Conducting Longitudinal Experiments with Behavioral Models in Repeated Stackelberg Security Games on Amazon Mechanical Turk

Debarun Kar, Fei Fang, Francesco Delle Fave, Nicole Sintov, and Milind Tambe

Department of Computer Science,
University of Southern California, Los Angeles, California - 90089. USA.
{dkar,feifang,dellefav,sintov,tambe}@usc.edu

Abstract. Recently, there has been an increase of interest in domains involving repeated interactions between defenders and adversaries. This has been modeled as a repeated Stackelberg Security Game (repeated SSG). Although different behavioral models have been proposed for the attackers in these games, human subjects experiments for testing these behavioral models in repeated SSGs have not been conducted previously. This paper presents the first "longitudinal study" – at least in the context of SSGs – of testing human behavior models in repeated SSG settings. We provide the following contributions in this paper. First, in order to test the behavioral models, we design a game that simulates the repeated interactions between the defender and the adversary and deploy it on Amazon Mechanical Turk (AMT). Human subjects are asked to participate in this repeated task in *rounds* of the game, with a break between consecutive rounds. Second, we develop several approaches to keep the human subjects motivated throughout the course of this longitudinal study so that they participate in all measurement occasions, thereby minimizing attrition. We provide results showing improvements of retention rate due to implementation of these approaches. Third, we propose a way of choosing representative payoffs that fit the real-world scenarios as conducting these experiments are extremely time-consuming and we can only conduct a limited number of such experiments[1].

Keywords: Game Theory, Human Behavior Models, Repeated Stackelberg Games, Longitudinal Experiments, Amazon Mechanical Turk

1 Introduction

Whereas previous real-world deployments of Stackelberg Security Games (SSGs) to protect airports, ports or flights have been one-shot game models [13], recent work has focused on domains involving repeated interactions between defenders and adversaries. These domains include security of wildlife (repeated interactions between rangers and

---

[1] This paper is based on the following AAMAS'15 full paper- "A Game of Thrones": When Human Behavior Models Compete in Repeated Stackelberg Security Games [5]. Unlike the main conference paper, this paper focuses on the methodological contributions of conducting human subjects experiments in repeated game settings.

poachers) [14], security of fisheries (repeated interactions between coast guard and illegal fishermen) [4], forest protection or drug interdiction, and are modeled via repeated SSGs. In a repeated SSG model, the defender periodically deploys new patrol strategies (in "rounds" of the game) and the adversary observes these strategies and acts accordingly. Recent studies based on human subject experiments (HSEs) on Amazon Mechanical Turk (AMT) have shown that modeling the adversary's bounded rationality using behavioral models such as quantal response (QR) [15] and subjective utility quantal response (SUQR) [11] leads to better performance when played against human subjects in single-shot games. However, there is no prior work on comparing these behavioral models and corresponding algorithms in *repeated* SSGs through HSEs. HSEs for repeated SSGs gives rise to several key challenges.

First, to test a particular behavioral model, we require the same set of participants to play each round of the game against the defender strategy computed for that model. The participants should also start with their earnings till the previous round to ensure similarity with the real-world scenario. This requires a new game design. Second, all the algorithms require collecting attack data of all the human subjects after each round to learn the parameters of the attackers' behavioral models and then calculate defender strategy for the next round. Once a particular round has been deployed and human subjects have been notified to play that round, the participants may play at different times and hence the completion of each round may take several hours or even a few days. Third, as the participants will need to play several rounds for one experiment, some players may leave the game before they finish all the rounds, raising a new challenge of ensuring the retention rate. Fourth, the time span of each experiment would be much longer than a single-shot SSG, thus restricting the total number of experiments that can be conducted and increasing the importance of identifying representative games for testing our models.

We fill these gaps by being the first to conduct *longitudinal experiments*[2] [7] with human subjects to test the effectiveness of existing behavioral models and algorithms on repeated SSGs. We also provide the following contributions in this paper. First, in order to test the behavioral models, we design a game that simulates the repeated interactions between the defender and the adversary and deploy it on Amazon Mechanical Turk (AMT). Human subjects are asked to participate in this repeated task in *rounds* of the game. Second, we develop several approaches to keep the human subjects motivated throughout the course of this longitudinal study so that they participate in all measurement occasions, thereby minimizing attrition which was a huge cause of concern for us. We provide results showing improvements of retention rate due to implementation of these approaches. Third, we propose a way of choosing representative payoffs that fit the real-world scenarios as conducting these experiments are extremely time-consuming and we can only conduct a limited number of such experiments.

