The Influence of Social Norms and Social Consciousness on Intention Reconciliation

Presented by Ian Cottingham and
Gabe Williamson

Citation

Grosz, B. J., S. Kraus, D. G. Sullivan, and S. Das (2002). The Influence of Social Norms and Social Consciousness on Intention Reconciliation, *Artificial Intelligence*, 142:147-177.

Motivation

- In the context of collaborative work, agents must be able to take on new tasks as the state space changes:
  - situations may arise where task swapping to team-related actions may conflict with other possibilities.
- Traditional agents have not way to reconcile conflicting intentions:
  - if an agent is preforming some task A and is presented with the option to preform some task B that precludes completing A, how can the agent make a reasonable decision;
  - various strategies could have a net-negative impact on local decision making.
- This kind of decision making by agents is necessary to model many "real world" simulations.

Social Consciousness

- Decision making can be influenced by domain-independent rules of behavior (norms) that control agent behavior in a group:
  - these rules of behavior can be exploited to allow agents to individually value decision outcomes in the broader group context.
- Socially conscious agents are still autonomous but are able to consider broader context when evaluating action payoffs:
  - rewards or penalties may be immediate or accumulate over time.
- Policies can be employed to motivate the agent both internally - increased individual reward - and externally - increased group reward.

Intention Reconciliation

- Resolving conflicting intentions requires that agents be able to weigh individual actions against global concerns.
  - Social consciousness can provide a mechanism for agents to make such decisions.
- Using social consciousness to resolve intention conflicts allows agent behavior to be studied over repeated interactions;
  - provides for a study of more realistic scenarios than have been previously studied when considering reconciling agent intentions.
- Social consciousness is used in this paper to study and simulate intention reconciliation in collaborative multi-agent systems.

SPIRE Framework

- Shared Plans Intention Reconciliation Experiments (SPIRE):
  - framework for modeling situations of agents working in teams to accomplish shared goals;
  - allows study of environmental factors, social commitment policies, utility functions on group outcomes;
  - used for experiments in this paper.
- Describes scenarios similar to those in economics and game-theory research:
  - many assumptions made by prior work not applicable as SPIRE agents consider different actions each iteration;
  - utility and task assignments can depend on other agents and past actions;
  - agents have imperfect knowledge when making decisions.

SPIRE Framework

- Tasks performed by individual agents over one time unit.
- Simulations divided into 'weeks' with a group of agents being assigned the same group activity each week;
  - different agents may work on different individual tasks;
  - tasks assignments are centrally scheduled - not the focus of study.
- Randomly selected agents chosen for a chance to perform outside tasks that conflict with group tasks:
  - choosing an outside offer results in defaulting on the group task;
  - if the task can be done by another agent in the time slot, it is reassigned;
  - replacement agent receives additional reward.
- Team incurs a cost when a agent defaults, cost is greater if no replacement is found (models real world).

Ranking-based Social Commitment

- Ranks agents based on past behavior by scoring them at the end of each round
  - $s(w+1) = \alpha s(w) - \rho 1d - \rho 2D$
  - α = decay factor, $\rho 1$ and $\rho 2$ are constants for defaulting with / without replacement, d = number of defaults
- Score reflects the number of times an agent defaults and the impact of a default diminishes over time;
  - score is decremented more for tasks that are not replaced.
- Scheduler will give N tasks per agent based in its score:
  - higher scoring agents will get higher scoring tasks;
  - round-robin assignment for equal scores.
- Difficult to estimate future income of agents due to dependency on actions.

Discount-based Social Commitment

- Uses agent reputation to discount the income an agent will receive for completing its portion of the task:
  - independent of the reputation of other agents;
  - allows the agent to estimate future income more accurately.
- Reputation is represented as a score that increases as the agent's reputation decreases.
- Score is computed using previous equation with negative values of ρ1 and ρ2:
  - group reputation is similarly maintained.
- All tasks randomly assigned with N tasks per agent randomly discounted based on the reputation of the agent and the group:
  - product of scaling scores to [0..1]

