

Murray, Cameron K.; Frijters, Paul; Vorster, Melissa

Working Paper

Give and You Shall Receive: The Emergence of Welfare-Reducing Reciprocity

IZA Discussion Papers, No. 9010

Provided in Cooperation with:

IZA – Institute of Labor Economics

Suggested Citation: Murray, Cameron K.; Frijters, Paul; Vorster, Melissa (2015) : Give and You Shall Receive: The Emergence of Welfare-Reducing Reciprocity, IZA Discussion Papers, No. 9010, Institute for the Study of Labor (IZA), Bonn

This Version is available at:

<https://hdl.handle.net/10419/110755>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

IZA DP No. 9010

**Give and You Shall Receive:
The Emergence of Welfare-Reducing Reciprocity**

Cameron K. Murray
Paul Frijters
Melissa Vorster

April 2015

Give and You Shall Receive: The Emergence of Welfare-Reducing Reciprocity

Cameron K. Murray

University of Queensland

Paul Frijters

*University of Queensland
and IZA*

Melissa Vorster

University of Queensland

Discussion Paper No. 9010

April 2015

IZA

P.O. Box 7240
53072 Bonn
Germany

Phone: +49-228-3894-0
Fax: +49-228-3894-180
E-mail: iza@iza.org

Any opinions expressed here are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but the institute itself takes no institutional policy positions. The IZA research network is committed to the IZA Guiding Principles of Research Integrity.

The Institute for the Study of Labor (IZA) in Bonn is a local and virtual international research center and a place of communication between science, politics and business. IZA is an independent nonprofit organization supported by Deutsche Post Foundation. The center is associated with the University of Bonn and offers a stimulating research environment through its international network, workshops and conferences, data service, project support, research visits and doctoral program. IZA engages in (i) original and internationally competitive research in all fields of labor economics, (ii) development of policy concepts, and (iii) dissemination of research results and concepts to the interested public.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ABSTRACT

Give and You Shall Receive: The Emergence of Welfare-Reducing Reciprocity

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 this reciprocal alliance behaviour. In each of the 25 rounds of our experiments, a player (the 'allocator') nominates one of three others 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. Males and business students are found to be more likely to form welfare reducing alliances. Random allocator rotation policies and low bonuses fail to significantly improve overall welfare: rotation policies significantly reduce the rate of formation of new alliances but do not lead to the breakdown of existing alliances, while low bonus policies are only found to be effective when alliances are well established. This points to the importance of the strength of existing alliances for the chances of institutional interventions curtailing welfare reducing reciprocity, i.e. 'back-scratching'.

JEL Classification: D71, D72, D63

Keywords: alliance formation, corruption, reciprocity, experiment

Corresponding author:

Cameron K. Murray
University of Queensland
Colin Clark Building, 39
St Lucia Brisbane Qld 4072
Australia
E-mail: ckmurray@gmail.com

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, is arguably a key element in the crisis (Roubini & Mihm, 2010; Levine, 2012; Matthews, 2014; Barth *et al.*, 2012).¹ For instance, Rohit Bansal and his superior Joe Jampietro, Goldman Sachs investment bankers who had previously worked at the NY Fed and Federal Deposit Insurance Corporation (FDIC), were alleged to have obtained market sensitive information from another former NY Fed employee, Jason Gross, that undermined regulatory oversight (Silver-Greenberg *et al.*, 2014).

Like most instances of favouritism, Bansal and Jampietro’s case does not cleanly fit the definition of corruption by Jain (2001) as “acts which utilise the power of public office for personal gain in a manner that contravenes stipulated rules” due to substantial grey areas in the interpretation of confidentiality rules. Yet it is a clear example of where a degree of discretion can give rise to *quid pro quo* alliances formed on bases other than merit. Such alliances can be payoff-increasing for the insiders, but come with efficiency losses that are hard to measure. Faccio *et al.* (2006) for instance, 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 the actual pay-backs to politicians. Mancur Olson (1982) nevertheless surmises that alliances can impose external costs as they redistribute gains towards themselves that can “exceed the amount redistributed by a huge multiple”, which raises the question of what institutions can reduce this kind of welfare-reducing reciprocity, or ‘back-scratching’.

Because costly alliances are hard to observe in real markets, we develop a new experimental design that captures their two key elements. First, alliance partners have to actively find each other and develop the trust necessary to sustain an alliance, and second, there is a conflict in payoffs between direct reciprocity within an alliance, and indirect reciprocity amongst a larger group. This design adds to existing experimental approaches to corruption and costly reciprocity.²

An influential prior design is the Repeated Bribery Game (RBG), pioneered by Abbink *et al.* (2002). The basic stage game of this experiment involves two players in the role of briber and 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 & Frank, 2010; Schikora, 2011b; van Veldhuizen, 2011).

A wide variety of interventions have been tried 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 &

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.

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

Frank, 2010), an outside informed monitor (Schikora, 2011b), payoff variations that mimic high non-bribery wages (van Veldhuizen, 2011; Armantier & Bolky, 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 (where the experimenter determines the possible alliance) does not fit situations where unobserved alliances already exist before policy interventions take place, nor where the composition of the alliances is completely indeterminate *ex ante*, i.e. where potential partners have to find each other rather than be paired by the experimenter.

In our new experimental design, in each of the 25 rounds of our baseline experiment, a player (the ‘allocator’) nominates one of three others as a co-worker (the ‘receiver’), which determines the group production that period to be the productivity of the receiver, where the productivity of all potential receivers is randomly drawn each round. Yet, the receiver also gets a bonus and becomes the allocator in the next round, allowing the receiver to return the favour independent of the productivity of their partner. Alliances thus form if two individuals keep choosing each other as receiver even when their productivities are lower than that of others in some rounds, creating a cost to every other individual in the 4-person experiment. 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 of alliance formation and the effect of policy intervention on both the emergence of new alliances and the break-up of existing ones.

Our first treatment is Rotation, which mimics a ‘staff rotation’ policy by introducing a degree of randomness in who makes the discretionary decision each round. This intervention is implemented in an additional 25 rounds with the same players, hence with possible existing alliances in place. Whilst such a policy reduces the payoffs of alliances and hence should reduce the rate of formation of new alliances, it is not clear whether they will help break up existing alliances.

In our second treatment, the bonus is reduced, again for an additional 25 rounds after the baseline experiment. The reduction in the bonus is sufficient to make a meritocratic strategy, wherein every allocator chooses the highest productivity player each round as receiver, more profitable than being in an alliance, as long as everyone else also plays meritocratic. We interpret this treatment as reducing the level of discretion, which in real organisations might take the form of having strongly enforced rules, or as Congleton (2014) explains, the bureaucratisation of individuals into “cogs in an organizational machine”. 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, 2002b), and hence we call this a Low Rent treatment.

In a separate set of treatments, we prime alliances by giving players stories about fictitious player friendships to read before the experiments. This priming is meant to increase expectations of reciprocity and thus allows us to see whether the effectiveness of institutions changes with the strength of the existing alliances.

As a prelude to our results, males and business students are found to be more likely to make choices leading to back-scratching alliances. In the baseline treatment 26% of rounds are identifiable as alliance-play, and we find that primed relationships do indeed more often form alliances, with 70% of rounds identifiable as alliance-plays. Primed alliances also change the anti-corruption treatment effects: Rotation is effective where players are anonymous and alliances emerged in the previous 25 rounds, but not in the ‘primed alliance’ set of treatments where subjects remain loyal to their pre-existing relationships despite rotation. Our Low Rent treatment however, is

somewhat effective in the primed condition, but not when anonymous, perhaps because in the primed experiments there is a general expectation of individuals to favour their pre-assigned partners, which means that there is less expectation of revenge if individuals deviate from their alliance when the payoffs make them unproductive (i.e., the partners ‘understand’ it is actually better for both not to play alliance).

Section 2 discusses the literature our paper adds to. Section 3 introduces the experiments, after which Section 4 presents and analyses the empirical results. In Section 5 we develop a simple theoretical framework to see whether the choices of the participants in the experiments can be interpreted as Nash-equilibria in player-types, where players have to choose before the game what type of player they will be and hence how they will react to the choices of others (e.g. strictly meritocratic or tit-for-tat reciprocal). We present simulations that suggest a reasonable fit between the experimental choices and this theoretical framework. Section 6 concludes.

2 Background

The paper falls within two related strands of experimental literature: the literature on corruption experiments that use the repeated bribery game (RBG), and the wider literature on coordination in the presence of negative externalities from which we discuss the experiments closest to ours.

As mentioned earlier, we mainly tie into the RBG experiments. Abbink & 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.³ 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 & 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, the higher the externality, the less bribes are 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 & 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, which suggests that in our experiments too we might expect men and economics students to form more stable alliances, even though in our design they have to find each other without knowing the other’s gender and other characteristics.

³Participants in India were on average 15 percentage points more likely to bribe than the participants in Australia, 9.9 percentage points more likely to bribe than participants in Indonesia, and 5.4 percentage points more likely than participants in Singapore.

A very different design is that by Greiner & 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.

The general literature on how coordination on an efficient outcome arises is vast, much of it using the public good games setup in which individuals have to contribute to a common project, with the efficient outcome being that everyone contributes maximally, though the optimising selfish strategy is to contribute nothing whilst hoping that others do contribute. Pro-social behaviour is then an individual decision in a one-shot game, though individuals can learn from others in repeated games. In a survey of the literature, Chaudhuri (2011) finds that individuals are more likely to coordinate on high contributions when their contributions are observed; or when there is the possibility of punishing non-contributors or rewarding high-contributors.

It is not clear whether the findings on coordination in public goods games should be expected to hold for our design. For instance, in our design punishment is also possible in the sense that allocators can decide who will become the receiver and who will not, but this punishment has ambiguous effects: just as much as other players can punish someone for not having played meritocratically in previous rounds, snubbing an alliance attempt might mean that the next time the snubbed person is the allocator he or she will attempt to form an alliance with someone else. This possibility of ‘forming an alliance with someone else’ is not a part of the public good game since the pro-social behaviour is towards the whole group, not a sub-set of it.

Regarding the role of information in public goods games, Choi *et al.* (2011) show that when a player can see that others have contributed in previous rounds, they are likely to contribute more themselves, pointing to the role of transparency in achieving efficient outcomes. In our setup too, the history of players’ decisions is known (yet not actively displayed), but only individuals who become allocators get to make decisions, meaning that those who are not chosen as allocators may never get a chance to ‘punish’ others who have not played pro-socially, making it unclear whether this public goods result would hold in our setting too.