These methodological contributions towards successfully conducting longitudinal experiments with human subjects on AMT led to the testing of several behavioral models and thus assess the performances of these models against human subjects. This, in turn, helped in identifying the strengths and weaknesses of existing models in repeated

---

[2] Whereas "longitudinal study" is often used to describe research that spans years – in which measurement occasions are conducted every X years – we use the term longitudinal study because our study included 5 measurement points with a single population.

Fig. 1. Game Interface for our simulated online repeated SSG (Reward, penalty and coverage probability for a selected cell are shown)

SSGs and allowed us to develop and test a new model called SHARP [5] which mitigates the weaknesses of existing models.

2   Game Interface Design

2.1   Game Overview

In our game, human subjects play the role of poachers looking to place a snare to hunt a hippopotamus in a protected wildlife park. The game interface is shown in Fig. 1. In the game, the portion of the park shown in the map is divided into a 5*5 grid, i.e. 25 distinct cells. The players may observe many hippos in some cells and fewer hippos in other cells. This is because the density of hippos may vary from one region of the park to another but note that the density in a particular region or cell remains constant, i.e. the number of animals in a particular region does not change over time. Hippos fit our description perfectly because they are territorial animals and stay very close to their starting locations. Also to remove any bias on part of the human subjects[3], we avoided popular species like tigers, elephants, etc. Also note that in the game, the movement of the hippos within a region is random and is currently not governed by any particular animal movement model.

In order to ensure that the participants play the game with the mindset similar to that of a poacher in the real-world, we first primed the participants before the start of the game with a background story about the hardships of a poacher's life and that if they successfully poached they can both feed their families and also earn money by selling some of the meat. This was in correlation with the real-world scenario.

---

[3] "Charismatic   Megafauna"   effect: `http://www.nature.com/news/` `charismatic-mammals-can-help-guide-conservation-1.14396`

Overlaid on the Google Maps view of the park is a heat-map, which represents the rangers' mixed strategy x — a cell i with higher coverage probability x_i is shown more in red, while a cell with lower coverage probability is shown more in green. As the subjects play the game, they are given the following detailed information: the reward if they poach successfully (R_i^a), the penalty if they get caught while poaching (P_i^a) and the coverage probability (x_i) for each target i. However, they do not know the pure strategy that will be played by the rangers in each round, which is drawn randomly from the mixed strategy x shown on the game interface. Thus, we model the real-world situation that poachers have knowledge of past pattern of ranger deployment but not the exact location of ranger patrols when they set out to lay snares. In our game, there were $M = 9$ rangers protecting this park, with each ranger protecting one grid cell. Therefore, at any point in time, only 9 out of 25 distinct regions in the park are protected. In this game, a player can only see his/her snare being placed in the park. This is in compliance with a simplified assumption of the real-world scenario, i.e. we assume that a particular poacher is unaware of the snares placed by other poachers and hence his decision to place a snare is unaffected by the other snares placed in the same or any other region.

2.2 Computation of Poacher Reward

In addition to animal density, which is strongly correlated with high-risk areas of poaching [10,9,3], distance is another important factor in poaching, e.g., recent snare-density studies have found that significant poaching happens within 5 kilometers of South Africa's Kruger National Park border [6]. Therefore, the reward obtained by a poacher in successfully placing a snare at target i is calculated by discounting the animal density by a factor of the distance traveled and is calculated as follows:

$$R_i^a = int(\phi_i - \zeta * \frac{D_i}{\max_j(D_j)})$$ (1)

Here, ϕ_i and D_i refer to the animal density at target i and the distance to target i from the poacher's starting location respectively. $int(y)$ rounds the value y to the closest integer value. The parameter ζ is the importance given to the distance factor in the reward computation and may vary based on the domain. Intuitively, the reward for successfully placing a snare in a region i near the starting location and which has animal density ϕ_i, is higher than the reward obtained by successfully placing a snare in a region with the same animal density but which is at a greater distance from the starting location as compared to i.

2.3 Non-zero sum game

In our games, the minimum and maximum animal density at each cell were 0 and 10 units respectively. The poacher received a flat penalty of -1 if he was caught at any target. Also in our game, when the poacher successfully poaches, he may obtain a reward that is less than the animal density (Eqn. 1), but the defender loses a value equal to that of the animal density, i.e., the game is non-zero-sum.

