Cooperation in an Uncertain and Dynamic World

Gallo, Edoardo and Riyanto, Yohanes E. and Roy, Nilanjan and Teh, Tat-How

30 December 2019

Cooperation in an Uncertain and Dynamic World*

Edoardo Gallo† Yohanes E. Riyanto‡ Nilanjan Roy§
Tat-How Teh¶

December 30, 2019

Abstract

We investigate how reputational uncertainty and the rate of change of the social environment interact to influence cooperation in social networks. Reputational uncertainty significantly decreases cooperation and welfare, induces more forgiveness toward defectors, and promotes opportunistic play. Compared to reputational uncertainty, a fast-changing social environment only causes a second-order qualitative increase in cooperation by making individuals more lenient in imposing a *network-punishment* (link removal). The interaction between reputational uncertainty and a fast-changing social environment induces more lenient strategies by reducing the frequency of *action-punishment* (retaliatory defection). Although neither of them affects the aggregate network metrics, their interaction decreases homophily among cooperators.

JEL Classification Numbers: C72, C73, C92, D81, D85

Keywords: Cooperation, experiments, prisoner's dilemma, uncertainty, repeated games, networks.

---

*We gratefully acknowledge the financial support from the Singapore Ministry of Education Academic Research Fund (AcRF) Tier 2 grant (MOE2014-T2-1-105). The paper benefited from discussions during presentations at the 2019 North American ESA Meeting in Los Angeles (USA) and the SEA 89th Annual Meetings in Fort Lauderdale (USA).

†Faculty of Economics, University of Cambridge, Sidgwick Avenue, Cambridge CB3 9DD, UK and Queens' College, CB3 9ET, UK.

‡Division of Economics, School of Social Sciences, Nanyang Technological University, 48 Nanyang Avenue, HSS 04-70, Singapore 639818.

§Department of Economics and Finance, College of Business, City University of Hong Kong, 83 Tat Chee Avenue, Kowloon Tong, Hong Kong.

¶Department of Economics, Faculty of Arts and Social Sciences, National University of Singapore, AS2 06-02 1 Arts Link, Singapore 117570.

1 Introduction

Understanding when and how repeated interactions lead to cooperation is fundamental to comprehend many economic phenomena. Seminal theoretical contributions have demonstrated that cooperative play is possible when individuals are sufficiently patient (e.g. Friedman, 1971; Fudenberg and Maskin, 1986), and several experimental papers have shown that people cooperate if the gains to cooperation are sufficiently large (e.g. Dal Bó, 2005; Dreber et al., 2008). Many of these studies, however, focus on a setup with two players with perfect monitoring of actions abstracting away from some of the features present in the real contexts where cooperation operates.

This paper examines experimentally the emergence and evolution of cooperation in groups of individuals who can decide with whom to interact and have access to information about their (potential) partners' reputation. A first dimension we investigate is the rate of change of the social-environment, defined as the frequency with which opportunities to change one's connections becomes available. For example, in an online world there are frequent opportunities to change connections, and this can be done in a snap of fingers, while offline these opportunities may be limited by spatial constraints or social norms. A second dimension we investigate is the presence of *uncertainty* on whether other people's observed action correspond to their intended action. In many contexts, it may be hard to verify someone's claim that they worked hard because performance also depends on luck, or, conversely, a negligent behavior may accidentally benefit someone. These two aspects have been investigated previously in isolation (e.g. Fudenberg et al., 2012; Rand et al., 2011), but our paper, to the best of our knowledge, is the first one to study them in a unified setting and explore their interaction.

In particular, in our experiment each subject is assigned to a group of 12 participants who play a repeated Prisoner's Dilemma (PD) game. The participants are connected by a network in which a link denotes a pair of subjects who play a PD game against each other in that round, and each subject is constrained to choose either to cooperate or defect with all their neighbors. At the beginning of a round, each subject has the opportunity to remove some of her existing links and (or) propose new links to others. The frequency with which opportunities to add (remove) links arises is one of our treatment variables and it captures the rate of change of the social environment. When making their decisions, subjects have access to information about the reputation of each of their connections in the form of their last five actions. A second treatment variable is *uncertainty* on the veracity of this reputational information.

Our first finding is that reputational uncertainty has a first order effect on cooperation and average welfare, while the rate of change of the social environment has, at best, a second order effect. Specifically, cooperation and average welfare significantly decrease with reputational uncertainty, while they only qualitatively increase in a fast-changing social environment.

Given that the action choices of players co-evolve with the interaction structure, there are two levers for individuals to punish defectors and enhance cooperation. *Network punishment* deters defectors with the threat of having their links with cooperators removed and (or) avoiding connections with defectors. The use of *network punishment* has been previously documented in studies involving endogenous updates to the interaction neighborhood without uncertainty (e.g. Rand et al., 2011). *Action punishment* is the standard threat of responding to a defector by engaging in retaliatory defection, which in this set-up is less targeted than the network punishment because not only the defector, but also the cooperative neighbors are affected by the retaliatory defection. By exploring how the use of network and action punishments vary with our treatment variables, we can shed light on the impact of reputational uncertainty on cooperative behavior.

Network punishment is less effective when there is reputational uncertainty because participants are more *lenient* and *forgiving* towards defectors, particularly in a fast-changing social environment. The concepts of leniency and forgiveness we employ in our setting are adaptations of the notions formulated in Fudenberg et al. (2012) for two-player games. In a fast-changing social environment, reputational uncertainty makes individuals employ lenient strategies – subjects are less likely to terminate an existing relationship with someone having a history of low cooperative play in an environment with uncertainty than without it. In the presence of reputational uncertainty, strategies are more forgiving as subjects are more likely to propose a link to someone with a history of low cooperative play independent of the rate of change of the social environment.

Our results point toward a behavioral asymmetry in the context of leniency and forgiveness with reputational uncertainty in a slow-changing social environment. Participants tend to react to a neighbor's defection by terminating the relationship, but they are more willing to give a second chance to potential neighbors with whom they are not connected yet.

Action punishment increases with reputational uncertainty in a fast-changing social environment. Specifically, given a proportion of defectors among their connections, a participant is more likely to switch from cooperation to defection in an uncertain and fast-changing social environment. A related finding is that reputational uncertainty promotes opportunistic play: participants have a higher tendency to defect after a history of highly cooperative play than in a setting without uncertainty because they can hide behind the uncertainty which makes it difficult for others to infer whether the defection was intentional or not.

Finally, we examine the impact of the rate of change of the social environment and reputational uncertainty on the emerging network structure. At aggregate level, there are no differences in network structure across treatments. Previous studies (e.g. Gallo and Yan, 2015; Gallo et al., 2019) have shown that reputational information and a fast-changing social environment increase cooperation by enabling the formation of clusters of cooperators who associate with each other. We find that reputational uncertainty undermines this process in a fast-changing social environment because defectors are more likely to remain embedded in cooperators' neighborhoods thereby creating an obstacle to the formation of the cooperative clusters that would have driven up the cooperation level and average welfare of the whole group.

This paper contributes to a growing literature that sits at the intersection of, first, the vast body of work investigating the emergence and sustenance of cooperation using variations of the PD and, second, research on how networks affect behavior and economic activity. These literature span different disciplines and it is not feasible to have an exhaustive review of all related contributions here. We aim, therefore, to position our paper in the context of the most relevant contributions with a bias toward the economic literature.

In experiments with an infinitely repeated two-player PD game under imperfect public monitoring, Aoyagi and Fréchette (2009) show that subjects cooperate in an environment with a continuous public signal, and that their payoff decreases with the level of noise in the public signal. Fudenberg et al. (2012) study a model of a repeated PD with reputational uncertainty in players' actions and observe that, compared to the setting with no uncertainty, individuals are slow to anger and fast to forgive. Specifically, with uncertainty subjects use strategies that, firstly, do not prescribe reverting to punishment after a single bad signal, and secondly, entail a return to cooperation after punishing the opponent. Aoyagi et al. (2019) find similar behavior with private monitoring.[1,2]

Beyond two-person games with uncertainty, Ambrus and Greiner (2012) study the effects of a costly punishment option on cooperation and social welfare in long, finitely repeated three-person public good contribution games. They show that the strong positive effects of increasing the severity of punishment as identified in Nikiforakis and Normann (2008) and Egas and Riedl (2008) do not hold when there is uncertainty. With five-person repeated public goods game with imperfect monitoring, Ambrus and Greiner (2017) further investigate the impact of democratic punishment, when members of a group decide by majority voting whether to inflict punishment on another member, relative to individual peer-to-peer punishment. Our contribution to this strand of work on cooperation in groups in an uncertain reputational environment is to introduce the possibility for individuals to choose whom to interact with given their (possibly imprecise) information on others' reputation.

When there is perfect monitoring of others' actions, recent research has demonstrated that the freedom to choose whom to interact with has a positive impact on efficiency in games of cooperation and coordination. Rand et al. (2011) and Wang et al. (2012) find that when subjects can update their links frequently, cooperation is maintained at a high level with participants preferentially removing links with defectors and forming new links with cooperators. Shirado et al. (2013) report that high cooperation is achieved at intermediate frequency of link updating.[3] Gallo and Yan (2015) further show that in this context knowledge about everybody's reputation is crucial to achieve a high level of cooperation.[4] Our study contributes to this growing literature

[1]Early experimental studies have explored the effect of manipulating players' information in repeated games in less standard ways. For example, Cason and Khan (1999) investigate a public goods game where information on others' contributions are presented with some delay. Bolton et al. (2005) study cooperation in an environment with limited information about reputation. Bereby-Meyer and Roth (2006) consider a repeated PD game with noisy payoffs, even though in their experiment players could monitor others' past actions perfectly.

[2]Related studies of repeated PD games under uncertainty include Rojas (2012), Rand et al. (2015), Embrey et al. (2013), and Arechar et al. (2017).

[3]When the frequency is too low, subjects choose to have many links, even if they attach to defectors. When it is too high, cooperators cannot detach from defectors as much as defectors re-attach and, hence, subjects resort to behavioral reciprocity and switch their behavior to defection.

[4]A substantial related literature studies how endogenous group formation affects play in PD, public

by introducing reputational uncertainty to investigate how it affects the effectiveness of the network formation mechanism in sustaining cooperation.

With regards to the role of punishment in inducing cooperation, starting with Fehr and Gächter (2000), a vast literature has demonstrated that the possibility of costly punishment facilitates increased cooperation in social dilemma situations. However, in the presence of noise, Ambrus and Greiner (2012) show that costly punishment is less effective in establishing cooperation. In our set-up, the only way to target punishment to a specific defector is by link removal, and therefore we investigate whether this type of punishment mechanism is less effective once we introduce reputational uncertainty.

An alternative form of punishment which has been shown to promote cooperation is ostracism (Cinyabuguma et al., 2005; Maier-Rigaud et al., 2010). Under ostracism, individuals can expel each other from the group using various voting schemes. The expelled players are then excluded from the benefits of any cooperation undertaken by those still remaining in the group. We can think of our set-up as allowing decentralized ostracism through the removal of links by single individuals in an environment with reputational uncertainty. Interestingly, we find that ostracism of defectors can occur even with a decentralized system, but only when there is no reputational uncertainty.

The remainder of the paper is organized as follows. Section 2 describes the design and procedures of the experiment. The research questions are listed in section 3. The results are presented in section 4, and the last section concludes.

2 Experimental Design and Procedures

In our experiment, each participant is assigned to a group of 12 who play a game that is repeated for 25 rounds with random termination thereafter. At the beginning of a round each participant has the chance to propose and/or remove links with some of the other group members. After the links have been updated, each participant plays a PD game with the other participants she is directly connected to. We now describe the protocol of the game in detail, which consists of three stages in each round.

