Economics Bulletin

Volume 30, Issue 3

Interdependent security experiments

Aric P. Shafran
California Polytechnic State University

Abstract

This paper analyzes the behavior of subjects in interdependent security experiments which exhibit strategic complementarity. In these experiments, subjects decide whether to pay to mitigate the risk of a loss, but the exact risk depends on the actions of other subjects. Two ranked equilibria exist, and the efficient equilibrium is for all agents to pay for the mitigation. Subjects in the interdependent security experiments rarely coordinate on the efficient equilibrium. Coordination is slightly more common in similar coordination games without the risk mitigation decision. The experiments also compare the effectiveness of two policies at inducing higher levels of mitigation.

This research was supported in part by funds provided by the Rocky Mountain Research Station, Forest Service, U.S. Department of Agriculture.
Citation: Aric P. Shafran, (2010) "Interdependent security experiments", *Economics Bulletin*, Vol. 30 no.3 pp. 1950-1962.
Submitted: Apr 01 2010. **Published:** July 27, 2010.

1 Introduction

Critical mass games are binary choice games with multiple equilibria where once a threshold of agents commit to an action, all other agents find it in their best interest to also take that action (Schelling, 1978). Examples from Schelling include urban segregation and faculty seminar participation. This idea has recently been applied to the problem of risk mitigation when there are risk interdependencies (Kunreuther and Heal, 2003). In interdependent security (IDS) games, agents decide whether or not to pay to mitigate the risk of a loss, but the exact risk of a loss depends on the actions of other agents. Kunreuther and Heal's leading example is the decision by airlines to screen passengers and bags for bombs. Other applications include computer security and wildfire protection decisions (Shafran, 2008).

When the risk interdependencies exhibit strategic complementarity, the games can have two or more Pareto-ranked pure strategy Nash equilibria.[1] The IDS games analyzed here have exactly two pure strategy equilibria, one equilibrium where no one invests in risk mitigation and a preferable equilibrium where all agents invest. Coordination failure occurs when agents choose not to invest in risk mitigation, and the less desirable equilibrium is achieved.

IDS games combine a traditional coordination game in which individuals unravel the strategic uncertainty regarding the actions of the other players with a decision based on individuals' personal preferences for risk mitigation. The risk mitigation aspect of the game adds a level of complexity to the game beyond that which has appeared in most coordination game experiments.[2] A goal of this paper is to investigate the effect of the additional complexity caused by the risk mitigation aspect of the game. To address this, behavior in IDS games is compared to that in a related but simpler game that does not depend on preferences for risk mitigation. The simpler game is closely related to many of Schelling's examples, and variations on it have been widely studied in experiments starting with Van Huyck et al. (1990).[3]

Subjects played in one of three variations of an IDS game, designed to test the effectiveness of two subsidy-based policies at inducing risk mitigation and leading to the preferable equilibrium. The first policy treatment, referred to as *symmetric subsidies*, offers a partial subsidy to each agent who chooses to mitigate the risk. The second policy treatment, referred to as *asymmetric subsidies*, makes a subset of the agents eligible for full subsidies. This policy is a novel approach to solving the problem of coordination failure, exploiting the risk interdependencies to induce mitigation among agents who do not receive a subsidy. By making use of the fact that mitigation is beneficial without a subsidy once enough agents are mitigating, this policy has the potential to lead to the preferable equilibrium in a more cost effective manner than the symmetric subsidies. In practice, asymmetric subsidies can be implemented through the use of "early bird" specials which reward the first agents to act or those who act by a set date.

The remainder of the paper is organized as follows. Section 2 develops a simple model of interdependent security upon which the experiments are based. Section 3 explains the experimental

[1]See Kunreuther et al. (2007) for an experimental study of IDS games with only one Nash equilibrium.

[2]Hess et al. (2007) also study IDS coordination games in an experimental setting. As in the experiments presented here, they find that coordination failure is common in simultaneous IDS games.

