificially inflate the incomes of less responsible subgroups, as it would under the RSCP.

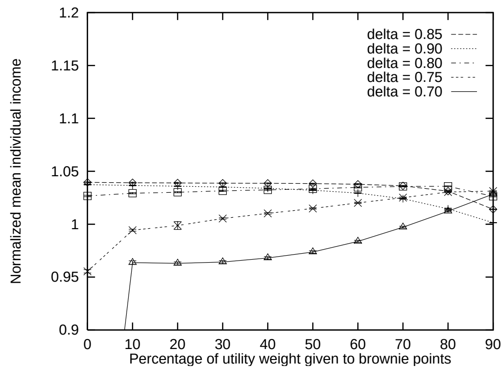

Figure 8. The impact of δ and *BPweight* on the normalized mean individual income of homogeneous groups of agents under the DSCP. To make the distinctions between the curves clearer, we have focused on a portion of the y-axis from [0.9:1.2].

5.2.1 Determining the baseline utility-function parameters for the DSCP

As in the RSCP experiments (Sect. 5.1.1), we began by conducting experiments to determine the baseline settings for δ—the parameter that agents use to compute FEI (Sect. 4.1)—and for *BPweight*—the parameter that agents use to weight brownie points (Sect. 4.5). We once again tested many different pairs of values for δ and *BPweight* to find the combination that maximizes mean individual income for homogeneous groups; Figure 8 displays the results of some of these experiments.

The tests show that agents operating under the DSCP maximize individual income when they use a δ value of 0.85 and no brownie points (*BPweight* = 0). This result confirms the hypothesis that agents operating in homogeneous groups do better when they give a moderate weight to future income, but it contradicts the hypothesis that such agents also do better with an intermediate amount of social consciousness (see the introduction to Section 5). Brownie points do not appear to offer the same benefit to agents operating under the DSCP as they do to agents operating under the RSCP, although they might still be useful to groups that use the DSCP under different environmental conditions.

Another important difference between the optimal RSCP and DSCP parameters is that the optimal δ value is higher under the DSCP, meaning that DSCP agents do better when they default less often. Agents cannot afford to default as often under the DSCP because there are three different possible losses from defaulting: (1) the group costs of defaulting, (2) income lost from tasks that are not completed because no replacement agent is available, and (3) income lost to discounting. In the RSCP, only the first two losses apply, allowing RSCP agents to default more often.

Unless otherwise stated, the settings of δ = 0.85 and *BPweight* = 0 were used throughout the experiments in this section. Because the optimal homogeneous-group results are obtained without brownie points, we did not conduct experiments involving heterogeneity in social consciousness like the ones discussed in Section 5.1.4.

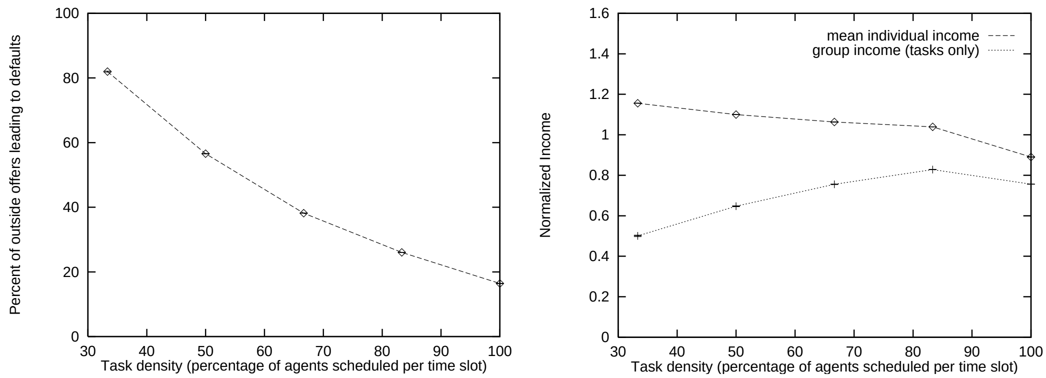

Figure 9. The impact of task density on the rate of defaulting (*left*) and on normalized individual and group incomes (*right*) when using the DSCP.