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.

The literature on who forms alliances with others, i.e. how in-groups and out-groups form and maintain, often follows the work of Tajfel (1970), who found that arbitrary allocation into groups

by the experimenter induced a striking degree of bias towards fellow group members, for instance on no other basis than that individuals are handed the same colour shirt or the same group label ('A or B'). Experiments of this type require priming of subjects with group identities, a process that has been repeatedly found to effectively invoke pro-sociality within groups or pairs (see Cameron *et al.* (2012) for a review).

Typically the priming occurs by arbitrary nomination into named identifiable groups by the experimenter (Tajfel, 1970; Chen & Li, 2009), using writing tasks (Yopyk & Prentice, 2005; McLeish & Oxoby, 2011), sharing information such as surnames or pictures of other players (Bohnet & Frey, 1999; Burnham, 2003; Charness & Gneezy, 2008; Andreoni & Petrie, 2004), simple messages or stories (Takács *et al.* , 2014), or in some cases using mutual gaze, touch or tapping rhythms in synchrony (Kurzban, 2001). We adopt the story-telling method in our groups of primed player pairs.

Summarising, the main innovation in our design is that we have an easy-to-implement dynamic setup in which alliances that have endogenous externalities arise even without priming, 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.

The external validity of experimental studies of questionable aspects of behaviour is a prime concern. On this point, Barr & 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 captures an aspect of real-world behaviour. Additionally, Armantier & Bolky (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.

Finally, approaching our design from a theoretical perspective is a challenge, with group formation theories generally characterised by multiple equilibria (Konishi & Ray, 2003; Nash *et al.* , 2012; Ray & Vohra, 2014). In simple terms, if it is beneficial to form a sub-group within a population in order to capture rents, then it will usually be beneficial to form a sub-group within the initial sub-group for the same reason. Backward induction immediately implies that rational selfish agents in our setup would always choose to 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 rationalise the finding that alliances do form, we appeal to strategic games wherein individuals choose to be a certain type of player before any of the 25 rounds, allowing us to see if, for instance, being a pure meritocrat would be a rational choice when all others are reciprocators. By picking a set of possible types, we use simulation analysis to explore the expected earnings from the interaction of strategy choices, the existence of Nash-equilibria in strategies and a trial-and-error mechanism for updating strategies.

3 Experimental Design and Research Questions

Subjects play in groups of four whose composition is fixed for the duration of the experiment, and are identifiable to each other by a unique coloured shape⁴. 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.

The payoff structure creates a conflict between maximising group and individual payoffs via back-scratching, with a ‘productivity number’ device that determines the group payoff in a round. 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 their decision. Each of the four players in the group receives a payoff in each round equal to the receiver’s productivity number. The receiver’s payoff that round includes the 25 ECU payment in addition. We call a choice meritocratic if the receiver has the highest productivity number. Table 1 lists player payoffs for the three possible choices available in a round in the situation that the ‘Circle’ is the allocator, the Triangle happens to have productivity 1 that round, Square has 2, and Pentagon has 3.

Table 1: Payoff structure for a stage game

	Circle	Triangle	Square	Pentagon
<i>Choice</i>	Allocator	Prod. 1	Prod. 2	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, which 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 individual payment to the allocating player from making a non-meritocratic choice. In all there are six potential two-player alliance combinations possible.

We limit the information provided to players each round in order for alliances to emerge under a *veil of ignorance*. Players not allocating in a round 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 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

⁴One might think it is better to be one shape than another, but no such pattern was found.

normally hidden nature of any individual favourable decisions.⁵ Images of decision, waiting and notification player screens are in Figure A.1 of the Appendix.

3.1 Treatments

Our two policy treatments mimic commonly prescribed anti-corruption measures, and we specifically test the efficacy of these treatments in breaking down alliances while deterring the formation of new alliances. As such we use a within-subject, or group, design in which the receiver in the final (25th) round of the baseline treatment becomes the allocator in the first round of a second policy 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. This treatment mimics policies of staff rotation in senior positions. To be clear, the person nominated by the allocator as receiver still gets the bonus (and can hence deduce that he was chosen by the allocator), but that receiver has a 50% chance of not being the allocator in the next round.

3.1.2 Low Rent

This treatment mimics policies designed to reduce the size of rents that can be allocated with discretion, and is achieved by reducing the bonus payment allocated each round from 25 ECU to 3 ECU.

3.2 Theoretical Effects of treatments on optimal choices

Selfish rational players should choose meritocratically. To nevertheless account for the logic of alliances, we here restrict ourselves to looking at two different types of players, neither of which is fully rational: the first is the meritocratic player type who chooses the highest productivity person as receiver in each round that they are the allocator. The second type is a tit-for-tat reciprocator who conditionally chooses the player who made them receiver in the previous round as long as that allocator displayed a willingness to reciprocate. A tit-for-tat reciprocator 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.

⁵In 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.

With these two types in mind, we can see how the 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, whilst the expected payoff of all other players is then 227.
2. An individual chooses to be meritocratic 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 meritocratic 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 game, there are two Nash-equilibria in these player types, i.e. where everyone plays TFT or where everyone plays meritocratic. In the TFT equilibrium, 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.⁶

Table 2 shows the expected payoffs for the baseline and the treatment scenarios, 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 in both these 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)

⁶It 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, leading to a payoff around 12 ECU higher than being a simple reciprocator too.

3.3 Anonymous or Primed

To mimic the situation where agents within institutions have prior relationships, we prime player pairs with bilateral expectations of friendship in a selection of experiment groups. Before the game, they had to read a brief recount of their imaginary ‘in-game’ relationships with a connected partner, effectively being told that one of the symbols was their friend that helped them in the past. The full description of this priming device is in the Appendix. A sample of hand-written participant notes in Table A.1 of the Appendix reveals that our priming method successfully manipulated some subjects’ expectations and behaviour towards reciprocity within an alliance, such as the remark by one primed person who wrote that “Hope the player information is correct with who trusts me to pay them so they repay me.” We expect these experiments to produce faster alliance formation, longer lasting alliances, and less responsiveness to anti-corruption policies than our anonymous groups in which all expectations of relationships are an emergent characteristic of game play.

3.4 Data

The experiments were conducted at the University of Queensland in Brisbane, Australia, in May and June 2013. In all 160 student subjects were recruited from the university student body using the ORSEE online recruitment system (Greiner, 2003), with experiments conducted in a university computer lab using CORAL software (Schaffner, 2013). The full experimental program, data and code is available at the research website⁷. 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 \$34, with the average payoff \$21 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 and written instructions provided to each subject for later reference. Subjects playing in primed groups received a single additional printed sheet with their priming device, which they were directed to read and consider prior to commencing the computerised part of the session. 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 payoff structure. At the completion of two sequential 25 round treatments, subjects answered a socio-demographic survey and were paid in cash upon departure via sealed envelopes.

Altogether 28 anonymous groups played the baseline, with 14 each continuing to either Rotation or Low Rent treatments. An additional 12 primed groups played in a second set of experiments, with each policy treatment implemented on 6 groups. A diagrammatic summary of the experimental setup is in Figure 1. Subjects were also given blank paper and pens in order to make notes during the experiment, which were collected upon completion in order to gain additional insight into the thought processes occurring during the game.

3.5 Research Questions

We define two choice variables of interest that we later use as measures of policy success and that are useful in terms of formulating the research questions: Alliance Initiation (AI) and Alliance

⁷<http://buddyexperiment.blogspot.com.au>

		Treatment sequence		Groups	Subjects
		25 rounds	25 rounds		
Anonymous	Baseline	Rotation		14	56
	Baseline	Low Rent		14	56
Primed	Baseline	Rotation		6	24
	Baseline	Low Rent		6	24

25ECU bonus per round
Productivity numbers from set {1,2,3}

Figure 1: Experiment setup

Reciprocation (AR). Alliance initiation is any non-meritocratic decision (allocation to a player with a productivity number less than 3) outside of a current alliance, and hence can be seen as an attempt at alliance formation. An alliance reciprocation is any non-meritocratic choice during an alliance period involving the same duo as the alliance initiation. An alliance is then a unbroken chain of choices involving an alliance initiation and at least one Alliance Reciprocation, ending when a different receiver is chosen than the duo in the alliance. These two alliance behaviour variables are roughly inverse to the frequency of meritocratic decisions, but they capture slightly different tendencies; AI captures active effort to form alliances, while AR captures willingness to sustain an alliance. We structure our enquiry around the following questions.

Question 1. Does implementing a rotation policy increase the frequency of meritocratic choices and total group payoffs?

This question is motivated by the success of Abbink’s (2004) experiment, and by the wide adoption of such policies in sensitive areas of large bureaucracies (World Bank, 1992). The intuition behind this hypothesis arises naturally from the expected payoffs summarised in Table 2, which shows that it is no longer a payoff-dominant strategy to play TFT, even for the alliance members that become receiver in the first round. If some individuals ‘re-set’ their strategy before the treatments (or after a few rounds) towards the payoff dominant equilibrium, then we should see more meritocratic play. Yet, since TFT play is still a Nash-equilibrium, strong expectations of alliance play on the part of others might prevent alliance members from switching strategies.

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

Like our first question, Question 2 arises due to the expected optimal strategy choices in response to the new payoff structure, making the meritocratic choice the payoff dominant Nash-equilibrium. Mainstream economic theories of rent-seeking also suggest that Low Rent environments are welfare enhancing due to lower rent-seeking activity (Lambsdorff, 2002a).

Question 3. Do primed groups have i) earlier Alliance Initiation, ii) more frequent Alliance Reciprocation, iii) longer lasting alliances, and iv) less responsiveness to anti-corruption policies?

Question 3 is motivated by the increase in certainty players will have about the strategy of the player they are now connected with via a primed relationship. Previous experimental evidence showed that priming relationships does increase cooperation within those relationships, and hence, we are likely to observe greater cooperating in alliance pairs than between the experimental group as a whole (Cameron *et al.*, 2012). Our expectation that our policy treatments will be less effective is informed by the general findings of cooperative games summarised in Chaudhuri’s (2011) review, whereby norms formed early in a repeated game are often persistent. Hence, our rotation policy may simply allow the continuation of two bilateral alliances based on the primed relationships, while our low discretion treatment may be insufficient to break the norm of alliance formation for fear of losing the outranking winning position of the alliance.