Decision Making in SPIRE

- Current Income (CI)
  - income(Task)
  - income(Outside Offer) - cost(Defaulting)
- Future Expected Income (FEI)
  - "estimate, based on the social-commitment policy and the agents score".
  - Week 1 estimated
  - Weeks 2-N estimated based on Week 1*discount factor

$$FEI(F) = \delta F + \delta^2 F + \ldots = (\frac{\delta}{1-\delta})F$$

RSCP FEI

- To Default, or Not to Default, that is the question.
- Estimate score (ranking) to estimate value of future assigned jobs
- Need to know other agents...
  - attitudes
  - offers
  - rankings
  - learning capabilities
- That's Hard! Let's simplify
  - default history + rankings -> relative performance & number of defaults

RSCP FEI 2

- Simplification (cont)
  - 4 agent equivalence classes
    - 1) Above/Default
      2) Below/Default
      3) Above/Stay
      4) Below/Stay
  - 3 possible results
    - Stays -> 2
    - defaults
      - replacement available -> between 2,3
      - none available -> 3

RSCP FEI 3

- Simplification (cont 2)
  - one-week income loss depends on 5 factors
    - Ranking last-week
    - Ranking this-week
    - Default last-week
    - Replacement Availble this-week
    - Total # Agents

DSCP FEI

- Expected value of it's tasks in each of 2 cases
  - Default/No Default
- Agents are able to infer GSF values, so they're provided
- Individual scores -> ISF values
- Multiplied by expected values and estimated GSF values

Brownie Points

- Agents private valuation of it's group reputation
- Not a social commitment policy
- Non-Monetary
- Two Cases
  - Default -> lose BP
  - Stay -> gain BP
- Gain/Loss is dependent on task value

Combining Factors

- CI+FEI=TEI (Total Estimated Income)
- TEI and BP are normalized to allow for comparison
  - divided by max(TEI_{def}, $TEI_{no\text{-}def}$)

$$U_{def} = (1 - BPweight) \times normTEI_{def} + BPweight \times normBP_{def}$$

$$U_{no-def} = (1 - BPweight) \times normTEI_{no-def} + BPweight \times normBP_{no-def}$$

Experimental Results

- Grouped by RSCP and DSCP
- examines the policy and the various possible inputs
- Homogeneous v. Heterogeneous agents