Figure 1 illustrates the three stages that constitute a round. In the initial round all

good and weak link games. Contributions include Ehrhart et al. (1999), Hauk and Nagel (2001), Riedl and Ule (2002), Coricelli et al. (2004), Cinyabuguma et al. (2005), Page et al. (2005), Güth et al. (2007), Ahn et al. (2009), Brandts et al. (2009), Charness and Yang (2014), Riedl et al. (2016), Yang et al. (2017). See Choi et al. (2016) for a comprehensive review.

participants are always linked with each other, but in any round thereafter (e.g. round T in Figure 1) they may not. In stage 1 of round T, x random pairs of participants, i.e. those drawn with dashed circles, may receive opportunities to propose and/or remove links with each other. Suppose i and j is a random pair picked at time T, each of them faces one of these choices:

- Remove the link if i and j were connected at round $T-1$. Deletion of a link is unilateral and does not require mutual consent.[5]

- Propose a link if i and j were not connected at round $T-1$. Formation of a link is bilateral and requires consent of the recipient of the proposal in Stage 2.

Stage 2 occurs if one and only one between i and j has proposed a link. Given that link formation requires mutual consent, in stage 2 the recipient of a link proposal has to decide whether to accept or reject it, as shown in the third row of Figure 1. Notice that if two participants send a link proposal to one another in Stage 1, then the link is automatically formed and there is no prompt.

In stage 3, participants play the PD game shown at the bottom right of Figure 1 with each one of their neighbors, subject to the restriction that they have to choose either action A or action B for all neighbors.[6] Choosing a single action simultaneously toward all neighbors is a standard restriction employed in the literature to allow for network effects, otherwise the network game becomes a collection of independent PD games between pairs of players in the group.[7] If a participant has no neighbors, then she has no action to choose.

Once all individuals make their choices, they see the points they receive from the interaction with each of the other participants. They receive zero points in this stage from participants with whom they are not connected to. Subjects are also reminded of their chosen action and the action chosen by their neighbors.

Appendix B shows screenshots for each of the three stages. For each participant, the left area of the interface displays a visualization of their links to their neighbors in

[5]This is a natural assumption in the context of a relationship that is not legally binding and is common in the economics literature (e.g. Jackson and Wolinsky, 1996).

[6]Action A is color-coded in green and action B is color-coded in blue. The colors are used as a visual aid to distinguish the actions and have no meaning. Action A and B corresponds to, respectively, cooperation (C) and defection (D). We used neutral labels A and B in our experiment to avoid the experimenter-demand effect. In Figure 1 we denote A as C and B as D.

[7]See for example Morris (2000), Charness et al. (2014), Riedl et al. (2016), and Yang et al. (2017).

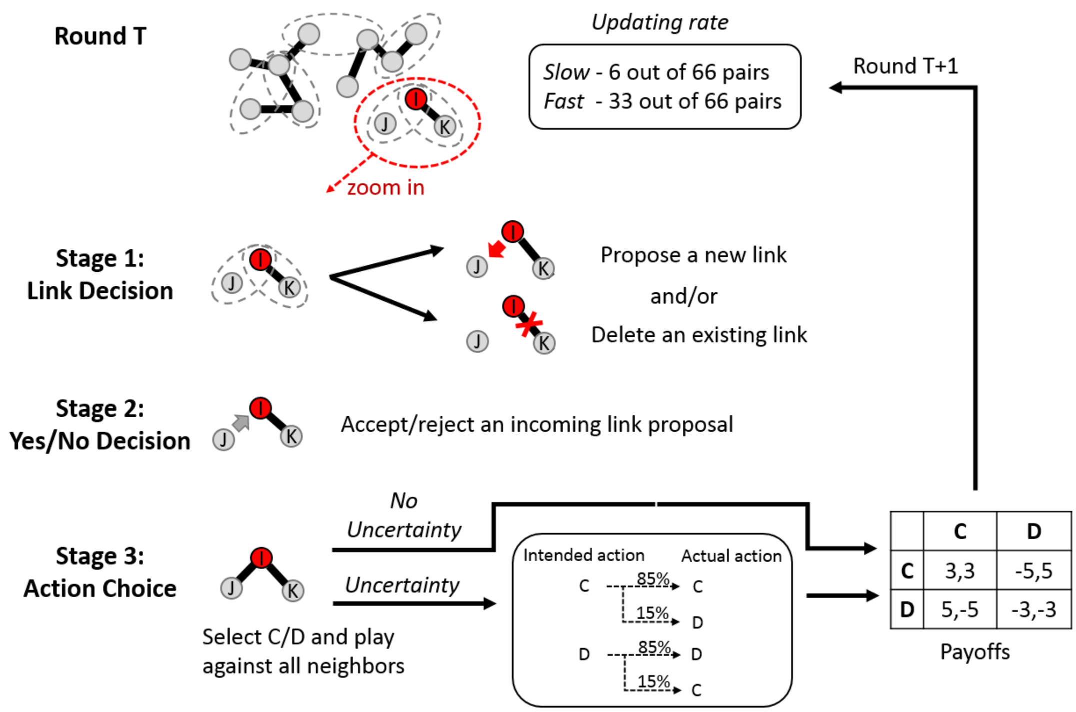

Figure 1: Schema of the experimental design.

which the nodes are colored in either green or blue depending on whether their action in the last round was to cooperate or defect respectively. The central part of the interface shows the list of current neighbors with their past 5 actions, and, at the end of stage 3, the points obtained in the PD game with each of the neighbors. The right area of the interface has the list of other participants who are not currently neighbors. The bottom part prompts the participant to make the relevant decision in each stage, e.g. the participants with whom she can propose/remove a link in stage 1 and the choice of action in stage 3.

Each group of 12 participants plays at least 25 identical rounds of this game. After round 25 there is a 50% chance of termination to reduce any endgame effects. Participants' earnings are based on 6 randomly selected rounds. For each participant, two other participants are randomly picked for each of the selected rounds and the sum of the points obtained in the interaction with these 12 randomly picked participants constitutes the total sum of points earned by the participant in this part. Notice that it is possible to pick an unconnected pair of participants and, in that case, the earnings from that interaction are equal to zero.[8]

Following the completion of the first part of the experiment, participants take a Holt and Laury (2002) multiple price list risk-elicitation test, and fill in a socio-demographic questionnaire, which includes the standard interpersonal trust question taken from the World Values Survey (WVS).[9]

The boxes with round edges in Figure 1 illustrate the two treatment dimensions in the experiment, namely the *updating rate* (*slow* and *fast*) and *whether uncertainty is present*. We conduct 4 treatments and use a 2×2 between-subject design to test the impact of uncertainty under two different rates of link formation/removal. Under *no reputational uncertainty* treatments, there is no uncertainty so participants receive payoffs according to their own action and their neighbors' actions. Under *reputational uncertainty* treatments, there is a commonly known 15% probability that an intended action is changed to the opposite action. Subjects are informed of their intended action as well as their implemented one. However, they are only notified of the neighbors' actual actions, and not what their intended one was. In terms of the updating rate dimension, we consider a *slow* $x = 6$ condition in which only 9% of pairings potentially change in each round, and a *fast* $x = 33$ condition in which 50% of pairings potentially change in each round.

The data were gathered in 16 experimental sessions conducted at the Nanyang Technological University (NTU) in Singapore, and the experiment involved 384 undergraduate students from various majors. No subject participated in more than one session and we used a between-subject design so that each subject participated in one and only one treatment. There are 4 sessions for each of the 4 treatments. Each session consists of 2 groups (networks) of size 12 for a total of 8 independent observations at the network level per treatment. The choice of a group size of 12 ensures that the network is large enough to allow the emergence of interesting structural features while at the same time it is small enough to collect a sufficient number of independent observations per

[8]This random selection of pairs for payment, independently of whether a connection exists or not, ensures that there are uniform incentives throughout the experiment in forming connections so the payment system does not introduce biases in the emerging network structure. For instance, if we had excluded unconnected pairs from the random selection for payment, then participants would have incentives to form just one link with a cooperator to ensure that a specific pairing is picked. This choice is counter to studies that pay the cumulative number of points participants have earned, which may lead to satisficing in the later rounds and therefore lower incentives to change the network.

[9]The interpersonal trust question is: *"Generally speaking, would you say that most people can be trusted or that you need to be very careful in dealing with people?"*.

treatment. The experiment was programmed with the software z-Tree (Fischbacher, 2007). Upon arrival, subjects were seated at visually isolated computer workstations. Instructions were read aloud and subjects also received a copy of the instructions.[10] Participants remained completely anonymous throughout the experiment.

The sessions lasted approximately two hours and participants earned on average S$13.73 in addition to the S$2 show-up fee. Out of the 384 participants, 53% are female, 62.4% are Singaporeans, and the average age is 21 years. About 23% of the participants have had some prior exposure to Game Theory, exactly half of the participants had participated in at least one laboratory experiment previously. From the WVS interpersonal trust question, 22% believe that others can be trusted. In the post-experiment questionnaire almost all participants state that they had no problem in understanding the experiment.[11]

3 Research Questions

This section lays out several research questions that motivate our study. Given the complexity of this repeated game on an endogenously evolving network, there is a multiplicity of equilibria rendering theory alone insufficient to formulate precise predictions. However, for each question we provide a discussion of what we should expect to find based on what we know from simpler theoretical set-ups and prior experimental work.

As our experimental set-up involves repeated interactions without a known end-point, insights from experimental repeated games suggest that there will be some degree of cooperation (Dal Bó and Fréchette, 2018). Starting from two-person cooperation games and public good games without any network structure, the experimental evidence is inconclusive for the aggregate level impact of uncertainty on overall cooperation, although the majority of the studies find a negative effect. Aoyagi and Fréchette (2009) and Aoyagi et al. (2019) find that the presence of uncertainty results in strictly lower cooperation. In contrast, Rojas (2012) reports that cooperation increases or decreases when uncertainty is introduced depending on the treatment. In finitely re-

[10] A sample copy of the instructions for one of the treatments is provided in Appendix B.

[11] The descriptive statistics of the participants in each of the four treatments are provided in Appendix A.

peated public good games, Ambrus and Greiner (2012) and Ambrus and Greiner (2017) report that although uncertainty (in the observation of other group members' contribution behavior) leads to lower contributions, it is statistically significant for only a subset of their treatments. Different from their setup, in our framework we have reputational uncertainty caused by the presence of uncertainty on whether other people's observed action correspond to their intended action.

Previous experimental studies investigating a PD game in a setting with more than two players and with an endogenous interaction structure have found that increasing the proportion of links a participant can update in each round increases cooperation. Rand et al. (2011) find that a fast-changing social environment encourages subjects to break links with defectors and form new links with cooperators. Wang et al. (2012) confirm this result by showing that there is a monotonic relationship between the level of cooperation and how fast the interaction neighborhood can be updated by investigating several values of the proportion of links that a participant can change in each round.

We thus are interested in investigating how the two variables, that is, the presence (or the absence) of reputational uncertainty, and the speed with which the interaction neighborhood can be updated influence the level of cooperation and aggregate welfare.

Question 1 (**Cooperation and Welfare**): *Does the level of cooperation and average welfare decrease with reputational uncertainty and increase in a fast-changing social environment?*

Given that in our set-up mutual cooperation between two individuals is the only way to produce social surplus, we expect that the presence of reputational uncertainty will have a negative effect on welfare, and, in agreement with the studies mentioned above, the level of cooperation and average welfare will increase in the *fast* treatment compared to the *slow* one.

Next, *network punishment* is a primary tool to maintain cooperation when individuals can choose their interaction neighborhood (e.g. Rand et al., 2011). It entails removing links with defectors and/or refusing to accept a proposal to interact from someone with a history of low cooperation. Given the importance of network punishment, it is useful to define the following two notions to analyze how it varies depending on our treatment variables.