Question 4. Can socio-demographic characteristics predict individual alliance choices, AI and AR?

Question 4 is informed by previous studies revealing biases associated with particular demographic traits, such as gender and profession (Frank & Schulze, 2000; Ones & Putterman, 2007; Lamsdorff & Frank, 2011; Chaudhuri, 2011), and offers some insight into what sort of personal qualities to look at when trying to select new members of a group that are not corruption-prone.

4 Empirical analysis

4.1 Summary statistics

Table 3 provides basic descriptive statistics of the experiment outcomes and player characteristics. In all, we have reasonable amounts of variation in gender, political views, wealth, and between local and international students. In the baseline treatment, 74% of groups saw an alliance formed at some point⁸, and of those, with the first alliance formed on average in round 8, and with 29% of all baseline rounds spent in an alliance. We see significant variation in many indicators including individuals earning (ranging from 53 to 384 ECU), alliance lengths (ranging from 2 to 25 rounds), and share of rounds in an alliance (ranging from 0 to 1). Together these basic metrics indicate that back-scratching alliances within larger groups can easily form in experimental settings of repeated games featuring discretionary decisions.

We can note that primed groups had fewer AI choices (0.14 compared to 0.28 share of rounds) and more AR choices (0.38 compared to 0.17 share of rounds), indicating that alliances were indeed both more rapidly formed and sustained for longer on average: when alliances form and are stable there are fewer rounds to initiate an alliance.

To examine Olson’s (1982) claim that groups will steal from the economic pie even if the cost to others is greater than their gain, we calculate a measure of *gains minus losses* for a group as the difference between the maximum group payoff under meritocratic play and the realised group payoff. Group losses are 59ECU on average (74ECU for primed groups), or 6.4% (8.0%) of the total expected group payoffs under the meritocratic outcome. If we take out the effect of bonuses on payoffs and only focus on the productivity payoffs, then the losses become 20% (25%) of total expected group payoffs from production.

⁸We define an alliance as a period of exclusive dealing between two players where in the first round, and at least one other, the allocation choice is non-meritocratic (meaning the payment receiver had a productivity number in that round of less than 3).

Table 3: Descriptive statistics for Anonymous and Primed groups

	Anonymous			Primed		
	Mean	Max	Min	Mean	Max	Min
Share Meritocratic rounds	0.60	1	0.24	0.48	1	0.24
Share Alliance Initiation	0.23	0.60	0	0.14	0.48	0
Share Alliance Reciprocation	0.17	0.72	0	0.38	0.72	0
Share Rounds in alliance	0.28	1	0	0.62	1	0
Mean alliance length [#]	8	25	2	13	25	3
Round of first alliance [#]	8	24	1	6	24	1
Number of alliances/group	0.9	3	0	1.2	2	0
Group earnings [!]	868	925	825	847	925	809
Individual earnings [!]	217	384	53	213	376	46
Equality (group Gini)	0.15	0.39	0.01	0.20	0.39	0.03
Gains minus losses	-59	0	-124	-74	0	-116
Age	21	34	18	22	39	18
Group male share	0.53	1	0	0.52	1	0.25
Group business student share	0.32	1	0	0.35	1	0
Happiness (1= V. Unhappy ... 5=V. Happy)	3.89	5	1	3.95	5	2
Political (1= Left ...10= Right)	5.63	10	1	5.65	10	1
Family wealth (1= Wealthy ... 3= Poor)	1.92	3	1	1.71	3	1
Native english speaker (1= Yes)	0.60			0.65		
International student (1=Yes)	0.44			0.48		
Club members (1= Yes)	0.66			0.67		
People v skills (1=People...5=Skills)	2.72	5	1	2.38	5	1

[#] For groups with any alliance

[!] Except Low Rent treatment

Figure 2 shows the smoothed distribution of group payoffs across the baseline and treatments. A bimodal distribution appears to be present, particularly for anonymous groups; a result of groups ending up in one of the two equilibria, either alliance play or meritocratic play. A striking feature here is the degree to which primed groups fail to display increased payoffs in the Rotation treatment compared to the anonymous groups, indeed suggesting that Rotation couldn't break existing alliances.

It is worth noting how the interactions of player choices creates rather complex outcomes. Charts of individual payoffs by round for all groups are in Figure A.3 of the Appendix, though a sample of payoff charts capturing unexpected complex interactions is in Figure 3.⁹ The 4 payoff lines completely characterise all the choices in an experiment, with the person being chosen as receiver in a round showing a jump in the payoff line. In the left panel, blue and green players form an initial alliance (alternating the jump in payoff reflecting the bonus), only for green to renege and form an alliance with orange for about 10 rounds, after which green again teams up with blue

⁹From left to right, the relevant treatments belonging to the graphs are: Anonymous Low Rent, Primed Rotation, Primed Rotation

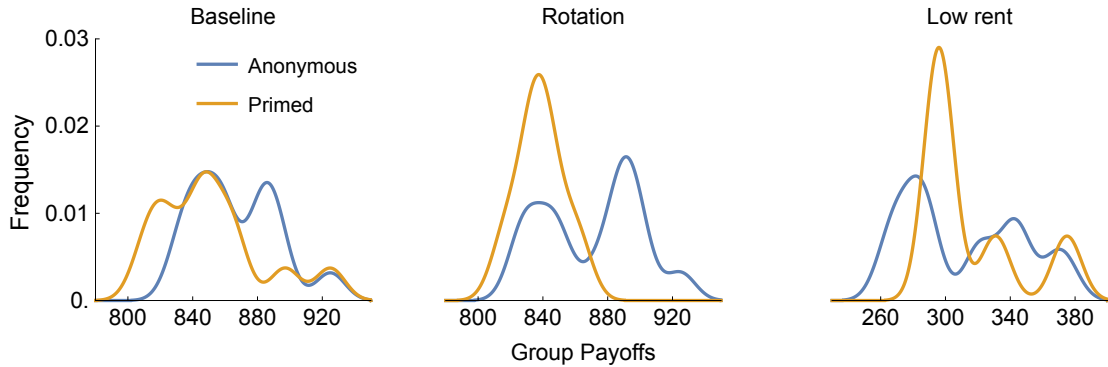


Figure 2: Smoothed frequency distribution of group payoffs (Gaussian kernel, bandwidth 9)

till round 25. In rounds 26-50 (a Low Rent treatment), orange and red eventually team up for most of the rounds.

A strange situation occurred in the centre panel, where green and red formed an alliance till round 13 when red chose blue as the receiver who promptly formed an alliance with green for the rest of the 25 round treatment. In rounds 26-50 (a Rotation treatment), blue and green maintained an alliance throughout most of the rounds, with occasionally red and orange being randomly chosen as allocator who then chose the other some of the time.

In the right panel, blue reneges on a prior alliance with orange in round 9 by choosing green, who promptly formed an alliance with red for the rest of the 25 rounds. In rounds 26-50 (a Rotation treatment), the two alliances (red-green and orange-blue) remained loyal to each other and hence only lost the allocator rights in some rounds due to random rotation.

These examples demonstrate that allocating meritocratically outside of the alliance can be punished either by 1) being replaced from your alliance with an outside player (as in the centre and left panels), or 2) giving power to an alliance of the two outsiders (as in the right panel). Fully capturing this kind of play in a theoretical model would be 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 A.1 in the Appendix gives more examples of ‘strange’ strategies as noted by some of the participants.

4.2 Treatment effects

We test for changes in the distribution of key outcome metrics between our baseline and policy treatments and report the results in Table 4. The primary outcome of interest is the change in the number of meritocratic rounds which ultimately determines the size of the group payoff.

In the anonymous Rotation groups, meritocratic choice frequency increases slightly to 0.62 from 0.60 in the baseline, whilst the Low Rent treatment decreases it to 0.56, though neither change is statistically significant. For primed groups, Rotation reduced meritocratic play to 0.42 from 0.47 whilst Low Rent increased it up to 0.57, again neither change being significant. Yet, the number of realised alliances (out of a maximum of 6 different pairs) per group decreased for anonymous groups under Rotation (1.1 to 0.8), but increased for primed groups (1.2 to 1.5) under Rotation,




Figure 3: Individual payoffs by round with complex interactions

showing that under Rotation the two primed pairs typically choose each other when given the chance via random rotation.

Additionally, Alliance Reciprocation increased from 0.14 to 0.24 for anonymous groups in the Low Rent treatment (if there was an alliance, which was less likely, it lasted longer), but fell for primed groups under Low Rents, from 0.44 to 0.35 (slightly more meritocratic choices). Together these measures indicate that Rotation policies are more likely to result in two alliance pairs in primed groups than in anonymous groups, while Low Rent policies do reduce the likelihood of alliance initiation in anonymous groups, but the alliances that then do form last longer, leading to no net gain.

Recall that all of our groups play two consecutive treatments during continuous game play, and that the first round allocator in the policy treatments received the payment in the final round of the baseline treatment. These averages are hence capturing a ‘policy plus learning’ effect with this setup, meaning that we can only identify the relative policy effects. To this end, the final three rows in Table 4 shows whether the combined effect is significantly different in the primed groups. We find that for the Rotation treatment the change in Meritocratic choices is 0.04 in Anonymous groups, and -0.07 for Primed groups ($p = 0.06$), while the change in Alliance Initiation choices is 0.01 in Anonymous groups, and 0.20 for Primed groups ($p = 0.02$), indicating that Primed groups are more likely to form two alliance pairs under Rotation than anonymous groups resulting in fewer meritocratic choices and lower groups payoffs. In contrast, the Low Rent treatment resulted in a -0.07 change in meritocratic choices in Anonymous groups and a 0.11 change in Primed groups ($p = 0.02$), while also creating a 0.16 change in AR choices for Anonymous groups and a -0.13 change for Primed groups ($p = 0.03$), suggesting that Primed groups responded to the Low Rent policy in a way that increased their group outcomes, while Anonymous groups did not. To visualise these relative differences, Figure A.4 of the Appendix plots these three choice types by round across all treatments, with these difference in the direction of treatment effects visible by the change in the trend lines for AI and AR choices between anonymous and primed groups in each treatment.