[3]See Van Huyck et al. (1991), Van Huyck et al. (1997), Cooper et al. (1992), Straub (1995), Schmidt et al. (2003), and Devetag (2003) for additional examples of coordination game experiments.

design. Section 4 presents the results. Section 5 concludes.

2 A Simple Model of Interdependent Security

Suppose that an agent with wealth Y faces some probability of a loss L. The agent chooses between strategy A (no mitigation) and strategy B (mitigation). Strategy B imposes a certain cost c in return for a lower probability of experiencing the loss L. The exact probability of a loss depends on the agent's choice as well as the choices of the other $N-1$ agents in the game. Let $n \in 0, 1, ..., N-1$ denote the number of other agents who choose B. Let $P(i,n)$ denote the probability that an agent incurs a loss where $i \in \{A,B\}$. Assume $P(i,n)$ has the following form:

$$P(i,n) \quad = \quad \begin{cases} P_0 - \alpha n & \text{if } i = A \\ P_1 - \beta n & \text{if } i = B \end{cases} .$$

Assume that $P_0 > P_1$ and $0 < \alpha < \beta$ to reflect the benefit of risk mitigation. The payoff to an agent who mitigates is:

$$Y - c - P(B,n)L \tag{1}$$

The payoff to an agent who does not mitigate is:

$$Y - P(A,n)L \tag{2}$$

An agent's best response is to mitigate if:

$$[P(A,n) - P(B,n)]L > c \tag{3}$$

The games studied in this paper have the characteristic that $[P(A,0) - P(B,0)]L < c$ and $[P(A,N-1) - P(B,N-1)]L > c$. Thus, it is a Nash equilibrium for all agents to mitigate and for no agents to mitigate. Kunreuther and Heal (2003) provide several examples of IDS games that have this feature. Their leading example is airline security. In their model, airlines choose whether or not to screen checked bags for bombs. Screening bags incurs a cost but reduces the probability that a bomb gets on a plane. As is common practice in the industry, airlines only screen bags checked directly with them. Bags transferred from other airlines are not re-screened. Airlines are therefore exposed to risk by partner airlines which elect not to screen bags.

Equilibrium selection criteria are necessary to predict which of the two equilibria will occur. Three selection criteria are considered here, each of which has been to shown to play a role in equilibrium selection in coordination game experiments. Harsanyi and Selten (1988) define an equilibrium as payoff dominant if every agent earns a higher payoff at that equilibrium compared to all other equilibria. In order for the mitigation equilibrium to be payoff dominant, it must be the case that $Y - c - [P_1 - \beta(N-1)]L > Y - P_0 L$, or $[P_0 - P_1 + \beta(N-1)]L > c$. The condition (from the previous paragraph) that $[P(A,N-1) - P(B,N-1)]L > c$ implies that $[P_0 - P_1 + \beta(N-1) - \alpha(N-1)]L > c$. Since $\alpha(N-1)$ is positive, the condition $[P_0 - P_1 + \beta(N-1)]L > c$ is always

Table I: Parameter Values

Treatment	Y	L	c	N	P_0	α	P_1	β
Baseline	\$0.30	\$0.24	\$0.04	7	0.4	0.02	0.38	0.05
Symmetric Subsidies	\$0.30	\$0.24	**\$0.02**	7	0.4	0.02	0.38	0.05
Asymmetric Subsidies	\$0.30	\$0.24	\$0.04	**5**	**0.36**	0.02	**0.28**	0.05

met when $[P(A,N-1)-P(B,N-1)]L > c$. Thus, the equilibrium where everyone mitigates is the payoff dominant equilibrium.

An equilibrium is risk dominant if agents are playing strategies that impose the least risk on them given their strategic uncertainty about the behavior of the other players (Harsanyi and Selten, 1988). If it is optimal to choose A when each opponent plays A with probability at least 0.5, then the non-mitigation equilibrium is risk dominant. If it is optimal to choose B when each opponent plays B with probability at least 0.5, then the mitigation equilibrium is risk dominant. The risk dominant equilibrium will depend on the specific values in the game.

