The back-scratching game

Cameron K. Murray*, Paul Frijters, Melissa Vorster

November 19, 2020

ACCEPTED FOR PUBLICATION VERSION
(PUBLISHED IN JOURNAL OF ECONOMIC BEHAVIOR AND ORGANIZATION, 2017)

Abstract

We develop a new experiment to study the emergence of welfare-reducing bilateral alliances within larger groups, and the effectiveness of institutional interventions to curtail ‘back-scratching’. In each of the 25 rounds of our experiments, a player (the ‘allocator’) nominates one of three others in a group as a co-worker (the ‘receiver’), which determines the group production that period to be the productivity of the receiver (which varies by round), but also gives the receiver a bonus and makes them the allocator in the next round. Alliances then form if two individuals keep choosing each other even when their productivities are lower than that of others, causing efficiency losses. We study whether particular interventions reduce the rate of alliance formation. Random allocator rotation policies were found to be ineffective, whether implemented prior to, or following, a baseline treatment. Low bonuses did significantly reduce the advent of alliances. There hence seems to be some hope in preventing back-scratching via reducing the gain that can be made by alliances.

Keywords: Alliance formation, Corruption, Reciprocity, Experiment

JEL Classification: D71, D72, D63

1 Introduction

Lehman Brothers filed for Chapter 11 bankruptcy protection on September 15, 2008, triggering a collapse in financial markets. A lack of diligence by regulatory authorities in the United States, whose impartiality was undermined by the ‘revolving-door’ of personnel between private banks, lobbyists and regulators, was arguably a key element in the crisis (Roubini and Mihm, 2010; Levine, 2012; Matthews, 2014; Barth et al., 2012).¹ For instance, Rohit Bansal, a Goldman Sachs investment banker who had previously worked at the NY Fed, pleaded guilty to obtaining market sensitive information from a NY Fed employee, Jason Gross, that undermined regulatory oversight (Silver-Greenberg et al., 2014; Protess and Eavis, 2015).

Because alliances like the one between Rohit Bansal and Jason Gross are hard to observe and conduct policy experiments on, we test the effect of institutional designs to combat back-scratching in a new lab experiment that captures both the issue of how specific alliances emerge, as well as how they can be broken up. Whilst the existing literature has focussed on the circumstances

*Corresponding Author: University of Queensland, Level 6, Colin Clark Building (39), St Lucia Brisbane Qld 4072. Email: ckmurray@gmail.com

¹Leaked recordings by Federal Reserve Bank of New York (NY Fed) whistle-blower Carmen Segarra provide evidence as to the degree of favouritism regularly shown between regulators and the regulated.

in which matched pairs are prepared to cause large negative effects on society, we focus on the circumstances where individuals must first find their alliance partners and earn their trust. Our design adds to existing experimental approaches to corruption and costly reciprocity.²

In our baseline experiment, in each of the 25 rounds the allocator nominates one of the three others in the group as a co-worker (the ‘receiver’), which determines the group payoff that period to be the productivity of the receiver, and where the productivity of all potential receivers is randomly drawn each round (from the set $\{1,2,3\}$). The receiver also gets a bonus payment of 25 Experimental Currency Units (ECU) and becomes the allocator in the next round, allowing the receiver to return the favour of allocating the bonus payment, independent of the productivity of their partner. Alliances form if two individuals keep choosing each other as the receiver even when their productivities are lower than that of others, creating a cost to others in the group. A choice to allocate the highest productivity receiver in a round maximises the group payoff, which we label a ‘meritocratic’ choice, the frequency of which is the main variable of interest. The key innovation with regards to the existing literature is that it is not given before the experiment who will be harmed, who will form alliances, and who will benefit from alliances, allowing us to look at the dynamics and social norms involved in alliance formation, and the effect of policy intervention on both the emergence of new alliances and the break-up of existing ones.

Our first policy treatment is Rotation, which mimics a staff rotation policy by introducing a degree of randomness in who makes the discretionary decision each round by reducing the probability of the previous round receiver being the new allocator to 0.5. Rotation policies are common in sensitive areas of business and regulation. For example, professional bodies require accountants to rotate out of auditing roles after two years to ensure they do not form alliances. Similar policies could be implemented in areas of financial regulation that are prone to alliance formation. This treatment reduces potential long-run payoffs of alliances, and allows for countervailing alliances to form, and hence should reduce the rate of formation of new alliances, but it is not clear whether it will help break up existing alliances.

In our second policy treatment, the bonus payment each round is reduced to 3 ECU, an amount that is sufficient to make a meritocratic strategy, whereby every allocator chooses the highest productivity player each round as receiver, more profitable than being in an alliance, as long as everyone else is meritocratic. We interpret this treatment as reducing the level of individual discretion, which might in reality take the form of having strongly enforced rules, or as Congleton (2014) explains, the bureaucratisation of individuals into “cogs in an organizational machine”. Ohashi (2009) studies an example of such successful rule systems, where the introduction of qualifying rules for government tenders in Japan reduced the discretion of public officials in choosing who could bid for contracts, and reduced procurement costs by 8%. This treatment is also informed by rent-seeking theory which suggests that the size of rents available is a primary determinant of the efficiency cost of rent-seeking activity (Lambsdorff, 2002), and hence we call this a Low Rent treatment.

Each of the 70 groups of four subjects play a sequence of two treatments, allowing us to vary whether the implemented policy treatments come first or second, hence enabling us to see if particular treatments are useful in preventing alliances from forming and/or breaking up alliances that have formed earlier.

In all baseline treatments, 38% of rounds are identifiable as alliance-play. The Low Rent treat-

²Our paper adds to the experimental literature on coordination with negative externalities, including Frank and Schulze (2000); Schulze and Frank (2003); Abbink (2004); Abbink and Hennig-Schmidt (2006); Armantier and Boly (2008); Schikora (2011b); Büchner et al. (2008); Schikora (2011a); Lambsdorff and Frank (2011); van Veldhuizen (2011); Barr et al. (2009); Barr and Serra (2009).

ment had 14% more meritocratic play, and fewer alliances, if it was the first treatment in the sequence, but had only small effects when it was preceded by a baseline treatment. In contrast, the Rotation treatment provided no substantial increase in meritocratic choices regardless of the sequence and even resulted in more overall alliance combinations being observed in each group.

Our notion of back-scratching is wider than just corruption: like most instances of favouritism, Bansal and Gross’s case does not cleanly fit a definition of corruption of “acts which utilise the power of public office for personal gain in a manner that contravenes stipulated rules” (Jain, 2001). Because the public office holder in this case did not receive a personal gain, but merely provided a favour to another person by capitalising on a loose interpretation of the confidentiality rules, it is a clear example of how a degree of discretion within rule systems can give rise to implicit *quid pro quo* alliances formed on bases other than merit. Such alliances can be payoff-increasing for the insiders, yet come with efficiency losses that are hard to measure. For instance, Faccio et al. (2006) find that companies are more likely to be bailed out with public money when they have political connections, though the authors were not able to observe any favours returned to politicians.

2 Background

The closest experiment to ours is the repeated bribery game (RBG) that was pioneered by Abbink et al. (2002). Each stage in the RBG involves two players who have been assigned by the experimenter into the role of briber or public official, where the briber first makes a choice of how much to offer as a bribe (if any), after which the official makes the choice to accept the bribe (or not), and then makes an allocation decision that may favour the briber. A public official’s allocation choice in favour of the briber triggers a negative externality determined by the experimenter, either by subtracting earnings from other subjects in the experimental session (Abbink et al., 2002; Abbink, 2004), or by reducing the size of a charity donation by the experimenter (Lambsdorff and Frank, 2010; Schikora, 2011b; van Veldhuizen, 2011).