In order to separate learning from policy effects, we can look at whether the choices in the experiments converge at some point before the 25 rounds of play, which would allow us to compare later rounds in the baseline with later rounds in the treatment experiments. To this end, Figure 4 shows the proportion of alliances in the 6 types of experiments (baseline plus 2

Table 4: Treatment effects for Anonymous and Primed groups

	Anonymous			Primed		
	Baseline	Rotation	Low Rent	Baseline	Rotation	Low Rent
Share Meritocratic rounds	0.60	0.62	0.56	0.47*	0.42	0.57
Share Alliance Initiation	0.23	0.31	0.12**	0.08***	0.30**	0.08
Share Alliance Reciprocation	0.14	0.06	0.24	0.44***	0.28	0.35
Share Rounds in alliance	0.26	0.11**	0.48	0.70***	0.43**	0.67
Mean alliance length	6.2	3.8**	18.8***	14.8***	6.4***	16.8
Round of first alliance [#]	8.4	10.3	3.3**	6.5	6.7	4.6
Number of alliances/group	1.1	0.8	0.6	1.2	1.5	1.0
Share groups any alliance	0.75	0.5	0.64	0.92	1	0.83
Mean group earnings [!]	867	870	312	851	838	315
Mean individual earnings [!]	217	218	78	213	210	79
Equality (group Gini) [!]	0.15	0.15	0.08	0.27***	0.18	0.09
Gains minus losses	-58	-55	-63	-74	-87	-60
<i>Treatment change (M rounds)</i>		0.04	-0.07		-0.07*	0.11*
<i>Treatment change (AI rounds)</i>		0.01	-0.09		0.20**	0.01
<i>Treatment change (AR rounds)</i>		-0.06	0.16		-0.13	-0.13**

Statistical comparisons of distributions of treatment measures are: Anon Rotation v Anon Baseline, Anon Low Rent v Anon Baseline, Primed Baseline v Anon Baseline, Primed Rotation v Primed Baseline, Primed Low Rent v Primed Baseline.

Statistical comparisons of distributions for *treatment changes* are: Primed Rotation v Anon Rotation, Primed Low Rent v Anon Low Rent.

p values from Wilcoxon signed-rank test * < 0.10, ** < 0.05, *** < 0.01.

[#] Of groups with any alliance.

[!] No statistical tests with low discretion treatment because of changed payoff structure.

treatments, with anonymous and primed groups). Whilst there are clear trends in many of these six groups in the early periods, there are no significant trends in the anonymous groups from round 12 onwards, which we take as an indication that learning effects have been incorporated by that round.¹⁰

To isolate the policy effects from group composition and other effects, we analyse two success variables. The first is the relative change in the frequency of meritocratic choices across the whole 25 rounds of the baseline and policy treatment, calculated as $(M_{2j} - M_{1j})/M_{1j}$. The second removes the learning effects by using only rounds 12 to 23, which reflects the period of stable alliances¹¹, calculated as $(M12_{2j} - M12_{1j})/M12_{1j}$. We have 40 observations on these two success variables S_j , which we then model as

$$S_j = X'_j(\alpha_j + \beta_j X'_j) + \epsilon_j$$

where S_j is either $(M_{2j} - M_{1j})/M_{1j}$ or $(M12_{2j} - M12_{1j})/M12_{1j}$; and X' represents a matrix of

¹⁰We also tried using the interval 15-24 rounds, or 10-25, which qualitatively made no difference to the analyses.

¹¹Alliances were observed to be significantly more likely to break down in the final two rounds of every treatment.

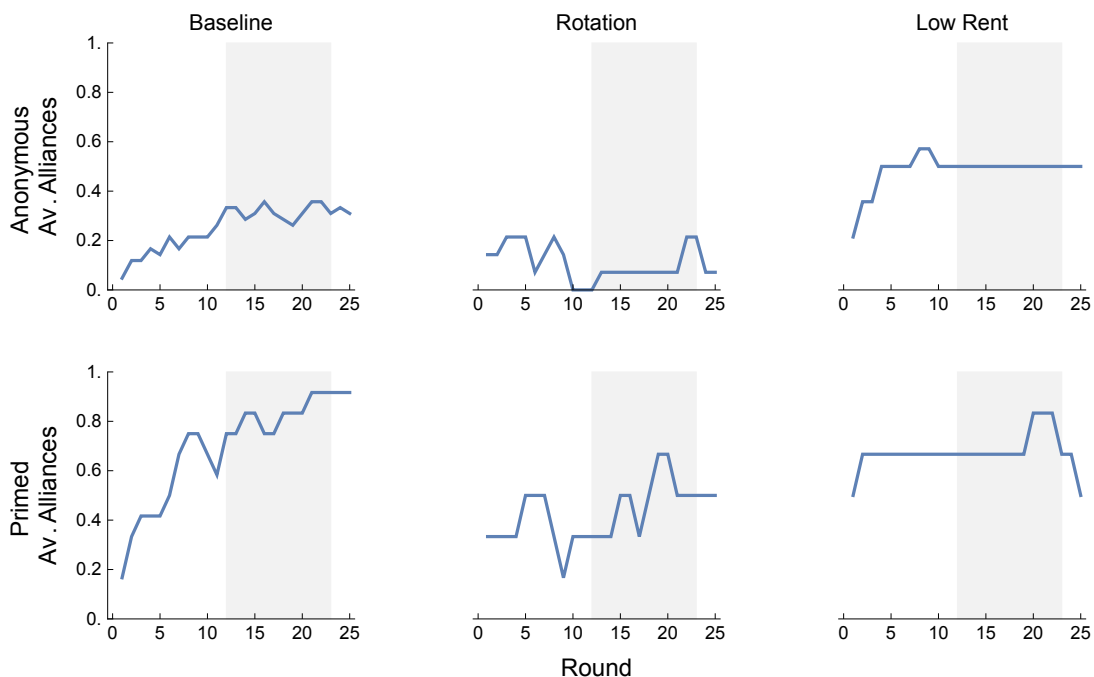


Figure 4: Establishment of stable norms in treatments

group-specific variables, such as the number of alliance reciprocation decisions in the baseline, the policy treatment and a control for the primed groups, together with a group level error term. The interaction effects between these variables are of particular interest. The results are in Table 5.

Table 5 goes through different specifications with each having a limited number of variables. The closest fit is found in column 4, which shows that the primed group has a 0.34 lower ratio of meritocratic plays than in the baseline rounds, but a 0.14 higher meritocracy ratio ($=-0.34-0.19+0.67$) in the Low Rent treatment. The Low Rent intervention in the anonymous group (the default) reduces the meritocratic ratio by 0.19, significantly different from the positive effect amongst the primed groups. None of the other treatments affects the ratio $(M_{2j} - M_{1j})/M_{1j}$ significantly.

The bottom panel shows that the same results hold qualitatively for the ratio $(M_{12_{2j}} - M_{12_{1j}})/M_{12_{1j}}$ as well: Rotation has no significant effects; primed groups play less meritocratic; Low Rent policies reduce meritocratic play in the anonymous groups; and amongst the primed groups Low Rent leads to relatively more meritocratic play.

We can summarise this as:

Result 1: Anti-corruption policy success is conditional on existing relationships

This result means that neither of our first two research questions can be answered with confidence. Rotation policies were not successful in breaking up primed pairs and merely allowed for sequential ‘reigns’ of the two existing pairs in a group. Amongst anonymous groups, rotation policy similarly had no net significant effect. Low Rent policies do turn out to be relatively

Table 5: Policy success models

	$(M_{2j} - M_{1j})/M_{1j}$						
	1	2	3	4	5	6	7
AR_{1j}	0.02**		0.01	0.02**		0.02	0.00
Primed (Pri)	-0.01		0.28	-0.34*		0.00	0.02
Rotation	-0.06	0.11	0.07		0.04	0.04	
LR	-0.08	-0.09		-0.19*	-0.17		-0.17
Pri x R		-0.21	-0.58**				
Pri x LR		0.52***		0.67***			
AR_{1j} x R					0.00	-0.02	
AR_{1j} x LR					0.04***		0.03*
Adjusted R^2	0.05	0.16	0.18	0.24	0.12	0.07	0.11
	$(M_{12_{2j}} - M_{12_{1j}})/M_{12_{1j}}$						
	1	2	3	4	5	6	7
$AR_{12_{1j}}$	0.09**		0.07**	0.09***		0.08**	0.04
Primed	-0.09		0.24	-0.58*		-0.09	-0.06
Rotation	-0.13	0.14	0.02		-0.06	-0.05	
LR	-0.15	-0.06		-0.28**	-0.23		-0.22
Pri x R		-0.19	-0.72*				
Pri x LR		0.73**		0.96**			
$AR_{12_{1j}}$ x R					0.05	-0.03	
$AR_{12_{1j}}$ x LR					0.11**		0.07
Adjusted R^2	0.12	0.08	0.18	0.24	0.15	0.12	0.15

Pri is primed group, R is Rotation treatment, LR is Low Rent treatment, AR_{1j} and $AR_{12_{1j}}$ are the Alliance Reciprocation frequency in a treatment.

more successful in situations where loyal alliances already exist, though they have no success in anonymous situations.

Result 2: Primed groups have longer lasting alliances

In line with our third research question, we can say that aligning player expectations with our priming device does result in faster alliances and longer lasting alliances, though it does not preclude the situation of players finding new alliances within the game outside of their primed alliances, as shown in the centre panel of Figure 3. Choosing personnel without existing loyal relationships is then arguably a prerequisite of establishing a group norm of meritocracy.

4.3 Individual characteristics

Another option available to combat back-scratching is to select group members on observable socio-demographic characteristics. But which ones are more likely to be associated with meritocracy?

We look at three individual level choices - meritocratic play, alliance initiation and alliance

reciprocation - across the all decisions in the experiments. There are 1,126 M choices across out of the 2000 opportunities to be meritocratic in all treatments, 411 AI choices (out of 1126 opportunities), and 465 AR choices (out of 697 opportunities). We model the decision variables $M_{ijt}, AI_{ijt}, AR_{ijt} \in \{0, 1\}$ as a linear probability model:

$$D_{ijt} = X'_{it}\alpha_D + v_t$$

where D_{ijt} denotes one of the three choice variables ($M_{ijt}, AI_{ijt}, AR_{ijt} \in \{0, 1\}$) by player i in group j in round t . X' represents a matrix of player-specific variables, α_D is the vector of coefficients, and v_t is an error term clustered by group, being $e_j + e_{ijt}$. We adopt a number of specifications to explore the relationship between socio-demographic variables and their impact on player choices in the game. Table A.2 in the Appendix shows the regression results.