Van Huyck et al. (1990) define the secure equilibrium as that in which agents are playing the strategy with the highest worst case payoff. Since both α and β are positive, the worst case payoff for both strategies occurs when no one else mitigates. Given the assumption that $[P(A,0)-P(B,0)]L < c$, choosing not to mitigate provides a higher worst case payoff. The equilibrium where no one mitigates is therefore the secure equilibrium.

3 Experimental Design

Subjects played a sequence of rounds which replicated one of three variations of the model from Section 2, a baseline game and two policy treatments that tested the effect of symmetric and asymmetric subsidies. In all three treatments, players received a payment of Y at the start of each round. There was a chance that they would lose L. Each player could choose between two strategies, A and B. Strategy B was the risk mitigation strategy. By choosing B, agents paid a small amount (c) to reduce the chance of a loss.

Table I shows how each parameter was set for each of the three treatments.[4] The symmetric subsidy treatment cuts the cost of mitigation in half from \$0.04 to \$0.02. The asymmetric subsidy treatment holds the cost of mitigation at \$0.04 but implicity allows two agents to mitigate for free, resulting in a game with five players instead of seven.

In all three treatments, the mitigation equilibrium is payoff dominant while the non-mitigation equilibrium is the secure equilibrium. By increasing the expected payoffs at the mitigation equilibrium, the symmetric subsidy strengthens the payoff dominance of that equilibrium compared to the baseline game. Both subsidy treatments reduce the risk associated with mitigating, the symmetric subsidy by increasing the expected payoffs from mitigating and the asymmetric subsidy

[4]Five of the fifteen cohorts did not play the final 20 rounds discussed below (three cohorts of the baseline game and two cohorts of the symmetric subsidy treatment). To hold expected session earnings constant, these cohorts had all dollar values scaled up by 2.5 times with the other parameters the same.

Table II: Summary of Treatments - IDS Game

Treatment	Choice	Cost	Number Choosing B						
			0	1	2	3	4	5	6
Baseline	A	\$0.00	40%	38%	36%	34%	32%	30%	28%
	B	\$0.04	38%	33%	28%	23%	18%	13%	8%
Symmetric Subsidies	A	\$0.00	40%	38%	36%	34%	32%	30%	28%
	B	\$0.02	38%	33%	28%	23%	18%	13%	8%
Asymmetric Subsidies	A	\$0.00	36%	34%	32%	30%	28%		
	B	\$0.04	28%	23%	18%	13%	8%		

Table III: Summary of Treatments - Expected Payoffs (In Dollars)

Treatment	Choice	Number Choosing B						
		0	1	2	3	4	5	6
Baseline	A	0.204	0.209	0.214	0.218	0.223	0.228	0.233
	B	0.169	0.181	0.193	0.205	0.217	0.229	0.241
Symmetric Subsidies	A	0.204	0.209	0.214	0.218	0.223	0.228	0.233
	B	0.189	0.201	0.213	0.225	0.237	0.249	0.261
Asymmetric Subsidies	A	0.214	0.218	0.223	0.228	0.233		
	B	0.193	0.205	0.217	0.229	0.241		

by reducing the probability of the worst outcomes to 0. As a result, the mitigation equilibrium is less "risky" with respect to the strategic uncertainty of the game under both subsidy treatments compared to the baseline game.

Because of the risk interdependencies, the exact probability of a loss depended on two factors: whether the player chose B and how many other players chose B. A table showed the players the exact chance of a loss for every possible outcome. Table II shows the information provided to the players for each of the three treatments. The percentages in the table indicate the probability that a player incurred a loss given their choice and the number of other players who chose B. The expected payoffs (in dollars) for the three treatments are shown in Table III, although players were not given this information.

After all choices were made each round, a random lottery number determined whether or not subjects who chose A and B incurred a loss. Players then received full information on the outcome of the round. Players were told how many other players chose B, the probability of a loss that corresponded with the outcome, the random lottery number drawn for the round, and whether they lost or not. From the information given, they could also determine whether they would have lost had they chosen the other strategy.

