A New Paradigm for the Study of Corruption in Different Cultures

Ya'akov (Kobi) Gal[1], Avi Rosenfeld[2], Sarit Kraus[3,4],
Michele Gelfand[4], Bo An[5], and Jun Lin[6]

[1] Department of Information Systems Engineering, Ben-Gurion University, Israel
[2] Department of Industrial Engineering, Jerusalem College of Technology, Israel
[3] Department of Computer Science, Bar-Ilan University, Israel
[4] Institute for Advanced Computer Studies, University of Maryland, USA
[5] Department of Engineering, Nanyang Technological University, Singapore
[6] Department of Engineering and Management, Beihang University, China
kobig@bgu.ac.il, rosenfa@jct.ac.il, sarit@cs.biu.ac.il,
mjgelfand@umd.edu, boan@ntu.edu.sg, linjun@buaa.edu.cn

Abstract. Corruption frequently occurs in many aspects of multi-party interaction between private agencies and government employees. Past works studying corruption in a lab context have explicitly included covert or illegal activities in participants' strategy space or have relied on surveys like the Corruption Perception Index (CPI). This paper studies corruption in ecologically realistic settings in which corruption is not suggested to the players a priori but evolves during repeated interaction. We ran studies involving hundreds of subjects in three countries: China, Israel, and the United States. Subjects interacted using a four-player board game in which three bidders compete to win contracts by submitting bids in repeated auctions, and a single auctioneer determines the winner of each auction. The winning bid was paid to an external "government" entity, and was not distributed among the players. The game logs were analyzed posthoc for cases in which the auctioneer was bribed to choose a bidder who did not submit the highest bid. We found that although China exhibited the highest corruption level of the three countries, there were surprisingly more cases of corruption in the U.S. than in Israel, despite the higher PCI in Israel as compared to the U.S. We also found that bribes in the U.S. were at times excessively high, resulting in bribing players not being able to complete their winning contracts. We were able to predict the occurrence of corruption in the game using machine learning. The significance of this work is in providing a novel paradigm for investigating covert activities in the lab without priming subjects, and it represents a first step in the design of intelligent agents for detecting and reducing corruption activities in such settings.

1 Introduction

Corruption – the abuse of entrusted power for private gain, is a global phenomenon that severely diminishes economic growth, education opportunities,

availability of health and welfare services, and undermines the ability of governments to implement needed policies [1]. Because corruption is inherently a covert activity and difficult to prove, it is notoriously hard to measure, and convictions are few and far between, especially in countries where corruption is often viewed as permissible. Thus most assessment tools have relied on surveys, such as the well-established Corruption Perception Index (CPI) which ranks countries based on their perceived levels of corruption, as determined by expert assessments and opinion surveys [8].

Lab experiments have studied corruption as three-player games involving a briber, a bribee, and a third party which is negatively affected by the bribe [9,2]. Other works have studied how wages and detection probabilities affect corruption levels when a public "official" can make a unilateral decision about how much money to divert from public funds, at a risk of being detected and punished [5,6]. In all of these works, illegal activities and bribes are primed or explicitly stated in players' strategy space, which does not reflect the secretive and dynamic nature in which corruption develops. A notable exception is the work by Falk and Fischbacher [4] in which illegal activities are framed as "take" actions, and are shown to reciprocate over time in complete information settings. They did not simulate a repeated task setting but asked subjects to provide a complete description of their strategy. Our study is more realistic in that we did not make references to any covert activities, participants could not observe such activities in the game and interactions among participants were conducted in real-time.

The goal of this paper is to study corruption in a realistic way as an evolving process in which participants engage in repeated economic activities. We used a board game in which three bidders compete to win contracts by submitting bids in repeated auctions, and a single auctioneer determines the winner of each auction. Bidders earn bonus points in the game if their bids are accepted. The auctioneer's score is constant and does not depend on their bids. The winning bid is paid to an external "government" entity. At given times in the bidding process, participants could exchange private messages with each other.

We played 276 games in Israel, the U.S. and China spanning hundreds of subjects. We analyzed the logs that were generated by the games posthoc in each country. We measured two types of corrupt activities, one in which the auctioneer did not pick the bidder with the highest bid, and one in which there was a distinct bribe that was transferred from a bidder to the auctioneer in return for getting chosen.