Many interventions have been tested in the RBG setting. A key finding by Abbink (2004) is that staff rotation is effective at reducing bribery and corrupt decisions, where rotation was implemented by randomly pairing bribers and public officials each round. The basic RBG design has since been augmented to allow for whistle-blowing (Schikora, 2011a; Lambsdorff and Frank, 2010), an outside informed monitor (Schikora, 2011b), payoff variations that mimic high non-bribery wages (van Veldhuizen, 2011; Armantier and Boly, 2008), and ‘citizen’ and probabilistic punishment (Abbink et al., 2002; Serra, 2012; Cameron et al., 2009). Abbink (2004) cautions though that the fixed briber-official setup of the RBG does not fit situations where unobserved alliances may exist before policy interventions take place, nor where potential alliance partners have to find each other rather than be paired by the experimenter.

Abbink and Serra’s (2012) survey of these and other corruption experiments report a large willingness to form alliances on the part of both briber and official. For instance, Cameron et al. (2009) found that even with the chance of punishment by a third player, 78–93% of players will offer a bribe, with 77–93% of bribed officials reciprocating by making an allocation decision that favours the briber but that has negative consequences for the third player. Whether a higher negative externality makes a difference to the actions of the official is uncertain: Abbink et al. (2002), Büchner et al. (2008) and Cameron et al. (2009), find it makes no difference in the RBG and its variations. Yet, Barr and Serra (2009) do find an effect. They implement a one-shot petty corruption experiment where externalities automatically come from accepting bribes, and fall on passive members of the experimental session not involved in making the

bribery decisions. In their setup higher externalities are related to fewer and smaller bribes, both offered and accepted. In our setup, where alliances are endogenous rather than fixed by the experimenter, the externalities come from alliance formation, which makes it unclear how greater externalities would affect alliances: the greater harm on others might make alliance partners more fearful of reprisals should they make someone outside of their alliance the allocator, potentially strengthening the alliance.

Lambsdorff and Frank (2011) extend the RBG design by having the briber be a potential whistleblower on the public official if that public official does not reciprocate on the bribe. Interestingly, they find that men and economics students are more corrupt in general as they are more likely to bribe and accept the bribe. Men and economics students also were more likely to punish the public official if their accepted bribe wasn't reciprocated.

A very different design is that by Greiner and Schneider (2015). In their study, subjects interact over 100 rounds in fixed groups. In the first stage of each round, subjects vote on who will be the dictator in the second stage and thereby gets to allocate group resources. Despite it being the case that there is no stable coalition if everyone were rational and selfish (because then every member should vote for themselves), they find strong and long-lasting coalitions, usually a minimum majority who get all the resources distributed amongst them by a dictator in that coalition, excluding the others. This setup mimics the emergence of large coalitions, like clans or ethnicities within larger societies, and also shows that laboratory participants quickly recognise the possibility of forming a reciprocal alliance with a subset of other players that benefits them at the exclusion of others. Yet, given that the size of the pie is fixed in their setting, there is no efficiency loss of such coordination and hence also no direct welfare implication.

Rand et al. (2011) had a local public goods design that included the possibility of players using information on the previous actions of others to infer who are likely to be contributors in future rounds. Their design had individuals pair up with several 'neighbours' whom they could send 50 units, with the neighbours receiving 100. Individuals could then choose new neighbours in later rounds, and were found to want to pair up with those individuals who indeed sent 50 units to their neighbours in previous rounds, indicating that individuals expect pro-social behaviour to be a persistent trait. This is very salient to our experiment where individuals looking to form an alliance will also be trying to team up with reciprocators contributing to the alliance, though in our case the question is how such reciprocation can be discouraged rather than encouraged, as there is a group loss.

Summarising, the main innovation in our design is that we have an easy-to-implement dynamic setup in which alliances that have endogenous externalities can arise, and in which we can see whether policy interventions can prevent alliances from emerging and/or break down existing alliances. The lack of enforceability or punishment of alliances in our design arguably corresponds well to real world favouritism which is often very hard to observe and can only be indirectly discouraged.

In light of the results of Lambsdorff and Frank (2011), we have an expectation that men and economics students are more likely to participate in back-scratching in our design. The implication any relationships between socio-demographic factors and back-scratching is that one available mechanism to reduce back-scratching is through the choice of new staff, which is particularly relevant to regulatory agencies.

In terms of external validity, Barr and Serra (2010) show that in an RBG game the degree of corruption by undergraduate subjects was predicted from the corruption index of their country of origin, suggesting that laboratory choices regarding bribery and corruption indeed capture an

aspect of real-world behaviour. Additionally, Armantier and Boly (2008) ran a field experiment in Burkina Faso where recruited exam-markers were unaware of their participation in a bribery experiment and were offered bribes folded in some exam papers. They found, consistent with their their laboratory results in Canada that had the same design, that higher wages for exam-markers decreased the prevalence of accepting bribes, showing a correspondence between the behavioural change following an institutional treatment in the lab with the same institutional treatment in the field.

3 Experimental Design and Research Questions

Subjects play in groups of four whose composition is fixed for the duration of the experiment and they are identifiable to each other by a unique coloured shape³. The game proceeds in rounds, of which there are 50 in total covering two treatments of 25 rounds. Each round, one player, the allocator, chooses which of the other three players to receive a 25 Experiment Currency Units (ECU) payment, with the first round allocator randomly chosen. The receiver of the payment in a round becomes the allocator for the next round, providing the potential for back-scratching to emerge.

Each round the three potential payment receivers are given a randomly shuffled productivity number, from the set $\{1, 2, 3\}$, which the allocator can observe before making a decision. The payoff for all players (including the allocator) in each round is equal to the chosen receiver's productivity number, while the chosen receiver also gets the 25 ECU payment in addition. We call a choice meritocratic if the receiver has the highest productivity number. Table 1 shows player payoffs for the three possible choices in a round where Circle is the allocator, Triangle happens to have productivity 1 that round, Square has 2, and Pentagon has 3.

Table 1: Possible player payoffs in a stage game

Circle's choice of receiver	Circle (Allocator)	Triangle (Prod. 1)	Square (Prod. 2)	Pentagon (Prod. 3)
Triangle	1	26	1	1
Square	2	2	27	2
Pentagon	3	3	3	28

With selfish rational players the equilibrium for both the one-shot game, and by backwards induction the repeated game, is what we call the meritocratic strategy of allocating to the highest productivity player. It is the institutional structure, whereby the payment receiver makes the allocation decision in the next round, that provides scope for a wide variety of cooperative outcomes in the repeated game. Since the receiver of the payment is able to discern whether the allocator has favoured them (i.e. their productivity number was less than 3), the allocating player can use a non-meritocratic choice as a costly signal to the receiving player of their intention to form an alliance; the cost arises from the reduced payment to the allocator. In all there are six potential two-player alliance combinations in each group.

