Influences on Intention Reconciliation

The Kiat's Meow -
Stosh Getzfrid, Janelle Rickstew, Travis Schreier

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.

Overview

"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. "

Key Terminology

- Intention reconciliation
- Social commitment policies
  - Ranking-based
  - Discount-based
- Social consciousness
- Shared Plans Intention Reconciliation Experiments (SPIRE) Framework

SPIRE

- Joins research on collaboration in MAS with research on rationality and resource-bounded reasoning
- Team of agents work together on group activities, each of which consists of doing a set of tasks
- Random agents are given outside offers
- Replacement agents
- Whole team incurs the cost if an agent defaults; cost is larger if no replacement agent is available
- Similar to Prisoner's Dilemma except agents' utilities and task assignments can depend on its own past actions and the actions of other agents

RSCP

- Each agent has a score that reflects how many times it has defaulted; this score is used to determine agent rankings.
- Higher ranked agents are assigned higher valued tasks.
- Impact of past weeks' defaults diminishes overtime (decay factor)

Agent's score is updated as follows:

$S(w+1) = \alpha(w) - p1(d) - p2(D)$

$p1 = 1 \qquad p2 = 5$

w: week number

α: constant decay factor of 0.5

d: # of defaults w/ replacement

D: # of defaults w/o replacement

DSCP

- Tasks are randomly assigned - agents with lower reputations do not receive tasks worth a lesser value (as in RSCP) but rather they receive a discounted amount of the value of the task they are assigned.
- Discount is determined by both the individual's reputation and the reputation of the group.
- Reputation represented using a numeric score that increases as the agent's reputation decreases(3)

$S(w+1) = \alpha(w) - p1(d) - p2(D)$

$p1 = -1 \qquad p2 = -1.1$

w: week number
α: constant decay factor of 0.5
d: # of defaults w/ replacement
D: # of defaults w/o replacement

Decision Making in SPIRE

- Agent must decide whether to default on task assigned – determined by utility of each option
- 3 Factors
  - Current Income (CI)
  - Future Expected Income (FEI)
  - Brownie Points (BP)

Decision Making in SPIRE – CI & FEI

- Current Income – income from task or outside offer & agent's share of group cost if it defaults
- Future Estimated Income – agent's estimate of income in future weeks based on social commitment policy & agent's score
  - Estimates impact of defaulting week 1, then calculates future weeks based on discount factor

Decision Making in SPIRE – Discounting

- Discounting can be seen in two different ways:
  - Agent's uncertainty about its own predictions
  - Assess present value of income (very similar to an interest rate)

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

Decision Making in SPIRE – Estimating Future Income under RSCP

- Two possible incomes for new position in rankings for if it defaults or not
  - Difficult to estimate because only know current and previous weeks ranking
- Makes assumptions based on previous week decision and whether ranking improved
  - Ex. If moved up in rankings, must mean agents ahead defaulted

Decision Making in SPIRE – Estimating Future Income under RSCP

- Uses assumption to create four classes:
  - 1st Agents above who will not default
  - 2nd Agents above who will default
  - 3rd Agents below who will not default
  - 4th Agents below who will default
- Adds itself to classes based on possible outcomes
  - No default, Default with replacement, Default with no replacement

Decision Making in SPIRE – Estimating Future Income under RSCP

- Finally, calculates estimated defaults and no defaults based on:
  - Previous rank
  - Current rank
  - Defaulted last week
  - Agent replacement
  - Total # of agents

Decision Making in SPIRE – Estimating Future Income under DSCP

- Compute new reputation score for both default & no default options, & expected value of discounted tasks
  - Uses Individual Scaling Function & Group Scaling Function

Decision Making in SPIRE – Brownie Points

- Brownie Points (BP) – Agent utility from being a good guy, considering for the good of the group, social consciousness
- BP only perceived by agent, not other agents
- Agents begin with identical # of BP
  - Default, lose BP
  - No Default, gain BP

Decision Making in SPIRE – Brownie Points

- Change in BP reflects values of tasks and offers
  - Ex. Agent defaults on low-value task, BP is reduced by less than if high-value task
  - Ex. Agent defaults for high-value offer, BP is reduced by less than if low-value offer
  - Increases work the same way
- Enables agents to do something less than optimal for individual, but benefits group more

