

# Is transparency an anti-corruption myth?

Cameron K. Murray<sup>a</sup>, Paul Frijters<sup>b</sup>, and Markus Schaffner<sup>c</sup>

<sup>a</sup>School of Economics, The University of Queensland

<sup>b</sup>Centre for Economic Performance, London School of Economics

<sup>c</sup>School of Economics and Finance, Queensland University of Technology

March 21, 2017

## Abstract

We look at the effect of transparency on the incidence of costly back-scratching in a laboratory setting by implementing player identification via photographs. In our experimental design players have an incentive to form bilateral alliances in which they favour their partner at the expense of others. We find no improvement in overall group payoffs from transparency. A plausible story that fits our results is that there may be two countervailing forces at play. First, more rapid alliance formation due to social cues from the photographs being used as a coordination device to facilitate faster alliance formation between some players. Second, shorter alliances due to prosocial forces at the group level. We draw out lessons for policy makers about the limits of transparency in curtailing “grey” types of corruption.

**JEL Classification:** C92, D7, D8

**Keywords:** Experiment, Alliance, Corruption, Transparency

## 1 *Quid pro quo* corruption

In recent decades many international organisations have adopted transparency in their charters with the expectation of reducing corruption. This includes the Extractive Industries Transparency Initiative (EITI), which commits resource companies to disclose all payments made to governments. Also, the 2005 United Nation Convention against Corruption calls on each State Party to “enhance transparency in its public administration”. The IMF includes the “transparency of government accounts” in its 1997 definition of good governance, and the EU 2011 Transparency Register was similarly developed to provide citizens with information about who is involved in European decision making. These initiatives and others build on a large development literature on corruption summarised by Tanzi (?), who concludes that “[i]n many countries, the lack of transparency in rules, laws and processes creates a fertile ground for corruption”. Kolstad and Wiig (?) re-emphasise that conclusion 10 years later, though they also draw attention to the importance of factors accompanying transparency, such as accountability.

Obtaining empirical evidence of the success of anti-corruption transparency policies is hampered by the strong incentive for concealment by corrupt parties. Many researchers have therefore

---

We thank Tim Frijters for extensive assistance in conducting these experiments.

turned to a more controlled laboratory setting to look for evidence of the effects of anti-corruption policies (see ? for a review of this literature). In this paper we study the importance of the visibility of back-scratching on corruption in the laboratory, where corruption entails a reciprocation of favours between parties that incurs a cost on others. To this end we use a new experimental game specifically designed to allow for mutual back-scratching at the expense of others, and then ascertain whether having all the players see pictures of all other players reduces the amount of backscratching observed, even in the absence of official punishment possibilities.<sup>1</sup>

Our study concerns a process somewhat outside the standard definition of corruption of “acts which utilise the power of public office for personal gain in a manner that contravenes stipulated rules” (?). This definition fits very much into a principal-agent view of corruption whereby a single decision maker controlling a public resource can be corrupted by an outsider. We instead focus on a more social definition of corruption that recognises that power is had by different people over time. We define back-scratching as the formation of reciprocal groups who trade favours over time for mutual gain at the expense of others.<sup>2</sup> This ‘grey-corruption’ view also differs from traditional notions of rent-seeking inasmuch as there is no lottery-type contest for economic rents controlled by a single principal (?). Instead we consider a long run game of *qui pro quo* between people who alternate in positions in power (?). This approach has a long history in the study of the political influence of interest groups (?????), and broadly incorporates the phenomena of nepotism, cronyism, and the revolving door of personnel between elite positions in government and private sector (???). This is almost invariably seen as unethical behaviour, but is not always seen as corruption from a legal standpoint.