To reflect the often hidden nature of back-scratching we limit the information provided to players each round. Players not allocating are able to see only their own productivity number, though they do know the distribution of all productivity numbers. After a round, players see only their

³One might think it is better to be one shape than another, but no such pattern was found.

new total earnings and not the specific earnings from the round, ensuring that only observant players who remembered their previous earnings balance, or had a productivity of 3, could infer whether or not the previous decision was meritocratic. We do this to reflect the normally hidden nature of favourable back-scratching decisions.⁴ Screenshots of the game are in Figure 3 of the Appendix.

3.1 Treatments

In addition to the Baseline treatment described above, we implement two policy treatments designed to mimic commonly prescribed anti-corruption measures. By having these treatments imposed either immediately, or after a Baseline, we test whether the treatments deter the formation of new alliances and/or break down alliances built up in a previous 25 rounds. As such, we use a within-subject, or group, design in which the receiver in the final (25th) round of the first treatment becomes the allocator in the first round of a second treatment.

3.1.1 Rotation

We reduce the ability for immediate reciprocation in this treatment by reducing the receiver’s probability of being the next allocator from 1 to 0.5, which allows all other players to have a 0.17 probability of being the next round allocator. To be clear, the person nominated by the allocator as receiver still gets the bonus and can hence deduce that they were chosen by the allocator.

This Rotation treatment mimics policies of staff rotation in senior positions, which are necessarily imperfect and do not occur on a ‘decision-by-decision’ basis (which would be implemented as a 0.25 probability of each player being the next round allocator). Our ‘mild’ rotation reflects rotation policies in practice that often rely on a small pool of candidates who make multiple decisions before being rotated out. It is expected that a fully random rotation policy each round in this design would almost completely avert alliance formation for players who had not yet played a Baseline treatment, as ‘pure’ rotation policies have been shown to achieve this in other corruption experiments (Abbink, 2004).

The basic idea of this design is to capture the various mechanisms observed in institutions which have the intent of introducing randomness. The jury system in law courts is one example of a very strong rotation policy, where each decision comes from newly rotated people from a large pool of candidates. Requirements by professional standards boards that ensure auditors do not work for the same company for more than a fixed period is an example of the type of imperfect rotation policy we have in mind. Similarly, following insider involvement in drug importation by Australia Customs officers, periodic staff rotation policies were recommended to combat the formation of alliances between officers and criminal gangs (Pezzullo, 2013), again, the type of imperfect rotation this experimental treatment represents.

3.1.2 Low Rent

In the Low Rent treatment we reduce the bonus payment allocated each round from 25 ECU to 3 ECU. The intention is to capture policies that reduce the size of the economics rents able to be allocated with discretion, such as via the introduction of fixed bureaucratic processes, or by taxation or sale of the private value of political decisions. A real world example is the practise of centralising staff purchasing arrangements with suppliers who meet set criteria, which avoids employee discretion of company purchases. Similarly, salaries may be determined rigidly

⁴In later work we change this element of the design and find no difference in choices from providing complete information to all players, so this turned out to be an innocuous part of the original design.

throughout an organisation to avoid small groups agreeing on salary increases they allocate to each other. For governments, policies that auction or sell at market values particular new rights, such as for mining, and property development, would reduce the value to recipients of any back-scratching favours given. Such systems, however, have the potential to simply elevate back-scratching to a higher level where greater discretion now resides.

3.2 Theoretical effects of treatments on optimal choices

Backward induction immediately implies that rational selfish agents in our experimental setup will always pick the most productive player in each round as the receiver, because this is optimal in the last round and thus also in each round before. To nevertheless account for the logic of alliances, we here demonstrate that even by looking at just two simple types of players from the universe of potential strategies, that alliances can greatly increase payoffs for their members, conditional upon the strategies of others.

Our two player types are: the meritocratic player type, who chooses the highest productivity person as receiver in each round that they are the allocator (regardless of whether the previous round was meritocratic or not); and the tit-for-tat reciprocator (TFT) whose strategy is as follows.

- If the previous round is not meritocratic, choose the previous round allocator.
- If the previous round is meritocratic and the two rounds prior allocator is me, choose the previous round allocator (continue an alliance).
- If the previous round is meritocratic and the two rounds prior allocator is not me (or it is the second round), choose productivity 2 player, or the productivity 1 player if the productivity 2 player is the previous allocator.
- If the first round, choose the productivity 2 player.

With these two types in mind, we can see how the two policy treatments change the incentives for being either a meritocrat or a tit-for-tat reciprocator (TFT), depending on the order of being able to make a choice in the game. The important situations in terms of a Nash-equilibrium are:

1. An individual chooses to be a meritocrat when all others choose to be a meritocrat. The first receiver's expected payoff is then 245 and the expected payoff of all other players is then 227.
2. An individual chooses to be a meritocrat when all others are TFT. Since choosing a type is only salient if a person gets to make a choice, we look at the individual being a meritocrat allocator in the first round. Their expected payoff is 51, and for the first receiver (who is TFT), the expected payoff is 375, as the receiver immediately switches to choosing someone else with whom a stable alliance is then formed.
3. An individual chooses to be a TFT when all others also choose this. The expected payoff for the first allocator is then 350, with 375 for the first receiver, and they immediately form an alliance lasting the whole treatment.
4. An individual chooses to be a TFT when all others are a meritocrat. The expected payoff to the first round allocator is then 219, while the first receiver (who is a meritocrat) can expect to get 245.

What this shows is that in the Baseline treatment, there are two Nash-equilibria in these player types, i.e. where everyone plays TFT or where everyone plays meritocratic. In the TFT equilib-

rium there is a winning alliance that gets formed in the first round and is sustained all 25 rounds, with the winning alliance members getting more than they would in the meritocratic equilibrium, meaning that the TFT equilibrium dominates the meritocratic equilibrium in expected payoffs for the winning alliance.⁵ This feature of multiple equilibria fits the general finding in group formation theories that study the stability of different possible coalitions (Konishi and Ray, 2003; Nash et al., 2012; Ray and Vohra, 2014).

Table 2 shows the expected payoffs for the Baseline and other treatments, for first round allocators, first round receivers, and other players. The Nash-equilibria are in bold font and show that the same equilibria hold in the Rotation and Low Rent treatment as in the Baseline, but that the two policy treatments the meritocratic equilibrium payoff dominates the TFT equilibrium, even for the winning coalition.

Table 2: Expected payoffs from TFT and meritocrat strategies in ECU

	Allocator Strategy	First receiver (others)			
		Meritocrat		TFT	
Baseline	Meritocrat	227,	245 (227)	51,	375 (201)
	TFT	219,	245 (214)	350,	375 (50)
Rotation	Meritocrat	227,	245 (227)	160,	245 (223)
	TFT	218,	240 (219)	214,	225 (194)
Low Rent	Meritocrat	93,	95 (93)	50,	90 (69)
	TFT	85,	88 (85)	86,	89 (50)

3.3 Implementation

The experiments were conducted at the University of Queensland in Brisbane, Australia, and in the University of New South Wales, Sydney, from May 2013 to February 2016. In all, 280 student subjects⁶ were recruited from the university student body using the ORSEE online recruitment system (Greiner, 2003), with experiments conducted in university computer labs using CORAL software (Schaffner, 2013). Each subject played 50 rounds in total and received their accumulated experimental currency earnings converted to Australian Dollars at a ratio of 20:1, earning between \$5 and \$36, with the average payoff \$20 for around a 1 hour experiment session, slightly above the minimum wage, with earnings inclusive of a show-up fee.