Decision Making in SPIRE – Combining Factors

- CI + FEI = Total Estimated Income (TEI)
- Values are normalized by dividing income by average total value of tasks received

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

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

Experiment Results – Basic Settings & Explanation

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

Experiment Results – Basic Settings & Explanation

- All agents can do all tasks
- # of task types & # of time slots chosen to model the work week of a Systems Admin Team

Experiment Results – Free-Rider Effect

- Less responsible agents take advantage of more responsible agents

RSCP Experiments – Baseline Utility Functions for RSCP

- Lower delta & BPweight led to more defaulting than was optimal
- Higher delta & BPweight led to not enough defaulting to take advantage of offers

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

RSCP Experiments – Different Task Densities

- Varied # of tasks scheduled for each time slot
  - Effects agents ability to replace other agents
- Hypothesis:
  - Agents default less often as task density increases
    - CONFIRMED
  - Increased task density leads to lower individual & group incomes
    - BUSTED; from 33.3% to 83.3% group task income increased
    - Defaulted tasks are still completed; outside offers can be taken for the good of the group

RSCP Experiments – Different Task Densities

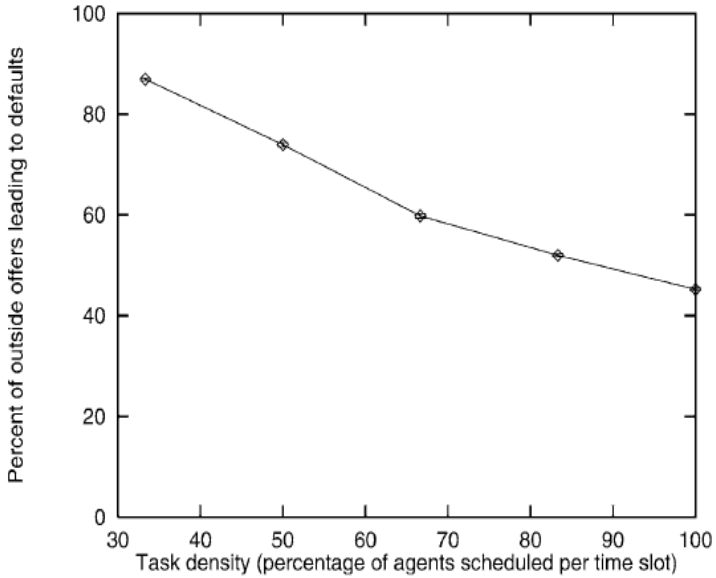

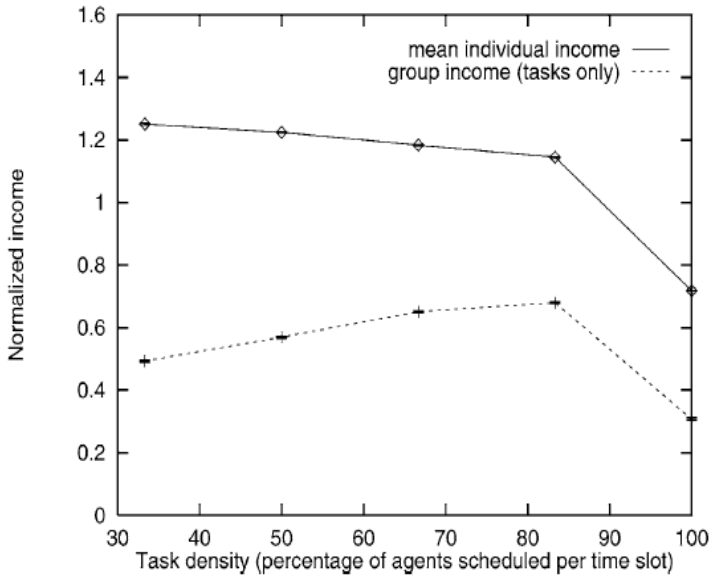

Task density	Offers with no replacement
33.3%	0.00%
50.0%	0.00%
66.7%	0.00%
83.3%	5.15%
100.0%	100.00%

RSCP Experiments – Different Outside-Offer Rates

- Varied # of outside offers as a % of total tasks
- Hypothesis:
  - As # of offers increases, rate of defaulting stays the same (# of defaults increases, but rate stays)
    - BUSTED; because available replacements are used up more quickly and thus no replacements, agents hesitant to default when no replacement
  - As # of offers increases, average individual income will increase, but group income decrease
    - CONFIRMED/BUSTED; 3 lowest offer rates confirmed but 50% and 70% disproved