To capture this back-scratching process in our experiment, in each of the 25 rounds of a treatment there is a discretionary decision made by a single player (the ‘allocator’) from a group of six subjects, about which of the other player will receive a payment (the ‘receiver’), worth 20 Experimental Currency Units (ECU). Non-allocator players are given a randomly drawn productivity number each round —a shuffled set of sequential integers from 1 to 5—which determines the group payoff for that round (an amount received by all players). Each subject’s payoff equals the chosen receiver’s productivity number that round, while the receiver’s payoff includes the discretionary payment for that round in addition. The ‘twist’ is that the receiver of the payment is the decision maker next round. Group payoff maximisation arises from allocating the payment to the highest productivity player each round, which we call a meritocratic strategy. But there is a strong incentive for individuals to form alliance pairs, and keep the allocator position within that alliance, at the cost of the other players. Players face the conflict of maximising individual earnings by forming an alliance pair, and maximising group earnings by playing meritocratically.

In this experimental setup we study the effect of transparency on alliances in a treatment where the six subjects in a group are shown photographs of each other, rather than being anonymously represented in the other player’s computers by coloured symbols, in order to provoke a cooperative response. In short, we lift the ‘veil of ignorance’ under which back-scratching alliances typically operate by mimicking possible transparency regulations, such as publishing the identities of decision-makers and their favoured partners.

While improved transparency might increase awareness of the social norm of meritocracy, and thus reduce the incidence of alliance formation, there is a counterveiling effect: the photos allow

<sup>1</sup>The game we use to study transparency is an extension of ? and sits in between the traditional experimental corruption literature centred around the repeated bribery game (RBG) of ?, and a number of literatures on coordination in repeated games such as the public goods game, and games of group formation and in-group bias (?????????).

<sup>2</sup>Loyalty to political connections, or political parties, is repeatedly observed in empirical analysis of lobbying and political donations, adding to the evidence that corruption is a relationship-based process (????).

potential partners to see each other, which can facilitate alliance formation if there is something visible that signals a willingness to coordinate. Seeming to be 'open for business' might become more salient when players can see each other in photos.

The results suggest that providing photographs seems to allow some players to use visible social cues to more rapidly form alliances, particularly business students and players rated as more beautiful by their group, confirming a number of previous findings (??). On balance however, the transparency treatment had no more meritocratic play because alliances lasted less long, and hence the same group payoffs eventuated. These results suggest a degree of caution regarding blanket calls for greater transparency as an anti-corruption policy, particularly where transparency is not a precursor to the option for punishment.

## 2 Background

The experimental literature on corruption, and the broader literature on the role of transparency on cooperation and group behaviour, inform our approach. ? pioneered an experimental design, the repeated bribery game (RBG), to look at the micro-level dynamics of corruption. In this design, repeated bribery opportunities are encountered by two subjects, one in role of a firm (a potential briber) and the other a public official, with the acceptance of a bribe being associated with a cost to other players in the experimental session. This setup has since been extensively modified to examine a number of elements in corruption, with the most relevant to our paper being ?, who implemented a *four eyes policy* in which the role of public official is given to two players who must jointly make the decision to accept or reject the bribe, and decide how to allocate payoffs between themselves and the potential briber. Rather than a reduction in bribery from this transparency policy, the 10 round *four eyes* treatment led to a higher frequency of corrupt transactions, even though the bribe was split between two public officials. Making corruption transparent to a third party in this case merely fostered that third party's involvement. This type of result is supported in a field study by ?, who found that grassroots monitoring is ineffective at reducing missing expenditures in construction projects.

Where our game differs substantially from the RBG is the nature of externalities arising from corruption. Typically the RBG employs one of two *ad hoc* approaches to generate a negative external cost. Either a small deduction is made from the earnings of other subjects in the same laboratory session (but who are not part of the matched briber-official pair), or a deduction is made from a charity donation by the experimenter (??). In either case there is relatively low salience of the externality insofar as the choice as to who suffers it remains detached from the choice whether or not to engage in corruption. A more common corruption situation involves direct discretionary choices about favouring a future alliance partner over identifiable others who suffer the cost, but could have been favoured instead. For example, a public officer choosing between contractors will, by explicitly choosing the winning contractor to be their alliance partner, implicitly choose the loser as well, and by doing so potentially forgo the chance of forming an alliance with the other contractor.