5.2.2 Behavior of the DSCP under different task densities

In the next set of DSCP experiments, we varied task density, which can be expressed as the percentage of agents who are scheduled to perform a task in each time slot. We maintained a constant percentage of discounted tasks as we increased task density. As in the RSCP version of these experiments (Sect. 5.1.2), we tested the following hypotheses: (1) agents will default less often as task density is increased, provided that the percentage of tasks affected by the social-commitment policy is held constant; and (2) increasing task density will lead to lower individual and group incomes. Figure 9 shows the results, which mirror the trends seen in the corresponding RSCP experiments.

As task density increases, the decreasing availability of replacements and the increasing number of discounted tasks leads agents to estimate larger losses from defaulting, and thus to default less often (Fig. 9, *left*). In addition, agents do worse as individuals as task density increases, while their group task income first rises and then falls (Fig. 9, *right*); the explanation of the RSCP task-density results (Sect. 5.1.2) and how they relate to our hypotheses also applies here.

5.2.3 Behavior of the DSCP under different outside-offer rates

The next set of DSCP experiments varied the outside-offer rate, considering cases in which the number of outside offers were 10, 20, 30, 50, and 75 percent of the total number of tasks. As in the RSCP version of these experiments (Section 5.1.3), we tested two hypotheses: (1) as the number of outside offers increases, the *rate* of defaulting will stay the same, but the absolute *number* of defaults will increase; and (2) as the number of offers increases, mean individual income will increase, but group income will decrease. The results are shown in Figure 10.

As we saw under the RSCP, agents default at a reduced rate as the outside-offer rate increases (Fig. 10, *left*), contradicting our first hypothesis, while the absolute *number* of defaults does increase slightly. The decreasing

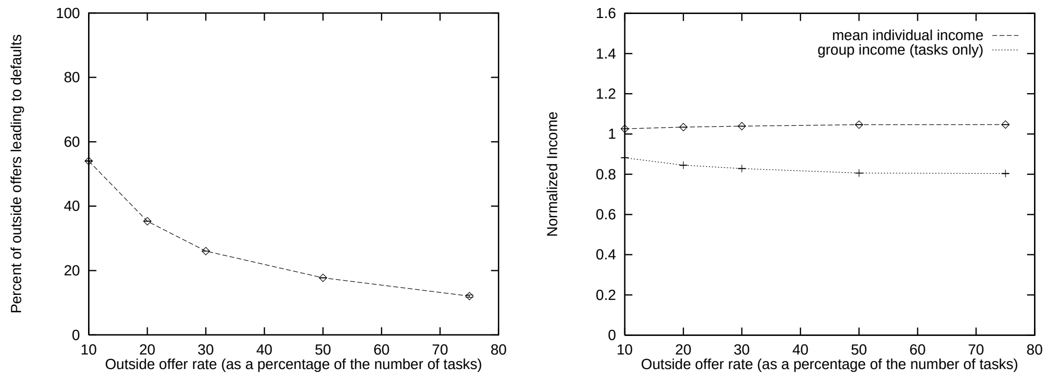

Figure 10. The impact of the outside-offer rate on the rate of defaulting (*left*) and on normalized individual and group incomes (*right*) when using the DSCP.

availability of replacement agents again plays a role (Table 3), but this factor is less significant than it was under the RSCP, because the lower overall rates of defaulting mean that the supply of replacement agents is exhausted less frequently. A more important factor involves the way in which an agent's tasks are discounted (see Sect. 3.2). An agent's reputation score, and thus its ISF discount factor, is affected by the *number* of times that it defaults, not by its rate of defaulting. Thus, to avoid increased estimated losses in FEI as the number of outside offers increases, an agent must default at a lower rate.

The income that individual agents receive (Fig. 10, *right*) is affected by three factors: (1) group costs, which increase as replacements become more scarce; (2) income that agents receive from outside offers, which increases as agents default on a larger number of tasks; and (3) discounts on task income, which increase as agents default on more tasks; these increased discounts are nonlinear because of the sigmoidal shape of the ISF function (Fig. 1, *left*). Across the range of offer rates, the extra offer income outweighs the other two factors, and the mean individual income increases slightly. For group task income, income from offers is not included, and the other two factors thus lead to lower group incomes. Both the increase in mean individual income and the decrease in group task income confirm our second hypothesis.