RSCP Experiments – Different Outside-Offer Rates

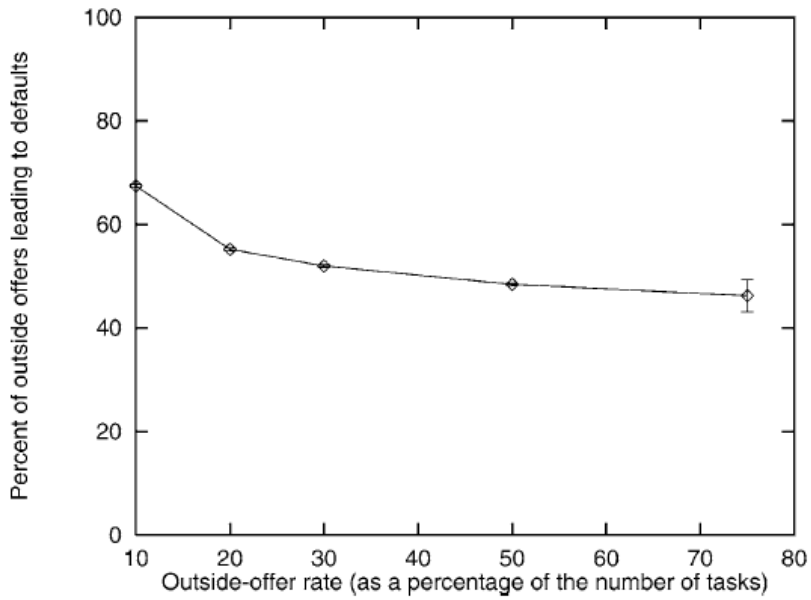

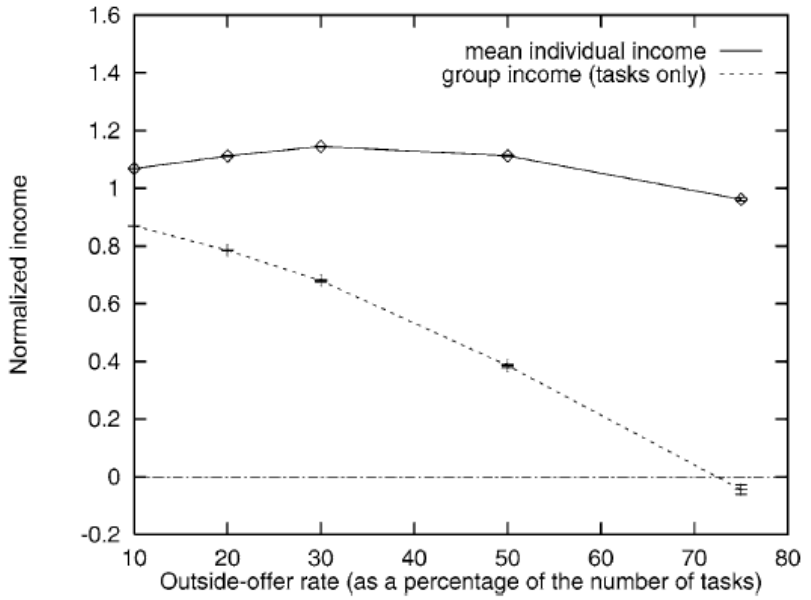

Outside-offer rate	Offers with no replacement
10%	0.03%
20%	0.91%
30%	5.15%
50%	25.45%
75%	47.81%

RSCP Experiments – Heterogeneity in Social Consciousness

- Varied % of agents who are socially conscious and those who are not
  - Those who are socially conscious use BP with the optimal BPweight of .1
- Hypothesis:
  - BP agents would do no better than no-BP agents, avoiding the free-rider effect
    - BUSTED; No-BP agents are seen to take advantage of BP agents
    - No-BP agents do worse as individuals in homogeneous but do better in heterogeneous

RSCP Experiments – Heterogeneity in Social Consciousness

- Was seen that agents who use BP default less often
  - As BP agents become larger portion of group, no-BP agents show small decrease in defaulting
- Agents benefit when more agents are socially conscious
- Found that no-BP do only slightly better as individuals