The results showed that corruption occurred in about 33% of games played in the U.S., 29% of games played in Israel and 56% in the games played in China. These results partially follow the CPI, in that China was the country that exhibited the largest amount of corruption. However, we measured a lower level of corruption in Israel as compared to the U.S. despite the higher CPI level in Israel. We also found that corruption benefited the auctioneer when compared to the case in which there was no corruption. However, while the bribing players benefited from this activity in Israel and in China, they did not benefit from bribery in the U.S. Further analysis revealed that bribes in the U.S. were often
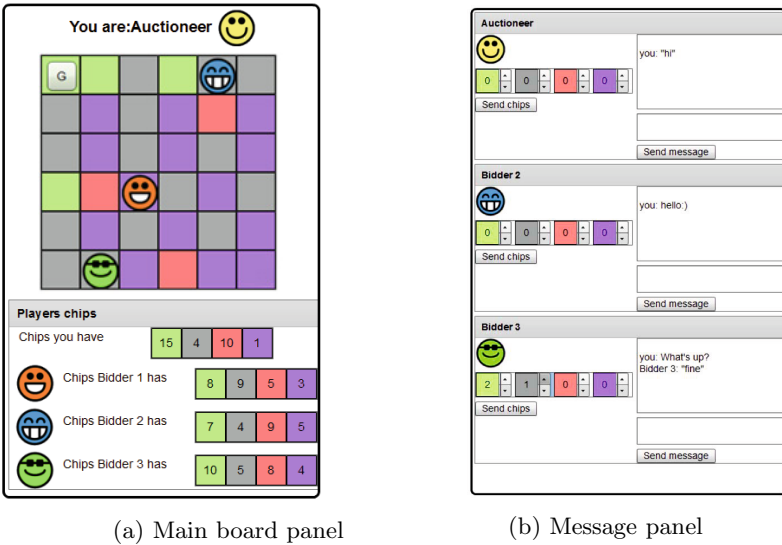
(a) Main board panel   (b) Message panel

Fig. 1. Snapshots of Olympic Game GUI

excessively large, which prevented the winning bidders from completing their tasks. In all countries, we found that there was a significant increase in the ratio of corruption in the last round of play. Additionally, we were able to predict the occurrence of corrupt activities in all countries using supervised learning approaches using information that was publicly available during the game.

The contributions of this paper are threefold. First, we provide a new empirical framework to study corruption in a lab setting without priming subjects. Second, we show that corruption is endemic to people's interaction in different countries in a way that is not predicted by the CPI. Third, we provide a predictive model for the occurrence of corruption in the lab.

2   Implementation Using Colored Trails

Our empirical study was based on a test-bed called Colored Trails [7], which we adapted to a multi-round contract allocation game called the Olympic Game. In the Olympic Game, a city is preparing to host the Summer Olympics by creating the necessary infrastructure such as hotels, transportation and restaurants. There are four participants in the game: three *bidders* and one *auctioneer*. The game comprises a board game of colored squares which includes players' icons and a goal square. Players are allocated a set of chips of colors chosen from the same palette as the board square. All players have full view of the board game and their allocated chip sets, however bidders cannot observe the chips of other players.

Figure 1 shows the main game panel from the point of view of the auctioneer. At the heart of the game is players' abilities to exchange resources and messages

with each other. The bidders send bids to the auctioneer in the form of chips. The auctioneer has full authority to make decisions about which bidder to choose for each project. The chips for the winning bid are automatically transferred to an external "government" and is not distributed among the players.

The Olympic Game is a variant of a family of board games that are analogous to task settings in the real world. Paths on the board from a bidder's position to the goal square represent carrying out a contract, with each square in the path corresponding to one of the tasks in the contract such as building a restaurant, hiring workers, and acquiring permits. To move an icon into an adjacent square a bidder player must turn in a chip of the same color as the square. Chips represent resources that can be used for bidding purposes and for completing the contracts. These chips directly translate to monetary units at the end of the game and are awarded to the subjects. The advantage of using the Colored Trails game is that it creates an environment which "abstracts away" domain specific details while still providing a task-like context in which decisions are made. The abstraction provided by Colored Trails is especially appropriate for investigating decision-making in different cultures because it avoids culturally-loaded contexts (e.g., religious connotations) that may confound people's behavior.

The Olympic Game is played repeatedly for an undisclosed number of rounds (four rounds in our study). In each round, a new contract is generated for a different project for the Olympic city. The goal for each round is to win the bid and advance to the goal square. Each round is played on a new board game, in which players' location on the board as well as their initial chip allocations may change. Players' roles do not change between rounds (i.e., the same auctioneer chooses the winning bid in all rounds).

The interaction in each round of the game proceeds in a sequence of phases with associated time limits. In the *strategy phase* (60 seconds), all players are given a chance to study the board and to think about possible strategies in the game. In the *exchange phase* (90 seconds), players may exchange resources and messages in free text with each other using the game interface. An example is shown in Figure 1 (bottom). The messages can be initiated by any of the participants, and are only visible to the message initiator and receiver. In the *bidding phase* (60 seconds), each of the bidders can choose to make a bid to win the contract for the round. The bid is the number of chips the bidder is willing to pay the government for the purpose of winning the contract. Bidders cannot see each others' bids. All of the bids are relaid to the auctioneer at the end of the bidding phase. At this point, the auctioneer can choose one bidder to win the contract. In the *execution phase* the bid of the winning bidder is deducted from the chips in its possession, and the winning bidder is advanced to the goal (automatically) given its available resources. In particular, if the winning bidder cannot get to the goal square, its icon is moved as close as possible to the square using its available chip set. Note that only the winning bidder is allowed to move towards the goal, and is deducted the chips used to move towards the goal.