5.2.4 Heterogeneity in the weight given to FEI under the DSCP

In this set of experiments using the DSCP, we fixed environmental factors and considered heterogeneous groups of agents who use different δ values to weight FEI. We considered cases in which none, 1/12, 4/12, 6/12,

Table 3. Effect of outside-offer rate on the average percentage of outside offers for which no replacement agent is available under the DSCP.

Outside-offer rate	Offers with no replacement
10%	0.16%
20%	0.94%
30%	1.98%
50%	3.79%
75%	5.56%

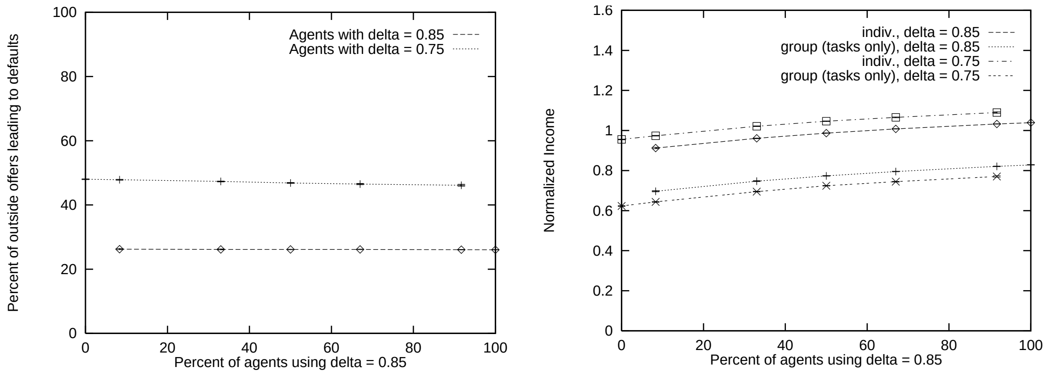

Figure 11. Outcomes of heterogeneous-group experiments in which agents use two different values of δ to weight future income under the DSCP. The rate of defaulting (*left*) and the normalized individual and group incomes (*right*) are shown for each subgroup as the percentage of agents placing a higher value on future income is increased.

8/12, 11/12, and all of the agents use the optimal homogeneous-group δ setting of 0.85 (the *higher-delta agents*), and the rest use δ = 0.75 (the *lower-delta agents*). As in the RSCP version of these experiments (Sect. 5.1.5), we hypothesized that the higher-delta agents would do better than the lower-delta agents, avoiding the free-rider effect seen in our earlier work. Figure 11 displays the results.

Agents who put a higher value on future expected income default less often than agents who do not value future income as much (Fig. 11, *left*), mirroring the corresponding RSCP results. As the percentage of more responsible agents increases, both subgroups default somewhat less often. The presence of a larger number of more responsible, higher-delta agents leads to slightly lower group discounts (i.e., the GSF factor is larger). As a result, the discounted tasks are worth more (making the current-income gains from outside offers less substantial), and agents estimate larger FEI losses. These two factors make all agents less likely to default.

Here again, the less responsible, lower-delta agents do worse as individuals in homogeneous contexts (comparing the individual-income results at the 0 and 100 percent points) and better as individuals in heterogeneous groups (Fig. 11, *right*). Like the lower-delta agents in the RSCP version of these experiments, these lower-delta agents are able to take advantage of the reduced group costs that their more responsible collaborators bring about, providing another example of the free-rider effect and contradicting our hypothesis; see Section 6 for a discussion of why these experiments yield different results than the preliminary experiments that formed the basis of our hypothesis. The presence of an increasing percentage of higher-delta agents also leads to lower group costs and better group reputations, and thus to higher incomes for both subgroups.