RSCP Experiments – Heterogeneity in Social Consciousness

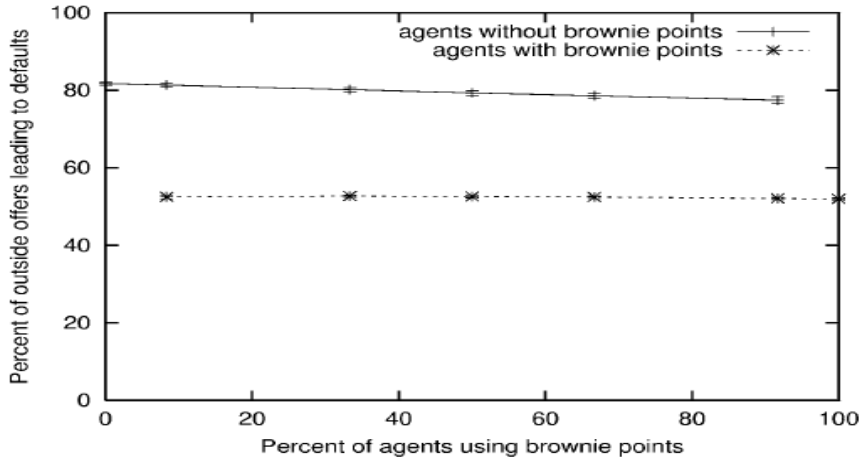

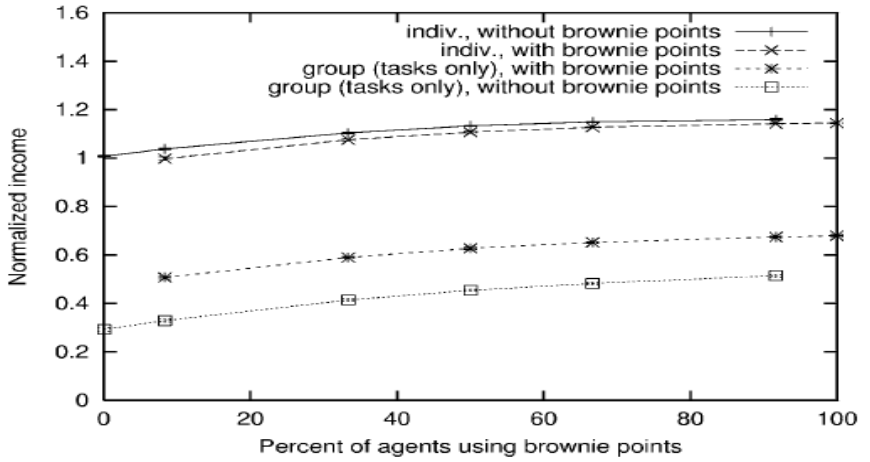

RSCP Experiments – Heterogeneity in Weight Given to FEI

- Varied % of agents who had an optimal delta value of .4, where other agents use a lower .3 value
- Hypothesis:
  - Higher delta agents do better than lower delta agents
    - BUSTED; Agents who put higher value on FEI default less often than agents who do not
    - Less responsible agents do worse as individuals in homogeneous but better in heterogeneous
    - Lower delta agents are able to take full advantage of reduced group costs