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

Baseline params: RSCP

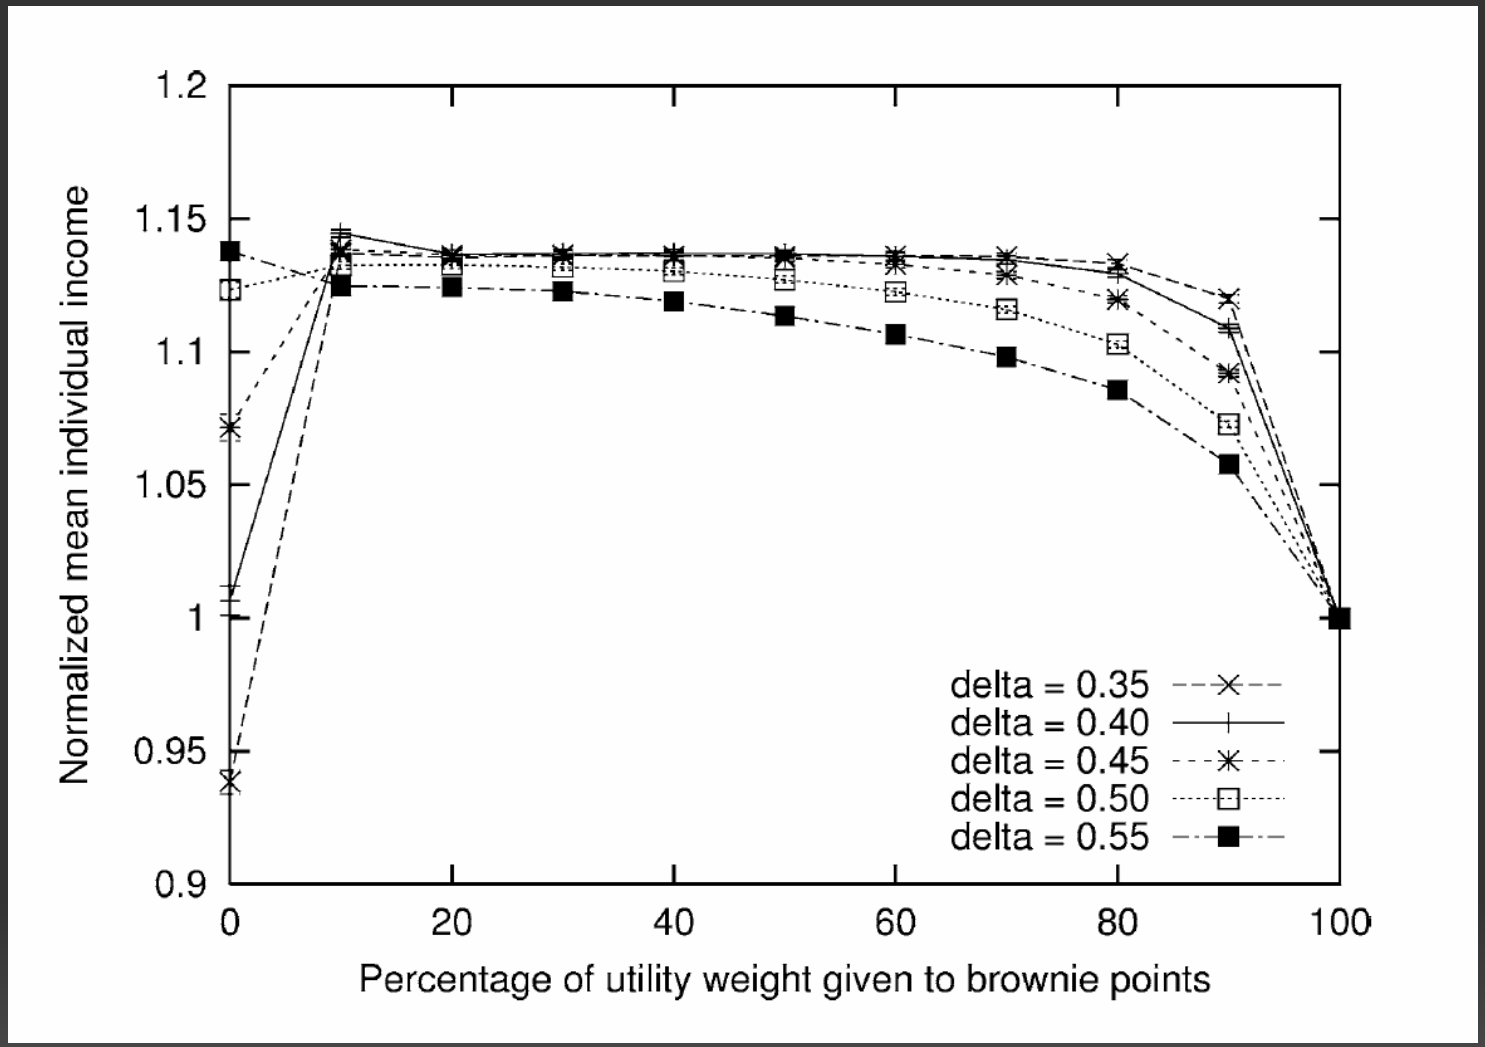

delta = 0.40
BPweight = 0.10
"Optimal mean individual incomes under 30% outside offer rate"

RSCP Task Density Overview

- Task Density: Tasks / Slot (% agents scheduled / slot)
- Effects
  - Replacement Availability
  - % of tasks affected by the RSCP -> possible task values
- % of score-assigned tasks is held constant as density is increased
- Hypotheses
  - agents will default less often as density is increased, provided that % tasks affected by RSCP is held constant
  - increasing task density -> lower individual and group incomes

RSCP: different outside-offer rates

- Varied outside offer rates
  - [10..20..70]% of total tasks
- Hypothesis
  - # OO increases -> default rate will be constant, total # defaults increases
  - # OO increases -> Individual Income increases, Group income decreases

Heterogeneity in Social Consciousness

- Differing BPweights (social consciousness factor)
- BPweight = 0.1 v BPweight = 0.0
- Hypothesis
  - BP agents would do better than no-BP agents, avoiding the free-rider effect.