6 Discussion

The two social-commitment policies provide different models for how agent societies might constrain the decision making of their members. One key difference between the two policies involves the accuracy with which agents can estimate their future losses from defaulting. Because the income of agents operating under the DSCP does not depend on their ranking with respect to other agents, they are able to make more accurate estimates of their future income. The RSCP, on the other hand, avoids the need for monetary punishments and the concomitant issue of what should be done with the income that agents lose to discounting. These differences are reflected in the results, as explained in more detail below.

As two different environmental factors—task density (Sect. 5.1.2 and 5.2.2) and outside-offer rate (Sect. 5.1.3 and 5.2.3)—are increased, agents under both policies decrease their rates of defaulting. However, the above-mentioned differences in predictive accuracy play a role in these results, as the DSCP agents, by reducing their rates of defaulting more sharply, incur smaller drops in individual and group task income than the RSCP agents. This confirms our hypothesis that agents operating under the DSCP would respond better to changes in the environment. On the other hand, it is frequently the case that the RSCP agents have higher individual and group incomes than the DSCP agents, because they avoid the losses in income from discounting that the DSCP agents incur. This contradicts our hypothesis that DSCP agents would achieve higher individual and group outcomes; in retrospect, we failed to fully anticipate the impact of the income losses from discounting when we formulated the hypothesis.

The accuracy of the future-income estimates under both social-commitment policies could be improved if estimates of future income from outside offers and losses from group costs were taken into account. Because agents fail to consider future offers, they need to use fairly low values of δ to achieve optimal rates of defaulting. To include future outside offers in their estimates, agents would need to learn about this aspect of their environment. More generally, because the optimal δ and *BPweight* values depend on factors like the outside-offer rate and task density, agents would benefit greatly from an ability to adapt their rates of defaulting in response to changing environmental conditions.

The impact of social consciousness in the form of brownie points seems mixed. For a given environment, both brownie points and the weight given to FEI by δ affect the rate at which agents default and thus the income that they receive. Under the RSCP, both δ and brownie points are needed to optimize the individual income of homogeneous groups operating in our baseline environment (Fig. 3). Under the DSCP, however, homogeneous groups of agents can optimize individual income without using brownie points by choosing a δ value of 0.85 (Fig. 8). However, the optimal choice of these parameters is highly dependent on the environment, and there may be situations in which brownie points could also be used to improve individual outcomes under the DSCP. In addition, the use of brownie points can often improve group outcomes without large sacrifices in individual gains (see Sect. 5.1.4).

Both social-commitment policies are susceptible to the free-rider effect in contexts in which different agents use different values of δ or *BPweight* (Sect. 5.1.4, 5.1.5, and 5.2.4). While we had seen this effect in earlier work that used high outside-offer rates, preliminary experiments led us to hypothesize that it would be avoided under a moderate offer rate [40]. However, these preliminary experiments were conducted using baseline settings for δ

and *BPweight* that were chosen to work well under the high offer rates. When the offer rate was lowered, the optimal parameter settings shifted, and the settings used by the more responsible agents happened to do better under these new conditions. In the work reported in this paper, the parameter settings used by the more responsible agents were optimized for use under the lower offer rate. By deviating only slightly from these optimal homogeneous-group settings, the less responsible agents do better by defaulting more, because their more responsible collaborators keep the group costs lower than they would be if all of the agents used the non-optimal settings.

In all three examples of the free-rider effect in this work, the less responsible agents do worse as individuals in homogeneous contexts (comparing the 0 and 100 percent points in the right-hand graphs in Figures 6, 7, and 11). This phenomenon is similar to the prisoner's dilemma from game theory [2], in which a strategy of defecting strictly dominates a strategy of cooperating, even though each agent's outcome is worse when everyone defects than when everyone cooperates. In situations like ours in which agents interact repeatedly, cooperating can become a stable strategy if each agent believes that defecting will terminate the interaction, resulting in a long-term loss that outweighs the short-term gain from defecting [29, 35]. The two social-commitment policies could in theory be changed to exact such a punishment when agents default more than a certain amount.