Each cohort played only one treatment. After playing 15 rounds of the original game, cohorts

Table IV: Summary of Treatments - Pure Coordination Game

Treatment	Choice	Cost	Number Choosing B						
			0	1	2	3	4	5	6
Baseline	A	\$0.00	40%	38%	36%	34%	32%	30%	28%
	B	\$0.00	54%	49%	44%	39%	34%	29%	24%
Symmetric Subsidies	A	\$0.00	40%	38%	36%	34%	32%	30%	28%
	B	\$0.00	47%	42%	37%	32%	27%	22%	17%
Asymmetric Subsidies	A	\$0.00	36%	34%	32%	30%	28%		
	B	\$0.00	44%	39%	34%	29%	24%		

played a new game with the same expected payoffs but where choice B was free.[5] The percentages corresponding to choice B were increased so that the expected value of choice B remained the same as in the previous game even though the choice was now free. The percentages corresponding to choice A remained unchanged. Table IV shows the new game for each of the three treatments.

In the first 15 rounds (the *IDS game*), subjects face strategic uncertainty combined with a risk mitigation choice. Conditional on the choices of the other subjects, strategy B always offers a smaller chance of a loss in return for a certain cost. After the strategic uncertainty is resolved, the preferred choices depends on an individual's value for risk mitigation. In rounds 11-25 (the *pure coordination game*), both choices are free, so individuals should prefer whichever strategy leads to the lowest probability of a loss. In this simpler game, subjects still face strategic uncertainty but the risk mitigation choice has been removed. The purpose of these rounds is to test if the simpler game leads to higher levels of coordination on the preferred equilibrium.

Because subjects played the IDS game first, there is a potential ordering effect in the pure coordination game rounds. After playing 10 rounds of the pure coordination game, cohorts played an additional ten rounds of the original IDS game. These last ten rounds will be used to test whether differences in behavior between the original game and the pure coordination game can be attributed to experience or to the differences between the games.[6]

Ninety-five participants were recruited from undergraduate classes at the University of Colorado. Five cohorts of seven people played the baseline game, five cohorts of seven played the symmetric policy treatment, and five cohorts of five played the asymmetric policy treatment. As an incentive to show up, students were offered a small amount of extra credit in one of their economics or business classes. In addition, students were told that they would earn between \$5 and \$12 for their participation in the experiments. A session lasted for 30 minutes, so participants should have expected to earn an hourly rate between \$10 and \$24. The average actual earnings in the thirty minute session were \$9.07.

The experiments took place in a computer lab on campus. All instructions were read aloud at the beginning of the session, and subjects could follow along from their computer terminal. The

[5]Five of the fifteen cohorts participated in sessions where only the initial 15 rounds were played.

[6]At the end of the session, players' risk preferences were measured using a procedure similar to that used in Holt and Laury (2002). These measures of risk aversion did not have a statistically significant effect on individual behavior.

full text of the instructions are available upon request from the author. Before the first round, subjects answered two questions about the game through their computer to demonstrate that they understood how payoffs were determined.

4 Results

Figure 1 shows the sequence of aggregate outcomes in the first fifteen rounds for each of the fifteen cohorts. The first column shows the results for the baseline game, the second column for the symmetric subsidy game, and the third column for the asymmetric subsidy game. It is clear from this figure that a higher percentage of subjects chose B in the symmetric subsidy treatment compared to the baseline game, although it is not as clear if there is a similar effect for the asymmetric subsidy treatment.

Most of the first fifteen rounds did not result in either of the predicted Nash equilibria, and play did not appear to converge toward either of the equilibria. The lack of convergence to either equilibrium can partly be explained by the fact that subjects condition future choices on the *ex post* lottery outcomes of previous rounds.[7] Subjects were likely to switch their choice immediately following a loss, regardless of whether it was the best choice in expectation. To demonstrate this, Table V shows the results of a probit where the dependent variable is 1 if the subject chose strategy B. Reported values are the marginal effects for a change in the variable with standard errors adjusted for correlation of observations across rounds for each subject.