Heterogeneity in Social Consciousness

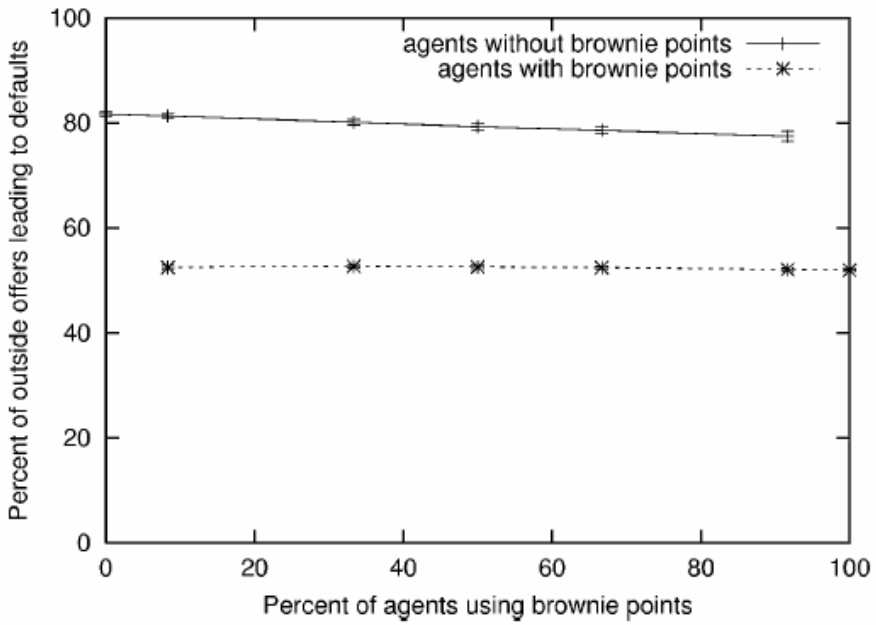

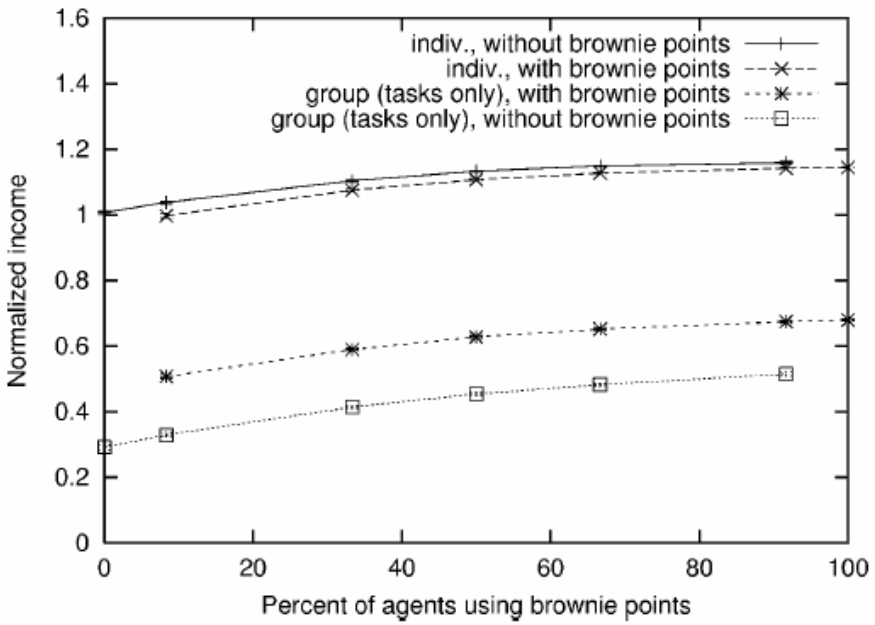

- No Brownie Points -> 80% agent default rate vs. 54% for BP agents
- BP Groups fared better, BP Agents fared worse
- BUT, free-rider effect is still observed
  - as %BP increases, overall welfare increases because BP agents reduce default costs
- Agent Designers can improve group outcomes without sacrificing individual gains w/ intermediate social consciousness