Table 1. Experiment Details

Average time taken per model per payoff structure (all 5 rounds)	Average time taken for a set of participants to play each round	Number of participants who played round 1 (minimum, maximum)	Average number of participants who played all 5 rounds	Average retention rate from round 2 to round 5
2.3 weeks	4 days	(42 , 49)	38	83.69%

3 Online Longitudinal Experiments

Longitudinal studies are typically conducted to observe and understand the effects of a particular set of variables over a period of time. Such studies can be conducted with a subject pool at an University lab or by recruiting participants in an online setting like AMT. Due to low costs, availability of a diverse subject pool and scalability, AMT has recently been a more favorable choice [12]. We tested a set of behavioral models (P-SUQR, P-BSUQR, P-RSUQR) and Maximin on AMT [5] by deploying the mixed strategy generated based on each of these models repeatedly over a set of five rounds. For each model, we recruited a new set of participants to eliminate any learning bias. Due to unavailability of data, the strategy shown for each first round was Maximin. We then learned the model parameters based on previous rounds' data, recomputed and re-deployed strategies, and asked the *same* players to play again in the subsequent rounds. For each model, all five rounds were deployed over a span of weeks. Such longitudinal studies on AMT are rare at AAMAS (in fact we are not aware of any); and certainly none have been conducted in the context of SSGs. Indeed, while the total time of engagement over our 20 experimental settings was 46 weeks, each setting required on average 2.3 weeks, a duration typically not seen in related research at AAMAS (See Table 1). Conducting such experiments on AMT gives rise to a set of challenges. Therefore this section highlights our methodological contributions to conduct such experiments.

Specifically, the challenges of longitudinal studies on AMT include: (i) ensuring that participants completely remember the game details and the procedures to play the game during each round of the experiment, as otherwise we may lose significant time and effort in collecting poor quality data, especially because each setting would take more than two weeks to be completed; (ii) minimizing attrition and maximizing retention, particularly for our study which required five separate measurement occasions; (iii) setting up a payment scheme to maximize participant retention across five rounds; and (iv) lengthy rounds due to participants' long latencies (in some instances forgetfulness) in responding to notifications.

To mitigate the above challenges, we took the following steps: (i) We set up proper validation and trial games in each round of the experiment, while not over-burdening the participants with many games and thus keeping their cognitive overload at a minimum. (ii) To be eligible for initial study enrollment, participants had to commit to completing all five rounds, i.e, remain in the study through completion. (iii) We allowed respondents sufficient time (4-5 days) to respond [8], as giving them an immediate deadline to finish a particular task can result in high attrition rates. (iv) We maintained persistent contact by sending repeated reminders [1]. (v) We set up the payment scheme to consistently

reward participation in each round plus offering a relatively high completion incentive at the end of the experiment. First, we discuss below how we set up the validation and trial games. Then, we discuss the implementation details of each of the above steps taken (in the order they were implemented) to reduce attrition rates and also provide results showing the improvements due to our approaches.

3.1   Validation and Trial Games

To enhance understanding of the game, participants were asked to play two trial games in round 1, with an option to view the instructions again after each game. After the trial games, they played one validation game, and finally the actual game. The validation game consisted of a cell with maximum animal density (=10) and the coverage probability of that cell was zero, while other cells had an animal density of 1 and non-zero but equal coverage probability. The participants were expected to select the target with the maximum animal density and zero coverage. Data from subjects who played the validation game incorrectly were discarded and they were not allowed to participate in future rounds of the experiment.

From second round onwards, participants were only asked to play one trial game and then the actual game. The trial game was kept in order to remind them of the game and its details and the playing procedures. Showing only the actual game without any trial games might have resulted in the participants not playing the game properly due to forgetfulness about the game details.

3.2   Payment Scheme

In our initial payment scheme, participants were paid a fixed 'base compensation' (=$0.50) for participation in each round of the experiment and a bonus amount based on the points earned (or lost) in each round by attacking a particular target region in the game. The participants started with an initial amount of $0.50 as the bonus amount in each round. For each point earned in a particular round (i.e., if they successfully poached), they received $0.10 as the bonus compensation. For each point lost (i.e., if they were captured by the ranger), $0.10 was deducted from their current amount. The bonus at the end of a round was *not* carried forward to the next round and was paid along with the fixed participation compensation for that round. For example, for an experiment with two rounds and $0.50 as the base compensation for each round, if a participant earned a reward point of 9 in the first round and got a penalty of 1 in the second round, (s)he was paid $(0.50+(0.50+9*0.10)) = $1.90 at the end of round 1 and $(0.50+(0.50-0.10*1)) = $0.90 at the end of round 2. However, this led to very high attrition rates, i.e., very few people returned to play in each round, thus making it difficult to compare the performances on various models on a varying number of participants for each model. This is shown in Fig. 2(a). Note that we had to abandon the experiments due to high attrition rates in round 5 for PSUQR in the first trial and rounds 4 and 5 for PSUQR in the second trial. This led us to implement a payment scheme which is discussed below.