Definition 1: Leniency. *Strategies are lenient if, conditional on receiving an opportunity to terminate a link, the deciding player does not terminate the existing relationship with a neighbor whose actual action was C at most once in the last five rounds (i.e, with a history of low cooperative play).*

Definition 2: Forgiveness. *Strategies are forgiving if, conditional on receiving an opportunity to propose a link, the deciding player proposes a link to someone whose actual action was C at most once in the last five rounds (i.e, with a history of low cooperative play).*

The notions of leniency and forgiveness in our setup are novel because they are defined to capture variations in link formation/removal rather than action punishment which is the focus in two-person PD games (Fudenberg et al., 2012; Embrey et al., 2013; Aoyagi et al., 2019). In these studies focusing on two-person PD games, strategies are defined to be lenient if they do not prescribe sure defection following a single defection by the opponent, and forgiving if they return to cooperation after having defected. These studies report that in the presence of reputational uncertainty subjects use strategies that are both lenient and forgiving. This is in contrast to what Dal Bó and Fréchette (2019) find in the perfect monitoring environment where grim-trigger is one of the most frequently observed strategies. Given the experimental results in two-player games, we expect the use of network punishment to decrease with reputational uncertainty and therefore strategies to be more lenient and forgiving in the reputational uncertainty treatments.

To the best of our knowledge, existing studies with endogenous interaction structure do not formally analyze how the leniency and forgiveness in players' strategies vary with the rate of change of the social environment. Intuitively, when updating is slow, there is a high potential cost of establishing or maintaining a link with a defector because the opportunities to break the link are harder to come by. When updating is fast, players can afford a more lenient and forgiving strategy, knowing that they do not have to wait long to have another opportunity to remove the links should they remain uncooperative. We thus expect that leniency and forgiveness are more prevalent in the fast-changing social environment condition.

Question 2 (**Leniency and Forgiveness**): *Are strategies more lenient and/or forgiving with reputational uncertainty and in a fast-changing social environment?*

Action punishment is the primary channel to ensure cooperative behavior in two-player games. It entails defecting as a response to a defection by the other player. The following definition extends this notion to our set-up with multiple players in a linked network.

Definition 3: Action Punishment. *The likelihood of choosing D as the intended action as a function of the proportion of neighbors who choose D.*

Because the same action is implemented against all linked neighbors, action punishment is less targeted than network punishment and therefore we would expect participants to opt for the latter to punish a defector. This is because choosing action punishment would imply that the player chooses D against all her linked neighbors even if some of these neighbors have been cooperative in the previous round. However, the presence of reputational uncertainty potentially enables an individual to punish a defector by defecting without necessarily losing her own reputation because others may infer that the defection was not intentional.

Question 3 (**Action Punishment**): *Is action punishment decreasing or increasing with reputational uncertainty?*

Another element that may undermine overall cooperation in the presence of reputational uncertainty is for cooperators to occasionally defect knowing that others would be unable to ascertain whether or not the defection is intentional. We define this behavior as follows.

Definition 4: Opportunistic Play. *The likelihood of choosing D as the intended action after a history of highly cooperative implemented actions.*

Previous experimental repeated games have found that the presence of reputational uncertainty can result in a higher share of players adopting non-cooperative strategies like "always defecting" or a "defection-inclined version of tit-for-tat" ((Fudenberg et al., 2012; Aoyagi et al., 2019)). While our set-up differs from these studies, it is reasonable to expect that opportunistic play, which can be interpreted as a type of non-cooperative strategy, should be more prevalent when there is reputational uncertainty.

Question 4 (**Opportunistic Play**): *Is the level of opportunistic play higher in the presence of reputational uncertainty?*

A natural question in our set-up is how the presence of reputational uncertainty affects the evolution of the overall network structure as captured by standard metrics such as average degree and clustering coefficient.[12]

Given the lack of prior studies that investigate reputational uncertainty in a network setting, it is difficult to formulate precise predictions on its impact on the overall structure of the network as there are opposite forces at play. If, for instance, we focus on average degree then reputational uncertainty may increase it because some links with defectors are not removed due to the uncertainty on whether a given defection was intentional. Average degree could, however, also decrease if participants become more wary in forming a link with someone who just defected, although the defection may have been unintentional. We therefore are interested in answering the following question.

Question 5 (**Aggregate Network Structure**): *Do aggregate network characteristics change with reputational uncertainty?.*

Aside from its potential impact on the overall network structure, reputational uncertainty may also cause a variation in the typical composition of the neighborhood of a participant. Previous experimental studies have shown homophily based on cooperativeness, i.e. cooperators tending to have more connections with other cooperators, increases when there is endogenous network formation (Wang et al., 2012), participants have complete information about the network structure (Gallo and Yan, 2015), and there is an opportunity to strengthen a connection (Gallo et al., 2019).

Whether the presence of faster interaction updating alone increases cooperativeness-based homophily may depend on the specific payoffs of the PD game. Wang et al. (2012) find that a faster updating increases homophily, but this result is possibly driven by incentives to over-connect because two individuals defecting on each other still receive positive payoffs in their set-up. Using the same payoffs as in our set-up, Gallo and Yan (2015) find little evidence of homophily based on cooperativeness in an endogenous network when participants do not have full information about the network. Accordingly, we expect a similar level of homophily in the slow- and fast-updating treatments. In the presence of reputational uncertainty, instead, cooperators may find it more difficult to recognize each other and connect so we expect that reputational

[12] Degree is the number of connections of an individual, while the (local) clustering coefficient captures the proportion of an individual's neighbor who are connected to each other. See section 4.2 for the mathematical definitions.

uncertainty will decrease cooperativeness-based homophily.

Question 6 (**Homophily**): *Does homophily among cooperators increase or decrease in the presence of reputational uncertainty and in a fast-changing social environment?*

4 Results

This part of the paper reports the results of our experiment. Section 4.1 examines overall cooperation and welfare. Section 4.2 investigates leniency and forgiveness in network punishment, and section 4.3 looks at action punishment. Finally, section 4.4 reports how our treatment variables affect the emerging network structure.

4.1 Overall Cooperation and Welfare

We start by analyzing how cooperation and welfare vary at the aggregate level across treatments. In the statistical analysis, we aggregate the data at the network level so that there are 8 independent observations for each treatment. Then, we apply the non-parametric Kruskal-Wallis test to compare across multiple groups to detect treatment effects, followed by the Dunn's (DT) test to explore differences between any two treatments.[13] In this section, we report p-values for the Dunn's test. The aggregate analysis in this section focuses on rounds 2-21 and we exclude the last four rounds to avoid any end-game effects.[14]

Figure 2 displays the evolution of the cooperation rate and average welfare over time. The cooperation rate is defined as the number of pairs with mutual (intended) cooperation at the end of a round, divided by the 66 possible pairs in a complete network with 12 nodes.[15] The average welfare is defined as the number of points obtained by participants from all interactions at the end of the round. After some

[13]The choice of a non-parametric test with a small sample per treatment after aggregation is rather conservative, and therefore any statistically significant finding denotes a sizable treatment effect.

[14]We focus on this time frame throughout, except in sections 4.2 and 4.3, where the analysis uses rounds 6-21 when we condition the analysis on the last five actions played by subjects.

[15]Our results remain the same if we use the fraction of intended cooperators as the measure of cooperation instead of cooperation rate. This alternate measure is defined as the ratio of participants in a group who choose C (as the intended action).

initial volatility, cooperation and average welfare in the different treatments tend to stabilize over time, aside from end-game effects in the very last rounds.[16]

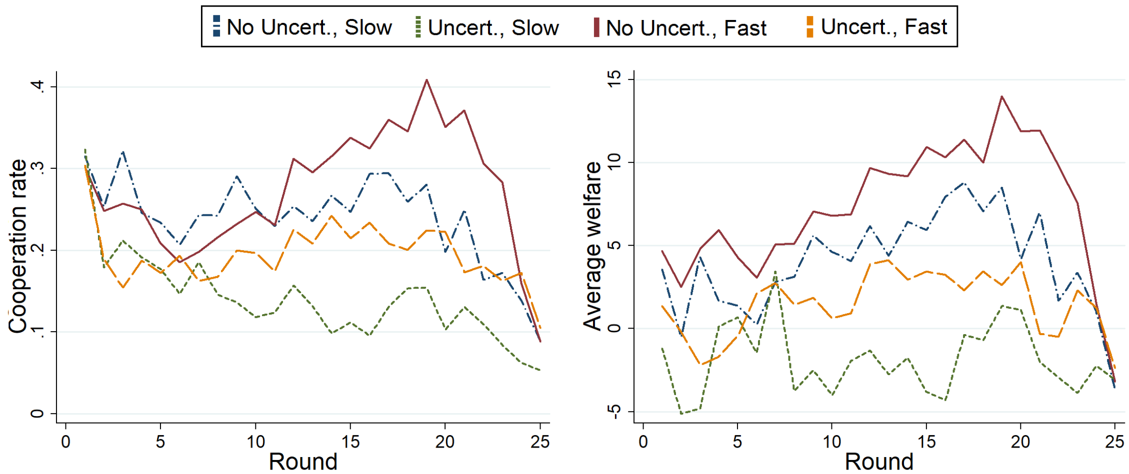

Figure 2: Cooperation rate and average welfare over time.

Reputational uncertainty leads to significantly lower cooperation and average welfare. In the *slow* condition, reputational uncertainty reduces cooperation rate from 0.24 to 0.14 ($p = 0.003$, DT) and average welfare from 3.85 to -1.91 ($p = 0.017$, DT). In the *fast* condition, reputational uncertainty reduces cooperation rate from 0.27 to 0.19 ($p = 0.024$, DT) and average welfare from 6.81 to 1.53 ($p = 0.028$, DT). The impact of reputational uncertainty is therefore large and significant, and also consistent with the majority of prior studies on PD games with two players.

The introduction of a fast-changing social environment increases cooperation and average welfare holding reputational uncertainty constant, but the effects are at best marginally significant. The p-values (DT) for differences in cooperation rate and average welfare between the *slow* and *fast* conditions are 0.345 and 0.182, respectively, without reputational uncertainty, and 0.071 and 0.132 with reputational uncertainty. The findings from previous studies on cooperation in a dynamic network showing that a fast-changing social environment increases cooperation (e.g. Rand et al., 2011; Wang et al., 2012) are therefore not robust to the introduction of reputational uncertainty. This is consistent with the recent evidence showing that the role of network formation

[16]End-game effects are common in games of cooperation (Gächter et al., 2008; Ambrus and Greiner, 2012).

is not robust to changes in the experimental environment (Rand et al., 2014) suggesting that its role is second order with respect to reputation. Our first result is therefore the following.

Result 1: *Reputational uncertainty significantly reduces cooperation and average welfare. A fast-changing social environment increases cooperation and welfare, but only qualitatively.*

Note that the effects we find are not driven by initial conditions. There is no significant difference for the cooperation rate between any two treatments in round 1 ($p > 0.1$ for all comparisons, DT). Because in all treatments the groups start, on average, from the same level of cooperation rate, it must be that reputational uncertainty and the dynamics of network evolution are the factors that yield the differences in the cooperation rate that emerge over time.

4.2 Leniency and Forgiveness in Network Punishment

We begin our investigation of leniency in network punishment by examining some aggregate statistics. Using observations from rounds 6-21, Figure 3 displays the average likelihood of removing an existing link, conditional on receiving an opportunity to do so. We consider three different histories of a neighbor's cooperative play: 0 or 1 C (*low*), 2 or 3 C (*medium*) and 4 or 5 C (*high*). Unsurprisingly, when the existing link is highly cooperative, there are very few occurrences of participants removing links in all treatments (around 2% of the total observations in each treatment). However, in instances of low or medium cooperative play of the existing link, a subject is less likely to remove a link when there is reputational uncertainty or in a fast-changing social environment.