RSCP Experiments – Heterogeneity in Social Consciousness

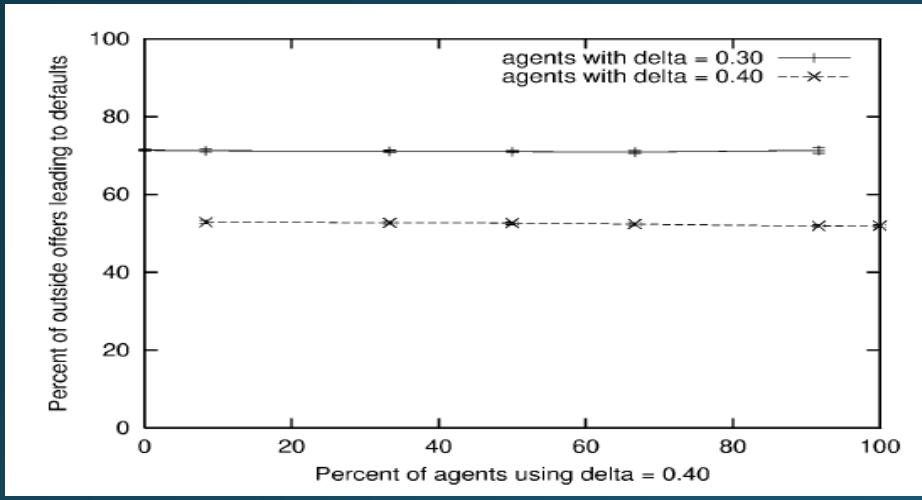

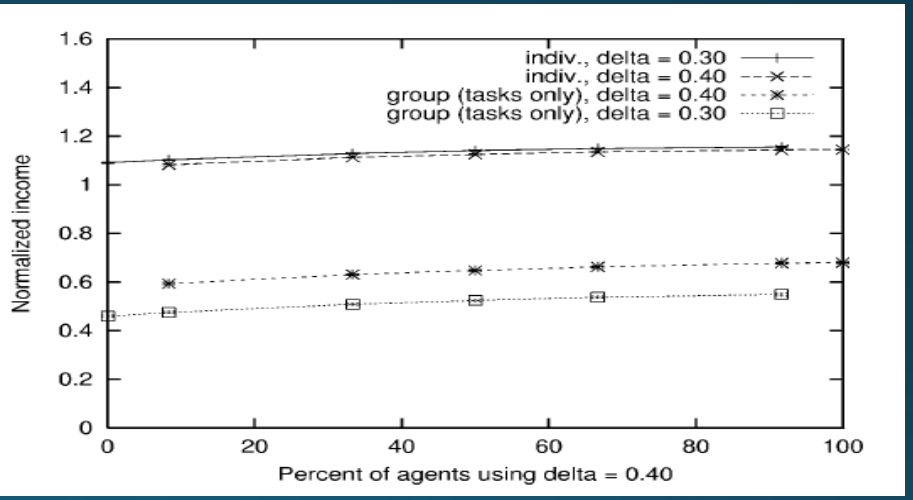

DSCP Experiments

1. Finding Baseline Utility Parameters
2. Comparing Different Task Densities
3. Comparing Different Outside-Offer Rates
4. Comparing Different δ Values for Estimating Future Incomes

- Normalized income by dividing it by total value of assigned tasks
  - Won't artificially inflate irresponsible agents' incomes

Exp. 1: Baseline Utility Parameters

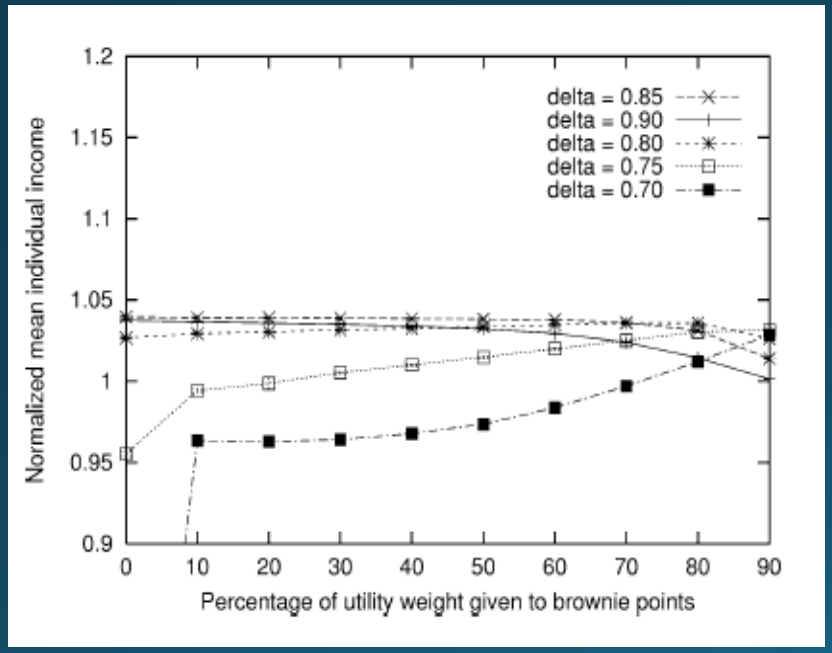

- Optimal δ = 0.85
  - Default less often (future income more important)
- Optimal BPweight = 0
  - Did not perform social consciousness experiment because of this