Heterogeneity in $Weight_{FEI}$ under RSCP

- Importance of future income
- delta = 0.4 or 0.3
- BPweight = 0.1
- Hypothesis
  - High Delta agents do better than low Delta agents
  - Free-Rider effect is avoided in High Delta situations

Heterogeneity in Weight$_{FEI}$ Results

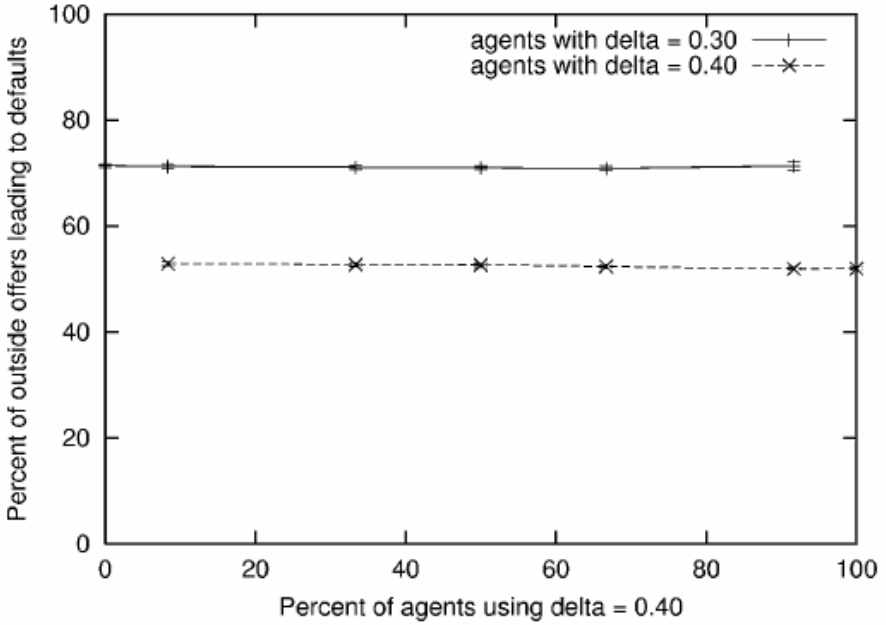

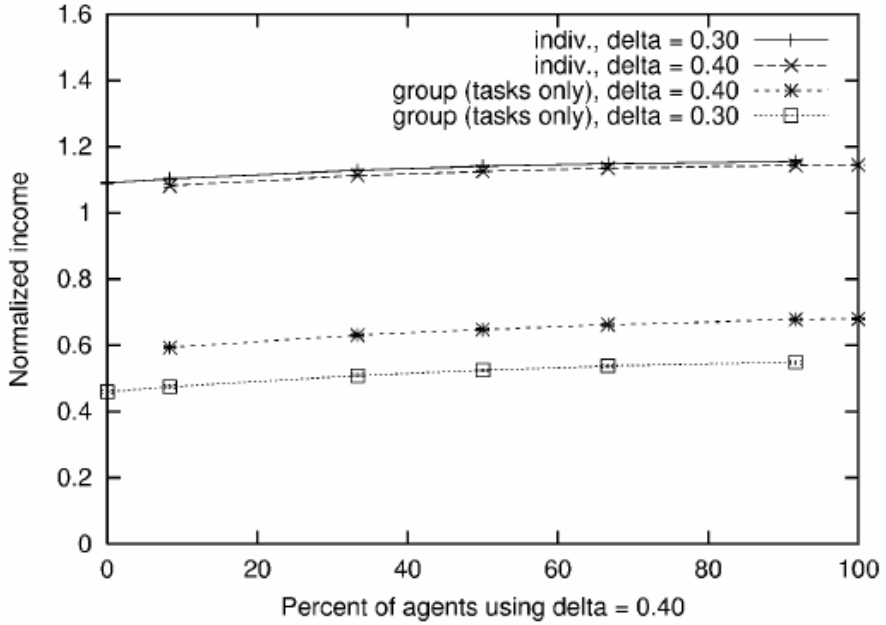

- default rates are as expected. Higher delta -> fewer defaults
- Less Responsibility -> more personal gain, less group gain
- More Responsibility -> slightly less personal gain, more group gain
- Free Rider effect remains