We first test the effect of reputational uncertainty on leniency. To do so, we run a random-effects logit regression with the dependent variable being the link removal decision (= 1 if the link is unilaterally removed). We conduct the regressions separately for the *slow* and *fast* conditions, and use the same three histories of a neighbor's cooperative play defined earlier. In all of these regressions we use observations from rounds 6-21 and cluster the standard errors at the individual level.

The independent variables include the dummy variable for *Reputational uncertainty* (= 1 if there is uncertainty), *Round*, and network metrics (*Degree* and *Local Clustering*).

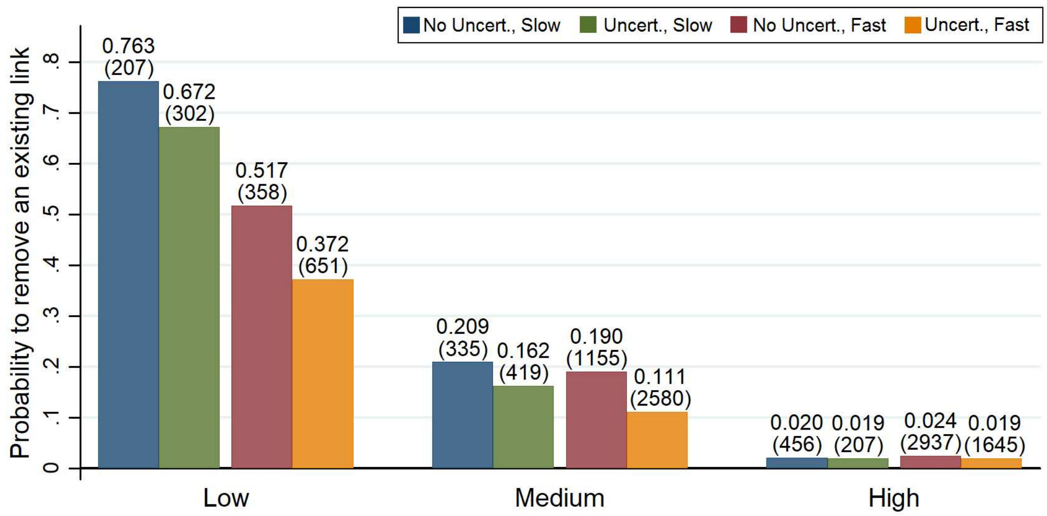

Figure 3: Average probability to remove an existing link, conditional on receiving an opportunity, divided by the last five rounds' history. *Low*: the actual action was C at most once in the last five rounds. *Medium*: C was the actual action two or three times. *High*: C was the actual action more than three times. The number of observations are in parentheses.

The degree d_i of a subject i is defined as the total number of links that i has at the end of the round. For each subject i, the local clustering coefficient measures how connected i's neighbors are, and it is formally defined as :

$$C_i = \frac{2\sum_{j,k\in N_i} g_{ij}g_{ik}g_{jk}}{d_i(d_i-1)}$$

where $g_{ij} = 1$ if a link exist between i and j and zero otherwise, $N_i = \{j|g_{ij} = 1\}$, and $d_i = |N_i|$.

We also include *Other Control Variables* consisting of *Age*, *Gender*, the measure of *Trust* obtained by asking the question taken from the World Values Survey (see section 2) to participants, the switching point in the Holt-Laury risk preference elicitation (*Risk Aversion*), an indicator variable capturing whether participants have taken a game theory course before (*Game Theory*) and an indicator variable representing whether subjects have participated in experiments before (*Prior Participation*).

Table 1 displays the regression results. In the *slow* condition, for any history, reputational uncertainty does not significantly affect the probability of removing an

Table 1: Random-effects logit regressions of the likelihood of participants removing existing links, conditional on receiving an opportunity and using last five rounds' history. Standard errors (clustered at the level of individual subject) are in parentheses.

	Slow			*Fast*		
	Low	*Medium*	*High*	*Low*	*Medium*	*High*
Reputational Uncertainty	-0.509	-0.414	-0.049	-0.920***	-0.857***	-0.124
	(0.414)	(0.267)	(0.663)	(0.344)	(0.203)	(0.329)
Round	0.041	-0.028	-0.018	-0.045**	-0.006	0.008
	(0.051)	(0.026)	(0.067)	(0.020)	(0.013)	(0.026)
Degree	0.16	-0.131*	-0.116	-0.061	-0.074**	-0.141*
	(0.120)	(0.067)	(0.216)	(0.048)	(0.036)	(0.076)
Local clustering	2.637*	1.185	-0.285	0.319	0.825**	0.808
	(1.467)	(0.797)	(1.905)	(0.626)	(0.353)	(0.685)
Constant	-1.563	-3.881**	-2.454	2.987	0.129	-2.757
	(2.852)	(1.737)	(4.092)	(2.180)	(1.065)	(2.073)
Other control variables	Yes	Yes	Yes	Yes	Yes	Yes
Observations	508	754	661	1009	3735	4582

Notes: Low- the actual action being C at most one time in last five rounds; *Medium*- C was the actual action two or three times; *High*- C was the actual action more than three times; *significant at 10%; **significant at 5% level; ***significant at 1% level; Other Control Variables include *Age*, *Gender*, *Trust*, *Risk Aversion*, *Game Theory*, *Prior Participation*.

existing link ($p > 0.1$). However, this probability is significantly lower with reputational uncertainty in the *fast* condition for low and medium cooperative play ($p = 0.007$ and $p < 0.001$). Thus, strategies employed by participants are lenient only when there are abundant opportunities to revise links, which happens when we have the *fast* condition. On the contrary, in the *slow* condition, instead, there is a high burden of maintaining a relationship with someone having a history of low cooperative play because the next opportunity to break the link might only arise in the distant future.

To formally show the point above which claims that participants are more lenient when updating is *fast*, we perform another set of random-effects logit regressions on link removal decisions, by including dummy variable for *Fast updating* (= 1 in the fast

Table 2: Random-effects logit regressions of the likelihood of participants removing existing links, conditional on receiving an opportunity and using last five rounds' history. Standard errors (clustered at the level of individual subject) are in parentheses.

	No Uncert.			*Uncertainty*		
	Low	*Medium*	*High*	*Low*	*Medium*	*High*
Fast	-1.002***	0.083	0.088	-1.551***	-0.590**	0.092
	(0.371)	(0.230)	(0.640)	(0.358)	(0.268)	(0.569)
Round	-0.043	-0.007	0.060**	-0.048**	-0.01	-0.087
	(0.027)	(0.015)	(0.026)	(0.022)	(0.016)	(0.053)
Degree	-0.018	0.034	-0.220***	-0.045	-0.169***	-0.133
	(0.071)	(0.044)	(0.082)	(0.053)	(0.046)	(0.121)
Local clustering	2.067***	0.513	1.734**	-0.146	1.175***	-0.479
	(0.802)	(0.432)	(0.784)	(0.744)	(0.446)	(1.720)
Constant	2.598	-2.401*	-4.551	1.919	-0.386	-1.296
	(1.994)	(1.238)	(3.897)	(2.259)	(1.508)	(2.692)
Other control variables	Yes	Yes	Yes	Yes	Yes	Yes
Observations	565	1490	3392	952	2999	1851

Notes: Low- the actual action being C at most one time in last five rounds; *Medium*- C was the actual action two or three times; *High*- C was the actual action more than three times; *significant at 10%; **significant at 5% level; ***significant at 1% level; Other Control Variables include *Age*, *Gender*, *Trust*, *Risk Aversion*, *Game Theory*, *Prior Participation*.

updating condition) while conducting the regression separately for the *no uncertainty* and *uncertainty* conditions. Table 2 displays the regression result, showing that faster updating makes individuals more lenient in maintaining a relationship. In the condition with no reputational uncertainty, faster updating significantly decreases the probability of removing an existing link for low cooperative play ($p = 0.007$). This probability is also lower with faster updating in the condition with reputational uncertainty, for low and medium cooperative play ($p < 0.001$ and $p = 0.028$).

Result 2: *(i) Leniency increases with reputational uncertainty only in a fast-changing social environment. (ii) Leniency is more prevalent in a fast-changing social environment, regardless of the presence of reputational uncertainty.*

Next, we focus on forgiveness in link formation decisions. Using observations from rounds 6-21, Figure 4 displays the average probability of sending a proposal (conditional on receiving an opportunity to do so) for the three different histories of the potential neighbor's cooperative play defined earlier, i.e. 0 or 1 C (*low*), 2 or 3 C (*medium*) and 4 or 5 C (*high*). As expected, the likelihood of sending a proposal is consistently high, close to 0.90, for recipients with a history of high cooperative play. In contrast, this likelihood is lower when a potential neighbor has a history of low or medium cooperative play. However, when reputational uncertainty is present, a subject becomes more willing to send a proposal to a potential neighbor who has a history of low and medium cooperative play, suggesting that the subject becomes more forgiving.

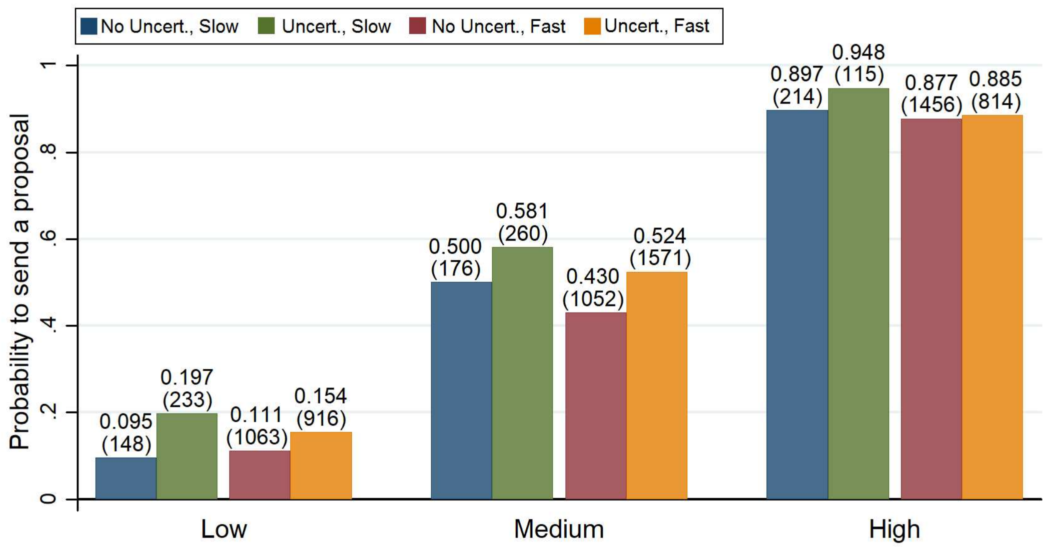

Figure 4: Average probability to send a proposal, conditional on receiving an opportunity, divided by the last five rounds' history. *Low*: the actual action was C at most once in the last five rounds. *Medium*: C was the actual action two or three times. *High*: C was the actual action more than three times. The number of observations are in parentheses.

Table 3 shows the results of a random-effects logit regression with the link proposal decision as the dependent variable (= 1 if the participant sends a proposal).[17] The regression results show that the presence of reputational uncertainty makes individuals more forgiving when starting a new relationship. In the *slow* condition, reputational

[17]Specifications and controls used for Tables 3 and 4 are the same as Table 1 and 2 respectively.

uncertainty increases the probability of sending a proposal to a potential neighbor who had a history of low cooperative play in the last five rounds ($p = 0.030$). This probability is also higher with reputational uncertainty in the *fast* condition for low and medium cooperative play of the potential neighbor ($p = 0.032$ and $p = 0.005$).

