Insuring Risk-Averse Agents

Greg Hines and Kate Larson

ggdhines@cs.uwaterloo.ca
klarson@cs.uwaterloo.ca
Cheriton School of Computer Science
University of Waterloo
Waterloo, Canada

Abstract. In this paper we explicitly model risk aversion in multiagent interactions. We propose an *insurance mechanism* that be can used by risk-averse agents to mitigate against risky outcomes and to improve their expected utility. Given a game, we show how to derive Pareto-optimal insurance policies, and determine whether or not the proposed insurance policy will change the underlying dynamics of the game (*i.e.*, the equilibrium). Experimental results indicate that our approach is both feasible and effective at reducing risk for agents.

1 Introduction

In almost every decision people make, risk is a factor. When negotiating a business contract, there is the risk of either side being unable to fulfill its obligations. When bidding for multiple items in an auction, there is the risk of winning too many or too few items. Even when using the Internet, there is the risk of congestion depending on the routing policy used. In most of these cases people are risk averse. The importance of the influence of risk aversion on peoples' decisions is reflected in the size of the insurance industry, a multi-trillion dollar business, [4] and the amount of research in economics relating to risk [17].

There is considerable research in multiagent systems on helping people make better decisions in settings such as those mentioned above. However, this research generally assumes that people are risk neutral [5, 15, 16]. Given the prevalence of risk aversion in the real world, we believe it is important to study how to manage the effects of risk and risk aversion in multiagent systems.

In this paper, we study non-cooperative multiagent systems. Our main contribution is an insurance mechanism that can be used in games to reduce agents' risk and increase their utility. We present a characterization of our mechanism that allows us to easily determine if the mechanism can be applied to any given game. Experimental results show that our mechanism is usable in many different situations and is scalable. The experimental results also examine how much risk aversion matters in different settings. We conclude with a discussion of related work and some promising areas for future work.

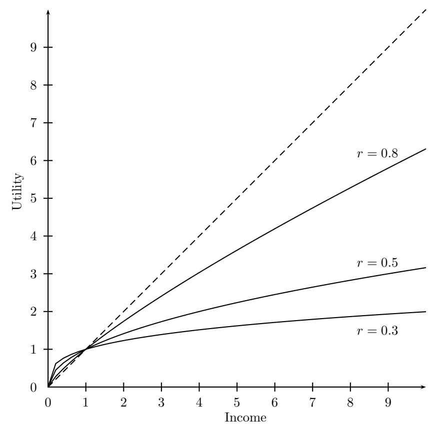

Fig. 1. A graphical depiction of income versus utility. The dashed line represents the utility of a risk-neutral agent while the solid lines represent the utilities of a risk averse agent for different values of r.

2 The Model and Background

In this section we introduce our model of risk aversion for multiagent systems, as well as define the key game-theoretic concepts used in this paper. For a more thorough introduction to game theory, we refer the reader to [8].

2.1 A Model of Risk Aversion

In this section we propose a model of risk aversion for a multiagent setting. The approach we take in modeling risk is motivated by models used in experimental economics [6]. If an agent is risk averse, then it dislikes uncertainty. For example, if given a choice between a lottery and a guaranteed payoff, a risk-averse agent will often prefer the guaranteed payoff, *even when the expected payoff from the lottery is higher.* In this paper, we model risk aversion by the concavity of an agent's utility function. Specifically, given an income I, if an agent's utility is of the form

$$U = I^r, \tag{1}$$

where $0 < r < 1$ is the *risk-attitude factor*, then the agent is risk averse. This utility function is depicted in Figure 1. For this paper we will assume that income is always greater than or equal to zero; we make this assumption since studies show that humans (and thus the agents designed to represent humans in different interactions) treat loss of income differently from equivalent gain of income [18].

Typically, a model of risk aversion requires a distinction between *income* and *utility* [11], therefore we generalize the notion of a game to reflect this distinction.

Definition 1 *Let N be a set of agents, $|N| = n$. An* income-based game *is defined as $G^I = \langle N, A, I_1, \ldots, I_n, u_1, \ldots, u_n \rangle$ where*

- *A_i is the set of possible actions for agent i, and $A = A_1 \times \ldots \times A_n$,*
- *$I_i : A \to \mathbb{R}$ specifies the* income *to agent i given joint action $a \in A$, and*
- *$u_i : \mathbb{R} \to \mathbb{R}$ is the utility function of agent i.*

We refer to the normal form (or matrix) representation of the income function as the *income matrix* and to the normal form representation of the utility function as the *utility matrix*.