AandLoss is a dummy variable that is 1 if the individual chose A and lost in the previous round. *BandLoss* is a dummy variable that is 1 if the individual chose B and lost in the previous round. *BandNoLoss* is a dummy variable that is 1 if the individual chose B and did not lose in the previous round. Coefficients on these three variables are relative to the omitted variable, individuals who chose A and did not lose in the previous round. The dummy variable *BestResponse* is set equal to 1 if strategy B was the best response to the choices of the other players in the previous round.[8] Subjects' earnings so far in the experiment are included to test for income effects in the later rounds. Other explanatory variables included are gender, year in school (with 0 representing first year college students), dummies for each of the subsidy treatments, and a dummy for business or economics majors.

Subjects choosing A were about 9% more likely to choose B following a loss, significant at the 5% level. To examine the effect of incurring a loss when choosing B, the difference between *BandLoss* and *BandNoLoss* indicates that subjects were about 7% less likely to choose B immediately following a loss (although this difference is not statistically significant). It is also noteworthy that subjects were no more likely to choose B immediately following rounds where B was the best response to the other players' choices. These results suggest that the salience of experiencing a loss plays an important part in individual decision-making and may inhibit convergence toward either equilibrium.

[7] A similar result is found in stochastic prisoner's dilemma games by Bereby-Meyer and Roth (2006).

[8] In Table V, the *BestResponse* variable assumes that subjects are risk neutral. Alternatively, the best response could be calculated using the estimates of risk aversion from the Holt and Laury (2002) test. The variable is not statistically significant using either approach.