In our setup, transparency provides additional information to those subjects who are suffering the negative externality from corruption, but who could be favoured instead. In contrast, the *four eyes* variation of the RBG does not offer transparency to those suffering the externality, nor can they be chosen to be part of a corrupt alliance. The incentive change from transparency in our setup is therefore quite different, though Schikora's (?) results suggest that knowing corrupt

alliances can be observed may serve to reinforce them; continuing to be corrupt even when being observed could signal a greater commitment to the alliance.

Our experiment differs from other cooperative games, such as dictator games, ultimatum games, and public goods games, in the sense of having more complex payoff conflicts, but transparency in these setups has been extensively studied and is hence informative. ? found that cooperative behaviour increases nearly twofold in a one-shot prisoner’s dilemma game when individuals can see each other in the laboratory, increasing further when they can also communicate, indicating that cooperation is not only reciprocity-based but also identifiability-based. ? ran standard dictator games that allowed sharing of a \$10 endowment with take-it-or-leave-it offer. In the anonymous treatment only 1 in 26 dictators offered a 50-50 division, yet when either recipients or dictators could see a picture of the other player 25% of dictators offered an equal split, demonstrating the importance of visibility for the social norm of equality to affect choices. However, no significant difference in the frequency of \$0 allocations between treatments was found, indicating that the effect of revealing visibility is conditional on the decision to already make a positive offer and cooperate to some degree. In a public goods game, where the level of information available on the group members was varied, ? found that providing photos of the other group members, and information as to their contributions, increased public good contributions by around 80%, though the effect of photos or information alone was insignificant.

Using a slightly different form of transparency, ? find an almost 50% increase in offers in the dictator game if the dictators are given the last name of the recipient, though they find a non-significant increase in offers in the ultimatum game where receivers could reject an offer, and where the initial offers were higher to begin with. In a group dictator game ? find that adding a set of ‘robot eyes’ to the screen of the dictator increased public good giving by 29%, again indicating a direct effect of being watched.

One would expect from these existing papers, particularly the two most closely aligned studies of ? and ?, that providing photos would increase pro-social behaviour. Yet, in the context of back-scratching it is not entirely clear what ‘pro-social’ behaviour actually is; cooperating with one other player increases payoffs for both ‘partners-in-crime’ at the cost of the anonymous others, while playing meritocratically represents a wider form of pro-sociality to the group as a whole.

Experiments by ? support the idea that transparency is not automatically effective at improving cooperation at the broadest level, and that information about others can provide cues that facilitate tacit coordination. When players in their experiment were faced with choosing a colour that matched the colour choice of another player, information about the gender, field of study, or a player’s additional choice of a university building, facilitated coordination. Similarly, in ultimatum games where photographs of players were revealed, men and attractive players were offered higher allocations (?), suggesting some underlying biases in preferences for cooperative partners. In terms of the strength of in-group favouritism when choosing team mates, ? conduct experiments where players choose team mates given a group signal (of their field of study), or a performance signal, based on a real effort task, and find a strong bias towards team mates choices being made on group signals rather than performance ones. In terms of our setup, these results suggest that revealing players identities through photographs may offer social cues about players that facilitate particular alliances.

Where we add to this considerable literature is our new experimental design capturing the process of corruption as costly back-scratching, and whereby players are able to choose alliance partners from a group, hence implicitly choosing who will bear the cost of the alliance. Players are also able to break and reform alliances with other subjects. The effect of transparency on choices in

our experiment is not clear cut; revealing player’s identities via photographs may generate social pressure from the group and encourage cooperation, but it may also provide social cues as to which other players are more likely to reciprocate and form an alliance.