We made three changes to our first method of compensation. First, we introduced a 'completion bonus' (=$2.50) for completing all the rounds of the experiment. Second,

like before, to motivate the subjects, the participants were incentivized based on the reward/penalty of the region they chose to attack. However now, while the base compensation was paid after each round was completed, the incentivizing bonus was carried forward from one round to the next and paid along with the completion bonus at the end of all the rounds of the experiment. This incentivizing bonus is called the 'performance bonus'. Third, the players now started with an initial base performance bonus amount of $1.50 in round 1 and they could win up to a maximum and a minimum amount in each round and hence at the end of all the rounds, based on how successful they were. We set the base compensation fee for each round to be only $0.50, thus resulting in a total base compensation fee of $2.50 over 5 rounds. However, the maximum amount they could potentially earn at the end of all the rounds from only the performance and completion bonus was as high as $7.60. The performance and completion bonus together at the end of all the rounds was much higher as compared to the total base compensation earned for playing all the 5 rounds. This ensured that majority of the participants remained motivated and returned to play all the rounds.

Taking the previous example of a two-round experiment where a participant earned a reward point of 9 in the first round and a penalty of 1 in the second round, (s)he was paid $0.50 at the end of round 1 (the participation bonus for round 1). (S)he also earned a performance bonus of $(1.50+9*0.1) = $2.40 in round 1 which was carried forward to round 2 and *not* paid at the end of round 1. Then at the end of round 2 she was paid $(0.50+(2.40-1*0.1)+2.50) = $5.30 (participation bonus for round 2 (=$0.50) + performance bonus at the end of round 2 (=$2.30) + completion bonus (=$2.50)). On an average, each participant was paid $7.60 upon completion for our five-round experiments. The effect of this payment scheme on participant retention rate can be seen in Fig.2(b).

3.3   Initial Study Enrollment

Even though the implementation of the new payment scheme saw an increase in retention rate as shown in Fig. 2(b), there was still a decrease in retention rates over rounds. Therefore, we implemented an approach where the participants had to commit to completing all five rounds before starting the first round of the game. They were asked to either 'agree' or 'disagree' to this commitment. On an average, 96% of the participants who enrolled in AMT for our study agreed to this commitment. These participants were then allowed to proceed towards playing the first round of the game. On the other hand, if they did not agree, they were thanked for their interest in our study, but not allowed to participate any further. The effect of this on the retention rate can be seen in Fig. 2(c). This clearly shows that a significant number of participants with prior commitment towards completing all the rounds of the experiment, returned and completed all the rounds.

3.4   Reminder Emails

Even though the implementation of the payment scheme and initial study enrollment procedures increased the retention rate as shown in Fig. 2(c), the retention rate still decreased over rounds for some of the experiments, even though at a slower rate. There-
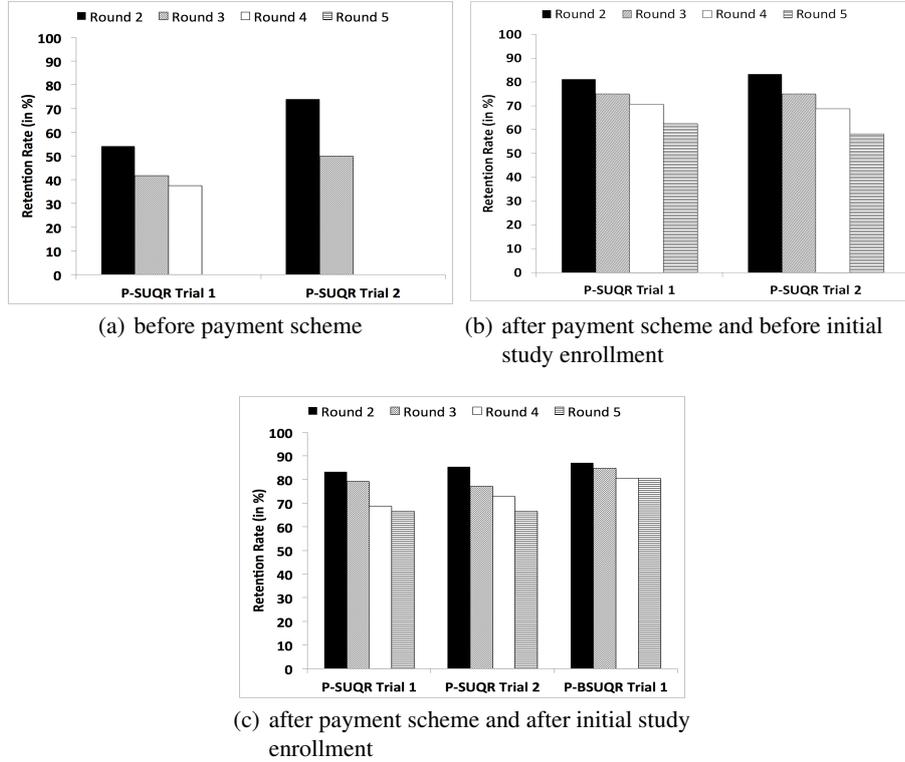
(a) before payment scheme