Each session began with a brief introductory talk (in Appendix). Subjects were provided printed instructions which they were directed to read prior to commencing the computerised part of the session, and blank paper and pens to keep notes. Prior to commencing both the Baseline and any treatment, subjects answered a series of hypotheticals to ensure they had a complete understanding of the experiment and the payoff structure. After their 50 rounds, subjects answered a socio-demographic survey and were paid in cash upon departure via sealed envelopes.

⁵It is easy to see that a simple reciprocity strategy whereby someone reciprocates to the last person who made them receiver, is not a Nash-equilibrium: if all others are simple reciprocators then being a meritocrat increases the payoff. Conditional on being involved in the first round, the meritocrat keeps allocating it to the highest productivity player each time they are allocator, and keeps being chosen as receiver in the subsequent round.

⁶Students from Brisbane (112) were in groups playing Baseline treatment first, followed by Low Rent or Rotation. Students from Sydney (168) played the other treatment sequences. Overall there was no significant difference in their tendency to play meritocratically, with 60% of rounds played meritocratically in both of Brisbane and Sydney (Mann-Whitney test of median values, 2,321 and p-value 0.89). The mean rounds of alliance play, a measure described in Section 3.4, was 26% for the Baseline treatments in Phase 1 in both Sydney and Brisbane.

The experiments involved 70 groups of 4 subjects, with Table 3 showing the sequences of the Baseline and policy treatments for different groups.

Table 3: Experiment setup

Treatment sequence		Groups	Subjects
First 25 rounds (Phase 1)	Second 25 rounds (Phase 2)		
Baseline	Baseline	14	56
Baseline	Low Rent	14	56
Low Rent	Baseline	14	56
Baseline	Rotation	14	56
Rotation	Baseline	14	56
Total		70	280

3.4 Research Questions

The main variable of interest is meritocratic play (MP). Meritocratic play is simply defined as a choice of the highest productivity player in that round. It is this choice that the policy treatments are intended to make more frequent. A number of other variables are also of interest. Alliance play (AP) we define as a sequence of plays where the same two players alternately choose each other to receive the bonus payment, starting with a non-meritocratic choice and involving at least one non-meritocratic choice by the other player. All rounds in such a sequence count as alliance plays, which hence can involve some meritocratic plays as well. We also look at the number of unique alliance player combinations and length of alliances as outcome variables. Together these measures complement the MP measure by allowing us to look at the whether the cause of low or high MP is a result of the establishment and maintenance of the same alliances, the breakdown of some alliances and the formation of new ones, or their complete breakdown. We also look at a number of other measures to help paint the picture of what happens in the experiments, such as the round of first alliance, average group and individual payoffs, and the group Gini coefficient. In addition, a variety of player socio-demographic factors are considered in terms of their potential relationship with alliance choices in the experiment. We structure our inquiry around the following research questions.

Question 1. Does the Rotation treatment increase the frequency of meritocratic choices compared to Baseline?

Question 2. Does a Low Rent treatment increase the frequency of meritocratic choices compared to Baseline?

Question 3. Do the Rotation or Low Rent treatments increase meritocratic choices more if they are implemented at the start (Phase 1) rather than following a Baseline?

Question 4. Do individual socio-demographic characteristics predict meritocratic play or alliance play?

4 Empirical analysis

4.1 Overview

In terms of average outcomes 60% of all rounds are meritocratic plays and 31% are alliance plays. Moreover, 86% of all groups showed at least one alliance play in their 50 rounds, meaning that the experimental design indeed gives rise to the back-scratching behaviour we wished to study. Before specifically addressing the main research questions, we first summarise and describe the behaviour observed in the 70 groups.

Table 4: Summary statistics

		Baseline			Rotation			Low Rent		
		All	Phase 1	Phase 2	All	Phase 1	Phase 2	All	Phase 1	Phase 2
Alliance behaviour	Share MP	0.57	0.59	0.56	0.62	0.61	0.61	0.64	0.73	0.56
	Share AP	0.38	0.26	0.51	0.11	0.12	0.11	0.31	0.13	0.48
	AP length	8.4	5.4	11.8	4.0	4.2	3.8	11.3	4.5	18.8
	1st AP	6.43	7.44	5.45	10.6	10.9	10.3	6.6	11.5	3.33
	Unique AP	1.06	1.1	1.02	0.75	0.71	0.79	0.64	0.64	0.64
Economic outcomes	Group pay.	864	866	863	870	870	870	324	336	312
	Ind. pay.	216	216	216	218	218	218	81	84	78
	Gini	0.19	0.15	0.21	0.15	0.15	0.15	0.06	0.04	0.08
	Net losses	61.4	59.3	62.1	54.7	54.6	54.9	51	38.6	63.4

Sparklines adjacent to variable group means are smoothed distributions. These help to show where bimodality has arisen due to learning of group norms.

AP length is the mean length of continuous AP in rounds, 1st AP is the round of first AP choice, and Unique AP is the mean number of unique alliance combinations in a group.

Yellow shading in Low Rent treatment indicates that the results cannot be directly compared to other treatments because of variation in the payoff structure.

In Table 4 are the mean outcomes for each treatment, which are also broken down according to the sequence, or phase, of the treatment occurring. The theoretical, and rational, maximum group payoff of 925 ECU over a treatment, which occurs from all players being fully meritocratic, occurred in only 5% of treatments. The mean group payoff in the Baseline treatment was just 864 ECU, meaning net losses were 61 ECU, or 6.6%. Using the reported alliance behaviour variables for all Baseline treatments, we can see that this arose because the average group had 38% of rounds in an alliance, with each group having one alliance pair, and this alliance lasting 8.4 rounds, first forming at round 6, on average.

A general pattern observable in Table 4 is that there are very minimal differences in the share of MP rounds in any particular treatment, save for Phase 1 Low Rent which had much higher meritocratic play than others. Additionally, the Phase 2 Baseline and Low Rent treatments have faster alliance formation, and longer alliances. The sparklines in Table 4 show a smoothed distribution of each outcome variable, making visible that, particularly in Phase 2 treatments, the distribution around the mean outcome is often bimodal. For example, the Low Rent Phase 2 treatment had 48% of rounds in an alliance. However, to arrive at this mean figure means overlooking at this came from seven groups with alliances for fewer than three rounds (five of those with no alliances), and the remaining seven groups having alliances lasting longer than 19 rounds (with three groups have alliances lasting the full 25 rounds).

We thus see in the data a tendency toward a creation of social norms within each groups, whereby the behaviour observed in early rounds predicts behaviour in the last rounds as subjects learn

the strategies of others. Across all 140 treatments played by the 70 groups the frequency of meritocratic (MP) choices in the first 10 rounds is highly predictive of the frequency in the last 10 rounds (For Phase 1 treatments Pearson correlation coefficient = 0.32 and $p = 0.01$, while for Phase 2 treatments Pearson correlation coefficient = 0.70 and $p = 0.00$), leading to bimodality in the distribution of many alliance variables. Figure 1 further demonstrates the learning process across groups in each treatment. Notably, the frequency of alliance (AP) decisions rises substantially over the 25 rounds, particularly in the Baseline condition (no shading), providing an indication that effects of policy treatments are likely to be conditional on the order in which they are played because of prior learning. For example, it can be seen in Figure 1 that the trend in MP and AP choices in the Low Rent treatment (darker shading) go in opposite directions depending on the order of this treatment. In such a setting, as in the type of real life settings we are trying to capture with this back-scratching game, deterring alliance formation, and breaking it down once formed, appear to be two distinct challenges.