Table 3: Random-effects logit regressions of the likelihood of participants sending proposal to a potential neighbor, conditional on receiving an opportunity and using last five rounds' history. Standard errors (clustered at the level of individual subject) are in parentheses.

	Slow			*Fast*		
	Low	*Medium*	*High*	*Low*	*Medium*	*High*
Reputational Uncertainty	1.929**	0.288	-0.088	0.794**	0.767***	0.565
	(0.891)	(0.360)	(1.819)	(0.371)	(0.272)	(0.413)
Round	-0.064	0.043	-0.124	-0.059*	-0.043***	-0.042*
	(0.101)	(0.038)	(0.103)	(0.031)	(0.016)	(0.022)
Degree	-0.135	-0.188*	-0.456	-0.152**	-0.088**	-0.175***
	(0.254)	(0.101)	(0.403)	(0.059)	(0.037)	(0.055)
Local Clustering	-5.952*	-1.066	-5.887	-1.315***	-0.566**	-0.892**
	(3.461)	(0.840)	(3.635)	(0.510)	(0.270)	(0.386)
Constant	-4.273	3.479	28.082**	-3.181	-1.856	3.342
	(9.080)	(2.753)	(11.647)	(2.226)	(1.607)	(2.559)
Other control variables	Yes	Yes	Yes	Yes	Yes	Yes
Observations	381	436	329	1979	2623	2270

Notes: Low- the actual action being C at most one time in last five rounds; *Medium*- C was the actual action two or three times; *High*- C was the actual action more than three times; *significant at 10% level, **significant at 5% level; ***significant at 1% level; Other Control Variables include *Age*, *Gender*, *Trust*, *Risk Aversion*, *Game Theory*, *Prior Participation*.

In Table 4, we test the effect of a fast-changing social environment on forgiveness by performing a similar random-effects logit regression on link proposal decisions (by replacing the treatment variable with *Fast updating*). For any history and both with and without reputational uncertainty, faster updating does not significantly affect the probability of sending link proposals ($p > 0.1$ for all comparisons).

Result 3: *(i) Forgiveness increases with reputational uncertainty, regardless of the*

Table 4: Random-effects logit regressions of the likelihood of participants sending proposal to a potential neighbor, conditional on receiving an opportunity and using last five rounds' history. Standard errors (clustered at the level of individual subject) are in parentheses.

	No Uncert.			*Uncertainty*		
	Low	*Medium*	*High*	*Low*	*Medium*	*High*
Fast	0.119	-0.531	-0.917	-0.265	-0.081	-1.131
	(0.553)	(0.356)	(0.540)	(0.528)	(0.296)	(0.723)
Round	-0.036	-0.015	-0.057**	-0.061	-0.035**	-0.018
	(0.040)	(0.023)	(0.024)	(0.039)	(0.018)	(0.038)
Degree	-0.182**	-0.169***	-0.167***	-0.111	-0.070	-0.090
	(0.082)	(0.050)	(0.061)	(0.073)	(0.045)	(0.089)
Local Clustering	-1.989***	-1.014***	-1.453***	-0.929	-0.508	-0.372
	(0.682)	(0.310)	(0.529)	(0.700)	(0.366)	(0.632)
Constant	-2.311	1.791	5.562**	-1.750	-1.512	6.842**
	(2.502)	(2.026)	(2.682)	(3.345)	(1.770)	(3.445)
Other control variables	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1211	1228	1670	1149	1831	929

Notes: Low- the actual action being C at most one time in last five rounds; *Medium*- C was the actual action two or three times; *High*- C was the actual action more than three times; *significant at 10% level, **significant at 5% level; ***significant at 1% level; Other Control Variables include *Age*, *Gender*, *Trust*, *Risk Aversion*, *Game Theory*, *Prior Participation*.

rate of change of the social environment. (ii) Forgiveness is unchanged if we compare fast and slow-changing social environments, regardless of the presence of reputational uncertainty.

4.3 Action Punishment and Opportunistic Play

Table 5 presents the results of random-effects logit regressions of the likelihood of participants choosing C as their intended action. We limit the analysis to the subsample of instances when a participant has selected C as the intended action in the previous round, and focus on rounds $2-21$.[18] The dependent variable is the intended action decision this round, which is equal to 1 if the intended action is C. It can be interpreted as the likelihood of continuing to choose C as the intended action.

Specifications (1)-(3) focus on the *slow* conditions, while (4)-(6) investigate the *fast* conditions. In the baselines (1) and (4), we only include the *Round* number, an *Reputational Uncertainty* dummy for the uncertainty condition and *Other Control Variables*. Specifications (2) and (5) also include the fraction of a participant's neighbors whose actual action was C in the last round (*Fraction of Cooperative Neighbors*) as well as the *Degree* and *Local Clustering* coefficient network metrics. Finally, in specifications (3) and (6) we include the interaction variable between *Reputational Uncertainty* and *Fraction of Cooperative Neighbors*.

The regression results show that the probability of continuing to choose C as the intended action in a round is significantly increasing ($p < 0.01$ for all specifications) in the fraction of cooperative neighbors and significantly decreasing in the number of connections ($p < 0.01$ for all specifications). This holds for all specifications in both the *slow* and *fast* conditions. Moreover, the interaction term is significant only in the *fast* condition, and therefore action punishment is decreasing with uncertainty in a fast-changing social environment, but it is unaffected by the uncertainty in a slow-changing one.

These results indicate that action and network punishments are (partial) substitutes, and can crowd each other out under certain circumstances. In the *slow* condition, network punishment is not very effective given that opportunities to update

[18] For this analysis, we focus on rounds 2-21 as we do not need to condition the analysis on the last five actions played by potential partners unlike in the previous analysis.

Table 5: Random-effects logit regressions of the likelihood of participants choosing C as their intended action this round, given that they chose C as their intended action last round. Standard errors (clustered at the level of individual subject) are in parentheses.

	Slow			*Fast*		
	(1)	(2)	(3)	(4)	(5)	(6)
Reputational Uncertainty	-0.763**	-0.231	0.526	-0.667***	-0.200	0.902**
	(0.304)	(0.319)	(0.515)	(0.238)	(0.241)	(0.441)
Round	0.023*	-0.027	-0.030	0.035***	0.011	0.006
	(0.013)	(0.018)	(0.018)	(0.012)	(0.012)	(0.012)
Frac. of Coop. Neighbors		3.921***	4.679***		2.754***	3.669***
		(0.396)	(0.560)		(0.326)	(0.474)
Frac. of Coop. Neighbors			-1.46			-1.708***
× Reputational Uncertainty			(0.776)			(0.608)
Degree		-0.168***	-0.167***		-0.137***	-0.143***
		(0.060)	(0.061)		(0.039)	(0.039)
Local Clustering		-0.036	-0.060		-0.026	-0.156
		(0.560)	(0.555)		(0.355)	(0.365)
Constant	0.172	0.359	0.157	2.453	1.590	1.017
	(1.933)	(2.124)	(2.125)	(1.517)	(1.522)	(1.540)
Other Control Variables	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1985	1985	1985	2335	2335	2335

Notes: *significant at 10% level, **significant at 5% level; ***significant at 1% level; Other Control Variables include *Age*, *Gender*, *Trust*, *Risk Aversion*, *Game Theory*, *Prior Participation.*

connections are infrequent so the usage of action punishment is unaffected by uncertainty because subjects have to resort to it as their main avenue to punish defectors even when past reputation is imperfect. In contrast, in the *fast* condition, network punishment is effective, and therefore the introduction of uncertainty decreases action punishment, which is less targeted and equally affected by the presence of imperfect reputational information.

Result 4: *Reputational uncertainty reduces the usage of action punishment in a fast-changing social environment. Action punishment is a substitute for network punishment when the latter is less effective in a slow-changing social environment.*

To examine the prevalence of opportunistic play, Figure 5 shows the likelihood that a subject chooses D, given a history of "highly cooperative" play defined as either at least 4 C actions in the last 5 rounds (left panel), or exactly 4 C actions in the last 5 rounds (right panel). The probability of defection increases with the introduction of reputational uncertainty in both slow and fast-changing social environments conditions.

In particular, we aggregate our data at the network level and compute the likelihood of choosing D for each type of cooperative history. We then use a Wilcoxon-Mann-Whitney (MW) test to evaluate the mean difference in this likelihood with and without reputational uncertainty, holding the rate of change of the social environment constant. When at least 4 out of the last 5 actions were C, the likelihood of choosing D increases with reputational uncertainty in both the *fast* ($p = 0.021$, MW) and *slow* ($p = 0.009$, MW) conditions. Similarly, when exactly 4 of the last actions were C, the likelihood of choosing D increases with reputational uncertainty in the fast-changing social environment ($p = 0.036$, MW), but only qualitatively in the slow-changing one ($p > 0.1$, MW).

In summary, the introduction of reputation uncertainty allows individuals who have a good reputation to "hide behind the veil of noise" to engage in opportunistic play. This rise in opportunistic play partly explains the drop in the overall cooperation level when reputational uncertainty is present.

Result 5: *Reputational uncertainty promotes opportunistic play.*

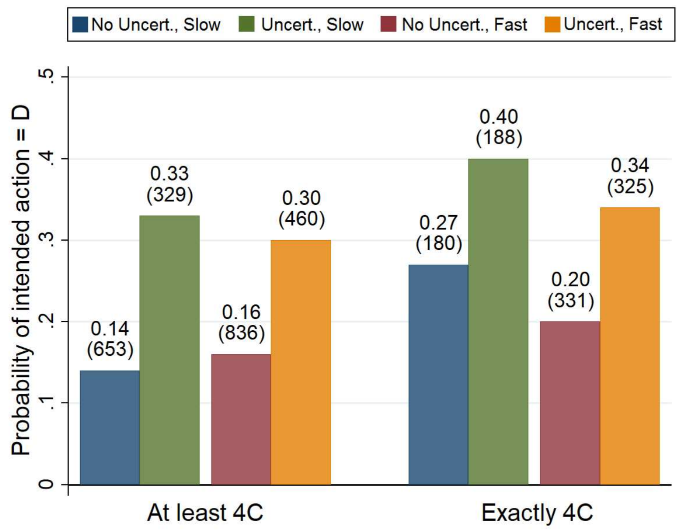

Figure 5: Probability of intended action being D given specific history of own actual actions in the last five rounds. Two histories are shown: one where a player had at least four occurrence of the actual action being C out of the last five rounds and the other where actual action was C exactly four times in the last five rounds. The number of observations are in parentheses.

4.4 Network Metrics and Homophily

Table 6 displays the average values for the following network metrics in each of our four treatments: *degree*, *local clustering* coefficient and *variance in degree*.[19] Similar to Section 4.1, we aggregate all observations for rounds 2-21, and apply the non-parametric Kruskal-Wallis test to compare the network metrics across multiple groups to detect treatment effects, followed by the Dunn's test to explore differences between any two treatments.

The presence of uncertainty does not affect any of the network metrics. Specifically, in the *slow* condition there is no difference in degree, clustering and variance in degree between the treatments with and without reputational uncertainty ($p > 0.1$ for all comparisons, DT). Similarly, there is no statistically significant difference for the same comparisons in the *fast* conditions ($p > 0.1$ for all comparisons, DT).

Result 6: *The aggregate network metrics remain unchanged whether or not repu-*

[19]For each round and each network group, we calculate the variance in the number of links. Subsequently, we take the average of the variance over rounds 2-21.

Table 6: Aggregate network metrics for rounds 2-21.

Treatment	Avg. degree	Avg. local clustering	Avg. variance in degree
(No Uncertainty, Slow)	7.49	0.71	2.89
(Uncertainty, Slow)	7.29	0.68	2.78
(No Uncertainty, Fast)	6.18	0.68	4.40
(Uncertainty, Fast)	6.73	0.71	4.85

tational uncertainty is present.