(b) after payment scheme and before initial study enrollment

(c) after payment scheme and after initial study enrollment

Fig. 2. Retention Rates for various models (a) before implementation of our payment scheme, and (b) after implementation of our payment scheme and before implementation of initial study enrollment procedure.

fore, we sent repeated reminders to the participants with clearly stated deadlines to ensure that they (i) do not forget to participate in the current round, and (ii) also remain motivated throughout the study. The construction of the email also turned out to be crucial in ensuring that more participants returned to play the next round. Specifically, the italicized portion in the sample email shown below proved effective. Results are shown in Figs. 3(a) and 3(b).

Sample Email:

Hi,

Thank you for participating in our experiment. Your base compensation for round 3 has been paid to you via AMT. Thank you also for your valuable comments and suggestions about the game and its strategies. We will definitely take those into account later on. Now, we would want you to participate in the 4th round of our experiment. Please follow the link below to participate: `http://cs-server.usc.edu:16568/gamelink/index.jsp`
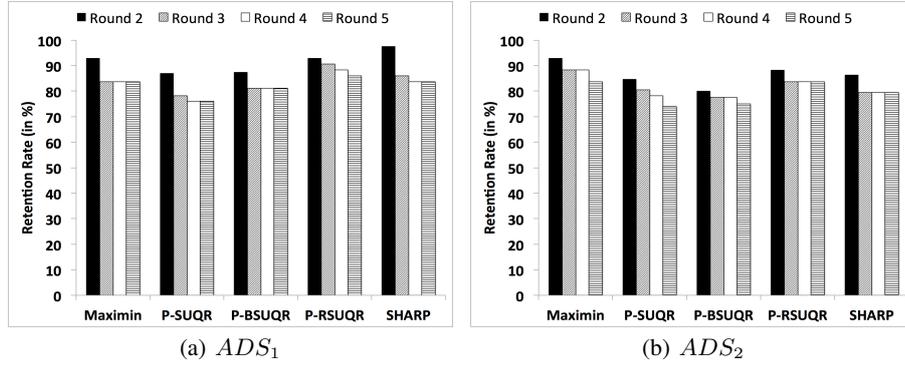
(a) ADS_1 (b) ADS_2

Fig. 3. Retention Rates for various models over 4 rounds, starting from round 2 to round 5, on (a) ADS_1 and (b) ADS_2 respectively.

In the first page, please read carefully the compensation details. You will be starting with the performance bonus that you earned in the last round. **The last date to participate in this round of our experiment is Wednesday (November 6 2014) 4 pm PST**. Please try to complete the experiment by the deadline because otherwise deployment of the next round gets delayed.

You are very important to the study and your continued participation is critical. Don't be discouraged if you got caught by a ranger in this round. The chance to play again and earn performance and completion bonuses are coming in a few days. We look forward to your continued participation.

Thank you.

3.5   Discussions

In this section, we discuss some of the key strengths and limitations of our approaches. (i) The implementation of each additional strategy (such as the payment scheme, email reminders, etc) required us to restart the experiment every time a new strategy was implemented. Also, since some of these approaches were not tested in isolation, but were only tested in combination with other approaches, we do not know how some of our strategies for maximizing retention in isolation affected retention rates (except for the first strategy, i.e. the new payment scheme we employed). So, if we only had implemented email reminders but did not give them sufficient time to respond, retention might be different and vice versa. (ii) It is possible to run parallel experiments with our game as we recruit completely different sets of participants for each of our 5-round experiment. We recruited different sets of players for each of our 5-round experiments because the knowledge accumulated by the subjects from a set of rounds of the game may affect the decision made by the same subjects if they were asked to play from round 1 in a different experiment. (iii) The problems addressed by the development of each of our strategy (for example, the choice of concrete metrics, payoffs and incentive
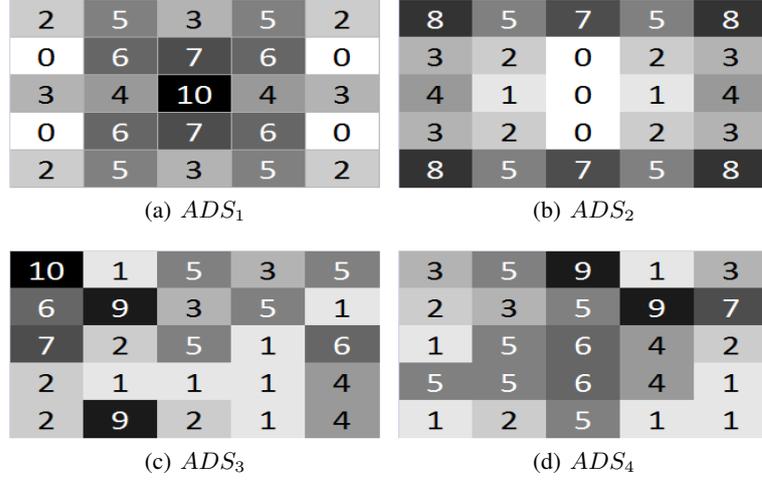
(a) ADS_1