Although the results of our heterogeneous-group experiments resemble the prisoner's dilemma, it is important to note that the less responsible agents outperform their more responsible collaborators by only small amounts. In addition, these small individual gains are accompanied by significantly lower group outcomes. Thus, agent designers may do well to pursue *epsilon-equilibria* [29]. For any epsilon greater than 0, an epsilon-equilibrium is a profile of strategies with the property that no agent has an alternative strategy that increases its payoff by more than epsilon. Applying this concept to our experiments would suggest that agent designers should chose the δ and *BPweight* values that lead to optimal outcomes in homogeneous contexts.

In addition, we believe that our heterogeneous-group results are largely a result of the fact that group costs are divided equally among the agents, while added income from an outside offer goes only to the defaulting agent. Given this imbalance, less responsible agents should always do better in heterogeneous contexts in which the more responsible agents use the parameter settings that are optimal for homogeneous groups. If group costs were not divided equally, the heterogeneous-group results might well be different. In the extreme case, the entire group cost could be borne by the defaulting agent; this would effectively prevent agents from defaulting. While this policy would increase group income, it would also decrease individual income significantly. Finding a balanced policy would require information about outside offers that may not be available when the policy is designed.

7 Related Work

The problem of intention reconciliation is relevant to Grosz et al.'s SharedPlans-based architecture for collaborative agents [17] and to the systems that have been built using this architecture. One system, called GigAgents, is designed to support joint human-computer collaborations [17]. The second system, called WebTrader [18], was developed to support the cooperative processes of buying and selling goods on the Web. WebTraders can sell their enterprise's goods to human or automated buyers, and they can buy items needed by their enterprises. SPIRE is general enough to be able to model agents from these and other domains in which intention reconciliation in collaborative contexts is needed.

Other researchers in multi-agent systems have examined similar decision-making scenarios in their work. Kalenka and Jennings [23] 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 three 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; and (3) they do not allow agents to default on their commitments. Kalenka and Jennings find that giving agents a degree of social consciousness tends to improve overall outcomes, but that it may cause individual performance to suffer in some cases; these results are similar to our own (see Sect. 5.1.1, Sect. 5.1.4, and Sect. 5.2.1).

Sen [36] also proposes 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, and, like Kalenka and Jennings, he considers decisions about whether to cooperate in the first place, not decisions about whether to default on existing commitments to collaborate. Sen's work also demonstrates the benefits of social consciousness, but he shows that less responsible agents may take advantage of their more responsible counterparts in some scenarios, a phenomenon that we have also seen (see Sect. 5.1.4, Sect. 5.1.5, and Sect. 5.2.4). However, Sen's results also demonstrate that the more responsible agents can do better in the long run by basing their decisions about whether to cooperate on an agent's past behavior.

Xuan and Lesser [44] present a negotiation framework that takes into account uncertainty about whether agents will honor their commitments. Similar to Sen, they focus on interactions between pairs of individual agents, and they also assume that the details of a commitment can be modified in a way that is satisfactory for both agents, whereas our work addresses situations in which this assumption fails to hold.

Hogg and Jennings [19] experiment with agents who make decisions based on a weighted sum of their own utility functions and those of their fellow team members, and they demonstrate the benefit of this type of social consciousness in certain resource-bounded contexts. The brownie-points model provides an alternative method of incorporating social consciousness that does not require agents to estimate the utility functions of their collaborators, and it is explicitly geared to the problem of intention reconciliation. Like Hogg and Jennings, we find that agents who have an intermediate level of social consciousness do better as individuals in at least some environments (see Sect. 5.1.1). However, Hogg and Jennings find that agents with a moderate amount of social consciousness also do better in the presence of purely selfish agents, while our experiments combining agents who use brownie points with those who do not reveal a small but statistically significant free-rider effect (see Sect. 5.1.4). Hogg and Jennings also consider the use of meta-level reasoning to control how much social reasoning the agents perform based on information about the state of the environment, as well as the use of learning to allow agents to improve their decision making over time.

The social-commitment policies that we have defined in our work differ from the “social laws” used in other multi-agent planning work [38]. Social laws provide constraints on agents that allow their actions to be coordinated; these laws constrain the ways agents *do actions* so that their activities do not negatively interact. In contrast, social-commitment policies concern *rational choice* and the ways in which a society can influence an individual’s decision making. As a result, social laws are by their nature domain-specific, whereas social-commitment policies affect decision making across domains and tasks. Rosenschein and Zlotkin [31] have presented mechanisms similar to social-commitment policies in the context of negotiation between agents.