At the end of each round, the score for each participant is computed as follows: One point is given for each chip in the possession of a participant at the end of

the round; a 100 point bonus is given to a winning bidder who is able to reach the goal. Otherwise, one point is deducted for every square in the path from the final position of the winning bidder and the goal square; a 5 point bonus is given to an auctioneer for choosing a winning bidder that is able to reach the goal square. There are several reasons for choosing these parameters. First, they reflect the fact that reaching the goal is the most important component in the game. Reaching the goal is analogous to completing the contract for which a winning bidder is chosen using the resources that are available to the bidder at the execution phase. The auctioneer is an employee of the government, and receives a constant salary for each round, as represented by the chips that are allocated to the auctioneer at the onset of the round. As a government employee, it receives only a nominal bonus for choosing a bidder that wins the contract.

3   Empirical Methodology, Setup, and Results

We ran 101 instances of the game in the U.S., 91 instances in Israel and 84 instances in China, totaling 276 participants. Our analysis based on those game instances in which one of the bidders was chosen (and in which corruption or bribery had a chance to occur according to our definition).[1] In all countries subjects were students enrolled in a university or college degree program. Each participant was given an identical 30 minute tutorial on the Olympic game which consisted of a self-guided presentation. Participation in the study was contingent on passing an on-line quiz about the game. There were at least 8 subjects in the lab at any given session (playing two consecutive games), and participants could not see others' terminals or communicate to each other in any way outside of the computer interface. All participants were paid a constant sum for their participation (about $10) plus a bonus that depended on their performance in the game (an additional $2-$7). The participants were randomly allocated to their respective roles in the game. All games comprised four rounds, but this information was not made public to the subjects.

We defined two classes of covert activities in the game. Corruption was defined as the case in which the auctioneer did not choose the highest bidder (or one of them in case of ties). Bribery was defined as a special case of corruption in which the auctioneer accepted resources from a bidder who was subsequently chosen to win the contract despite not submitting the highest bid. We hypothesized that the occurrence of corruption will closely follow the Corruption Perception Index level in each country. Thus, we expected China to exhibit the highest level of corruption and bribery in the game, followed by the Israel and the U.S. In addition, we hypothesized that the number of instances of corruption will increase with the number of rounds played in the same game. Lastly, we expected that corruption and bribery will result in diminished profit for the government, and increased revenue for the auctioneer and bidder. All results

---

[1] There was no bidder chosen for a contract only in a minority of cases, less than 8 games in each country.

Table 1. Corruption and bribery measures for all countries

	Bribery	Corruption
U.S.	19%	33%
Israel	14%	29%
China	**46%**	**56%**

reported in this section were verified for statistical significance in the $p < 0.05$ range using appropriate ANOVA, chi-square and t-tests.

To illustrate a corrupt exchange in one of the games, consider the following discourse fragment between the auctioneer (A) and bidders (B1, B2, B3) in one of the games played in the U.S: A to B1: "wanna send up some chips? B3 sent me some." B2 to A: "Do you send me some back?" A to B2: "I make the decision on the final choice of bidding, so no". B2 to A: "Or is this just a bribe? haha" In this game, B2 sent the auctioneer a bribe of 24 chips and was chosen as the winning bidder despite submitting the lowest bid.

Table 1 shows the percentage of corruption and bribery activities in the game. These results partially follow the Corruption Perception Index in that China was the country that exhibited the largest amount of corruption. However, we measured lower level of corruptions in Israel than in the U.S., despite the higher corruption level index in Israel.

There was no significant correlation between the number of rounds played and the occurrence of corruption in any of the countries. However, in all countries, the number of corruption cases increased significantly from round 1 to round 4. Specifically, the corruption ratio increased from 25% to 57% in the U.S.; from 21% to 46% in Israel; and from 37% to 67% in China.

Figure 2 compares the performance (measured by average score) of the government ("gvrn"), auctioneer ("acnr") and the winning bidder ("bdr") for those games that exhibit and do not exhibit corruption in the three countries. (The results for bribery activities followed the same trend and are not reported due to brevity concerns). As shown by the figure, the auctioneer significantly benefited from corruption in the U.S., but not in China and in Israel (there was no significant difference in the score to auctioneer with and without corruption in these countries). In addition, while winning players significantly benefited from corruption in Israel and in China, they did not benefit from bribery in the U.S.