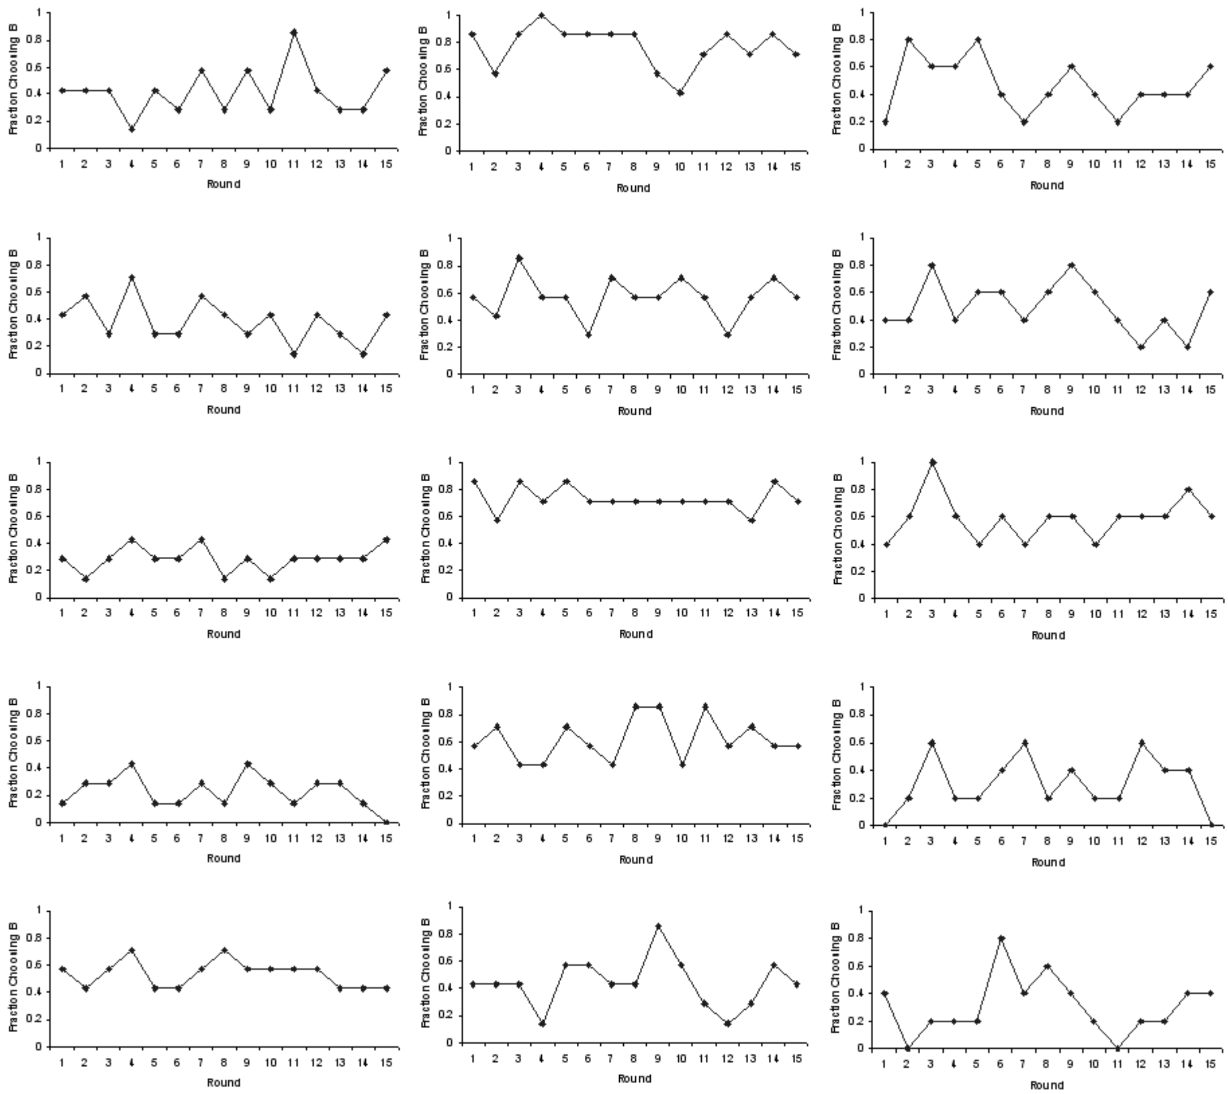

Figure 1: Aggregate Outcomes By Cohort

Table V: The Effect of Prior Outcomes

Dependent Variable: B=1, A=0	
AandLoss	0.092**
	(0.045)
BandLoss	0.255***
	(0.055)
BandNoLoss	0.321***
	(0.054)
BestResponse	-0.028
	(0.044)
Earnings	-0.019**
	(0.010)
SymmetricPolicy	0.213***
	(0.066)
AsymmetricPolicy	0.044
	(0.063)
Female	-0.027
	(0.054)
YearInSchool	0.011
	(0.027)
BusinessEcon	0.021
	(0.054)
Observations	1330

Standard errors in parentheses
Significant at the ** 5%; *** 1% level

Table VI: Pure Coordination Game vs. IDS Game

Dependent Variable: $diff$	
Rounds 16-25	-0.311***
	(0.090)
Rounds 26-35	-0.059
	(0.091)
Observations	340
Standard errors in parentheses	
*** significant at 1%	

A small income effect is apparent from Table V, with subjects 2% less likely to choose B for every dollar in accumulated earnings. Given the levels of earnings in the first 15 rounds, this effect is very small in magnitude. The symmetric policy has a strong positive effect on mitigation while the effect of the asymmetric policy is small and insignificant. Section 4.2 discusses these treatment effects in more detail.

4.1 The IDS Game vs. the Pure Coordination Game

The purpose of the pure coordination game rounds was to test if the reduced complexity of the game results in outcomes closer to the equilibrium predictions. Because subjects played these rounds after playing the IDS game, it is possible that differences observed between the pure coordination game and the IDS game are the result of experience and are not related to the change in the game. To distinguish between these two hypotheses, subjects played an additional ten rounds of the IDS game after playing the pure coordination game.