(b) ADS_2

(c) ADS_3

(d) ADS_4

Fig. 4. Animal density structures

mechanisms) does indeed lead to the development of some methodological contributions towards conducting such longitudinal experiments on crowdsourcing platforms like AMT. The principles followed are *not* only based on common sense, but previous literature [2] and also the results shown in Figs. 2(a)-2(c) and 3(a)-3(b) .

4 Payoff Structures

Selecting the appropriate set of payoff structures is critical as the total number of experiments that can be conducted is limited due to the time taken to conduct each experiment. The payoff structures used in our human subjects experiments vary in terms of the animal densities and hence the adversary rewards. We henceforth refer to payoff structures and animal density structures interchangeably in this paper. The total number of animals in all the payoffs we generate is the same (= 96). However, the variation in these payoffs is in the way that the animals are spread out in the park. In payoff structure 1, the animal density is concentrated towards the center of the park, whereas the animal density is higher towards the edges of the park in payoff structure 2. These represent scenarios that might happen in the real world. The animal density for both payoffs is symmetric, thus eliminating any bias due to the participant's starting point in the game. Contrary to the above, animals in the park may be randomly scattered without any particular orientation. So, we randomly generate two additional animal density structures (payoffs 3 and 4) and test our proposed model on these payoffs. To generate a random structure, one out of 25 cells was chosen uniformly at random and then an animal was allocated to it until the total number of animals in all the cells was 96, keeping in mind that if a cell total reached 10 (maximum animal density in a cell), that cell was not chosen again. Figs. 4(a), 4(b), 4(c) and 4(d) show heatmaps of these animal density structures, denoted as ADS_1, ADS_2, ADS_3 and ADS_4 respectively.

5   Results of Human Subjects Experiments

The methodological contributions discussed in this paper resulted in conducting successful human subjects experiments with existing adversary behavioral models. This also led to the identification of certain limitations in such models and hence the development of a novel human behavior model called SHARP for repeated SSG settings. Although the details of this model are discussed in detail in [5], we present here some of the performance results of SHARP against existing models in the human subjects experiments we conducted. The performance is measured in terms of the defender utilities per round and cumulative defender utilities over rounds. The interesting observation from the following results is that SHARP not only performs well in initial rounds of the game, unlike existing models which incur heavy losses, it also performs consistently well in the subsequent rounds.

5.1   Defender Utilities

In Figs. 5(a)- 5(d) we show actual defender utilities obtained over 5 rounds for P-SUQR, P-BSUQR, P-RSUQR, SHARP and Maximin on ADS_1, ADS_2, ADS_3 and ADS_4 respectively, with an average of 38 human subjects playing per round. In the plots, y-axis corresponds to defender utility and the models tested for each round is shown on the x-axis. For example, in Fig. 5(b), P-SUQR performs worst in round 2 with an utility of -5.26. In Fig. 5(b), we also show (inset) zoomed in results of the second round to highlight the difference in performance between Maximin (= -0.18) and SHARP (= -0.1). First round utilities for all models are same as Maximin strategy was played due to absence of data. All significance results reported below are computed with bootstrap t-test. Following are key observations from our experiments.

Heavy initial round losses: For all models except SHARP, there is statistically significant (p=0.05) loss in defender utility as compared to Maximin in second round. P-SUQR recovers from initial round losses and outperforms Maximin in rounds 3, 4 and 5 for ADS_1 (statistically significant at p=0.05), and in round 4 (statistically significant at p=0.15) and round 5 for ADS_2. P-RSUQR, which is a robust model, also outperforms Maximin in rounds 4 and 5 (statistically significant at p=0.05) for ADS_1 after initial round losses. Surprisingly, P-BSUQR, which is the basis for wildlife security application PAWS, performs worst on both payoffs over all rounds. Figs. 5(e)- 5(h) show cumulative defender utility over five rounds on ADS_1, ADS_2, ADS_3 and ADS_4 respectively. Observe that it takes five rounds for P-SUQR to recover from initial round losses and outperform Maximin in terms of cumulative defender utility for ADS_1 (Fig. 5(e)). None of the other models recover from the initial round losses on either payoffs in five rounds, thus highlighting the impact initial round losses have on model performance for a long period of time.