Sandholm et al. [32, 1, 33, 34] propose a mechanism built into contracts that allows agents to renege on their commitments by paying a predetermined cost known as a decommitment penalty. A decommitment penalty is another example of a social-commitment policy, one that is based solely on a given instance of defaulting. While SPIRE experiments could be run using this type of social-commitment policy, we have focused on longer-term policies that are based on an agent’s history of defaulting, and we have considered the use of internal measures of social consciousness in conjunction with externally imposed social-commitment policies. In addition, Sandholm et al. focus on contracts involving two or three non-collaborative agents[4], while we study the problem of intention reconciliation in the context of an ongoing collaboration of a large set of agents. Sandholm et al. find that allowing agents to decommit improves outcomes, a conclusion that matches our own results showing that agents do better when they are allowed to default (by using a *BPweight* value less than 1; see Sect. 5.1.1 and 5.2.1). Sandholm et al. are able to derive algorithms for optimizing contracts involving two or three agents [23], while SPIRE allows you to experimentally determine the optimal utility-function parameters for large groups of agents operating in a given environment (see Sect. 5.1.1 and Sect. 5.2.1).

Teague and Sonenberg [43] compare decommitment penalties with a measure of social consciousness that is based on brownie points but uses a slightly different set of formulas than the ones defined by Glass and Grosz [7]. Unlike the scenarios that we have examined, which focus on whether agents should maintain their team commitments in the face of subsequent offers, Teague and Sonenberg consider situations in which agents receive two rounds of offered tasks (some of which require collaboration) and must decide whether to accept an offer in the first round, as well as whether to default in the second round for the sake of a new, more valuable offer. In this

4. Sandholm et al. do consider scenarios involving many agents [1], but the contracts—and thus the decommitments—are still between two agents. When they consider contracts involving three agents [33], one of the agents is the contractor, and the other two agents are contractees who do not work as a team.

type of scenario, they find that both brownie points and decommitment penalties constrain defaulting, but that decommitment penalties generally lead to lower group outcomes because they can discourage agents from committing to collaborative tasks in the first round when they anticipate more valuable, second-round tasks. The social-commitment policies that we have proposed could also lead agents to reject initial task offers, but we have not focused on the type of scenario in which this effect could emerge. Teague and Sonenberg also propose an alternative mechanism in which agents maintain estimates of the reliability of other agents (an external parallel to the internal brownie-point measure) and use these estimates when deciding whether to collaborate.

There is a significant body of economics literature on rational choice and intention reconciliation. Iannaccone [22] examines social policies that alter individual utility functions to encourage group commitment. While these policies are similar in spirit to social-commitment policies, they are aimed at group formation, not at conflicting intentions. His approach is also completely penalty-based, and is not applicable to agents that face multiple decision points over time. Holländer [20] 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.

8 Conclusions

The SPIRE empirical framework has enabled us to simulate and study the process of intention reconciliation in collaborative contexts, examining the impact of environmental factors, social norms, and agent utility functions on individual and group outcomes in both homogeneous and heterogeneous groups. In this paper, we presented the results of experiments involving two different types of social norms. Agents operating under both sets of norms responded similarly to varying environmental factors, but the choice of social norms influenced the effectiveness of social consciousness and other aspects of the agents’ utility functions. Both sets of norms were shown to be susceptible to the free-rider effect, but the gains of the less responsible agents were minimal. In fact, our results suggest that agent designers can achieve epsilon-equilibria by choosing utility-function parameters that lead to optimal outcomes in homogeneous contexts. We plan to investigate additional social-commitment policies to determine if this agent-design principle holds more generally.

In other future work with SPIRE, we will experiment with policies in which group costs are not divided equally. As discussed in Section 6, we suspect that giving a larger share of these costs to the defaulting agent may allow agents who are more responsible to fare better in heterogeneous contexts. In addition, we are interested in exploring situations in which agents have more information about their collaborators (e.g., knowledge of how responsible they tend to be), and how this information can influence the agents’ decisions about whether to default. We also plan to explore the related possibility of agents modeling and adapting to the behavior of other agents. Das [10] has conducted preliminary investigations in this area, showing that...