Our next goal is to understand how the presence of reputational uncertainty affects homophily among cooperators, which is defined as the tendency for cooperators to connect to each other. Following Currarini et al. (2009), suppose there are 2 different categories of individuals $i = 1, 2$. Let N_i denote the number of category i individuals in the population and let $w_i = N_i/N$ be the relative fraction of i individuals in the population, where $N = N_1 + N_2$. Let s_i denote the average number of connections that individuals of category i have with individuals who are of the same category and let d_i be the average number of connections that category i individuals form with individuals of categories different than i. The *Homophily index* H_i is equal to $H_i = s_i/(s_i + d_i)$ and it captures the fraction of the links of individuals of category i that are with that same category.

A drawback of the homophily index is that it varies with population size so it is not possible to use it to evaluate the extent of homophily for categories of individuals belonging to groups of different sizes. To deal with this drawback, we define a profile (s_1, d_1, s_2, d_2) to satisfy the baseline homophily if $H_i = w_i$ and inbreeding homophily if $H_i > w_i$. The *Inbreeding Homophily index* (IH) of category i is defined as $IH_i = (H_i - w_i)/(1 - w_i)$, and it measures the amount of bias with respect to the baseline homophily as it relates to the maximum possible bias which is the denominator $1 - w_i$.[20]

We classify a subject to be a cooperator if they chose C as their first round action,

[20]It can be easily checked that we have inbreeding homophily for category i if and only if $IH_i > 0$. The index of inbreeding homophily is 0 if there is pure baseline homophily and it is 1 if a group completely inbreeds.

and a defector if they chose D. This classification has the advantage that it is exogenous to the evolution of cooperation in the group as it is fixed before participants receive any feedback about game play. The average number of cooperators range from 6.63 to 6.88 across treatments, and there are no significant differences across treatments ($p > 0.1$, DT).

Table 7 displays the results of a Wilcoxon-Mann-Whitney (MW) sign-rank test on within treatment differences in the fraction of intended cooperative choices, average degree, homophily index and inbreeding homophily index between the cooperators and the defectors. As one would expect if our categorization into these types is meaningful beyond the first round, cooperators have more cooperative actions compared to defectors ($p < 0.05$ for all treatments, MW). Moreover, when there is no reputational uncertainty, cooperators have a higher degree and IH compared to defectors ($p < 0.05$ for both treatments, MW). Once we introduce reputational uncertainty, cooperators have a significantly higher degree and IH than defectors only in the slow-changing social environment ($p < 0.05$, MW), while the difference is only marginally significant in the fast-changing condition ($p < 0.1$, MW).

Table 7: The results of the Wilcoxon-Mann-Whitney (MW) sign-rank test on within-treatment differences in the fraction of intended cooperative choices, average degree, inbreeding homophily index (IH), and homophily index (H) between type-C and type-D participants for rounds 2 to 21. We have $n = 8$ sample size per treatment after aggregation at the group (network) level. The second row denotes the p-value from the MW test

Variables	**Within-treatment differences btw C and D types** $\binom{\text{difference}}{p\text{-value}}$			
	(*No Uncert., Slow*)	(*Uncertainty, Slow*)	(*No Uncert., Fast*)	(*Uncertainty, Fast*)
Fraction of intended cooperative actions	0.209	0.233	0.211	0.242
	0.050	0.012	0.012	0.012
Average degree	0.852	0.857	1.497	0.673
	0.012	0.036	0.012	0.069
Inbreeding Homophily index	0.120	0.128	0.253	0.110
	0.012	0.050	0.012	0.069
Homophily index	0.166	0.182	0.273	0.183
	0.093	0.208	0.012	0.093

The introduction of reputational uncertainty significantly decreases the inbreeding homophily index (IH) from 0.253 to 0.110 and the homophily index (H) from 0.273 to 0.183 among cooperators in the fast-changing social environment. We aggregate the IH index for cooperators over rounds 2-21 and use a Wilcoxon-Mann-Whitney test to evaluate the difference in IH among cooperators with and without reputational uncertainty holding the rate of change of the social environment constant. We find that this difference is statistically significant in the *fast* ($p = 0.046$, MW) but not in the *slow* condition ($p = 0.115$, MW).

Result 7: *Homophily among cooperators decreases with reputational uncertainty, but only in a fast-changing social environment.*

This result is consistent with the conjecture we put forward in Section 3 when we discussed homophily among cooperators. There we argued that with reputational uncertainty, it becomes difficult for individuals to recognize the truly cooperative individuals, and therefore we should expect that cooperativeness-based homophily to decrease in the presence of reputational uncertainty.

5 Conclusion

We provide empirical evidence regarding cooperation in an uncertain world with individuals receiving opportunities to update their neighborhood. Our experiments demonstrate that the presence of noise, defined as uncertainty on whether other people's observed action correspond to their intended action, affects individual behavior along the network as well as the action dimension. We further show that the extent of the impact of noise depends fundamentally on the frequency at which individuals receive opportunities to propose and/or remove links with others. In doing so, we extend our knowledge on the characterization of behavior under noise with endogenous network formation.

Our study raises a number of questions that suggest interesting possibilities for follow-up studies as future research. First, we do not impose an explicit cost on forming, maintaining, or removing links in our experiment.[21] Exploring the effect of such

[21] Real-world relationships take time to form and maintenance of such links involves time and effort, both of which are costly.

cost might provide us with additional insights than mere opportunity costs of not interacting as in our setup. Second, there are very few studies that investigate the effect of allowing communication on cooperation in repeated games (Embrey et al., 2013; Arechar et al., 2017). It would be interesting to study the effect of communication in our framework. Third, noisy information is publicly observable in our setup, so our game has an imperfect public monitoring structure. However, it is possible that players do not even know what signals other players have observed about their own play, thereby making the monitoring structure private. There are few empirical studies on cooperation under private monitoring, e.g., Kayaba et al. (2016) and Aoyagi et al. (2019) investigate two-person games with such a monitoring structure. Our setup could be extended to include this type of challenging monitoring technology.

References

Ahn, T.-K., R. M. Isaac, and T. C. Salmon (2009). Coming and going: Experiments on endogenous group sizes for excludable public goods. *Journal of Public Economics 93*(1-2), 336–351.

Ambrus, A. and B. Greiner (2012). Imperfect public monitoring with costly punishment: An experimental study. *American Economic Review 102*(7), 3317–32.

Ambrus, A. and B. Greiner (2017). Democratic punishment in public good games with perfect and imperfect observability. *Duke ERID Working Paper* (183).

Aoyagi, M., V. Bhaskar, and G. R. Fréchette (2019). The impact of monitoring in infinitely repeated games: Perfect, public, and private. *American Economic Journal: Microeconomics 11*(1), 1–43.

Aoyagi, M. and G. Fréchette (2009). Collusion as public monitoring becomes noisy: Experimental evidence. *Journal of Economic Theory 144*(3), 1135–1165.

Arechar, A. A., A. Dreber, D. Fudenberg, and D. G. Rand (2017). "I'm just a soul whose intentions are good": The role of communication in noisy repeated games. *Games and Economic Behavior 104*, 726–743.

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*(4), 1029–1042.

Bolton, G. E., E. Katok, and A. Ockenfels (2005). Cooperation among strangers with limited information about reputation. *Journal of Public Economics 89*(8), 1457–1468.

Brandts, J., A. Riedl, and F. Van Winden (2009). Competitive rivalry, social disposition, and subjective well-being: An experiment. *Journal of Public Economics 93*(11-12), 1158–1167.

Cason, T. N. and F. U. Khan (1999). A laboratory study of voluntary public goods provision with imperfect monitoring and communication. *Journal of Development Economics 58*(2), 533–552.

Charness, G., F. Feri, M. A. Melendez-Jimenez, and M. Sutter (2014). Experimental games on networks: Underpinnings of behavior and equilibrium selection. *Econometrica 82*(5), 1615–1670.

Charness, G. and C.-L. Yang (2014). Starting small toward voluntary formation of efficient large groups in public goods provision. *Journal of Economic Behavior and Organization 102*, 119–132.

Choi, S., E. Gallo, and S. Kariv (2016). Networks in the laboratory. In Y. Bramoullé, A. Galeotti, and B. Rogers (Eds.), *The Oxford Handbook of the Economics of Networks*, pp. 440–475. Oxford University Press.

Cinyabuguma, M., T. Page, and L. Putterman (2005). Cooperation under the threat of expulsion in a public goods experiment. *Journal of Public Economics 89*(8), 1421–1435.

Coricelli, G., D. Fehr, and G. Fellner (2004). Partner selection in public goods experiments. *Journal of Conflict Resolution 48*(3), 356–378.

Currarini, S., M. O. Jackson, and P. Pin (2009). An economic model of friendship: Homophily, minorities, and segregation. *Econometrica 77*(4), 1003–1045.

Dal Bó, P. (2005). Cooperation under the shadow of the future: Experimental evidence from infinitely repeated games. *American Economic Review 95*(5), 1591–1604.

Dal Bó, P. and G. R. Fréchette (2018). On the determinants of cooperation in infinitely repeated games: A survey. *Journal of Economic Literature 56*(1), 60–114.

Dal Bó, P. and G. R. Fréchette (2019). Strategy choice in the infinitely repeated prisoners dilemma. *American Economic Review 109*(11), 3929–3952.

Dreber, A., D. G. Rand, D. Fudenberg, and M. A. Nowak (2008). Winners don't punish. *Nature 452*(7185), 348.

Egas, M. and A. Riedl (2008). The economics of altruistic punishment and the maintenance of cooperation. *Proceedings of the Royal Society B: Biological Sciences 275*(1637), 871–878.

Ehrhart, K.-M., C. Keser, et al. (1999). Mobility and cooperation: On the run. *mimeo*.

Embrey, M., G. R. Fréchette, and E. Stacchetti (2013). An experimental study of imperfect public monitoring: Efficiency versus renegotiation-proofness. *SSRN Working Paper*.

Fehr, E. and S. Gächter (2000). Cooperation and punishment in public goods experiments. *American Economic Review 90*(4), 980–994.

Fischbacher, U. (2007). z-tree: Zurich toolbox for ready-made economic experiments. *Experimental economics 10*(2), 171–178.

Friedman, J. W. (1971). A non-cooperative equilibrium for supergames. *Review of Economic Studies 38*(1), 1–12.

Fudenberg, D. and E. Maskin (1986). The folk theorem in repeated games with discounting or with incomplete information. *Econometrica 54*(3), 533–554.

Fudenberg, D., D. G. Rand, and A. Dreber (2012). Slow to anger and fast to forget: Leniency and forgiveness in an uncertain world. *American Economic Review 102*(2), 720–749.

Gächter, S., E. Renner, and M. Sefton (2008). The long-run benefits of punishment. *Science 322*(5907), 1510–1510.

Gallo, E., Y. E. Riyanto, T.-H. Teh, and N. Roy (2019). Strong links promote the emergence of cooperative elites. *Scientific Reports 9*(10857).

Gallo, E. and C. Yan (2015). The effects of reputational and social knowledge on cooperation. *Proceedings of the National Academy of Sciences 112*(12), 3647–3652.

Güth, W., M. V. Levati, M. Sutter, and E. Van Der Heijden (2007). Leading by example with and without exclusion power in voluntary contribution experiments. *Journal of Public Economics 91*(5-6), 1023–1042.

Hauk, E. and R. Nagel (2001). Choice of partners in multiple two-person prisoner's dilemma games: An experimental study. *Journal of Conflict Resolution 45*(6), 770–793.

Holt, C. A. and S. K. Laury (2002). Risk aversion and incentive effects. *American Economic Review 92*(5), 1644–1655.

Jackson, M. O. and A. Wolinsky (1996). A strategic model of social and economic networks. *Journal of Economic Theory 71*(1), 44–74.