Figure 1: Frequency of choice types across groups by round and treatment (no shading is Baseline)

To further explore the scope of alliance behaviours a sample of individual accumulated payoff charts capturing unexpected interactions in three groups is in Figure 2.⁷ These charts show individuals payoffs for player in a group, with the person being chosen as receiver in a round showing a jump in the payoff line equal to the bonus plus their productivity of that round, and alliances observable by alternating jumps in the payoffs of two players. In the left panel, blue and green players (dot-dashed and solid line respectively) form an initial alliance only for green to renege in round nine and immediately form an alliance with orange (dashed) for nine rounds, after which green again teams up with blue till round 25. In rounds 26-50, orange and red (dotted) eventually team up for most of the rounds.

A strange situation occurred in the group in the centre panel, where blue and orange (solid and dashed) formed an alliance, partially disrupted but regained, which lasted till round 20 when orange chose green (dot-dashed) as the receiver who promptly formed a new alliance with red (dotted). In the group in the right panel, blue and green (solid and dashed) form a loyal alliance for the whole Low Rent treatment, before blue allocates to orange (dot-dashed), promptly forming a new alliance that lasts six rounds. Orange reneges on this alliance, leaving green and red (dashed and dotted) to form a new alliance for the next sixteen rounds.

These examples demonstrate that allocating meritocratically outside of an alliance can be punished either by i) being replaced from an alliance with an outside player (as in the left and right

⁷Individual accumulated payoff charts for all 70 groups are in Figure 4 of the Appendix.

Figure 2: Individual payoffs by round with complex interactions

panels), or ii) giving power to an alliance of the two outsiders (as in the centre panel). Fully capturing this kind of play in a theoretical model is a challenge, as one might imagine, for the obvious reason that the players left outside an alliance almost certainly were surprised by the actions of the other players. Table 7 in the Appendix gives more examples of ‘strange’ strategies as noted by some of the experiment subjects.

One point of note is the first round of an alliance begins with a non-meritocratic choice, of which there are two choices. Either the player with productivity one is chosen, giving a lower payoff to the allocator, or the player with productivity two is chosen, giving the highest payoff to the allocator. In our experiment only 56% of first round AP choices were to the productivity two player. This conflicted with our expectations that it would be close to 100%. One reason could be that the longer no alliance persisted, the more costly the signal necessary to form an alliance. The mean round of first AP choice in all Phase 1 treatments was 8.6 (standard deviation = 6.33, min = 1, max = 24, and median = 8). Yet there is no correlation between the number of rounds taken and the first non-meritocratic choice that forms an alliance in Phase 1 (Pearson correlation coefficient = 0.09 and $p = 0.48$). The process of forming alliances therefore appears to be driven more by the choice of alliance partner than efficiency considerations.

4.2 Treatment effects

Table 5 summarises the main linear regressions of the policy treatments on the main outcome variable of the share of MP rounds in a treatment, in addition to the share of AP rounds, number of unique alliance combinations, and mean alliance length.

The first column for each dependent variable captures the raw average effects of the policy treatments, and the second column treats the order, or phase, of the each treatment is separately. We see that Rotation and Low Rent treatments have insignificant effects on meritocratic play when their average effects are taken into account. However, when the order of the treatment is considered, the Low Rent in Phase 1 shows a 14% increase in meritocratic choices, the only significant result in terms of our main outcome variable.

In terms of alliance play, the Rotation coefficients are difficult to interpret because alliances are

Table 5: Treatment effects

	MP share		AP share		Unique alliances		Alliance length	
Rotation	0.04		-0.27***		-0.31		-5.87***	
	(0.05)		(0.08)		(0.19)		(1.90)	
Rotation x 1st		0.02		-0.14		-0.38		-2.54
		(0.07)		(0.11)		(0.27)		(2.64)
Rotation x 2nd		0.07		-0.40***		-0.24		-9.19***
		(0.07)		(0.11)		(0.27)		(2.64)
Low Rent	0.07		-0.08		-0.42**		-1.12	
	(0.05)		(0.08)		(0.19)		(1.90)	
Low Rent x 1st		0.14**		-0.13		-0.45*		-3.28
		(0.07)		(0.11)		(0.27)		(2.64)
Low Rent x 2nd		0.00		-0.02		-0.38		1.03
		(0.07)		(0.11)		(0.27)		(2.64)
2nd Phase	-0.06	-0.04	0.22***	0.25***	-0.03	-0.07	5.6***	6.11***
	(0.04)	(0.05)	(0.06)	(0.08)	(0.15)	(0.19)	(1.47)	1.87
R ²	0.04	0.06	0.17	0.20	0.05	0.05	0.17	0.21
N	140	140	140	140	140	140	140	140

Linear models are fitted with standard errors in parentheses clustered by experiment group (of which there are 70 groups undertaking two treatments each), and resulting p-values represented by * < 0.10, ** < 0.05, *** < 0.01.

mechanically broken due to random selection.⁸ The direction and magnitude of the Low Rent coefficients in the AP model are consistent with (being a negative sign of) the corresponding MP models. The main outcome here is that the Phase 2 average had 25% more alliance play, again indicating learning behaviour.

While the number of unique alliances was expected to be higher in Rotation, as the randomness would allow the ‘outside’ alliance pair to form, this was not the case. Low Rent also appears to generate fewer alliances overall, with the effect coming mostly from the Phase 1 implementation. In terms of the average length of an alliance, again the Rotation treatment mechanically shortens this variable and is difficult to interpret, but the main effect of reducing alliance from Rotation appears to be coming when it is a Phase 2 treatment. Low Rent seems to have little effect in raw terms, or considering its treatment order separately.

As a robustness check, Table 8 in the Appendix presents these same regression results using only data from treatment rounds 11-23, to remove any the early round learning effects and later round ‘end-of-game’ effects. The overall direction and magnitude of the full-round results is supported. Additionally, the use of these select rounds allows for additional controls for early round play. The third column of Table 8 adds an additional control for the share of MP choices in the first ten rounds. The effect is to capture the variation that Low Rent Phase 1 term had captured, further supporting the notion that the Low Rent treatment’s effect comes from preventing the desire to form alliances in the first place, which is why it is mainly effective in Phase 1.

Overall, the effects of the two policy treatments in this experiment appear to be not particularly large or significant. The main result here is that the Low Rent treatment can increase the amount of meritocratic play, but only if implemented prior to any other treatment. The data suggest that

⁸A number of alternative measures to capture alliance play that continued in some form after rotation were investigated, but these required a long list of assumptions about the intended behaviour of participants, and would have still been difficult to interpret.

lack of effect of the Rotation treatment could be due to a number of factors. For example, even though it is more difficult to form an alliance in Rotation, our mild form of this treatment was perhaps not strict enough deter some participants from still attempting for form, and reform, alliances, which is a costly activity itself.

Table 6: Pair-wise treatment comparisons of MP and unique alliances.