### 3 Experimental Design and Research Questions

#### 3.1 Basic design

The design was first developed in ?. The baseline treatment consists of 6 subjects<sup>3</sup> whose composition is the same for the whole duration of the experiment, and who are identifiable to each other by a coloured shape on the screen. Each round one player, the allocator, chooses which of the other players to receive a 20 ECU payment, with the first round allocator randomly chosen. This payment represents a discretionary allocation of economic rent. 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 players not allocating are given a randomly shuffled productivity number, from the set {1, 2, 3, 4, 5} which the allocator can observe before making their decision. Each of the 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 payment in addition. We call a choice meritocratic if the payment is allocated to the player with the highest productivity number as it maximises the total group payoff that round. The payoff matrix for a single round is summarised in Table 1.

Table 1: One round payoff structure

<i>Choice</i>	Player 1 Allocator	Player 2 Prod. 1	Player 3 Prod. 2	Player 4 Prod. 3	Player 5 Prod. 4	Player 6 Prod. 5	Group Total Payoff
Player 2	1	21	1	1	1	1	26
Player 3	2	2	22	2	2	2	32
Player 4	3	3	3	23	3	3	38
Player 5	4	4	4	4	24	4	44
Player 6	5	5	5	5	5	25	50

The situation here shows a single round where Player 1 is the allocator. Payoffs are in ECU.

Randomising productivity numbers captures the idea that choosing an alliance partner to favour when making a discretionary decision often means forgoing the most efficient choice for the group, such as choosing the less talented person for a job, choosing a less efficient contractor, and so forth. Meritocratic choices are, in contrast, always the most efficient, or productive, choice for the group as a whole. The conflict between incentives to form an alliance pair or play meritocratically are made clear in Table 2. Under meritocratic play, all players make 208ECU

<sup>3</sup>We increase the group size for these experiments to ensure greater anonymity of subject choices under the baseline condition. Identifying players to each other with photographs enables subjects to determine who may have been an alliance pair in the anonymous treatment. A larger pool of players increases the difficulty of identifying any alliance pairs after the experiment. Earlier experiments showed that subjects excluded from alliances could be quite emotional about the experiment result, and we wanted to avoid any real life conflict after the completion of the experiment. Subjects in groups were also randomised in their seating position in the lab, and between treatments the location of players on screen was randomised.

in expectation, and the group payoff is maximised at 1,250ECU. However, the expected payoffs of players within a bilateral alliance from the first round are much higher, at 325ECU, while for those in that same group that are outside the alliance make just 75ECU each. In total, the group payoff declines by 24% to 950. In sum, there is a 56% increase in payoffs available from forming an alliance from the first round, through it reduces the earnings of the four players outside your alliance by 64%.

Table 2: Expected payoffs from successful alliance and meritocratic strategies in ECU

Strategy	In my alliance		Out of my alliance		Group	
	Payoff	Change	Payoff	Change	Payoff	Change
Meritocratic	208		208		1,250	
Alliance	325	56%	75	-64%	950	-24%

Meritocratic play maximises group payoff ( $20 \times 25 + 5 \times 25 = 1,250$ ), which is equally divided amongst players as their expectation prior to the first round. For alliance play, for players able to form an alliance by being the first round allocator or receiver is jointly ( $20 \times 25 + 3 \times 25 \times 2 = 650$ ), with an equal share for alliance partners being 325, with the first allocator getting 315 and the first receiver 335. For alliance play the players not in the alliance will each receive ( $3 \times 25 = 75$ ).

Under standard rationality assumptions, only meritocratic play is individually optimal. Any another strategy unravels with backwards induction; in the last round, profit-maximisation requires choosing the player with the highest productivity, independent of the history of the game. This makes it optimal to also be meritocratic in the penultimate round, and so on, till round one. Alliance-forming non-meritocratic behaviour only makes sense where expectations that others will reciprocate can be generated. It is these expectations that may be affected by improving the transparency of decisions being made by revealing the identity of players using photographs.