Kayaba, Y., H. Matsushima, and T. Toyama (2016). Accuracy and retaliation in repeated games with imperfect private monitoring: Experiments and theory. *Working Paper*.

Maier-Rigaud, F. P., P. Martinsson, and G. Staffiero (2010). Ostracism and the provision of a public good: Experimental evidence. *Journal of Economic Behavior and Organization 73*(3), 387–395.

Morris, S. (2000). Contagion. *Review of Economic Studies 67*(1), 57–78.

Nikiforakis, N. and H.-T. Normann (2008). A comparative statics analysis of punishment in public-good experiments. *Experimental Economics 11*(4), 358–369.

Page, T., L. Putterman, and B. Unel (2005). Voluntary association in public goods experiments: Reciprocity, mimicry and efficiency. *Economic Journal 115*(506), 1032–1053.

Rand, D. G., S. Arbesman, and N. A. Christakis (2011). Dynamic social networks promote cooperation in experiments with humans. *Proceedings of the National Academy of Sciences 108*(48), 19193–19198.

Rand, D. G., D. Fudenberg, and A. Dreber (2015). It's the thought that counts: The role of intentions in noisy repeated games. *Journal of Economic Behavior and Organization 116*, 481–499.

Rand, D. G., M. A. Nowak, J. H. Fowler, and N. A. Christakis (2014). Static network structure can stabilize human cooperation. *Proceedings of the National Academy of Sciences 111*(48), 17093–17098.

Riedl, A., I. M. Rohde, and M. Strobel (2016). Efficient coordination in weakest-link games. *Review of Economic Studies 83*(2), 737–767.

Riedl, A. and A. Ule (2002). Exclusion and cooperation in social network experiments. *mimeo*.

Rojas, C. (2012). The role of demand information and monitoring in tacit collusion. *The RAND Journal of Economics 43*(1), 78–109.

Shirado, H., F. Fu, J. H. Fowler, and N. A. Christakis (2013). Quality versus quantity of social ties in experimental cooperative networks. *Nature communications 4*, 2814.

Wang, J., S. Suri, and D. J. Watts (2012). Cooperation and assortativity with dynamic partner updating. *Proceedings of the National Academy of Sciences 109*(36), 14363–14368.

Yang, C.-L., M.-L. Xu, J. Meng, and F.-F. Tang (2017). Efficient large-size coordination via voluntary group formation: An experiment. *International Economic Review 58*(2), 651–668.

A Demographic variables, by treatment

Table A1. Average values of demographic variables.

Treatment	*(No Uncertainty, Slow)*	*(Uncertainty, Slow)*	*(No Uncertainty, Fast)*	*(Uncertainty, Fast)*
Age	20.91	21.31	20.84	21.01
Gender (Female)	57%	45%	59%	49%
Risk Aversion	5.42	5.62	5.91	6.03
Game Theory	25%	24%	23%	22%
Prior Participation	63%	46%	46%	47%
Trust	25%	17%	28%	20%

Risk Aversion shows the score from Holt and Laury (2002) risk-elicitation test, in which a score of 0 indicates minimal risk aversion; *Trust* shows the percentage of participants indicating "yes" in the standard interpersonal trust question taken from the World Values Survey (WVS).

Table A1 summarizes the main socio-demographic characteristics of our participants. All 384 participants are students from various majors, ranging from humanities, social sciences, economics and business, engineering and sciences, of Nanyang Technological University (NTU) in Singapore. We verified that there are no significant differences ($p > 0.10$) across the four treatments in terms of these socio-demographic characteristics.

B Experimental Instructions for *(No Uncertainty, Fast)* treatment

Welcome to all of you! You are now taking part in an interactive study on decision making. Please pay attention to the information provided here and make your decisions carefully. If at any time you have questions to ask, please raise your hand and we will attend to you in private.

Please note that **unauthorized communication is prohibited**. Failure to adhere to this rule would force us to stop this study and you may be held liable for the cost incurred in this experiment.

Your participation in this study is voluntary. You will receive SGD **2** show-up fee for participating in this study. You may decide to leave the study at any time. Unfortunately, if you withdraw before you complete the study, we can only pay you for the decisions that you have made up to the time of withdrawal, which could be substantially less than you will earn if you complete the entire study.

The amount of your earnings from this study depends on the decisions you and others make and also on chance. At the end of this session, your earnings will be paid to you privately and in cash. They will be contained in an envelope (indicated with your unique user ID).

General Information

Each of you will be given a unique user ID and it will be clearly stated on your computer screen. At the end of the study, you will be asked to fill in your user ID and other information pertaining to your earnings from this study in the payment receipt. **Please fill in the correct user ID to make sure that you will get the correct amount of payment**.

Rest assured that your **anonymity will be preserved** throughout the study. You will never be aware of the personal identities of other participants during or after the study. Similarly, other participants will also never be aware of your personal identities during or after the study. You will only be identified by your user ID in our data collection. All information collected will **strictly be kept confidential** for the sole purpose of this study.

Specific Information

The total duration of this study is approximately 2 hours.

Your total earnings = earnings from Part I + earnings from Part II + show up fee (S$2)

All incentives will be denominated in **experimental dollars** (expressed as **ECU**). The exact conversion rate will be detailed later.

Part I

You and 11 other participants will engage in several rounds of interactions. Each round is the same and consists of 3 Stages.

In Stages 1 and 2, you can form and delete links with any of the other participants, which we will explain shortly. If you are linked with another participant then this participant is a "neighbor".

In Stage 3, you choose either action A or action B, and the choice of action A or B applies to all your neighbors. Action A is color-coded in green and action B is color-coded in blue. The colors are only a visual aid to distinguish the actions and have no meaning. You get points for the action you choose and the action each of the other participants chooses, in the following way:

You get 0 points if you are not linked with another participant regardless of your choice of action.

The number of points you get depends on the actions you and your neighbor choose, according to the table below:

		Neighbor	
		A	B
You	A	3	-5
	B	5	-3

This is what the table says:

- If you choose A and the neighbor chooses A, you get 3 points
- If you choose A and the neighbor chooses B, you get -5 points
- If you choose B and the neighbor chooses A, you get 5 points
- If you choose B and the neighbor chooses B, you get -3 points

At the end of each round you will see a summary of the number of points you get with each of the other participants.

You will play 25 rounds of this game for sure. After round 25, there is a 50% chance that the game will terminate in the following round. In other words, in every round after round 25 the experimenter will "flip a coin" and if the outcome is "Heads" then the game will terminate.

Stage 1 - Link Decisions

In Stage 1, you will have the opportunity to (i) form a link with another player who is not linked to you or (ii) to delete an existing link with one of your partner(s). However, you will not have the complete freedom to form or to delete links. In particular, 33 random pairs of players (out of the 66 possible pairs of players in this network) will be selected to be updated such that:

(i) If you and another player that is selected are not linked currently, you will need to decide whether to propose a new link between you two.

(ii) If you and another player that is selected are linked currently, you will need to decide if you would like to delete the link.

For illustration purposes, the screenshots from Round 4 are shown (see also the attachment (stage 1).

<u>Top-left of the screen</u>

You will see the following figure that visualizes your current neighbors.

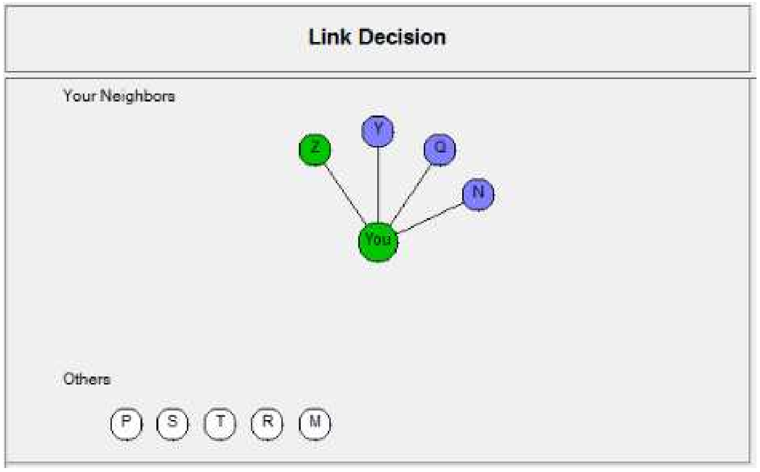

The neighbors you have at the start of each round are the same as the neighbors you had at the end of the previous round. A link means that if you are linked to participant Z then participant Z is linked to you.

Each circle denotes a participant with an ID. The "You" circle is always positioned at the center. For example, above you are linked with Z, Y, Q, and N. The color of the circles represents the last actual action (A or B) that has been chosen by each participant.

Note that in Round 1 all participants are connected to each other.

<u>Center of the screen</u>

This table allows you to make your decisions to **delete** links in Stage 1. The first row of the table ("Your Last 5 Actions") reminds you of the actions you have chosen in the last 5 rounds. For example, the sequence below means that you chose action A in the last round (the **leftmost**), action A two rounds ago (the **second-leftmost** slot), action A three rounds ago (the **middle**), and action B four rounds ago (the **second rightmost**), and action B five rounds ago (the **rightmost**).

For the pairs of players selected to be updated, a box will be shown next to the ID of the participant. You can delete a link to the neighbor by ticking the corresponding box under the "Unlink" column. There is no limit to how many boxes among the shown boxes you can tick, and you can also choose not to delete any link by leaving the shown boxes unticked.

The table has 3 columns, from left to right:

- **"Neighbor"**: lists the IDs of the participants who are currently linked with you
- **"Past actions"**: lists the last 5 actual actions by each of your neighbors
- **"Unlink"**: allows you to delete your link to the neighbors that have been randomly selected by the system. For example, ticking the box in the last row will delete your link to participant Z

The color-codings of the other participants' actions are the same as before. The green letter A denotes the choice of action A. The blue letter B denotes the choice of action B. The "-" denotes that the participant did not have any neighbor in that round, and therefore the participant did not choose an action.

Right of the screen

This table allows you to make your decisions to **form** links in Stage 1.

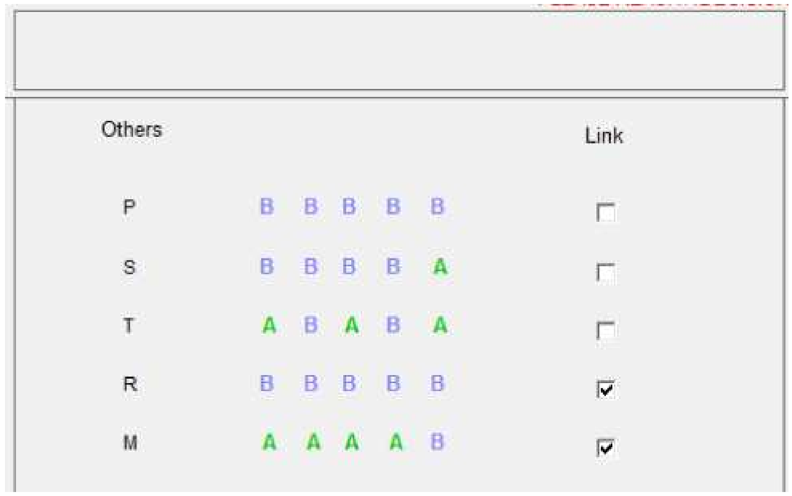

For the pairs of players selected to be updated, a box will be shown next to the ID of the participant. You can propose a link to the participant by ticking the corresponding box under the "Link" column. There is no limit to how many boxes among the shown boxes you can tick, and you can also choose not to propose any link by leaving all the boxes unticked.