		Baseline		Rotation		Low Rent	
		Phase 1	Phase 2	Phase 1	Phase 2	Phase 1	Phase 2
Base.	Phase 1	42	-0.07	0.38	0.31	0.45*	0.45*
	Phase 2	-0.06 (28,28)	42	0.25	0.24	0.38	0.38
Rot.	Phase 1	0.02 (14,42)	0.05 (14,28)	14	-0.07	0.07	0.07
	Phase 2	0.03 (14,42)	0.07 (14,28)	0.01 (14,14)	14	0.14	0.14
LR	Phase 1	0.14*** (14,42)	0.17** (14,28)	0.12* (14,14)	0.11* (14,14)	14	0.00
	Phase 2	-0.04 (14,42)	0.00 (14,42)	-0.06 (14,14)	-0.07 (14,14)	-0.17** (14,14)	14

Lower triangle (blue) is the difference in treatment mean of meritocratic share of rounds (row minus column).

Upper triangle (orange) is the difference in treatment mean of number of unique alliances per group (row minus column).

Diagonal is the number of total treatments of that type played.

Number of unique group-treatment observations are below in parentheses (row, column), and apply in reverse to the upper-triangle. Mann-Whitney test p-value reported for the comparison of group share of meritocratic decisions and unique alliance between treatments using between group-treatment data * < 0.10, ** < 0.05, *** < 0.01.

As a robustness check, Table 8 in Appendix B shows the results of running the same regression models using data from rounds 11 to 23 of each treatment group. This provides a comparison that does not depend on play in early rounds when subjects are still learning the set-up of the game, or on play in the last two rounds, when final round effects may be present. The results overall are consistent with those in Table 5 with the main significant result being the additional MP share in Low Rent Phase 1. We also run an additional model that controls for the effect of MP share in rounds 1 to 10 on MP share in later rounds, with the results in the third column. This control wipes out the treatment effects, supporting the view that the effect of Low Rent played in Phase 1 is primarily because it deters alliance formation.

To provide further clarity on relationships between treatment outcomes, Table 6 summarises the full suite of pair-wise combinations of the MP and unique alliance variables (which are less affected by the mechanical game changes in Rotation) for all treatment-phase combinations. Here too the only significant result in terms of meritocratic play is from Phase 1 Low Rent, with 14% more MP choices on average than the Phase 1 Baseline, but also having significantly more than Rotation as well. Compared to Baseline, Low Rent also appears to generate fewer unique alliance pairings when undertaken in Phase 1.

4.3 Socio-demographics and alliance behaviour

A final consideration is whether the types of choices made by individual players is related to their socio-demographic characteristics, with our expectation from the literature being that males and business students are more likely to form alliances. Table 9 in the Appendix shows the coefficient estimates from a series of regressions of MP or AP choices made each round, dependent on a number of characteristics of the player making the choice that round. There are 3,500 observations that arise from considering each choice in every round played by each group. Standard errors are clustered by experiment group to account for the non-independence of each observation within a group due to the interaction of individual choices and the game-play of others. These results are a wash. There are no clear and significant directional effects from any of the measured socio-demographic characteristics of the student participants.

4.4 Summary of main results

Answer 1. The mean difference in meritocratic play between all Rotation treatment groups and all Baseline treatment groups was only 0.04, which was small and insignificant, implying that a mild Rotation treatment does not in general increase meritocratic play in our experimental set-up.

Answer 2. The mean difference in meritocratic play between all Low Rent treatment groups and all Baseline treatment groups was 0.07. This effect is somewhat larger than that of Rotation but still insignificant ($p = 0.05$), implying that Low Rent does not in general increase meritocratic play in our set-up.

Answer 3. Low Rent did not significantly increase meritocratic play overall, but when implemented in Phase 1 it increased meritocratic play by 0.14 ($p = 0.07$). The effect on alliance play was also higher in Phase 1 compared to Phase 2, yet not significantly so. In contrast, Rotation significantly reduced alliance play when implemented in Phase 2 (by -0.40) but not in Phase 1, which is logical as Rotation mechanically reduces alliance play in the situation where there is already alliance play between a particular player pair who randomly lose the chance to sustain that alliance (which is much more prevalent in Phase 2). However, the lack of a significant increase in meritocratic play in Phase 2 Rotation indicates that the mechanical effect of breaking one particular alliance did not lead to increased meritocratic play and overall efficiency gains for the group. We interpret these results to say that Low Rent does help prevent the formation of alliances when implemented immediately rather than after the alliances have formed, whilst Rotation mechanically reduces alliance play if it has formed, but has no large effects on meritocratic play in any Phase and is hence not particularly successful at breaking up alliances once formed.

Answer 4. Our discrete choice analyses of the effect of individual characteristics on meritocratic or alliance play drew a complete blank in terms of significance. We hence cannot find any individual characteristic that significantly relates to meritocratic or alliance play, against our expectation that, for instance, men would be more likely to form alliances than women.

5 Discussion and conclusions

In this paper we introduced a new experimental design aimed at producing welfare-reducing reciprocity, or back-scratching, in the laboratory. In our baseline setup, 38% of rounds could be identified as part of an alliance, while 86% of all experiment groups had at least one alliance, indicating that the design closely represents situations where discretionary choices allow for

the emergence of costly alliances amongst partners, or small groups, who act in their joint interest. Our policy treatments mimicked two common types of anti-corruption procedures, namely staff rotation, and reducing the level of individual discretion to allocate economic rents. The Rotation policy, whereby the next round allocator is the previous receiver with only a 50% chance, reduced alliance rounds from 38% to 11%, yet the share of meritocratic decisions was relatively unchanged, showing that this effect is primarily due to Rotation breaking alliances rather than reduced intentions to form and reform alliances. This mild form of rotation is clearly not as strict seen in other setups in the literature, such as Abbink (2004), but we feel it is closer to the imperfect rotation policies available in many institutional environments, such as banking and financial regulation, which as we earlier mentioned are prone to favouritism due to their high stakes.

The Low Rent treatment did slightly (but not significantly) increase the frequency of meritocratic rounds and reduce alliance rounds, but this is due solely to the large and significant effect of 14% higher meritocratic play when played as the first treatment, probably due to subjects having no prior alliances or expectations of cooperative strategies as the payoff for alliance formation is lower than the meritocratic payoff.

Overall our results are not as strong as expected. The *prima facie* policy conclusions from these experiments are that existing alliances matter in terms of the success of institutional interventions aiming to curtail back-scratching. Ad hoc rotation of staff from a pool of people, who may even already have established alliances, is unlikely to deter them from taking advantage of their temporary position to favour their alliance partners. This points to a significant trade-off in the choice of personnel for regulatory agencies between greater industry experience and being free of previous alliances, with our motivating example of personal alliances in financial regulation demonstrating clearly this trade-off. Rotating staff from regulators abroad instead of the local regulated industry may help ensure a level of experience in new staff without bringing loyalties with them.

In addition, our lack of any observed relationship between socio-demographic factors and alliance choices indicates that policies that focus on selecting the ‘right’ people for a group or organisation will not be effective in terms of preventing or disrupting backscratching.

Having smaller economic rents able to allocated with discretion led to more meritocratic play amongst experiment groups who were yet to observe alliance-forming strategies of others, suggesting that when designing new regulatory institutions, reducing individual discretion at the start is a key way to prevent the formation of costly back-scratching alliances. We also hypothesise that perhaps low rents can reduce alliance behaviour where strong loyalties exist because weaker loyalties could partly persist due to the fear of losing the dominant position and subsequently being punished by a former alliance partner teaming up with a new partner.

Our design can be extended in many directions, such as increasing the transparency of choices and player identities, varying group size, introducing other allocation rules such as voting, and introducing the possibility of punishment.