Performance of SHARP against other models: SHARP consistently outperforms (statistically significant at p=0.05) all the models over all rounds (Figs. 5(a)- 5(d)), most notably in initial rounds (round 2) and ends up with significantly high cumulative utility at the end of all rounds (Figs. 5(e)- 5(h)).

Therefore, our results on extensive human subjects experiments on repeated SSGs show SHARP's ability to perform well throughout, including the important initial rounds.
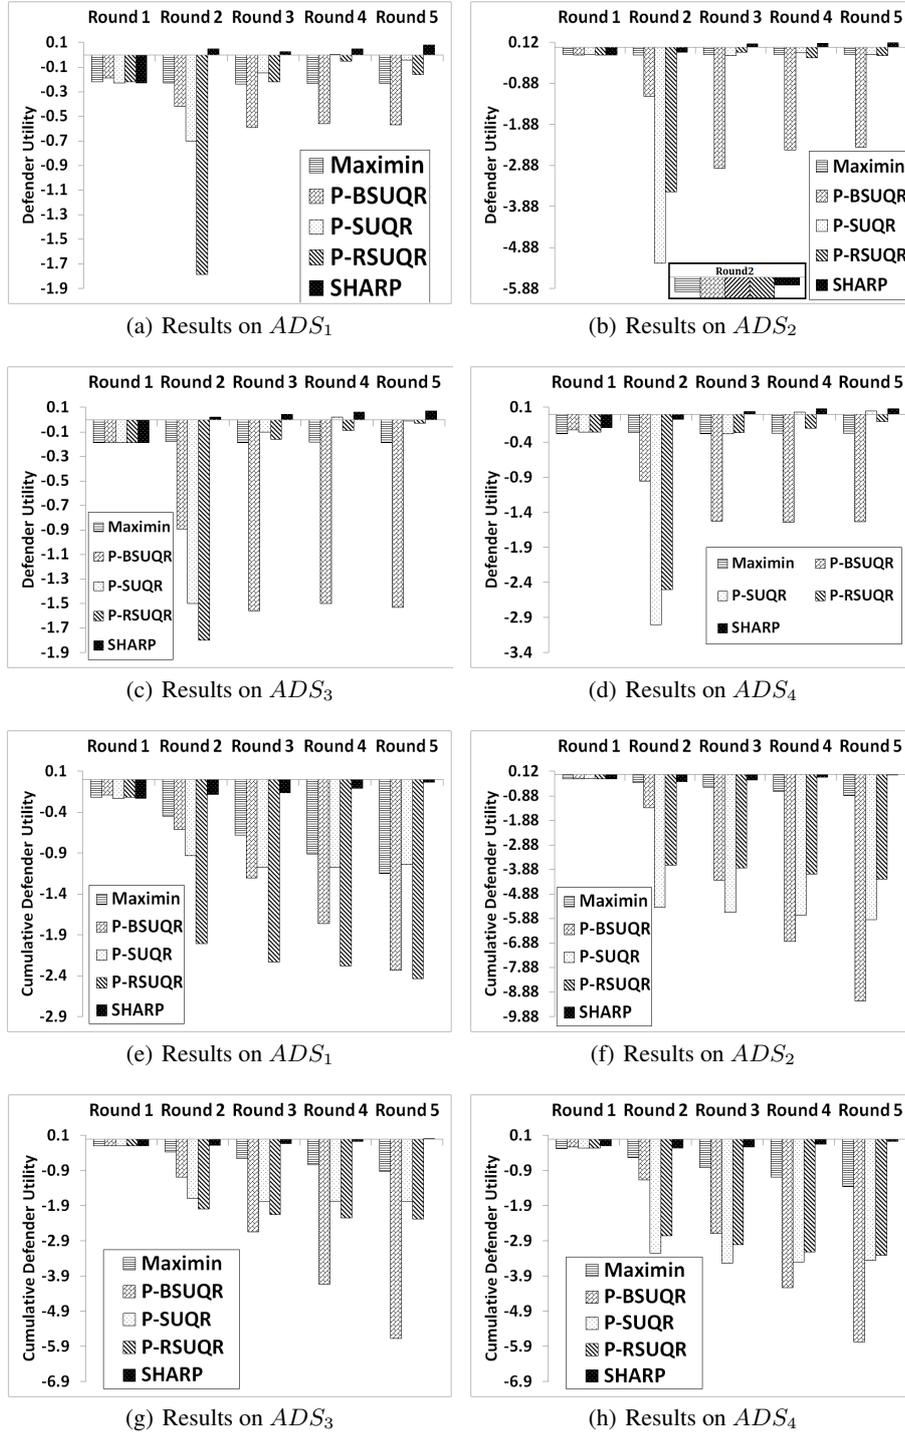
(a) Results on ADS_1