Agents interact by following strategies, that is, by selecting actions to play according to some distribution. We are particularly interested in *correlated strategies*.

Definition 2 *A correlated strategy $\sigma_A = \{\sigma_A(a) | a \in A\}$ is any probability distribution over A. The conditional correlated strategy $\sigma_{A_{-i}}(a_{-i}|a_i)$ is the probability of the joint action (a_{-i}, a_i) according to σ_A given the action a_i where $a_{-i} = (a_1, \ldots, a_{i-1}, a_{i+1}, \ldots, a_n)$.*

In an income-based game, agents try to maximize utility, not income. That is, agents are trying to maximize their expected utility, given by

$$\sum_a \sigma_A(a) u_i(I_i(a)). \tag{2}$$

In this paper, we are interested in analyzing situations involving equilibrium play. Specifically, we are interested in studying situations where agents are playing a correlated equilibrium.

Definition 3 *A correlated strategy σ^*_A is a correlated equilibrium if for every agent i and every $a_i \in A_i$,*

$$\begin{aligned} &\sum_{a_{-i} \in A_{-i}} \sigma^*_{A_{-i}}(a_{-i}|a_i) u_i(I_i(a_i, a_{-i})) \\ &\geq \sum_{a_{-i} \in A_{-i}} \sigma^*_{A_{-i}}(a_{-i}|a_i) u_i(I_i(a'_i, a_{-i})), \end{aligned} \tag{3}$$

for all $a'_i \in A_i$.

We choose correlated equilibria as our solution concept for several reasons. First, it is a generalization of Nash equilibria, and so the techniques we propose in this paper can also be applied when a Nash equilibrium is chosen as the solution concept. Second, correlated equilibria can be "less risky" than Nash equilibria in that the correlation can help the agents in reducing the risk of mis-coordinating. Thus, if we can show that our approach is effective when risk is already being mitigated to some extent, then we believe that we can extrapolate our results and argue that our approach is broadly applicable and widely effective.

To illustrate the effects of risk, we consider an example of agents determing routing policies.

Example: *Consider two agents (Row and Column) on the Internet, each deciding on a routing policy to use. There is a public route available with a bandwidth of 200 kb/sec that is shared evenly if both agents decide to use it. Each agent also has a private route available with a bandwidth of 75 kb/sec. Suppose that both agents' utilities are given by Equation 1 with* $r = 0.5$*. Figure 2 shows the income (measured in bandwidth) and utility matrices of the game. Note that the unit of income is in kb/sec while utility has no units. If agents are playing the correlated strategy* $(\sigma_A((B, L)), \sigma_A((T, R))) = (0.5, 0.5)$*, the expected income for both agents is 137.5 kb/sec and the expected utility is 11.4. However, if the agents are guaranteed a bandwidth of 137.5 kb/sec, their utilities would increase to 11.7.*

	L	R
T	75,75	75,200
B	200,75	100,100

	L	R
T	8.7,8.7	8.7,14.1
B	14.1, 8.7	10,10

Fig. 2. *Income (left) and utility (right) matrices for a routing problem.*

2.2 Characterization of the Set of Correlated Equilibria

Since correlated equilibria play a central role in our paper, we wish to derive a formal characterization of the set of correlated equilibria in an income-based game. (For simplicity, the rest of the paper will consider games with only two agents each with two actions. Our results can be generalized to an arbitrary number of agents and actions.) Given the income and utility matrices for a general income-based game shown in Figure 3, suppose Agent 1 is trying to determine whether or not to follow a given correlated strategy σ_A (assuming Agent 2 does as well).

	L	R
T	a,b	c,d
B	e,f	g,h

	L	R
T	a^r,b^r	c^r,d^r
B	e^r,f^r	g^r,h^r

Fig. 3. Income matrix (left) and utility matrix (right) for a general game.