DSCP: Baseline Utility Parameters

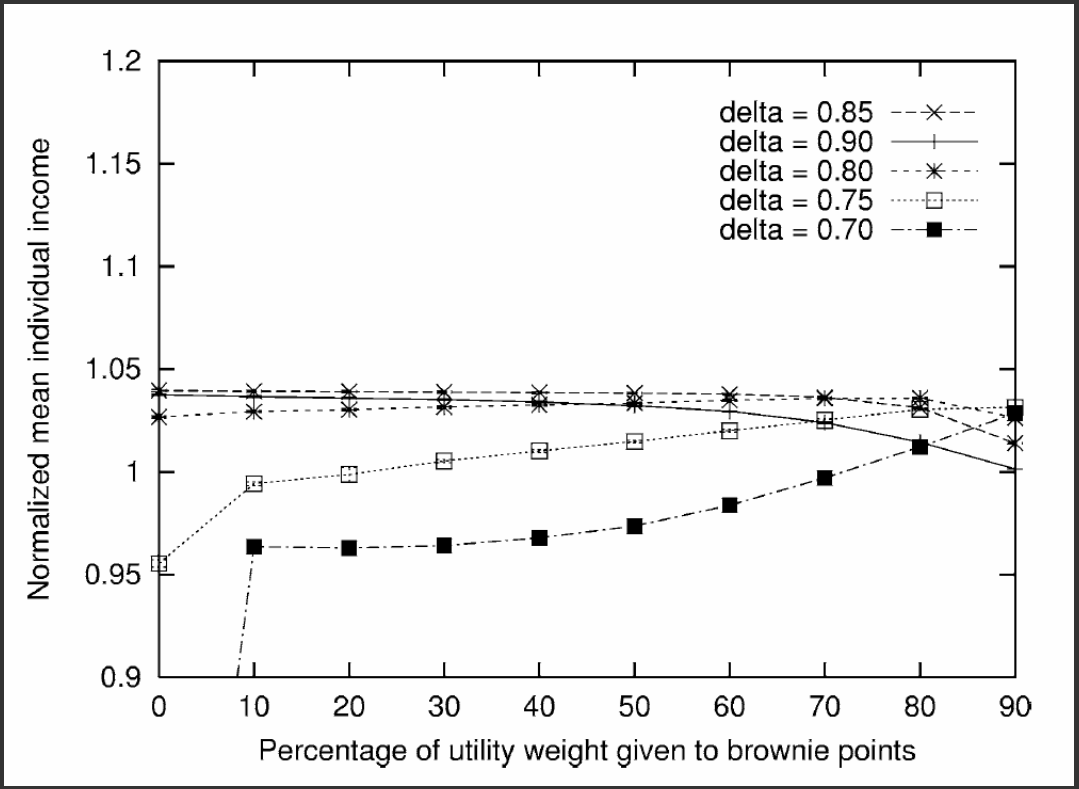

- delta=0.85, BPweight=0
- no social consciousness?
  - no explanation
- .85 is much > than RSCP delta of 0.4
  - less incentive to default
  - 3 possible losses (group,discount,no-replacement)

DSCP: Task Densities

- Hypotheses
  - agents will default less often as density is increased, provided that % tasks affected by RSCP is held constant
  - increasing task density -> lower individual and group incomes