To explain this result, we measured the average number of chips sent by the winning bidder, and the number of times the winning bidder defaulted on its contract (was not able to complete its task) when corruption existed. We found that the average number of chips sent in China (11.79) was highest, followed by the U.S. (7.14) and Israel (4.2). In addition, we found that there were 6 instances of failed contracts following corruption in the U.S. as opposed to 3 instances in China and only 1 instance in Israel. Such failed contracts occurred because the winning bidder did not retain enough chips to be able to complete its task.

Fig. 2. Average scores for participants in the 3 countries

Table 2. Modeling corruption overall and across each of the three countries with all attributes and exclusively with the StateProfit value

	All – Acc.	Recall	China – Acc.	Recall	Israel – Acc.	Recall	U.S. – Acc.	Recall
Corr. (all)	77.54%	0.75	72.62%	0.83	75.82%	0.37	74.26%	0.59
Corr. (profit)	**80.07%**	**0.78**	**77.38%**	**0.89**	**79.12%**	**0.48**	**76.23%**	0.56

4   A Machine Learning Model of Corruption

One goal of this study was to quantify and predict when corruption occurs. In constructing our models, we considered all game specific data and general demographic attributes. Game specific information included: the initial position of each bidder, the round number, the board configuration, and the distance of that bidder from the goal. We also considered the known outcome of the game: how much the state profited from the winning goal. Demographic information included the bidders' age, sex, and country. We intentionally did not consider private information relating to messages and chip exchanges or identity of winning bidders. We employed a standard decsision tree classifier using ten-fold cross validation.

We first considered a cross-cultural model for corruption, the results of which are found in left-most section of Table 2. In constructing these results, we observed that the decision tree for all of the data in predicting corruption had the following rule: If StateProfit $<= 15$, then a series of three rules involving the distance of Bidder1 from the goal, the Auctioneer's gender, and Bidder2's age, otherwise, if StateProfit > 21 then there is no corruption, otherwise again a series of rules involving Bidder2's age and Bidder1's distance from goal.

We found that the country of origin attribute did not constitute a main attribute within the decision trees. This was noted from the absence of this attribution from the output of the decision trees, implying that the rules primarily based on StateProfit are independent of the 3 countries we considered.

However, as the previous section documented, we did observe that differences exist across cultures. Thus, we postulated that explicitly creating decision trees for each culture might yield additional insights, the results of which are found in columns 3–8 of Table 2. These results show, as was the case previously, that the accuracy of corruption models exclusively with the StateProfit attribute were more accurate and usually yielded higher recall than those with all attributes (these instances are bolded within the table). This again confirms the significance of this attribute.

5   Conclusions and Future Work

This paper presented a new empirical framework to investigate corruption. This novel setting allows us to study people's actions without biasing their actions through explicitly including corruption as a factor within the game, or merely noting their subjective feelings about corruption. This work has implications for both the social and computational sciences. It presents a new tool for studying corruption in the lab. We do not claim that corrupt behavior in the game predicts the occurrence of corruption in real life. We *do* claim that covert activities occur in a more realistic manner than in traditional studies. Our test-bed can be used to study the effects of different mechanisms and policies on the evolution of corruption over time. It opens the door for designing computer agents (whether by researchers or students [3]) that play the game and attempt to reduce corruption by adopting different strategies, based on machine learning methods.

References

1. United nations handbook on practical anti-corruption measures for prosecutors and investigators (2012)
2. Abbink, K., Irlenbusch, B., Renner, E.: An experimental bribery game. Journal of Law, Economics, and Organization 18(2), 428–454 (2002)
3. Chalamish, M., Sarne, D., Lin, R.: The effectiveness of peer-designed agents in agent-based simulations. Multiagent and Grid Systems 8(4), 349–372 (2012)
4. Falk, A., Fischbacher, U.: "Crime" in the lab-detecting social interaction. European Economic Review 46(4), 859–869 (2002)
5. Fehr, E., Gächter, S., Kirchsteiger, G.: Reciprocity as a contract enforcement device: Experimental evidence. Econometrica 65(4), 833–860 (1997)
6. Fehr, E., Kirchsteiger, G., Riedl, A.: Does fairness prevent market clearing? an experimental investigation. The Quarterly Journal of Economics 108(2), 437–459 (1993)
7. Gal, Y., Grosz, B., Kraus, S., Pfeffer, A., Shieber, S.: Agent decision-making in open mixed networks. Artificial Intelligence 174(18), 1460–1480 (2010)
8. Transparency international. Corruption perceptions index (2011)
9. Treisman, D.: The causes of corruption: a cross-national study. Journal of Public Economics 76(3), 399–457 (2000)