The expected utility for Agent 1 from playing T when following σ_A is

$$a^r \sigma_{A_2}(L|T) + c^r \sigma_{A_2}(R|T), \tag{4}$$

and the expected utility for Agent 1 for instead playing B is

$$e^r \sigma_{A_2}(L|T) + g^r \sigma_{A_2}(R|T). \tag{5}$$

Therefore, for Agent 1 to be willing to follow σ_A, we require that

$$a^r \sigma_{A_2}(L|T) + c^r \sigma_{A_2}(R|T) \geq e^r \sigma_{A_2}(L|T) + g^r \sigma_{A_2}(R|T)$$

or

$$(a^r - e^r)\sigma_A(TL) + (c^r - g^r)\sigma_A(TR) \geq 0. \tag{6}$$

Similar constraints can be created for every possible agent-action combination. The set of all such constraints completely defines the set of correlated equilibria for a game.

3 An Insurance Policy for Risk-Averse Agents

In this section we present an insurance policy mechanism for risk-averse multiagent systems. By buying insurance for certain outcomes, risk-averse agents are able to reduce their risk and increase their expected utility. We study the use of such a mechanism in a setting where agents are playing a correlated equilibrium.

The basic idea is for agents to be both buyers and sellers of insurance. For example, Agent i can buy coverage for the joint action a by selling coverage for another joint action, a'. If the outcome of the game is a, agent i *receives* income from another agent. If the outcome of the game is a', agent i *gives* some of its income to another agent. The unit cost of buying insurance (and the unit revenue from selling insurance) is set by a price vector $p = \{p(a)|a \in A\}$. The price vector is set by some third party; it is reasonable to assume that the correlating device of the correlated equilibrium also sets p.

3.1 Creating the Insurance Policy

Given a correlated equilibrium σ_A^* and a price vector p, each agent must determine how much insurance to buy and sell for each joint action a, *i.e.*, its demand $d_i(a)$. If $d_i(a) > 0$ then agent i wishes to buy insurance coverage for the joint action a and if $d_i(a) < 0$, agent i wishes to sell insurance coverage for the joint action a. Since agents are utility maximizers, the demand for each joint action can be computed by solving the constraint maximization problem:

$$\max_{d_i} \sum_{a \in A} \sigma_A^*(a) u_i(I_i(a) + d_i(a)), \tag{7}$$

$$\text{s.t.} \sum_{a \in A} p(a) \cdot (I_i(a) + d_i(a)) = \sum_{a \in A} p(a) \cdot (I_i(a)), \tag{8}$$

where $p(a)$ is the unit price for buying or selling insurance for the joint action a. The RHS of Equation 8 is agent i's budget given p, or the maximum amount of insurance it can possibly buy. Thus, agent i is simply trying to determine which insurance to buy ($d(a) > 0$) or to sell ($d(a) < 0$)), to maximize its expected utility while being constrained by its budget. For simplicity, let

$$x_i(a) = I_i(a) + d_i(a). \tag{9}$$

Let

$$R_i(a) = \frac{\partial \sigma_A(a) u_i(x_i(a))}{\partial x_i(a)} \cdot \frac{1}{p(a)}, \tag{10}$$

be the ratio of marginal utility compared to cost for buying insurance coverage for the joint action a. Agent i's utility is maximized when

$$R_i(a) = R_i(a'), \tag{11}$$

for all $a, a' \in A$. If Equation 11 is not satisfied, for example if $R_i(a) > R_i(a')$, then agent i would increase its utility by buying more coverage for the joint action a and buying less (or selling more) coverage for the joint action a'. For our utility function, this gives

$$x_i(a) = \sqrt[r-1]{\frac{\sigma(a')}{\sigma(a)} \frac{p(a)}{p(a')}} x_i(a'). \tag{12}$$

Equation 12 can be substituted into agent i's budget constraint (Equation 8) to determine its overall demand.