References

- Abbink, K., 2004. Staff rotation as an anti-corruption policy: an experimental study. *European Journal of Political Economy* 20, 887–906.
- Abbink, K., Hennig-Schmidt, H., 2006. Neutral versus loaded instructions in a bribery experiment. *Experimental Economics* 9, 103–121.
- Abbink, K., Irlenbusch, B., Renner, E., 2002. An experimental bribery game. *Journal of Law, Economics, & Organization* 18, 428–454. doi:[10.2307/3555050](https://doi.org/10.2307/3555050).
- Abbink, K., Serra, D., 2012. Anticorruption policies: Lessons from the lab, in: Serra, D., Wantchekon, L. (Eds.), *New Advances in Experimental Research on Corruption*. Emerald Group Publishing Limited. Book Series: Research in Experimental Economics. chapter 4, pp. 77–115.
- Armantier, O., Boly, A., 2008. Can corruption be studied in the lab? comparing a field and a lab experiment. *CIRANO - Scientific Publications* 26.
- Barr, A., Lindelow, M., Serneels, P., 2009. Corruption in public service delivery: An experimental analysis. *Journal of Economic Behavior & Organization* 72, 225–239. doi:<http://dx.doi.org/10.1016/j.jebo.2009.07.006>.
- Barr, A., Serra, D., 2009. The effects of externalities and framing on bribery in a petty corruption experiment. *Experimental Economics* 12, 488–503.
- Barr, A., Serra, D., 2010. Corruption and culture: An experimental analysis. *Journal of Public Economics* 94, 862–869. doi:<http://dx.doi.org/10.1016/j.jpubeco.2010.07.006>.
- Barth, J.R., Caprio, G., Levine, R., 2012. *Guardians of Finance: Making regulators work for us*. MIT Press.
- Büchner, S., Freytag, A., González, L., Güth, W., 2008. Bribery and public procurement: an experimental study. *Public Choice* 137, 103–117.
- Cameron, L., Chaudhuri, A., Erkal, N., Gangadharan, L., 2009. Propensities to engage in and punish corrupt behavior: Experimental evidence from Australia, India, Indonesia and Singapore. *Journal of Public Economics* 93, 843–851. doi:<http://dx.doi.org/10.1016/j.jpubeco.2009.03.004>.
- Congleton, R.D., 2014. Rent seeking and organizational governance: Limiting losses from intra-organizational conflict. Available at SSRN 2444756 .
- Faccio, M., Masulis, R.W., McConnell, J.J., 2006. Political connections and corporate bailouts. *The Journal of Finance* 61, 2597–2635. doi:[10.1111/j.1540-6261.2006.01000.x](https://doi.org/10.1111/j.1540-6261.2006.01000.x).
- Frank, B., Schulze, G.G., 2000. Does economics make citizens corrupt? *Journal of Economic Behavior & Organization* 43, 101–113. doi:[http://dx.doi.org/10.1016/S0167-2681\(00\)00111-6](http://dx.doi.org/10.1016/S0167-2681(00)00111-6).
- Greiner, B., 2003. An online recruitment system for economic experiments, in: Kurt Kremer, V.M. (Ed.), *Forschung und wissenschaftliches Rechnen. GWDG Bericht 63*, Göttingen: Ges. für Wiss. Datenverarbeitung.
- Greiner, B., Schneider, P., 2015. Campaigns, coalition formation, and coalition stability in a repeated dictator election game. *UNSW Working Paper Series* .

- Jain, A.K., 2001. Corruption: A review. *Journal of Economic Surveys* 15, 71–121. doi:10.1111/1467-6419.00133.
- Konishi, H., Ray, D., 2003. Coalition formation as a dynamic process. *Journal of Economic Theory* 110, 1–41. doi:http://dx.doi.org/10.1016/S0022-0531(03)00004-8.
- Lambsdorff, J.G., 2002. Making corrupt deals: contracting in the shadow of the law. *Journal of Economic Behavior & Organization* 48, 221–241.
- Lambsdorff, J.G., Frank, B., 2010. Bribing versus gift-giving -an experiment. *Journal of Economic Psychology* 31, 347–357.
- Lambsdorff, J.G., Frank, B., 2011. Corrupt reciprocity - experimental evidence on a men’s game. *International Review of Law and Economics* 31, 116–125.
- Levine, R., 2012. The governance of financial regulation: Reform lessons from the recent crisis. *International Review of Finance* 12, 39–56. doi:10.1111/j.1468-2443.2011.01133.x.
- Matthews, D., 2014. Whistleblower’s tapes suggest the fed was protecting goldman sachs from the inside. *vox.com* .
- Nash, J.F., Nagel, R., Ockenfels, A., Selten, R., 2012. The agencies method for coalition formation in experimental games. *Proceedings of the National Academy of Sciences* 109, 20358–20363.
- Ohashi, H., 2009. Effects of transparency in procurement practices on government expenditure: A case study of municipal public works. *Review of Industrial Organization* 34, 267–285. doi:http://dx.doi.org/10.1007/s11151-009-9208-1.
- Pezzullo, M., 2013. Statement on integrity issues. Technical Report. Senate Legal and Constitutional Affairs Legislation Committee Consideration of Additional Estimate for 2012-13.
- Protest, B., Eavis, P., 2015. Ex-goldman banker and fed employee will plead guilty in document leak.
- Rand, D.G., Arbesmand, S., Christakis, N.A., 2011. Dynamic social networks promote cooperation in experiments with humans. *Proceedings of the National Academy of Sciences* 108, 19193–19198.
- Ray, D., Vohra, R., 2014. Coalition formation, in: Young, P., Zamir, S. (Eds.), *Handbook of Game Theory*. North-Holland. volume 4.
- Roubini, N., Mihm, S., 2010. *Crisis economics: A crash course in the future of finance*. Penguin.
- Schaffner, M., 2013. Programming for Experimental Economics: Introducing CORAL - a lightweight framework for experimental economic experiments. *QuBE Working Papers* 16. Queensland University of Technology.
- Schikora, J.T., 2011a. Bringing good and bad Whistle-blowers to the Lab. Munich Discussion Paper 2011-4. Department of Economics University of Munich.
- Schikora, J.T., 2011b. Bringing the Four-Eyes-Principle to the Lab. Munich Discussion Paper 2011-3. Department of Economics University of Munich.
- Schulze, G.G., Frank, B., 2003. Deterrence versus intrinsic motivation: Experimental evidence on the determinants of corruptibility. *Economics of Governance* 4, 143–160.
- Serra, D., 2012. Combining top-down and bottom-up accountability: evidence from a bribery experiment. *Journal of Law, Economics, and Organization* 28, 569–587.

- Silver-Greenberg, J., Protesse, B., Eavis, P., 2014. New scrutiny of Goldman's ties to the New York Fed after a leak. New York Times November 19.
- van Veldhuizen, R., 2011. Bribery and the Fair Salary Hypothesis in the Lab. Working Paper. Tinbergen Institute.

Instructions

These instructions were read aloud in each session before subjects began the experiment.

Welcome to this experiment. Thank you all for participating.

Please quietly read the information sheet in front of you, and complete your name, date, and sign the Consent Form.