First, we see that males are indeed 9% less likely to make meritocratic choices which is significant at the 10% level), 5% more likely to initiate, and 19% more likely to reciprocate alliances. Business students show a similar pattern, with 9% fewer meritocratic choices and 10% more likely to initiate alliances. In terms of predicting meritocratic choices, each additional year of age makes one 2% less likely to play meritocratic, a small but significant negative effect. Somewhat unexpectedly, higher family wealth was not associated with less meritocratic choices or more alliance initiation: if anything, it was associated with less alliance reciprocation.

When we look at the effects of social opinions in the middle panel, we first see very low explanatory power¹², which disappears as soon as we add the demographics back in, so there is no indication of separate contributions from social views and opinions. This can be summarised as:

Result 3. Men and business students are more alliance-prone

This result adds further evidence to a growing body of literature that looks at the socio-demographic determinants of cooperation in general, and corruption in particular. Males have a higher tendency to form alliances, as do business students, suggesting that back-scratching is likely to be more prevalent in institutions and sectors of the economy dominated by men with a financial or business education.

Looking at the effects of the policies in these choice models, we see the same effects earlier observed: primed pairs are more likely to reciprocate in their alliances, and Low Rent policies seem to have positive effects amongst primed groups, but not in anonymous groups. This suggests that the fear of ‘missing out’ is a much bigger factor in anonymous groups than in primed groups where individuals are more secure in their partnering. Support for this contention also comes from participant’s hand-written notes, with many commentators noting the risk involved in attempting for form alliances due to conflicting expectations. Notes to the effect of “I hope my friend trusts me” were common. A list of such notes about strategies and friendships are Table A.1 of the Appendix. In terms of policy interventions pertaining to existing institutions, it is perhaps the clearest message of the paper: with strong crystalised pairings, one can more easily coordinate on a new equilibrium wherein everyone abandons alliance play altogether, whereas with uncertain pairings the introduction of additional uncertainty of lower rents does not diminish the fear that others may form alliances at all.

¹²Apart from the AR coefficient for the Skills v people variable, which may represent an opinion based on the experience within the experiment itself, whereby those in alliances answered the survey that people matter more than skills.

5 Theoretical analysis

The main experimental results do not correspond to that of selfish rational agents under which meritocracy is the only Nash equilibrium. Yet, the reality is a murky mess of alliance seekers, back-scratchers, back-stabbers, and meritocrats: in the baseline experiments, stable alliances only emerge after round 10 or so, with the first alliance initiation typically taking till round 8. At best, players seem to be using a trial-and-error process of moving towards an equilibrium in strategies, or types, that would represent the establishment of stable expectations.

To capture some of this behaviour, we want to see here whether the main results can be rationalised if we presume that there is a small set of types that individuals belong to; a type which is chosen prior to any of the 25 rounds of the experiment. We restrict the set to three strategies with one round of memory and two strategies relying on two previous rounds of play to inform the choice. The three one round strategies we consider are

1. Meritocratic (M) - makes the meritocratic choice every time.
2. Reciprocator (R) - allocates to the previous allocator.
3. Loyal partner (LP) - allocates exclusively to one player.

The two round prior strategies allow strategic interaction and the development of social norms.

1. Tit-for-tat reciprocator (TFT), who is a conditional alliance reciprocator.
 - (a) If the previous round is meritocratic
 - i. And two rounds prior allocator is not me (or it is the second round)
 - A. If productivity 2 player is not previous allocator, then productivity 2 player (this is the first round default)
 - B. If productivity 2 player was previous allocator, then productivity 1 player
 - ii. And two rounds prior allocator is me, then chooses previous round allocator (reciprocate)
 - (b) If the previous round is not meritocratic
 - i. Chooses the previous round allocator (reciprocate)
2. Social Norm follower (SN), who is a conditional group cooperator.
 - (a) If the previous round is meritocratic
 - i. And two rounds prior allocator is not me, then chooses productivity 3 player (this is the first and second round default)
 - ii. And two rounds prior allocator is me, then chooses previous round allocator (reciprocates)
 - (b) If the previous round is not meritocratic
 - i. Chooses the previous round allocator (reciprocates)

Table 6: Best response to type expectations

		Baseline					Rotation					Low Rent				
Alloc. strat	Expect.	M	R	SN	TFT	LP	M	R	SN	TFT	LP	M	R	SN	TFT	LP
	M	231	216	231	60	64	231	212	231	165	178	94	75	94	52	52
	R	225	205	316	179	159	225	206	256	205	186	88	69	88	66	63
	SN	231	210	231	108	110	231	210	231	175	186	94	69	94	57	58
	TFT	225	210	328	207	131	225	206	262	205	183	86	69	88	69	60
	LP	225	203	316	164	206	225	206	251	194	207	88	69	88	64	69

These five strategies were selected for this analysis following participant feedback, observed prevalence of conditional cooperation strategies in repeated public goods games¹³ (Chaudhuri, 2011), and from players’ decisions in the experiments: the meritocratic strategy is consistent with 60% of all experiment decisions, while the reciprocator strategy is consistent with 21% of decisions. From the two round strategies, SN is consistent with 61% of decisions, and the TFT with 35% of decisions.

5.1 Strategy equilibria

We calculate the expected payoffs to a first-round allocator who is one of these five types, conditional upon the expectation of a single type for all others, by simulating the 25 rounds of a treatment 10,000 times with that mix of player types, which allows us to look for pure symmetric strategy equilibria. Table 6 documents the results of this exercise, presenting the expected payoff for the row type conditional on the type of others in the columns.

In Table 6 the best-response types for the allocator are shaded, and of these the symmetric pure Nash equilibria in types occur along the diagonals and are denoted in bold. The equilibria in the baseline treatment are meritocratic (M), loyal partner (LP), and tit-for-tat reciprocator (TFT). The logic of why it is not an equilibrium for all players to be the reciprocal type is that if I believe that others will play strategy R, then my best response is to play M.¹⁴ A similar logic holds if I expect all others to play SN, for which the best response is TFT.

Note here that the existence of multiple equilibria depending on expectations aligns very closely with the story so far in that early rounds of play create a persistent social norm, which in the parlance of this exercise translates to ‘early play aligns expectations about strategies’, hence allowing groups to settle on an equilibrium in strategies by way of the reasoning “if that’s how everyone else is going to play, then I’ll do the same”. The bimodal distribution of experimental outcomes in Figure 2 closely match the distribution of group and player payoffs from the meritocratic and alliance equilibria. Hence, this analysis leads us to believe that the emergence of groups norms is the product of aligning strategy expectations.

Can this exercise provide insights into the divergent policy treatment effects? The key change to equilibria in the policy treatments is that the SN strategy can potentially be rationalised as

¹³Whereas TFT reciprocates in order to form an alliance conditional on the other players reciprocating, and SN allocates meritocratically, conditional on others allocating meritocratically, or else reciprocates.

¹⁴Obviously this relies on prior reasoning, as observing someone playing strategy R will be indistinguishable from playing M, SN, or TFT if the reciprocate to the highest productivity player.

an equilibrium under a Low Rent policy, though here the strategy choices are like to collapse back to the meritocratic equilibrium. How a group moves from an alliance equilibrium of either LP or TFT to SN or M in response to a policy change is not clear, and our experimental evidence suggests that only groups on the LP equilibrium are more likely to move to an SN or M equilibrium under the Low Rent treatment. Clearly though, the three equilibria in strategies in the baseline are not themselves affected by the policy treatments.

5.2 Evolutionary strategy fitness

To account for the existence of a mechanism allowing groups to move from one equilibrium to another, we simulate the evolutionary competition between strategies as follows. We take a random selection of four out of our five strategies to play a 25 round selection stage. Selection occurs after each stage, whereby the player with the lowest payoff updates their strategy to one of the other four available strategies, randomly chosen, and the group plays another 25 round stage. We look at 10,000 simulations of a sequence of 50 selection stages.¹⁵ Under this selection mechanism, no strategy can become fully extinct, as it has a 0.2 chance of being randomly chosen as a replacement strategy in any selection stage, and hence a minimum limit of the fitness measure (how widespread the strategy is in the population) in this selection process is 0.05, and the maximum is 0.5.¹⁶ What this means is that all strategies are potentially available when we introduce a policy change in this evolutionary simulation after a sequence of 50 baseline selection stages, and hence can look at the direction of change in the fitness of each strategy under each of our two treatments.

In Figure 5 we plot the results of these simulations. The top left panel shows the relative fitness of LP, and while it is relatively fit compared to the other strategies under the baseline conditions, neither the Low Rent or Rotation policy, significantly alters its fitness. For the Low Rent treatment the first few selection stages after the policy change do see a slight decline in the fitness of LP and TFTR being matched by an increase in M and SN; a change that is not persistent, though in keeping with the direction of policy treatment effects in the primed groups in which the Low Rent treatment appeared to generate a shift towards meritocratic play.

The bottom left panel shows that under the Low Rent policy a TFT strategy maintains a high level of relative fitness, while the Rotation policy clearly reduces the fitness of the strategy, a reduction matched by an increase in the fitness of M and SN strategies, which again supports our main experimental results. If players in our anonymous experimental groups rarely settle into a LP strategy, they are likely to be following an alliance-forming strategy that is indiscriminate in partner choice, such as TFT. It makes sense that the policy treatments led to different results for primed groups (which are more likely to find the LP equilibrium), and anonymous groups (which are more likely to find the TFT equilibrium). In all, the consistency of these changes in evolutionary fitness of strategies with the direction of our experimental results suggests that in a trial-and-error process of strategy selection is a useful way of approaching this problem.

Result 4. Policy success may depend on the composition of strategies within a group

Analysing collective behaviour in this environment of potential back-scratching as a repeated interaction of strategies shows a high degree of consistency with the experimental results. Not

¹⁵The results are robust to the number of rounds in each selection stage, and equally apply with 5 or 25 rounds. After just ten selection stages the fitness of each strategy shows stability.

¹⁶A one fifth chance of being randomly chosen by the lowest one of four players generates the minimum of 0.05, while any fitter strategy at best can sustain two out of four players who form an alliance in the group.