Example: *Continuing the example from Section 2, suppose the insurance price vector*

$$\{p((B, L)) = 1, p((T, R)) = 2\} \tag{13}$$

is announced and we wish to determine Row's optimal insurance coverage. In this case, Row's budget will be 350 kb/sec. Assuming $r = 0.5$*, Equation 12 simplifies to*

$$x_{Row}((B, L)) = \sqrt{2} x_{Row}((T, R)). \tag{14}$$

Substituting this into Equation 8, we get

$$\sqrt{2} x_{Row}((T, R)) + 2 x_{Row}((T, R)) = 350, \tag{15}$$

$$x_{Row}((T, R)) = 102.5. \tag{16}$$

Similarly, we find $x_{Row}((B, L)) = 150.0$*. Therefore, agent Row wishes to purchase insurance coverage of 27.5 kb/sec for the outcome* (T, R) *by selling 50 kb/sec of insurance coverage for the outcome* (B, L)*.*

Since, given p, each agent can determine its demand, the next challenge is to find an appropriate p. The insurance price vector should be set with several goals in mind. First, the resulting insurance policy should be budget balanced, *i.e.*, no external source of funding is required and the insurance policy does not make a profit. Second, the insurance policy should also be Pareto optimal, *i.e.*, no agent's expected utility can be increased without decreasing another's.

The insurance policy can be guaranteed to be budget balanced by choosing a p that results in supply equaling demand, *i.e.*, for all $a \in A$

$$\sum_i d_i(a) = 0. \tag{17}$$

To find which price vectors result in supply equaling demand, note that with our insurance mechanism, agents can only trade coverage and not create it. Therefore, our insurance mechanism is an example of an *exchange market* [8]. Arrow and Debreu proved that for every exchange economy, there exists a price vector which results in supply equaling demand. With respect to our insurance mechanism, Arrow and Debreu's theorem is as follows:

Theorem 1 *[2] For a given game G and correlated equilibrium σ^*_A, there exists some price vector p^* such that*

$$\sum_i d_i(a) = 0, \tag{18}$$

for every $a \in A$, assuming that u_i is continuous, strictly concave and strongly monotone.[1] *Such an d is known as a competitive equilibrium.*

It is straightforward to check that the utility function in Equation 1 satisfies all the required conditions in Theorem 1. Theorem 1 also requires that agents are *price-takers* – that is, each agent is unable to influence the price of the insurance policy. If an agent's demand (or supply) of insurance for a joint action is only a small fraction of the overall supply (or demand) then we can reasonably assume that the agent is a price-taker. Thus, our approach will work for games where there are many agents. However, if there are only a few agents, then they may be able to influence prices, and an alternative approach might be necessary. We propose that the price-setter also guarantees that the market will clear: that is, the third party promises to meet any extra demand and buy any extra supply.

Now that we have determined that p^* exists, we would like to know if it results in a Pareto optimal allocation. To do so, we rely on the following result [8].

Theorem 2 (First Fund. Thm. of Welfare Economics) *Any competitive equilibrium will always result in a Pareto optimal allocation.*

While for most exchange markets there are multiple p^*, with our utility function p^* is unique [8]. Since p^* is unique, this implies that p^* is also social-welfare maximizing.

At the same time, there is the complication that since the insurance policy will change agents' incomes and utilities, σ^*_A, the correlated equilibrium in the original game, may not remain an equilibrium once the insurance policy is in place. That is, since agents' utilities will have changed, $\sigma^*_{A_i}$ may no longer be a best response to $\sigma^*_{A_{-i}}$ and agents may wish to play other strategies. In this case, the insurance policy will no longer work since agents may now demand more coverage for certain joint actions or be willing to supply less coverage for other joint actions. Therefore, we are interested in finding a correlated equilibrium that will still be one after the insurance policy is in place. We call such an equilibrium an *insurable equilibrium*. Furthermore, for a given game G we would also like to provide a test to determine if any insurable equilibria exist.

To determine whether σ^*_A is an insurable equilibrium, we start by characterizing the set of all Pareto optimal allocations. The set of Pareto optimal allocations of income can be determined by solving the following constraint maximization problem:

$$\max_{x_i} u_i(x_i) \tag{19}$$

$$\text{s.t. } u_{-i}(I - x_i) = \bar{u} \tag{20}$$

[1] A utility function is strongly monotone if $u(y) > u(x)$ if $y \geq x$ and $y \neq x$.