The experiments that we have conducted with SPIRE have made a number of simplifying assumptions, including that: (1) all agents are initially available at all times; (2) all agents have the ability to perform all task types; (3) each of the tasks in the group activity can be performed by an individual agent; and (4) each task can be performed in a single time slot. We plan to investigate scenarios that eliminate these simplifications. SPIRE is already able to handle agents with differing availabilities and capabilities, although some scheduling issues may still need to be addressed. More substantial modifications to the current system will be needed to accommodate tasks that require multiple agents or span multiple time slots.

We are also interested in further investigating the interaction of social-commitment policies with the social consciousness provided by brownie points. In our experiments using the DSCP, brownie points were not needed to maximize outcomes. We suspect that this result may stem from the more accurate estimates of future income that are possible under the DSCP, but more experiments are needed to confirm this intuition. Regardless of what these experiments reveal, brownie points will still be useful to agents who are unable to make accurate FEI estimates, as well as in agent societies that track the reliability of other agents, because the impact of defaulting in such contexts will extend beyond the current collaboration.

Because intention reconciliation in realistic multi-agent contexts is an extremely complex problem, a system like SPIRE is essential for obtaining the insights needed to design collaboration-capable agents [17]. 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] Andersson, M.R. and Sandholm, T.W. 1998. Leveled commitment contracting among myopic individually rational agents. In: *Proc. of ICMAS-98*, pp. 26-33.

[2] Axelrod, R. 1984. *The Evolution of Cooperation*. Basic Books, New York, NY.

[3] Axelrod, R. and Dion, D. 1988. The further evolution of cooperation. *Science*, 242:1385-1390.

[4] Axelrod, R. and D'Ambrosio, L. 1994. Annotated bibliography on the evolution of cooperation. Unpublished.

[5] Benoit, J.-P. and Krishna, V. 1998. The folk theorems for repeated games: a synthesis. Unpublished.

[6] Boutilier, C. 1999. Multiagent systems: challenges and opportunities for decision-theoretic planning. *AI Magazine*, 20(4):35-43.

[7] Bratman, M.E. 1987. *Intention, Plans, and Practical Reason*. Harvard University Press, Cambridge, MA.

[8] Bratman, M.E., Israel, D., and Pollack, M.E. 1988. Plans and resource-bounded practical reasoning. *Computational Intelligence,* 4(4):349-355.

[9] Castelfranchi, C. 1998. Modeling social action for AI agents. *Artificial Intelligence*, 103:157-182.

[10] Das, S.

[11] Decker, K. and Li, J. 1998. Coordinated hospital patient scheduling. In: *Proc. of ICMAS-98*, pp. 104-111.

[12] Doyle, J. 1991. Rational belief revision. In: *Proc. of the Second Intl. Conference on Knowledge Representation and Reasoning*, pp. 163-174.

[13] Glass, A. and Grosz, B.J. 1999. Socially conscious decision-making. In: *Proc. of Agents 2000*, pp. 217-224

[14] Goldman, C. and Kraus, S. 2001. The value of temptation. Working paper, Bar-Ilan University.

[15] Grosz, B.J. and Kraus, S. 1996. Collaborative plans for complex group action. *Artificial Intelligence*, 86(2):269-357.

[16] Grosz, B.J. and Kraus, S. 1999. The Evolution of SharedPlans. In Wooldridge, M. and Rao, A., editors, *Foundations and Theories of Rational Agency*. Kluwer Academic Publishers, The Netherlands, pp. 227-262.

[17] Grosz, B.J., Hunsberger, L., and Kraus, S. 1999. Planning and acting together. *AI Magazine*, 20(4):23-34.

[18] Hadad, M. and Kraus, S. 1999. SharedPlans in electronic commerce. In Klusch, M., ed., *Intelligent Information Agents*, pp. 204-231. Springer Verlag.