Exp. 2: Task Densities

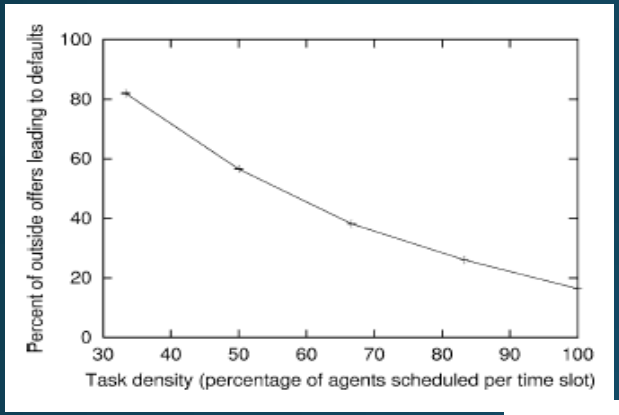

- Similar results to RSCP
- As task density increases:
  - Agents default less often
  - Individual income decreases
  - Group income increases until 83% ensity and then decreases

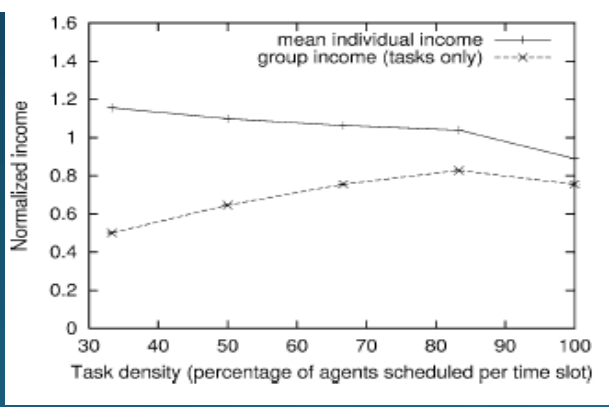

Exp. 3: Outside-Offers

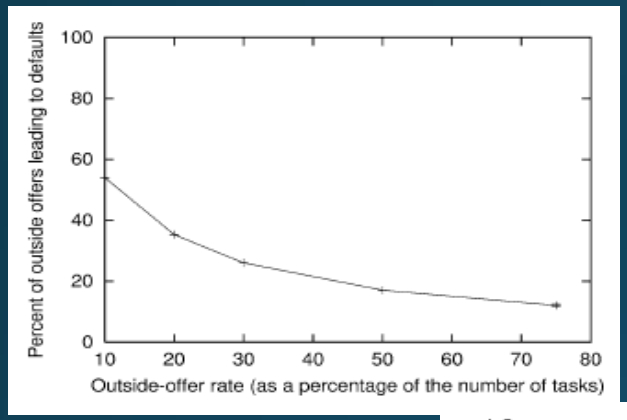

- As outside-offer rate increases:
  - Agents default less often
  - Individual income fairly constant
  - Group income slightly decreases