	L	R
T	$(a+b)l$, $(a+b)(1-l)$	$(c+d)l$,$(c+d)(1-l)$
B	$(e+f)l$, $(e+f)(1-l)$	$(g+h)l$,$(g+h)(1-l)$

Table 1. The set of Pareto optimal allocation games where $0 \leq l \leq 1$.

for some fixed utility $\bar{u}$. This maximization problem can be solved using the Lagrangian method to find the constraint

$$\frac{x_i(a)}{x_i(a')} = \frac{I(a)}{I(a')}. \tag{21}$$

For brevity, we omit the derivation.[2] Therefore, the set of Pareto optimal allocations is given by

$$x_1 = \{I_{a_1} \cdot l, I_{a_2} \cdot l, \ldots\}, \tag{22}$$

$$x_2 = \{I_{a_1} \cdot (1-l), I_{a_2} \cdot (1-l), \ldots\} \tag{23}$$

for $0 \leq l \leq 1$. The resulting set of *Pareto optimal allocation games* is shown in Table 1.

Using the same reasoning that we used to find the correlated equilibrium constraint in Equation 6, we can find the analogous constraint for the Pareto optimal allocation game:

$$\begin{aligned} &\{[(a+b)l]^r - [(e+f)l]^r\}\,\sigma(TL) \\ &+ \{[(c+d)l]^r - [(g+h)l]^r\}\,\sigma(TR) \geq 0. \end{aligned} \tag{24}$$

Note that l^r cancels out in Equation 24 leaving the simplified constraint

$$\begin{aligned} &[(a+b)^r - (e+f)^r]\sigma_A(TL) \\ &\qquad + [(c+d)^r - (g+h)^r]\,\sigma_A(TR) \geq 0. \end{aligned} \tag{25}$$

Since Equation 25 does not depend on l, this constraint must hold for every Pareto optimal allocation, including the one created by p^*. Therefore, to determine if σ_A^* is an insurable equilibrium, we simply check whether it satisfies the general set of constraints for correlated equilibria in a Pareto optimal allocation game. By incorporating a convex programming method similar to that of Papadimitriou and Roughgarden [10], our approach can easily determine whether any insurable equilibria exist for a given income-based game.

4 Experimental Evaluation

In this section we describe the experimental evaluation of our insurance mechanism. In particular, we study the extent of its applicability and the amount that it improves the utility of the participating agents.

[2] Note that if the agents did not all have the same value for r, this characterization would not be possible.

	Distribution	
Game Size	Uniform	Bi-modal Gaussian
2 agents, $\|A_i\| = 2$	5%	5%
3 agents, $\|A_i\| = 3$	95%	94%
4 agents, $\|A_i\| = 4$	91%	100%

Table 2. The percentage of insurable equilibria in randomly-generated games of different sizes, generated by two different distributions.

4.1 Experimental Setup

We conducted our experiments using games with 2 agents with 2 actions per agent, 3 agents with 3 actions per agent, and 4 agents with 4 actions per agent. For each game, the income values were drawn randomly from one of two distributions: the first was the uniform distribution over $[0, 10]$, and the second distribution was a bi-modal Gaussian with $N_1(10, 3)$ and $N_2(100, 3)$. This second distribution was used to generate high-risk games. For all experiments we set the risk-attitude factor $r = 0.6$. Experimental evidence suggests that this value often captures humans' risk-attitudes [6], without being too extreme in either direction. We repeated each experiment 100 times, and all results reported are averaged over these repetitions. A linear program was used to determine if an equilibrium was insurable. We then used the tâtonnement process to find the resulting insurance prices [8].

In each of our experiments, there were three things we studied. First, we were interested in determining what percentage of games actually had insurable equilibria. This measurement allows us to determine the applicability of our approach. Second, for insurable equilibria, we were interested in understanding how effective our insurance mechanism was at reducing risk. For an agent to completely remove risk from a game, it must receive the same income for every possible outcome, so for a given game, we define u_i^{Orig} to be the expected utility to agent i in the original (non-insured) game, u_i^{Ins} to be the expected utility in the insured game, and u_i^{RF} to be the expected utility if the game was made to be risk-free. We define the *insurance effectiveness* for agent i (IE_i) as

$$IE_i = \frac{u_i^{Ins} - u_i^{Orig}}{u_i^{RF} - u_i^{Orig}} \times 100\%. \tag{26}$$

Finally, we were interested in understanding the underlying cost of risk aversion in terms of utility loss for an agent i. We define the *cost of risk aversion* of agent i ($CORA_i$) as

$$CORA_i = \frac{u_i^{RF} - u_i^{Orig}}{u_i^{RF}} \times 100\%. \tag{27}$$

4.2 Results

Table 2 presents our findings on how often insurable equilibria exist. We present our findings from games generated from both distributions. In general, we found that for games with more than two agents, insurable equilibria were very common, and over