DSCP: Task Densities Results

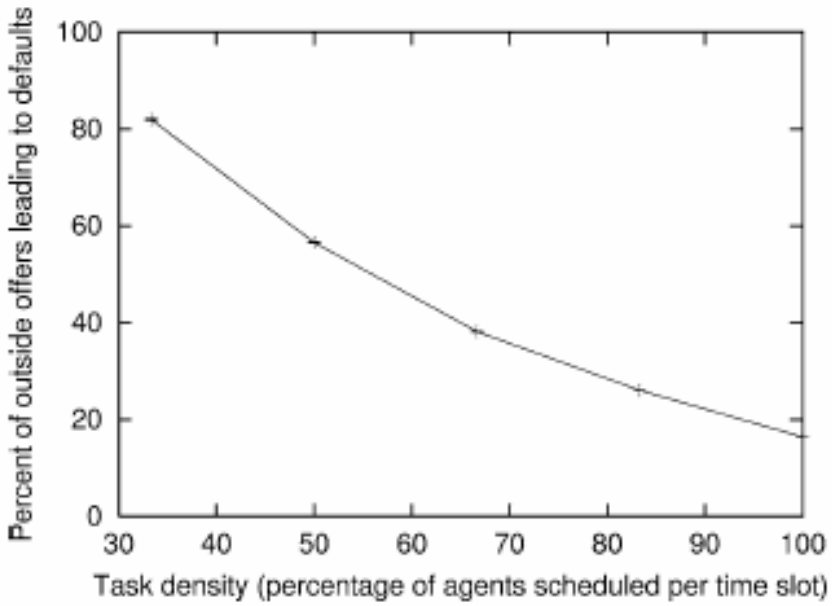

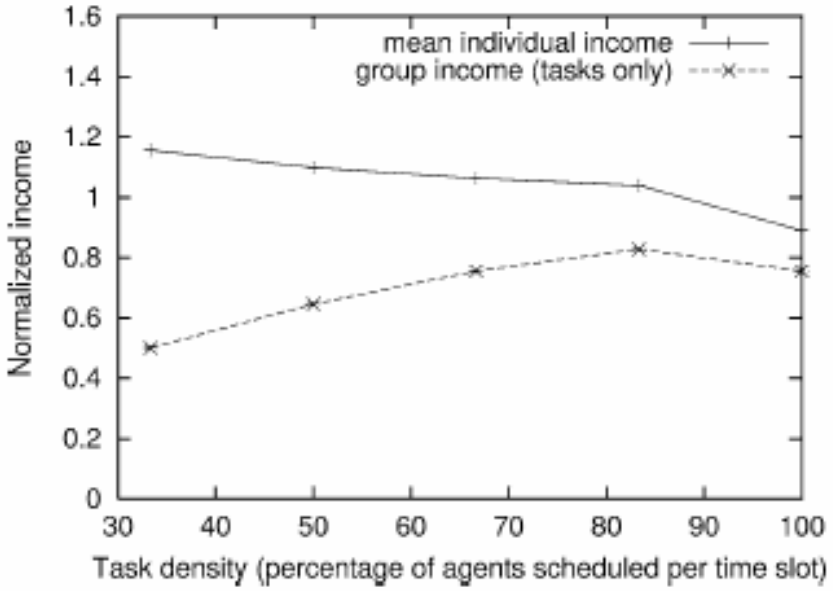

Fig. 9. The impact of task density on the rate of defaulting (left) and on normalized individual and group incomes (right) when using the DSCP.

- Same results as RSCP

DSCP: Outside Offer rate

- Varied outside offer rates
  - [10..20..70]% of total tasks
- Hypothesis
  - # OO increases -> default rate will be constant, total # defaults increases
  - # OO increases -> Individual Income increases, Group income decreases