The table has 3 columns, from left to right:

- **"Others"**: lists the IDs of the participants who are not linked with you
- **"Past actions"**: lists the last 5 actual actions by each of these participants
- **"Link"**: allows you to propose a link to the participants randomly selected by the system you are not linked with. For example, ticking the box in the last row will propose a link to participant M

Bottom-right of the screen

Note that the participants whom you can propose or delete links with are decided randomly by the computer. The text reminds you about the decisions you need to make and lists the participants whom you can propose or delete links with. You can propose new links and/or delete existing links, or you can choose not to propose and/or delete any links. In either case, you need to click the "Submit" button to confirm your decisions.

Stage 2 - Yes/No Decision

In case you decided not to link to another participant in Stage 1 but this other participant has opted to send you a link proposal, then in Stage 2 we ask you to confirm your earlier decision not to link to this other participant. You will be shown the following table where you need to respond to link proposals in Stage 2.

Please refer to the attachment (Stage 2) for illustration purpose.

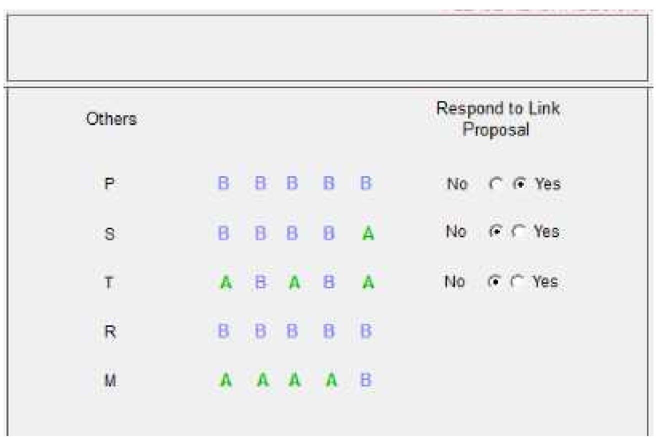

As in Stage 1, the table has 3 columns, from left to right:

- **"Others"**: lists the IDs of the participants who are not linked with you
- **"Past actions"**: lists the last 5 actual actions by each of these participants
- **"Respond to link proposal"**: allows you to accept or reject a link proposal from a participant that is not currently linked with you

For example, look at participants P, S and T: under the "Respond to link proposals" column a "No" and a "Yes" buttons appear. This means that P, S and T have sent you a link proposal. For each proposal, you need to click on “Yes” to accept it or “No” to reject it. In the example shown, you have chosen to reject the proposals from S and T, and accept the proposal from P.

Note that if you have sent a proposal to connect to a participant in Stage 1 and that participant has done the same then the link is automatically formed without the need of further approval in Stage 2, and hence no button will be shown.

Stage 3 - Action Choice

In Stage 3, you choose either action A or action B. Please refer to the attachment (Stage 3) for illustration purpose.

Bottom-right of the screen

This is where you choose your action in Stage 3 by clicking either the A or B button. You must pick an action in order to proceed.

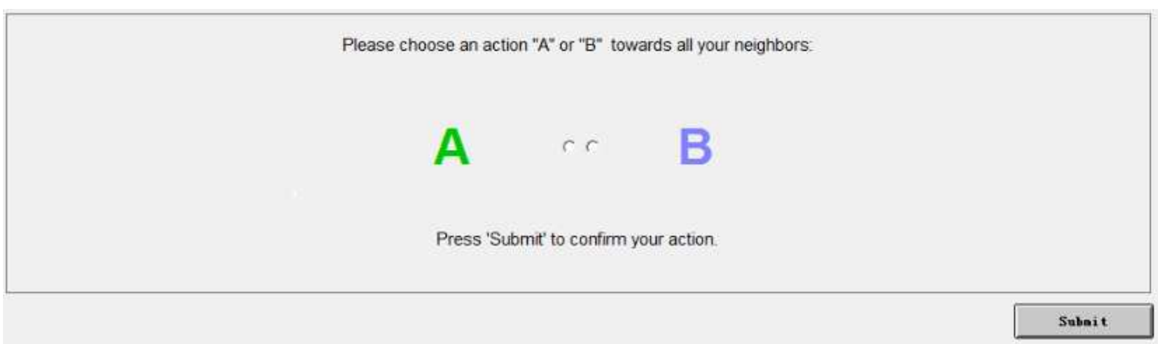

Note that it is possible that you are not linked with any participant (i.e. you have no neighbor) as the result of yours and others' earlier decisions. In such case, you will instead see the message reminding you that you do not have any neighbor and so you do not need to choose an action.

Action Outcome

After Stage 3, you will see a summary of the points you got with each of the other participants in the round.

The action you have chosen (**intended action**) during Stage 3 may or may not be implemented correctly. Specifically, there is a **15% chance** that your intended action A is implemented as **actual action** B and a **15% chance** that your intended action B is implemented as **actual action** A. Note that other players only observe your implemented (**actual**) action and **not** your **intended** action.

Please refer to the attachment (Action Outcome) for illustration purpose.

The table shows your neighbors and your neighbors' **actual actions** in this round. The first row of this table reminds you of your intended action and the actual action implemented. For example, in the table below your intended action was A and your actual action was A. Note that the table is updated with yours and the others' decisions in the previous Stages.

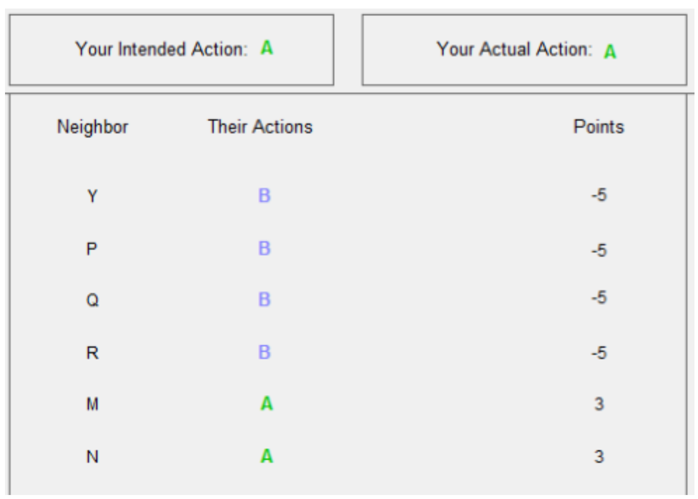

Your Intended Action: A

Your Actual Action: A

Neighbor	Their Actions	Points
Y	B	-5
P	B	-5
Q	B	-5
R	B	-5
M	A	3
N	A	3

The table has 3 columns, from left to right:

- **"Neighbor"**: lists the IDs of the participants who are currently linked with you
- **"Action"**: lists the actual action by each of your neighbors in this round
- **"Points"**: lists the points you gained from your and each neighbor's actual action

If your intended action is implemented correctly, you will see the following box.

If your intended action is **not** implemented correctly, you will see the following box. You need to click the "Continue" button to move to Stage 1 of the next round.

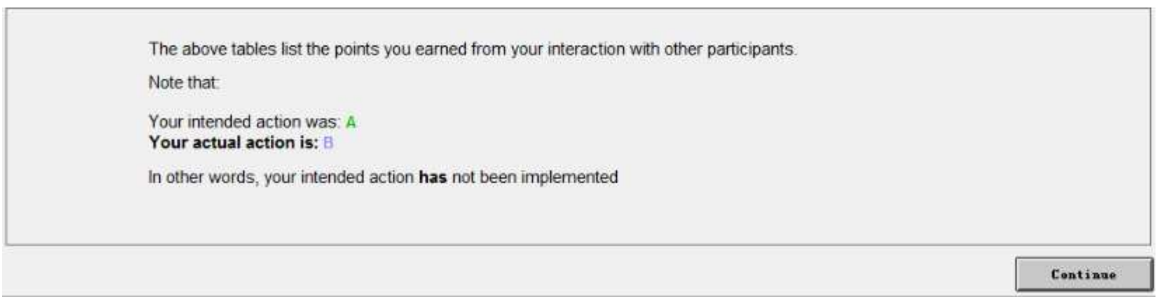

Earnings in Part I

At the end of the Experiment, we will randomly select **6** rounds for payment. In each of these **6** rounds, we will randomly pick **2** of your potential link pairs (There are 11 potential link pairs in total). Note that a participant who was not linked with you can be picked, and in that case you will get 0 points for the interaction with that participant.

To determine your earnings, we sum the number of points you got with each of the picked participants in each of the **6** rounds. The exchange rate from ECU to SGD is:

1.5 ECU = 1 SGD

Part II

In this part of the study you will be asked to make a series of choices. How much you receive will depend partly on chance and partly on the choices you make. The decision problems are not designed to test you. What we want to know is what choices you would make in them.

For each line in the table you will see in this stage, please indicate whether you prefer option A or option B. There will be a total of 10 lines in the table but just one line will be randomly selected for payment. You do not know which line will be paid when you make your choices. Hence you should pay attention to the choice you make in every line.

After you have completed all your choices, the computer will randomly generate a number, which determines which line is going to be paid out.

Your **earnings for the selected line depend on which option you chose**: If you chose **option A** in that line, you will receive **20 ECU**. If you chose **option B** in that line, you will receive **either 60 ECU or 0**. To determine your earnings in the case you chose option B, there will be a second random draw. The computer will randomly determine if your payoff is 0 or 60, with the chances set by the computer as they are stated in Option B.

Earning in Part II

In Part II, the conversion rate from experimental dollars (ECU) to Singaporean dollars is:

20 ECU = 1 SGD

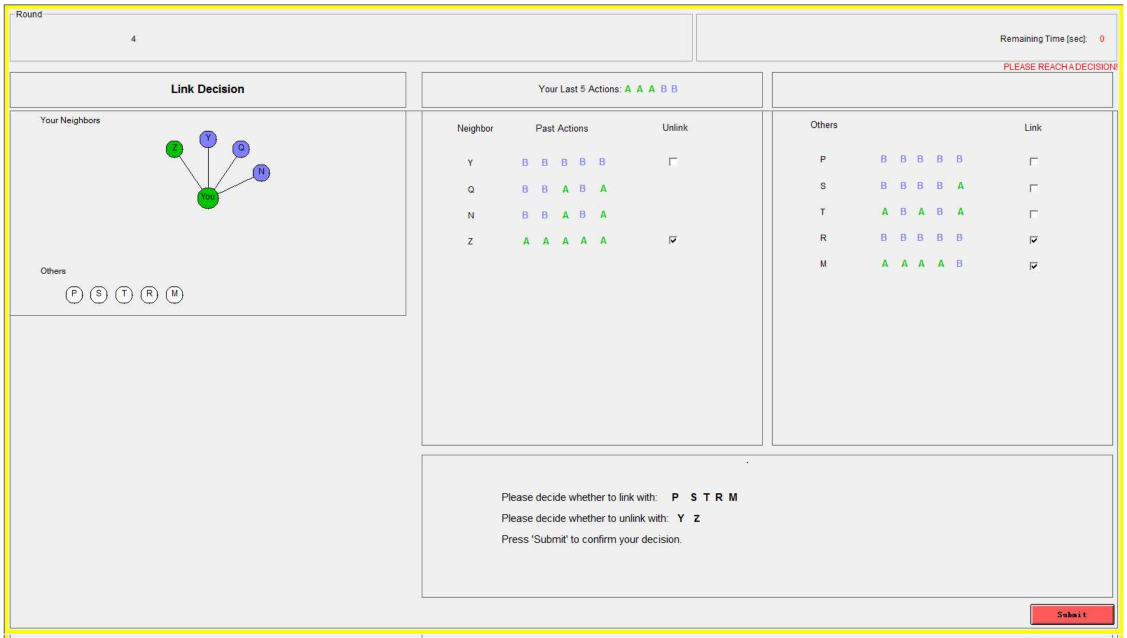

Attachment (Stage 1)

Attachment (Stage 2)

Attachment (Stage 3)

Attachment (Action Outcome)