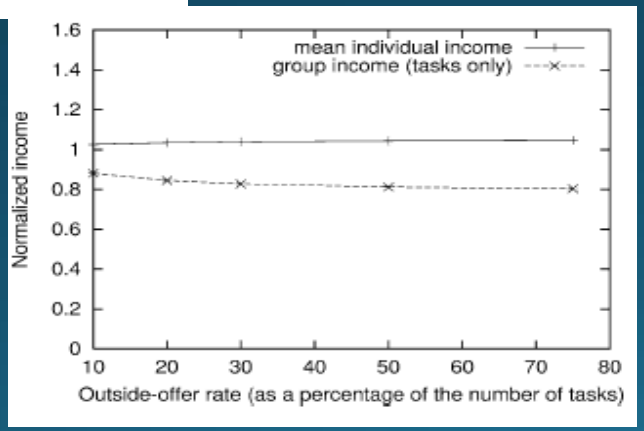

Exp. 4: Weight of FEI

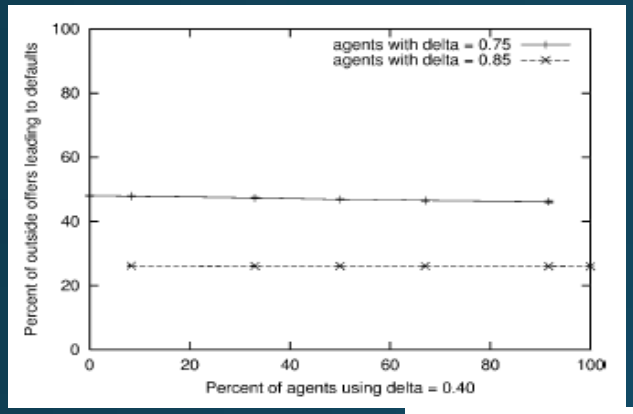

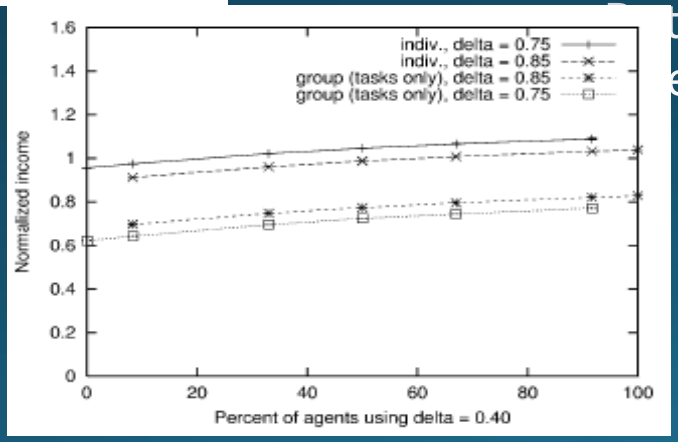

- Two types of agents:
  1. Use $\delta = 0.85$ (optimal)
  2. Use $\delta = 0.75$
- As number of agents using $\delta = 0.85$ increases:
  - Both types default at constant rates

Differences between DSCP and RSCP

- Agents under DSCP are more accurately able to estimate their future losses due to defaulting
- DSCP includes monetary punishments for irresponsible agents (discounting income from tasks)
  - Income lost to this discounting is just lost (not redistributed to others)
  - RSCP avoids this problem
- Agents under RSCP require Brownie Points in addition to δ to optimize individual incomes
  - Not needed in DSCP

Differences between DSCP and RSCP

- Agents under DSCP respond to environmental changes better
  - Reduce rates of defaulting more sharply when task density and outside-offer rates increase
  - Incur smaller drops in individual and group income
- Group income was higher under DSCP, but lower average individual incomes
  - Agents default less often → higher group incomes
  - Discounting of incomes leads to lost income → lower individual incomes

Improvements

- Optimal δ and BPweight depends on environment
  - Give agents the ability to alter rates of defaulting
- Free-riders can take advantage of heterogeneous situations
  - Remove individuals from group if they defect more than X times
  - Defaulting (defection) will not strictly dominate cooperating as it does in Prisoner's Dilemma
- Costs are divided evenly among individuals in a group
  - Have larger portion of cost attributed to defaulting agents

Conclusions

- Under either RSCP or DSCP, agents default fewer times as both task density and outside-offer rates increase
- Agents default fewer times under DSCP than RSCP due to the increased importance on estimated future income
- More irresponsible agents can take advantage of responsible agents in mixed groups to slightly improve their incomes, but groups of only irresponsible agents result in worse group income

Future Work

- Test when group costs are not divided evenly
  - Defaulting agents take on more of the group cost
- Test when agents have more information about the responsibility of other agents
  - Possibly have agents model others
- Explore conflicts that arise when agents are committed to multiple group activities

Related Works

- Kalenka and Jennings – "socially responsible" decision-making principles
  - Domain-dependent and not decision-theoretic
  - Whether agents choose to help each other, not decision to default
  - Agents cannot default
- They found that giving agents a degree of social consciousness tends to improve overall outcome but may cause individual performance to suffer

Related Works

- Sen – decision-making strategies that encourage cooperation among self-interested agents
  - Focuses more on interactions between pairs rather than a team
  - Whether to cooperate in the first place, not decision to default
- He found that less responsible agents may take advantage of their more responsible counterparts, but that more responsible agents tend to do better in the long run because they evaluate an agent's past behavior before deciding to cooperate.

Questions?