Out of 100 rounds of the pure coordination game, equilibrium was reached ten times, a higher percentage (10%) than in the first set of rounds of the IDS game (about 3%). When subjects returned to the IDS game at the end of the session, equilibrium was reached 4 times (4%), about the same frequency as in the first set of IDS rounds. To further investigate this, a variable $diff = min(n_b, N - n_b)$ was constructed representing the difference between actual outcomes and the closest equilibrium outcome, where n_b is the number of subjects who chose B. Table VI shows the results of a regression of $diff$ on a dummy variable for the pure coordination game and a dummy variable for the later risk mitigation rounds with fixed cohort effects and allowing for first order autocorrelation of the error terms.

The pure coordination game coefficient is negative and significant at the 1% level, indicating that the number choosing B in the pure coordination game is closer to one of the two equilibria compared to the original IDS game. In contrast, there is no significant difference between the final IDS rounds and the original IDS rounds. Outcomes moved closer to one of the two equilibria when subjects played the simpler pure coordination game, then moved back when subjects returned to the original game. This leads to the conclusion that the additional complexity of the IDS game partly explains the lack of convergence to either equilibria and that the changed behavior in the

Table VII: Treatment Effects

Dependent Variable: Percentage of Subjects Choosing B			
	(1)	(2)	(3)
Symmetric Policy	0.257***	0.299***	0.257***
	(0.031)	(0.034)	(0.071)
Asymmetric Policy	0.062**	0.088**	0.062
	(0.031)	(0.038)	(0.071)
Observations	225	225	225

Standard errors in parentheses
** significant at 5%; *** significant at 1%

pure coordination game is a result of the reduced complexity and not the effect of experience.

4.2 The Effect of Symmetric and Asymmetric Subsidies

To analyze the effect of the two subsidy treatments, Table VII reports estimation results of the effect of each treatment on the percent choosing B, using outcomes in the first 15 rounds. The first column is an OLS regression, ignoring any autocorrelation within cohorts or heteroscedacity across cohorts. Column 2 allows for first order autocorrelation of the error terms as well as groupwise heteroscedacity. Column 3 is a random effects model, in which the error is made up of a cohort-specific error term and a random error term, still allowing for first order autocorrelation of the random error term.

It is clear from Table VII that the symmetric subsidy treatment has a strong, positive, statistically significant effect on the percentage of subjects choosing B in each round. About 26% more subjects (almost two more subjects) choose B in each round compared to the baseline game. From Columns 1 and 2, there is some evidence that a statistically significant higher percentage of subjects chose B in the asymmetric subsidy treatment compared to the baseline game. However, we cannot conclude whether this higher percentage reflects the effect of the treatment or is simply a random cohort effect. Furthermore, even if there is a treatment effect, it is small in magnitude compared to the effect of the symmetric treatment. Between 6% and 9% more subjects choose B compared to the baseline game.

Although the symmetric subsidy treatment is clearly more effective at inducing higher percentages of individuals to mitigate the risk, that does not necessarily imply that it is the more cost effective solution. To analyze cost effectiveness, the two freebies in the asymmetric subsidy treatment are also included in the mitigation totals. Looking at the overall levels of mitigation and overall costs of the two policy treatments, the symmetric subsidy treatment results in slightly higher levels of mitigation for a slightly higher cost. For the symmetric treatment, B was chosen 329 times in the first fifteen rounds, costing \$6.58 in subsidies. The cost of the subsidies in asymmetric subsidy treatment is fixed at \$0.08 per round, resulting in a total cost of \$6.00. Altogether, choice B was chosen 162 times in addition to the 150 implicit B choices corresponding to the full subsidies. Thus, the symmetric subsidy treatment resulted in a 5.4% increase in mitigation for a 9.7% increase in cost.

The number of subjects mitigating per dollar spent is slightly higher under the asymmetric subsidy treatment. The mean level of mitigation per dollar spent is 51.86 under the asymmetric subsidy scheme compared to 50 under the symmetric subsidy scheme. When the two full subsidies are included in the mitigation totals, the asymmetric subsidy outperforms the symmetric subsidy in cost effectiveness, although the difference is not statistically significant.