Game Size	IE	CORA
2 agents, $\|A_i\| = 2$	25%	4.5 %
3 agents, $\|A_i\| = 3$	72%	8.2 %
4 agents, $\|A_i\| = 4$	76.8%	12 %

Table 3. Average insurance effectiveness (IE) and costs of risk aversion ($CORA$) in games generated using the bi-modal Gaussian distribution. For IE, values closer to 100% show that the insurance mechanism is improving the expected utility for an agent (if $IE = 100\%$ then the optimal utility is achieved). For $CORA$, a value of 4.5%, for example, indicates that risk aversion leads to a 4.5% decrease in utility, compared to a risk-neutral approach.

90% of all games generated had insurable equilibria. For the randomly generated two-agent, two-action games, we found that insurable equilibria were quite rare, occurring in only 5% of games. While at first glance this was disappointing, upon further investigation of the two-agent, two-action games, we noticed that most of these games had a single pure-strategy Nash equilibrium. For such games there is no risk in mis-coordinating, and thus no need for an insurance policy.

Table 3 presents our findings for the IE and $CORA$ measurements for different sizes of games, drawn from the bimodal distribution. The results presented are averaged over all games where there was an insurable equilibrium, and over all agents in those games. We make two important observations. First, as the game increases in size, the impact that risk has (as measured by $CORA$) also increases. Second, as CORA increases, so does the effectiveness of our insurance mechanism (as measured by IE). When games were generated using the uniform distribution (results not presented), we observed that there was less overall risk. In particular, the $CORA$ measurement was never greater than 3% on average, and thus, overall, the insurance effectiveness was also quite low. Given these results, we conclude that when risk is an important factor in a game, our insurance mechanism is highly effective. When there is little risk, however, it provides only minimal advantage.

5 Related Work

The standard insurance model assumes an initial level of wealth with some probability of an accident, *i.e.*, some loss of wealth, represented by a probability density function [1, 12]. Someone interested in buying insurance decides on the type of coverage they want: the maximum coverage, the deductible, the level of coinsurance, *etc.* The insurer then decides on the premium to charge for that particular insurance policy. Research in insurance has examined questions such as determining the optimal policy to buy and the optimal premium to charge [1, 12]. Other work has dealt with the effects of asymmetric information, moral hazards, and adverse selection [14].

Game theory has been used to a limited degree in the study of insurance; for example, in analyzing the actions of insurance companies in competitive markets [14]. However, this work ignores any strategic interaction between insurance buyers by assuming the insurance companies can supply any and all requested insurance policies with non-negative returns. Arrow and Raviv have both used decision theory in determining the optimal actions of both buyer and seller [1, 12]. Since their work focuses on the actions of an isolated buyer and seller, this is more an application of decision theory than game theory.

The insurance research closest to our work is the study of *reciprocal reinsurance*: the exchange of risk between insurance companies [3]. Borch considered the problem as an n-person coalitional game and was able to solve it for $n = 2$. This work made several assumptions that our work does not, such as assuming that the probabilities of the different outcomes are independent and companies have additional outside money to use. The goal is for the two companies, A and B, to reach a deal where A pays B to cover a specific amount of A's risk (or vice-versa). Borch presented this as a bargaining problem, where companies had to find an amount to be paid and the risk to be covered, and suggested the Nash bargaining solution as the desired outcome.

There has been limited work on risk in multiagent systems. Exceptions include work by Lam *et al.*, which proposed an insurance scheme for agents trying to obtain resources [7]. In their work insurance premiums were paid to specific *insurance agents* who, in return, guaranteed that necessary resources were always available. The effects of risk aversion have been studied more often in auction design. Page, for example, studied the problem of optimal auction design with both risk-averse buyers and a risk-averse seller [9] while the effects of risk aversion in sequential auctions was studied by Robu and La Poutré [13].

6 Conclusion

In this paper we commenced a formal study of risk and risk-aversion in multiagent systems. We presented a mechanism that allows agents to buy and sell *insurance* in order to protect themselves against undesirable outcomes. We described how to derive Pareto-optimal insurance policies, and provided a characterization of *insurable equilibria* for two-player games. Experimental results indicated that when risk is prevalent in the agents' interactions, our insurance mechanism effectively mitigates the risk and improves the expected utility of all agents.