You will spend the next hour or so playing a computer game with other people in this room. At the end of the experimental game the computer will prompt you to complete a socio-economic survey. I would appreciate honest answers. Be assured that the answers will not be linked to your name in any way. Once you leave this room there is no way to trace you to the participant number assigned to your computer terminal.

In the experiment you will earn real money. Cash Australian dollars. At the end, your earning in the experimental game will be paid to you in cash at a conversion rate of 20:1. For example if you earn 500 experimental dollars, you will earn \$25, which I will bring in an envelope to your seat once complete. This amount includes your \$5 show-up fee, which is the minimum earnings. You are hence you are playing to increase your earnings above this amount.

If at any point in time you have questions about the rules of the experimental game, please first go back and reread the instruction sections on the computer. If you still are uncertain, please raise your hand and I will come and assist you.

If you encounter any problems or error messages on the computer during the session, please raise your hand.

I will remind you that you are playing a game with some of the other people in this room. Decisions made in the game require other players to be ready. Please do not wait to click continue on any screen in the game. Simply do so after you have understood the information on the screen and followed any instructions.

There are no right or wrong choices in this game. You are in an interactive game with other players. Think strategically and make decisions that you think could increase your earnings over the whole game.

Finally, the sheet of paper and pencil in front of you is to make notes should you desire during the game. Written on the paper is your participant number. Please now check that this matches the number entered in the text field on your screens. If it does not match, please correct it on your computer now.

When the computer prompts you that your experiment is over please write the cash amount shown on your screen, your name, and sign the form in front of you. Then simply wait quietly in your seat for me to arrive with your payment.

If you would like to hear about the research findings that arise from this experiment you can leave your email address on the table by the printer on your way out.

Are there any questions before we begin.

Figure 3: Clockwise from top left: Allocator decision screen, non-allocator (observer) wait screen, notification of decision for receiver of the payment, notification of decision for non-receivers

Figure 4: Accumulated player payoffs by group

Table 7: Examples of written notes by subjects

<i>My friend and I are greedy except when they are 1 and others are 3.</i>
<i>Not going to get it back until I get a 3 again</i>
<i>Hope the player information is correct with who trusts me to pay them so they repay me.</i>
<i>Never use productivity number 1 even if it's a mate. Always go for highest productivity.</i>
<i>Check to see if he repays me, try to get repayment continuous b/w the two. Same tactic for round 2</i>
<i>Seems best to continuously allocate to the player with the highest productivity number</i>
<i>Round 20 didn't get from hexagon. Then they didn't give to me when I was 3 despite me giving it to them when they were 2 and others were 3. But I took it they were 1.</i>
<i>It's always good to have support. If they support you, you should support them</i>
<i>Choose red circle every time and hopefully they realise they are best off choosing me.</i>
<i>For one note that this game would be played very differently if people were face to face.</i>
<i>Secondly, I am pretty sure a cultural bias exists, where Asians (the majority) would exhibit more group consciousness. My guess is however that you are testing the relationship between leadership and group consciousness? (altruism / egotism)</i>
<i>Strategy: pick one player and always allocate to that person. What if I get isolated?</i>
<i>Isolate the other two?</i>

Table 8: Treatment effects using rounds 11-23 only

	MP share		AP share		Unique alliances		Alliance length	
Rotation	0.02		-0.30***		-0.20		-3.60***	
	(0.06)		(0.09)		(0.14)		(1.19)	
Rotation x 1st	-0.01	-0.02		-0.13		-0.10		-1.37
	(0.08)	(0.06)		(0.13)		(0.20)		(1.66)
Rotation x 2nd	0.06	0.03		-0.47***		-0.31		-5.83***
	(0.08)	(0.06)		(0.13)		(0.20)		(1.66)
Low Rent	0.06		-0.10		-0.27*		-1.12	
	(0.06)		(0.09)		(0.14)		(1.19)	
Low Rent x 1st	0.13*	0.05		-0.16		-0.31		-1.84
	(0.08)	(0.06)		(0.13)		(0.20)		(1.66)
Low Rent x 2nd	-0.02	-0.03		-0.05		-0.24		-0.40
	(0.08)	(0.06)		(0.13)		(0.20)		(1.66)
2nd Phase	-0.05	-0.03	-0.01	0.21***	0.25***	0.04	0.07	2.89***
	(0.04)	(0.06)	(0.05)	(0.07)	(0.09)	(0.11)	(0.14)	(0.92)
Share MP (rounds 1-10)			0.79***					3.50***
			(0.11)					(1.17)
R ²	0.02	0.04	0.35	0.14	0.17	0.04	0.05	0.14
N	140	140	140	140	140	140	140	140

Linear models are fitted with standard errors in parentheses clustered by experiment group (of which there are 70 groups undertaking two treatments each), and resulting p-values represented by * < 0.10, ** < 0.05, *** < 0.01.

Table 9: Individual socio-demographics and choices

	Socio-Dem.		S-D & treats.	
	MP	AP	MP	AP
Age	0.00 (0.01)	0.00 (0.01)	0.00 (0.01)	0.00 (0.01)
Gender	0.03 (0.04)	0.01 (0.05)	0.00 (0.04)	-0.00 (0.04)
Marital	-0.01 (0.04)	0.04 (0.05)	-0.01 (0.04)	0.03 (0.05)
Inter. student	-0.05 (0.04)	-0.06 (0.05)	-0.05 (0.04)	-0.05 (0.05)
Bus. student	-0.06 (0.07)	0.01 (0.08)	-0.05 (0.07)	0.00 (0.07)
Happiness	-0.01 (0.02)	0.04 (0.03)	-0.01 (0.02)	0.03 (0.03)
Family wealth	0.01 (0.03)	0.03 (0.04)	-0.01 (0.03)	0.02 (0.04)
Politics	0.01 (0.01)	-0.01 (0.01)	0.01 (0.01)	-0.01 (0.01)
Skills vs people	-0.00 (0.02)	-0.02 (0.02)	-0.00 (0.02)	-0.02 (0.02)
Clubs	-0.01 (0.04)	-0.05 (0.05)	-0.02 (0.04)	-0.04 (0.05)
Leadership	-0.00 (0.04)	0.06 (0.05)	0.00 (0.04)	0.05 (0.05)
Rotation			0.03 (0.08)	-0.14 (0.10)
Low rent			0.14* (0.09)	-0.13 (0.10)
2nd treat			-0.03 (0.07)	0.24*** (0.08)
Rotation x 2nd treat			0.02 (0.16)	-0.24 (0.18)
Low rent x 2nd treat			-0.16 (0.16)	0.09 (0.18)
Sydney control	-0.01 (0.04)	0.07 (0.05)	0.01 (0.08)	-0.01 (0.09)
R ²	0.01	0.03	0.02	0.13
N	3,500	3,500	3,500	3,500
# MP or AP choices	2,090	1,095	2,090	1,095

Linear models are fitted with standard errors clustered by experiment group (of which there are 70), and resulting p-values represented by * < 0.10, ** < 0.05, *** < 0.01.

Gender = 1 for male, 0 for female. Skills vs people is 1-5 where “success in life is mostly due to...” skills is 5, and people is 1. Politics is 1-10 where 1 is Left and 10 is Right. Clubs is a [1,0] dummy of student club memberships, and Leadership is a [1,0] dummy for organisational leadership positions. Happiness is a 1-5 scale with 5 as overall very happy. Marital is a [1,0] dummy of romantic relationship or marriage. Family wealth is 1-3 with 3 being wealthier than peers.