5 Conclusions

This paper compared interdependent security games to coordination games without a risk mitigation decision. The IDS games played at the beginning of each session involved a costly low-risk choice and a free high-risk choice. As a result, a player's attitudes toward risk come into play in addition to the strategic uncertainty that is present in all coordination games. Equilibrium outcomes were uncommon in these rounds. Playing a game with the same expected payoffs but without the risk mitigation decision resulted in outcomes closer to an equilibrium, although not necessarily the efficient equilibrium. Returning to the original game moved outcomes away from equilibrium, indicating that the added complexity of the original game is in part responsible for the lack of convergence to an equilibrium.

The paper also tested the effectiveness of two policies at inducing mitigation in IDS games. The results suggest that the financial incentive offered under the symmetric subsidy policy has a significant positive impact on the number of subjects who mitigate although it does not lead to the efficient level of mitigation in most cases. It is less clear whether the asymmetric subsidy induces un-subsidized subjects to mitigate via the risk interdependencies. A small effect is observed which is not always statistically significant. In terms of cost effectiveness, both policies are about the same with the traditional subsidy leading to slightly higher levels of mitigation for a slightly higher cost.

It is clear from the data that subjects react to *ex post* outcomes in previous rounds when making their decisions. After a loss occurs, agents are significantly more likely to mitigate if they were not already mitigating. One explanation for this is that the importance of reducing risk becomes clear only when the consequences of not reducing risk are experienced. Further research is necessary on the impact of prior outcomes on behavior in IDS games.

References

Bereby-Meyer, Y. and A. E. Roth (2006) "The speed of learning in noisy games: Partial reinforcement and the sustainability of cooperation" *American Economic Review* **96**, 1029–42.

Cooper, R., D. V. DeJong, R. Forsythe, and T. W. Ross (1992) "Communication in coordination games" *Quarterly Journal of Economics* **107**, 739–71.

Devetag, G. (2003) "Coordination and information in critical mass games: An experimental study" *Experimental Economics* **6**, 53–73.

Harsanyi, J. C. and R. Selten (1988) *A general theory of equilibrium selection in games*, MIT Press.

Hess, R. O., C. A. Holt, and A. M. Smith (2007) "Coordination of strategic responses to security threats: Laboratory evidence" *Experimental Economics* **10**, 235–50.

Holt, C. A. and S. K. Laury (2002) "Risk aversion and incentive effects" *American Economic Review* **92**, 1644–55.

Kunreuther, H. and G. Heal (2003) "Interdependent security" *Journal of Risk and Uncertainty* **26**, 231–49.

Kunreuther, H., G. Silvasi, E. T. Bradlow, and D. Small (2007) "Deterministic and stochastic prisoner's dilemma games: Experiments in interdependent security" *National Bureau of Economic Research, Technical Working Paper 341*.

Schelling, T. C. (1978) *Micromotives and Macrobehavior*, W. W. Norton & Company.

Schmidt, D., R. Shupp, J. M. Walker, and E. Ostrom (2003) "Playing safe in coordination games: the roles of risk dominance, payoff dominance, and history of play" *Games and Economic Behavior* **42**, 281–99.

Shafran, A. P. (2008) "Risk externalities and the problem of wildfire risk" *Journal of Urban Economics* **64**, 488–95.

Straub, P. G. (1995) "Risk dominance and coordination failures in static games" *The Quarterly Review of Economics and Finance* **35**, 339–63.

Van Huyck, J. B., R. C. Battalio, and R. O. Beil (1991) "Strategic uncertainty, equilibrium selection, and coordination failure in average opinion games" *Quarterly Journal of Economics* **106**, 885–910.

Van Huyck, J. B., R. C. Battalio, and R. O. Beil (1990) "Tacit coordination games, strategic uncertainty, and coordination failure" *American Economic Review* **80**, 234–48.

Van Huyck, J. B., J. P. Cook, and R. C. Battalio (1997) "Adaptive behavior and coordination failure" *Journal of Economic Behavior and Organization* **32**, 483–503.