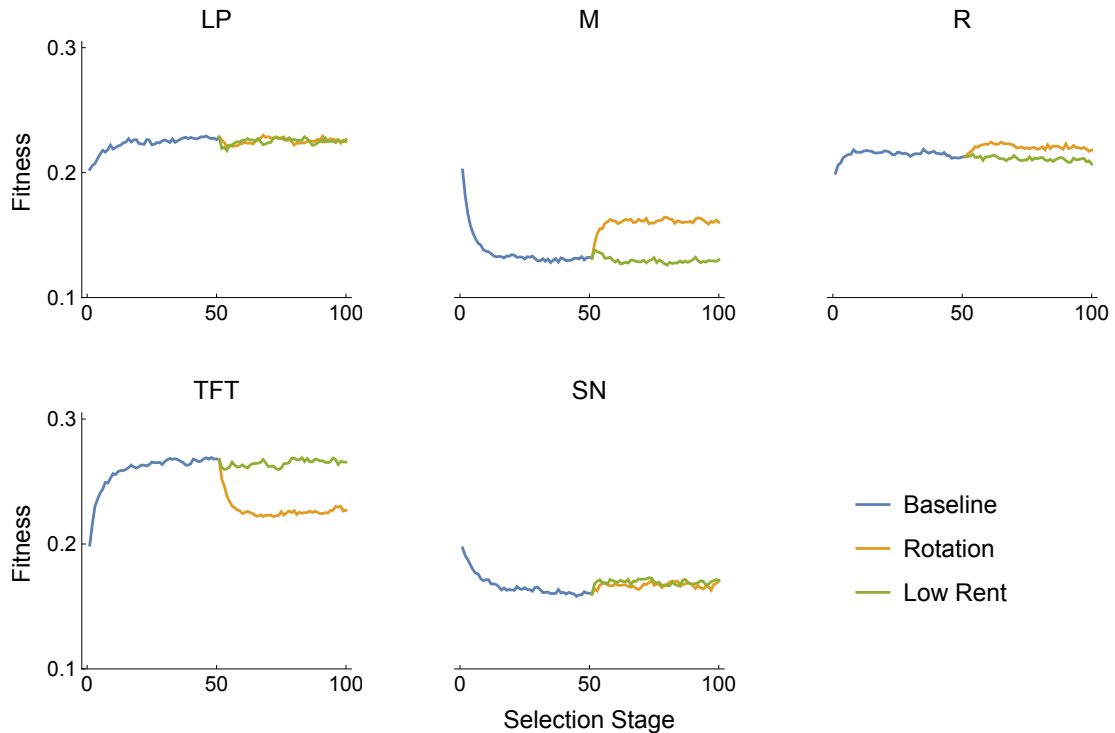


Figure 5: Evolutionary strategy fitness

only do multiple equilibria exist in the static sense, which reflect potential meritocratic and alliance-forming norms, but implementing a trial-and-error process of players finding, or shifting between equilibria reveals that the direction of policy effects will be conditional on prior strategies or norms. In particular, our simulation result that a Rotation policy will be more effective in groups where players are anonymous to their alliance partners, consistent with the actual difference we find when we compare the primed and anonymous groups.

6 Discussion and conclusions

In this paper we introduced a new experimental design aimed at producing welfare-reducing reciprocity, or back-scratching, in the laboratory. The basic design involves three alliance-related choices: alliance initiation, whereby a player not yet in an alliance chooses a receiver non-meritocratically; alliance reciprocation, where a player reciprocates non-meritocratically to the previous round's allocator and hence is seeking to sustain an alliance; and meritocratic, where players choose meritocratically that round. We looked at the frequency of these choices and the characteristics of who makes them, as well as the effect of policy interventions.

In our baseline setup, 26% of rounds could be identified as part of an alliance, defined as a sequence of at least two reciprocating rounds following an alliance initiation and involving an alliance reciprocation. These alliances were established by the 8th round and lasted for 6 rounds on average. The primed group baseline treatment had 70% of rounds as part of an alliance,

and much longer lasting alliances of 15 rounds on average. Alliance reciprocation increased from 14 to 44% of rounds while meritocratic rounds decreased from 60 to 47% on average. Males were 20% less likely to play meritocratic, and business students 21% less likely, suggesting that these groups quickly recognised the opportunities to gain from alliance formation. We were thus able to generate significant back-scratching in the laboratory even within groups of anonymous student participants.

Our staff rotation policy, whereby the next round allocator is the previous receiver with only a 50% chance, reduced alliance rounds from 26% to 11% in the anonymous groups. Yet, the share of meritocratic decisions was unchanged showing that the entire effect is due to rotation policy breaking alliances rather than reduced intentions to form and reform alliances. The increase in alliance initiation attempts from 23% to 31% of rounds, coupled with the decrease in the number of alliances per group (from 1.1 to 0.8) and mean alliance length (6.2 to 3.8 rounds on average), show that attempts to reinitiate alliances were more frequent but less successful, and that players attempted to form new alliances rather than remain loyal to their partner.

For the primed groups this policy led to a slight reduction in meritocratic rounds (from 47% to 42%), with a significant increase in alliance (re)initiation from 8% to 31%, and an increase in the number of alliances per group, showing rotation didn't prevent members from the two established alliances pairs in each group reforming their alliances whenever possible.

The Low Rent treatment did not increase the frequency of meritocratic rounds in the anonymous groups, and even significantly decreased it when we compare the stable periods of play, apparently because alliances that were already formed persisted, increasing the average alliance length from 6.2 to 18.8 rounds. In the primed groups, there was an increase in meritocratic choices (from 47% to 57% of rounds).

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. Rotating staff from a pool of people with already 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. Having low rents from alliance formation helps to induce more meritocratic play amongst groups with strong existing alliances, though not with weak existing alliances, perhaps because with weak alliances the role of the fear of losing the dominant position and then being punished is stronger.

Our design can be extended in many directions, such as increasing the transparency of previous choices, varying the groups size, varying the cultural background of participants, and introducing the possibility of punishment. An open question is to how to nudge groups towards coordinating on the meritocratic equilibrium in the first place and hence preventing alliance formation from ever happening.

References

- Abbink, Klaus. 2004. Staff rotation as an anti-corruption policy: an experimental study. *European Journal of Political Economy*, **20**(4), 887–906.
- Abbink, Klaus, & Hennig-Schmidt, Heike. 2006. Neutral versus loaded instructions in a bribery experiment. *Experimental Economics*, **9**(2), 103–121.
- Abbink, Klaus, & Serra, Danila. 2012. Anticorruption Policies: Lessons from the Lab. *Chap. 4, pages 77–115 of: Serra, Danila, & Wantchekon, Leonard (eds), New Advances in Experimental Research on Corruption*. Book Series: Research in Experimental Economics. Emerald Group Publishing Limited.
- Abbink, Klaus, Irlenbusch, Bernd, & Renner, Elke. 2002. An Experimental Bribery Game. *Journal of Law, Economics, & Organization*, **18**(2), 428–454.
- Andreoni, James, & Petrie, Ragan. 2004. Public goods experiments without confidentiality: a glimpse into fund-raising. *Journal of Public Economics*, **88**(7–8), 1605–1623.
- Armantier, Olivier, & Bolky, Amadou. 2008. Can Corruption be Studied in the Lab? Comparing a Field and a Lab Experiment. *CIRANO - Scientific Publications*, **26**.
- Barr, Abigail, & Serra, Danila. 2009. The effects of externalities and framing on bribery in a petty corruption experiment. *Experimental Economics*, **12**(4), 488–503.
- Barr, Abigail, & Serra, Danila. 2010. Corruption and culture: An experimental analysis. *Journal of Public Economics*, **94**(11–12), 862–869.
- Barr, Abigail, Lindelow, Magnus, & Serneels, Pieter. 2009. Corruption in public service delivery: An experimental analysis. *Journal of Economic Behavior & Organization*, **72**(1), 225–239.
- Barth, James R, Caprio, Gerard, & Levine, Ross. 2012. *Guardians of Finance: Making regulators work for us*. MIT Press.
- Bohnet, Iris, & Frey, Bruno S. 1999. Social Distance and other-regarding Behavior in Dictator Games: Comment. *The American Economic Review*, **89**(1), 335–339.
- Büchner, Susanne, Freytag, Andreas, González, LuisG., & Güth, Werner. 2008. Bribery and public procurement: an experimental study. *Public Choice*, **137**(1-2), 103–117.
- Burnham, Terence C. 2003. Engineering altruism: a theoretical and experimental investigation of anonymity and gift giving. *Journal of Economic Behavior & Organization*, **50**(1), 133–144.
- Cameron, C. Daryl, Brown-Iannuzzi, Jazmin L., & Payne, B. Keith. 2012. Sequential Priming Measures of Implicit Social Cognition: A Meta-Analysis of Associations With Behavior and Explicit Attitudes. *Personality and Social Psychology Review*, **16**(4), 330–350.
- Cameron, Lisa, Chaudhuri, Ananish, Erkal, Nisvan, & Gangadharan, Lata. 2009. Propensities to engage in and punish corrupt behavior: Experimental evidence from Australia, India, Indonesia and Singapore. *Journal of Public Economics*, **93**(7–8), 843–851.
- Charness, Gary, & Gneezy, Uri. 2008. What’s in a name? Anonymity and social distance in dictator and ultimatum games. *Journal of Economic Behavior & Organization*, **68**(1), 29–35.
- Chaudhuri, Ananish. 2011. Sustaining cooperation in laboratory public goods experiments: a selective survey of the literature. *Experimental Economics*, **14**(1), 47–83.