DSCP: Outside Offer rate

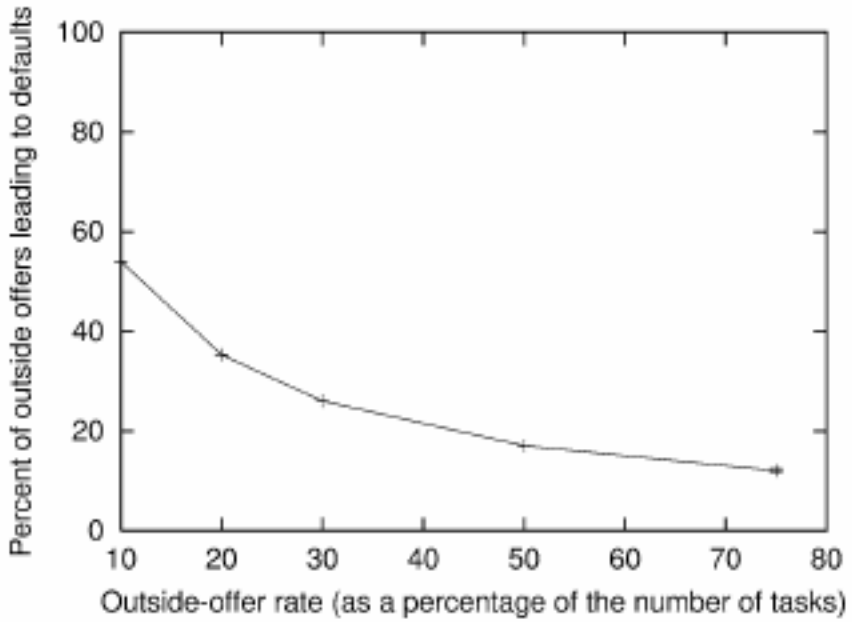

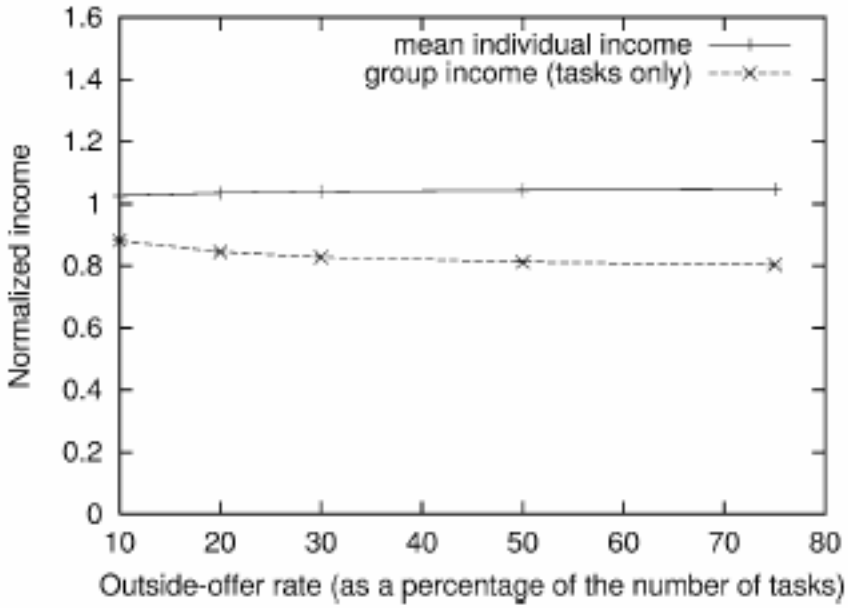

Fig. 10. The impact of the outside-offer rate on the rate of defaulting (left) and on normalized individual and group incomes (right) when using the DSCP.

- Absolute # increases slightly, but rate decreases
  - reputation is decreased by number, not rate he defaults
- 3 factors: Group Costs, Outside Income, task discount
  - Outside income (barely) outweighs other costs
  - But, group income decreases

Heterogeneity in $Weight_{FEI}$ under DSCP

- Importance of future income
- delta = 0.85 or 0.75
- BPweight = 0.1
- Hypothesis
  - High Delta agents do better than low Delta agents
  - Free-Rider effect is avoided in High Delta situations

Conclusions

- Future income estimation has a large impact on the evaluation of acting when in conflict:
  - estimations for both DSCP and RSCP would be improved by taking into account future income and losses from group costs were considered;
  - would require greater aware of environmental conditions and would allow agents to dynamically adapt default rates.
- Both policies are susceptible to the free-rider effect:
  - caused by parameters that allowed higher default rates from less responsible agents due to lowered group costs from more responsible agents.
- Social norms do influence the accuracy of an agent's response to changing environmental factors.

Contributions

- Work joins research in collaboration and resource-bounded reasoning.
- Demonstrates the roles that social consciousness can play in resolving decision conflicts resulting from changes in environmental factors:
  - useful for modeling real world and complex scenarios.
- Presents a framework for modeling agent interaction in collaborative environments:
  - support for usefulness of the model and its applicability in the decision domain.
- Demonstrates that improved decision making in cases of high uncertainty can be achieved by applying social consciousness in agent design.

Questions?