(b) Results on ADS_2

(c) Results on ADS_3

(d) Results on ADS_4

(e) Results on ADS_1

(f) Results on ADS_2

(g) Results on ADS_3

(h) Results on ADS_4

Fig. 5. (a), (b), (c) and (d): Defender utilities for various models on ADS_1, ADS_2, ADS_3 and ADS_4 respectively; (e), (f), (g) and (h): Cumulative defender utilities for various models on ADS_1, ADS_2, ADS_3 and ADS_4 respectively.

6   Conclusions

In this paper, we discussed about an online simulation game that we developed for conducting longitudinal human subjects experiments to test the performances of various behavioral models in repeated SSGs. We deployed this game on AMT and faced several challenges while conducting these experiments, each of which lasted several weeks. Therefore, we present our methodological contributions towards conducting these longitudinal experiments successfully on AMT and experimental results of the effectiveness of our approaches. Specifically, our contributions include: (i) developing an empirically supported payment scheme to reduce the attrition rate of participants for our repeated task; (ii) sending carefully worded reminders periodically to maximize participation in all measurement occasions; and (iii) generate a set of representative payoff structures to test our behavioral models. These contributions allowed us to successfully test a suite of existing behavioral models and identify their strengths and weaknesses in repeated SSGs and finally develop and test a novel human behavior model called SHARP which outperforms existing approaches [5].

7   Acknowledgements

References

1. Cotter, R.B., Burke, J.D., Stouthamer-Loeber, M., Loeber, R.: Contacting participants for follow-up: how much effort is required to retain participants in longitudinal studies? Evaluation and Program Planning 28(1), 15–21 (2005)
2. Estrada, M., Woodcock, A., Schultz, P.W.: Tailored panel management: A theory-based approach to building and maintaining participant commitment to a longitudinal study. In: Evaluation Review (2014)
3. Hamisi, M.: Identification and mapping risk areas for zebra poaching: A case of Tarangire National Park, Tanzania. Thesis, ITC (2008)
4. Haskell, W., Kar, D., Fang, F., Tambe, M., Cheung, S., Denicola, E.: Robust protection of fisheries with compass. In: Innovative Applications of Artificial Intelligence (IAAI) (2014)
5. Kar, D., Fang, F., Fave, F.D., Sintov, N., Tambe, M.: "A Game of Thrones": When human behavior models compete in repeated stackelberg security games. In: International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2015) (2015)
6. Lemieux, A.M.: Situational Crime Prevention of Poaching (Crime Science Series). Routledge (2014)
7. Mao, A., Parkes, D., Chen, Y., Procaccia, A.D., Gajos, K.Z., Zhang, H.: Turkserver: Enabling synchronous and longitudinal online experiments. In AAAI HCOMP Workshop (2012)
8. Menard, S.W.: Handbook of longitudinal research: Design, measurement, and analysis. Academic Press (2008)
9. Montesh, M.: Rhino poaching: A new form of organised crime. Tech. rep., College of Law Research and Innovation Committee of the University of South Africa (2013)
10. Moreto, W.: To Conserve and Protect: Examining Law Enforcement Ranger Culture and Operations in Queen Elizabeth National Park, Uganda. Thesis, Rutgers (2013)

11. Nguyen, T.H., Yang, R., Azaria, A., Kraus, S., Tambe, M.: Analyzing the effectiveness of adversary modeling in security games. In AAAI (2013)
12. Paolacci, G., Chandler, J., Ipeirotis, P.G.: Running experiments on amazon mechanical turk. Judgment and Decision Making pp. 411–419 (2010)
13. Tambe, M.: Security and Game Theory: Algorithms, Deployed Systems, Lessons Learned. Cambridge University Press, New York, NY (2011)
14. Yang, R., Ford, B., Tambe, M., Lemieux, A.: Adaptive resource allocation for wildlife protection against illegal poachers. In: International Conference on Autonomous Agents and Multiagent Systems (AAMAS) (2014)
15. Yang, R., Kiekintveld, C., Ordonez, F., Tambe, M., John, R.: Improving resource allocation strategy against human adversaries in security games. In IJCAI (2011)