- Chen, Yan, & Li, Sherry Xin. 2009. Group Identity and Social Preferences. *The American Economic Review*, **99**(1), 431–457.
- Choi, Syngjoo, Gale, Douglas, Kariv, Shachar, & Palfrey, Thomas. 2011. Network architecture, salience and coordination. *Games and Economic Behavior*, **73**(1), 76–90.
- Congleton, Roger D. 2014. Rent Seeking and Organizational Governance: Limiting Losses from Intra-Organizational Conflict. Available at SSRN 2444756.
- Faccio, Mara, Masulis, Ronald W., & McConnell, John J. 2006. Political Connections and Corporate Bailouts. *The Journal of Finance*, **61**(6), 2597–2635.
- Frank, Björn, & Schulze, Günther G. 2000. Does economics make citizens corrupt? *Journal of Economic Behavior & Organization*, **43**(1), 101–113.
- Greiner, Ben. 2003. An Online Recruitment System for Economic Experiments. In: Kurt Kremer, Volker Macho (ed), *Forschung und wissenschaftliches Rechnen*. GWDG Bericht 63, Göttingen: Ges. für Wiss. Datenverarbeitung.
- Greiner, Ben, & Schneider, Patrick. 2015. Campaigns, Coalition Formation, and Coalition Stability in a Repeated Dictator Election Game. *UNSW Working Paper Series*.
- Jain, Arvind K. 2001. Corruption: A Review. *Journal of Economic Surveys*, **15**(1), 71–121.
- Konishi, Hideo, & Ray, Debraj. 2003. Coalition formation as a dynamic process. *Journal of Economic Theory*, **110**(1), 1–41.
- Kurzban, Robert. 2001. The Social Psychophysics of Cooperation: Nonverbal Communication in a Public Goods Game. *Journal of Nonverbal Behavior*, **25**(4), 241–259.
- Lambsdorff, Johann Graf. 2002a. Corruption and rent-seeking. *Public Choice*, **113**(1-2), 97–125.
- Lambsdorff, Johann Graf. 2002b. Making corrupt deals: contracting in the shadow of the law. *Journal of Economic Behavior & Organization*, **48**(3), 221–241.
- Lambsdorff, Johann Graf, & Frank, Björn. 2010. Bribing versus gift-giving -An experiment. *Journal of Economic Psychology*, **31**(3), 347–357.
- Lambsdorff, Johann Graf, & Frank, Björn. 2011. Corrupt reciprocity - Experimental evidence on a men's game. *International Review of Law and Economics*, **31**(2), 116–125.
- Levine, Ross. 2012. The Governance of Financial Regulation: Reform Lessons from the Recent Crisis. *International Review of Finance*, **12**(1), 39–56.
- Matthews, Dylan. 2014. Whistleblower's tapes suggest the Fed was protecting Goldman Sachs from the inside. *vox.com*.
- McLeish, Kendra N., & Oxoby, Robert J. 2011. Social interactions and the salience of social identity. *Journal of Economic Psychology*, **32**(1), 172–178.
- Nash, John F., Nagel, Rosemarie, Ockenfels, Axel, & Selten, Reinhard. 2012. The agencies method for coalition formation in experimental games. *Proceedings of the National Academy of Sciences*, **109**(50), 20358–20363.
- Olson, Mancur. 1982. *The Rise and Decline of Nations: Economic Growth, Stagflation, and Social Rigidities*. Yale Univeristy Press.

- Ones, Umut, & Putterman, Louis. 2007. The ecology of collective action: A public goods and sanctions experiment with controlled group formation. *Journal of Economic Behavior & Organization*, **62**(4), 495–521.
- Rand, David G., Arbesmand, Samuel, & Christakis, Nicholas A. 2011. Dynamic social networks promote cooperation in experiments with humans. *Proceedings of the National Academy of Sciences*, **108**(48), 19193–19198.
- Ray, Debraj, & Vohra, Rajiv. 2014. Coalition Formation. In: Young, Peyton, & Zamir, Shmuel (eds), *Handbook of Game Theory*, vol. 4. North-Holland.
- Roubini, Nouriel, & Mihm, Stephen. 2010. *Crisis economics: A crash course in the future of finance*. Penguin.
- Schaffner, Markus. 2013. *Programming for Experimental Economics: Introducing CORAL - a lightweight framework for experimental economic experiments*. QuBE Working Papers 16. Queensland University of Technology.
- Schikora, Jan Theodor. 2011a. *Bringing good and bad Whistle-blowers to the Lab*. Munich Discussion Paper 2011-4. Department of Economics University of Munich.
- Schikora, Jan Theodor. 2011b. *Bringing the Four-Eyes-Principle to the Lab*. Munich Discussion Paper 2011-3. Department of Economics University of Munich.
- Schulze, Günther G., & Frank, Björn. 2003. Deterrence versus intrinsic motivation: Experimental evidence on the determinants of corruptibility. *Economics of Governance*, **4**(2), 143–160.
- Serra, Danila. 2012. Combining top-down and bottom-up accountability: evidence from a bribery experiment. *Journal of Law, Economics, and Organization*, **28**(3), 569–587.
- Silver-Greenberg, Jessica, Protess, Ben, & Eavis, Peter. 2014. New Scrutiny of Goldman’s Ties to the New York Fed After a Leak. *New York Times*, **November 19**.
- Tajfel, Henri. 1970. Experiments in intergroup discrimination. *Scientific American*, **223**(5), 96–102.
- Takács, Károly, Squazzoniy, Flaminio, Bravo, Giangiacomo, & Castellaniy, Marco. 2014. Employer networks, priming, and discrimination in hiring: An experiment. *Analytical Sociology: Actions and Networks*, 373.
- van Veldhuizen, Roel. 2011. *Bribery and the Fair Salary Hypothesis in the Lab*. Working Paper. Tinbergen Institute.
- World Bank. 1992. *Governance and Development*. Tech. rept. World Bank.
- Yopyk, Darren J. A., & Prentice, Deborah A. 2005. Am I an Athlete or a Student? Identity Salience and Stereotype Threat in Student–Athletes. *Basic and Applied Social Psychology*, **27**(4), 329–336.

Appendix

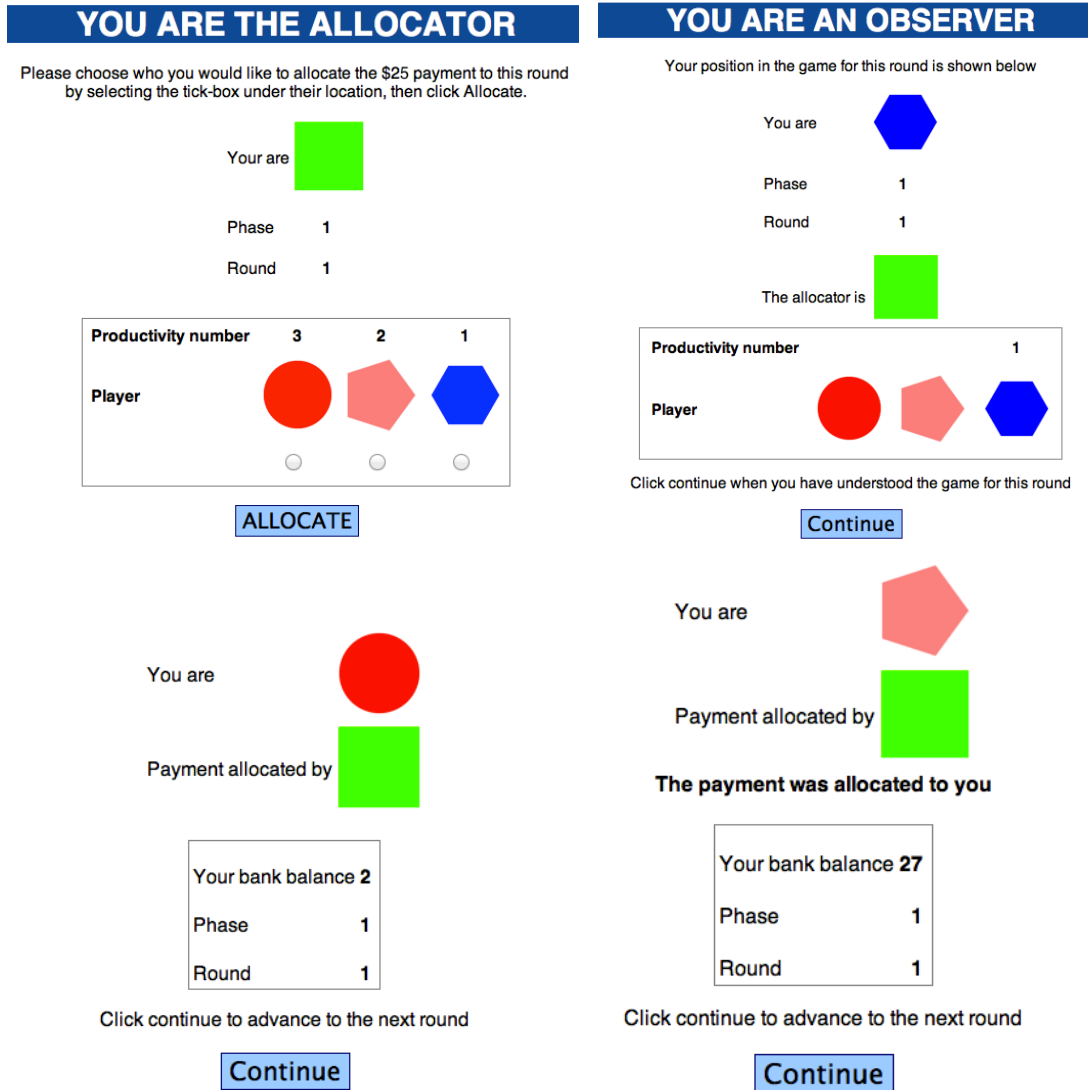


Figure A.1: Clockwise from top left: Allocator decision screen, non-allocator wait screen, notification of decision for non-receivers, notification of decision for receiver

Player Information

Not all players in the game are equal.

In the game you are friends with the player whose symbol is



The two of you always helped each other in the past. This player found you a job with their uncle when you finished school. You are each always ready to repay the favour of the other, because you know that helping each other out is what friends do.

You have no history with other players in the game - they have never done you any favours, nor have you done any for them.

Your friend is also being reminded of your shared history together.

Remember to play the game to earn as much money as you can.