[19] Hogg, L. and Jennings, N.R. 2001. Socially intelligent reasoning for autonomous agents. *IEEE Transactions on Systems, Man, and Cybernetics — Part A*, 31(5):381-399.

[20] Holländer, H. 1990. A social exchange approach to voluntary cooperation. *American Economic Review*, 80(5):1157-1167.

[21] 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.

[22] Iannaccone, L.R. 1992. Sacrifice and stigma: reducing free-riding in cults, communes, and other collectives. *Journal of Political Economy*, 100(2):271-291.

[23] 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.

[24] 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.

[25] Kraus, S., Wilkenfeld, J., and Zlotkin, G. 1995. Multiagent negotiation under time constraints. *Artificial Intelligence*, 75(2):297-345.

[26] Kreps, D. and Wilson, R. 1982. Reputation and imperfect information. *Journal of Economic Theory*, 27:253-279.

[27] Levesque, H., Cohen, P. and Nunes, J. 1990. On acting together. In: *Proc. of AAAI-90*, pp. 94-99.

[28] Luce, R.D. and Raiffa, H. 1957. *Games and Decisions*. John Wiley and Sons, New York, NY.

[29] Osborne, J. and Rubinstein, A. 1994. *A Course in Game Theory*. MIT Press, Cambridge, MA.

[30] Pollack, M.E. and Horty, J.F. 1999. There's more to life than making plans: plan management in dynamic, multi-agent environments. *AI Magazine*, 20(4):71-84.

[31] Rosenschein, J.S. and Zlotkin, G. 1994. *Rules of Encounter: Designing Conventions for Automated Negotiation among Computers*. MIT Press, Cambridge, MA.

[32] Sandholm, T.W. and Lesser, V.R. 1996. Advantages of a leveled commitment contracting protocol. In: *Proc. of AAAI-96*, pp. 126-133.

[33] Sandholm, T.W., Sikka, S., and Norden, S. 1999. Algorithms for optimizing leveled commitment contracts. In *Proc. of IJCAI-99*, pp. 535-540.

[34] Sandholm, T. and Zhou, Y. 2000. Surplus equivalence of leveled commitment contracts. In: *Proc. of ICMAS-00*, pp. 247-254.

[35] Selten, R. and Stoecke, R. 1986. End behavior in sequences of finite prisoner's dilemma supergames. In: *Journal of Economic Behavior and Organization*, 7:47-70.

[36] Sen, S. 1996. Reciprocity: a foundational principle for promoting cooperative behavior among self-interested agents. In: *Proc. of ICMAS-96*, pp. 322-329.

[37] Sen, S., Haynes, T., and Arora, N. 1997. Satisfying user preferences while negotiating meetings. *Intl. Journal on Human-Computer Studies*, 47(3):407-427.

[38] Shoham, Y. and Tennenholtz, M. 1992. On the synthesis of useful social laws for artificial agent societies. In: *Proc. of AAAI-92*, pp. 276-281.

[39] 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.

[40] Sullivan, D., Grosz, B., and Kraus, S. 2000. Intention reconciliation by collaborative agents. In: *Proc. of ICMAS-00*, pp. 293-300.

[41] Sycara, J. and Zeng, D. 1996. Coordination of multiple intelligent software agents. *Intl. Journal of Intelligent and Cooperative Information Systems*, 5:181-211.

[42] Tambe, M. 1997. Towards flexible teamwork. *Journal of Artificial Intelligence Research*, 7: 83-124.

[43] Teague, V. and Sonenberg, L. 2001. Investigating commitment flexibility in multi-agent contracts. In Parsons, S., Gymtrasiewicz, P., and Wooldridge, M.J., eds., *Game Theory and Decision Theory in Agent-Based Systems*, Kluwer.

[44] Xuan, P. and Lesser, V. 1999. Incorporating uncertainty in agent commitments. In: *Proc. of ATAL-99*, pp. 221-234.

[45] Yoon, K.P. and Hwang, C.-L. 1995. *Multiple Attribute Decision-Making: An Introduction.* Sage University Papers, Series on Quantitative Applications in the Social Sciences, vol. 104.