There are many interesting open challenges related to our insurance mechanism. First, we would like to study our mechanism in a 2-stage game model. This might allow for a more generalized equilibrium model, and also be useful in a repeated game model. The repeated game model could be used to study a non-equilibrium setting; a non-equilibrium model might involve relaxing the balanced-budget requirement and using a *targeted optimality* approach, where we would optimize the insurance mechanism for specific types of agents and games. Studying repeated games may also allow the use of credit and savings to reduce risk, and it would be interesting to compare the advantages of an insurance mechanism against a credit-and-savings mechanism. Secondly, we are interested in trying to apply our insurance mechanism to other models of multiagent systems such as cooperative games and collaborative multiagent systems. Thirdly, we would like to investigate whether other models of risk aversion are more useful. We are specifically interested in how cumulative prospect theory and loss aversion could be used in multiagent systems [18]. It would also be interesting to compare the advantages and disadvantages of the core and competitive equilibrium as different solution concepts. Finally, we would like to implement our insurance mechanism in real life settings. In such settings, agents may be unaware of their own degree of risk aversion or may choose to lie about it. As a result, there would be a need to use preference elicitation and mechanism design.

References

1. Kenneth Arrow. *Essays in the Theory of Risk-Bearing*. North-Holland, 1971.
2. Kenneth Arrow and Gérard Debreu. Existence of an equilibrium for a competitive economy. *Econometrica*, 22(3):265–290, 1954.
3. Karl Borch. *The Mathematical Theory of Insurance*. Lexington Books, 1974.
4. Swiss Reinsurance Company. World insurance in 2006: Premiums came back to "life", 2006.
5. Vincent Conitzer and Tuomas Sandholm. AWESOME: A general multiagent learning algorithm that converges in self-play and learns a best response against stationary opponents. *Machine Learning*, 67(1-2):23–43, 2006.
6. Jacob K. Goeree, Charles A. Holt, and Thomas R. Palfrey. Risk averse behavior in generalized matching pennies games. *Games and Economic Behavior*, 45(1):97–113, October 2003.
7. Yuk-Hei Lam, Zili Zhang, and Kok-Leong Ong. *AI 2005: Advances in Artificial Intelligence*, chapter Insurance Services in Multi-agent Systems, pages 664–673. Springer, 2005.
8. Andreu Mas-Colell, Michael Whinston, and Jerry R. Green. *Microeconomic Theory*. Oxford University Press, 1995.
9. Frank H. Page. Optimal auction design with risk aversion and correlated information. Technical report, Tilburg University, 1994.
10. Christos H. Papadimitriou and Tim Roughgarden. Computing correlated equilibria in multiplayer games. *Journal of the ACM*, 55(3), 2005.
11. Matthew Rabin. Risk aversion and expected-utility theory: A calibration theorem. *Econometrica*, 68:1281–1292, 2000.
12. Artur Raviv. The design of an optimal insurance policy. *The American Economic Review*, 69:84–96, 1979.
13. Valentin Robu and Han La Poutré. Designing bidding strategies in sequential auctions for risk averse agents: A theoretical and experimental investigation. In *Proceedings of the 9th Workshop on Agent Mediated Electronic Commerce*, pages 76–89, Honolulu, U.S.A., 2007.
14. Michael Rothschild and Joseph Stiglitz. Equilibrium in competitive insurance markets: An essay on the economics of imperfect information. *The Quarterly Journal of Economics*, 90:629–649, 1976.
15. Ola Rozenfeld and Moshe Tennenholtz. Routing mediators. In *Proceedings of IJCAI-07*, pages 1488–1493, Hyderabad, India, 2007.
16. Yoav Shoham, Rob Powers, and Trond Grenager. If multiagent learning is the answer, what is the question? *Artificial Intelligence*, 171(7):365–377, 2007.
17. Joseph Stiglitz. Nobel lecture. "`http://nobelprize.org/nobel_prizes/economics/laureates/2001/stiglitz-lecture.pdf`", 2001.
18. Amos Tversky and Daniel Kahneman. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, 5(4):297–323, October 1992.