Figure A.2: Instructions given to players in Friendship condition to simulate ongoing social relationship and manipulate expectations

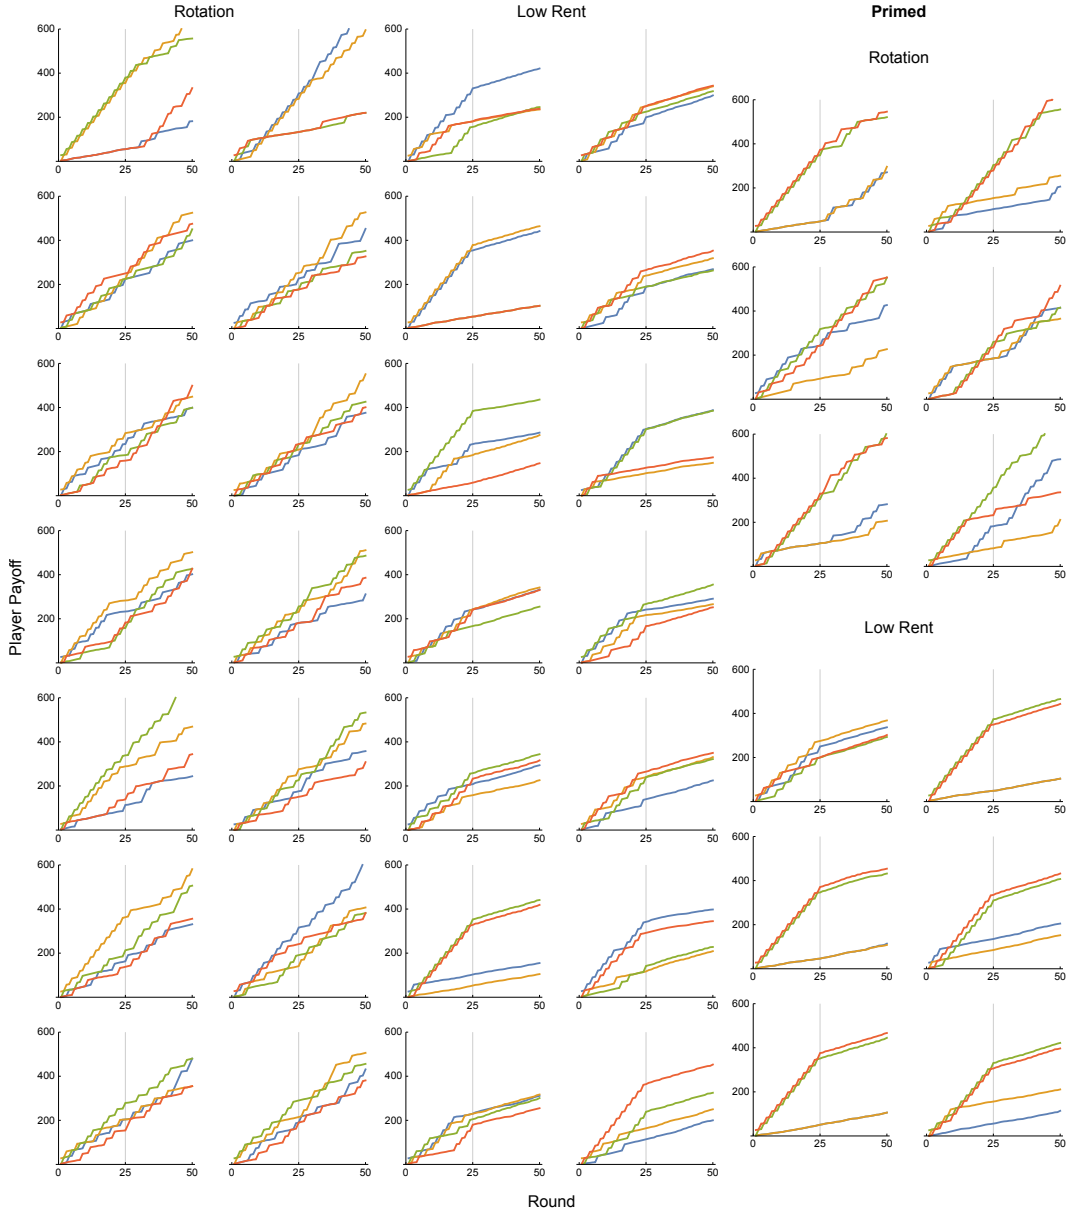


Figure A.3: Accumulated player payoffs by group, round and treatment

Table A.1: Examples of participant's written comments

*My friend and I are greedy except when they are 1 and others are 3.
Not going to get it back until I get a 3 again*

*Hope the player information is correct with who trusts me to pay them so they repay me.
Never use productivity number 1 even if it's a mate. Always go for highest productivity.
Check to see if he repays me, try to get repayment continuous b/w the two. Same tactic for round 2*

*Seems best to continuously allocate to the player with the highest productivity number
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.*

It's always good to have support. If they support you, you should support them

Choose red circle every time and hopefully they realise they are best off choosing me.

*For one note that this game would be played very differently if people were face to face.
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)*

*Strategy: pick one player and always allocate to that person. What if I get isolated?
Isolate the other two?*

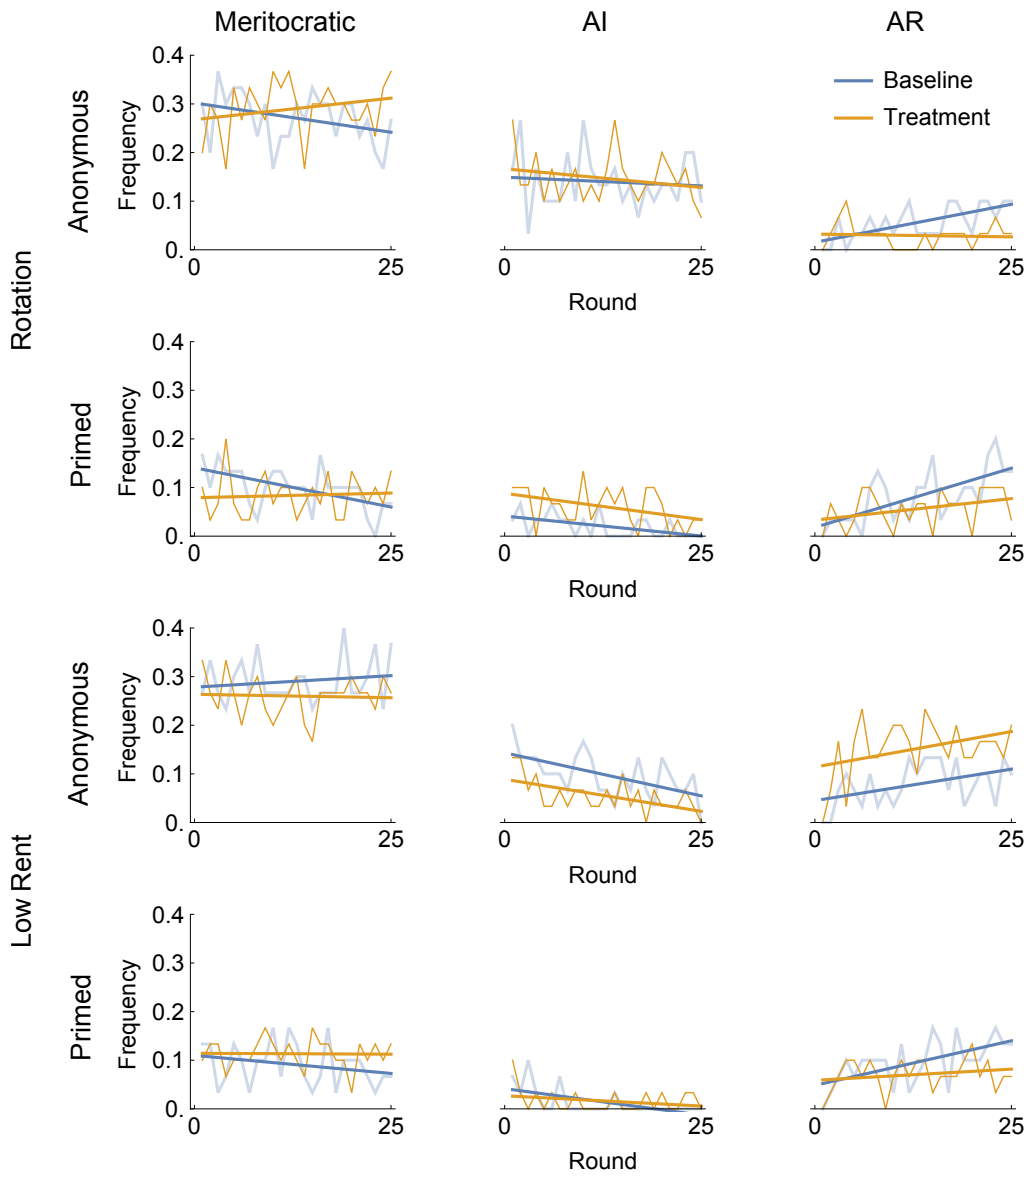


Figure A.4: Choices by round and treatment including fitted linear model

Table A.2: Choice models

	Demographics			Social			Treatments			Combined		
	M	AI	AR	M	AI	AR	M	AI	AR	M	AI	AR
Age	-0.02*	0.01	0.01							-0.01	0.01	0.00
Gender	-0.09*	0.05	0.19*							-0.08	0.06	0.07
Relationship	-0.05	0.01	0.10							-0.04	0.01	0.08
Inter. student	0.02	0.03	-0.10							0.01	0.04	-0.08
Bus. student	-0.09	0.10	0.11							-0.09	0.11	0.05
Happiness	-0.02	-0.00	0.05							-0.03	0.01	0.06
Family wealth	0.05	-0.01	-0.11							0.04	-0.01	-0.05
Politics				-0.00	0.01	-0.02				-0.01	0.00	-0.02
Skills v people				0.03	0.01	-0.09*				0.01	0.01	-0.03
Clubs				0.01	0.00	-0.01				0.01	-0.00	-0.01
Leadership				-0.01	-0.03	0.18				0.02	-0.05	0.08
Rotation							0.02	0.02	-0.23	0.02	0.02	-0.21
Low Rent							-0.04	-0.11	0.33***	-0.03	-0.10	0.27**
Primed							-0.13*	-0.09	0.46***	-0.07	-0.10	0.40***
Rotation x Pri							-0.07	0.21	-0.01	-0.10	0.22	-0.01
Low Rent x Pri							0.14	0.09	-0.39*	0.11	0.06	-0.33
N	2000	1318	697	2000	1318	697	2000	1318	697	2000	1318	697
R^2	0.03	0.02	0.11	0.00	0.00	0.06	0.02	0.02	0.31	0.04	0.04	0.30

p values are * < 0.10, ** < 0.05, *** < 0.01 with standard errors clustered by experiment group.

Gender = 1 for male. People skill is 1-5 where mostly skills is 5, and mostly people is 1. Politics is 1-10 where 1 is Left and 10 is Right.