## 3.2 Treatments

Ten groups of 6 subjects play either two baseline treatments, or one baseline then one transparency treatment, or one treatment then one baseline, with 30 groups in total, playing 20 transparency treatments and 40 baseline. In the transparency treatment, each subject sees a photo of each of the others in their group instead of a coloured shape. The first round of a treatment involves a new random draw of the first round allocator, as well as randomising player positions on the screen so that subject cannot tell who played what in a previous treatment.

## 3.3 Procedures

The experiments were conducted between April and July 2013 with 180 students participants recruited from the Queensland University of Technology in Brisbane, Australia, using ORSEE (?). They took place in university computer labs, using iPads running CORAL software (?). Photos were taken immediately prior to the transparency treatment using the iPad camera. Each subject played two treatments of 25 rounds, and received their accumulated experiment ECU earnings converted to Australian Dollars at a ratio of 20:1, making the average payoff \$AUD 18 for an average of 50 minutes of play, which is above the minimum wage.

After receiving experimental instructions, subjects went through a set of hypotheticals to ensure they understood the instructions and the basic dynamics of the game. At the end of the experiment subjects completed a socio-demographic survey, which included questions about whether they recognised other subjects in their group, and if so, what their relationship is, as well as being asked to rate attractiveness of other subjects and answer a series of questions about their strategic considerations during the experiment.

### 3.4 Research questions

We are mainly concerned with choices involved in the dynamics of alliance formation and maintenance. We define a number of choice types to that end. First, a meritocratic (M) choice is any allocation of the payment to the player with the highest productivity number in that round. An alliance initiation (AI) is defined as a choice to allocate the payment in a round to a player with a productivity number less than the maximum, and where those players are not already in an alliance. This choice variable captures intentions to form alliances through costly signals, or favours, that a player hopes will be reciprocated in future rounds. An alliance reciprocation (AR) choice is a non-meritocratic reciprocal allocation of the payment in the round immediately following an alliance initiation, or any other non-meritocratic decision within an alliance period. AR choices capture a slightly different element in alliance formation, which is the willingness of players to respond to the choices of others, or in general terms, to follow social norms being set by others after an alliance is formed by being loyal to that alliance. A round is classified as in an alliance (IA) if it forms part of a period of exclusive dealing between two players in which the first decision is an AI, and at least one other is non-meritocratic (being an AR). We look closely at the effect of the experiment treatments on this suite of choice types, along with other basic descriptive measures of group outcomes, in order to answer the following research questions.

**Question 1:** Does identifying players increase meritocratic play and increase group payoffs?

Implementing a policy of increased transparency in corruption and cooperation games has previously found to be ineffective (??), yet in other cooperative games, photos of players has increased overall cooperation (????). Because our experimental design captures a different type of cooperation, whether these results can be replicated in our back-scratching setting, where the choice of alliance formation is directly linked to the choice of those who suffer a negative externality, is of interest. Our expectation from the literature is that there will be a decrease in alliance formation with transparency, and hence an increase in meritocratic play and group payoffs, representing an efficiency dividend.

**Question 2:** Who are more likely to form alliances, and do social cues from identification facilitate alliance formation between these players?

A secondary question is whether there are general demographic determinants of alliance choices, and if so, whether these are important in transparency. The cooperation literature suggests that social cues from improved information, such as photographs, can be used as a coordination device (??). In our setup this would mean an increase in the likelihood of back-scratching alliances between certain player types in the transparency treatment compared to the baseline. Even if alliance formation declines on the whole from transparent, an observed increase in the incidence of particular types of alliances will condition any policy guidance offered by these experiments.

## 4 Results

### 4.1 Descriptive statistics

We present in Table 3 selected summary statistics, encompassing patterns of play, payoffs and player characteristics, and in the Figure A.1 of the Appendix we plot individual player earnings for each group by round, visually displaying the complexity of actual alliance behaviour.<sup>4</sup> On average, groups earned 173ECU (or 14%) less than the maximum possible. Meritocratic choices made up a little over half of all choices, and 32% of rounds were in an alliance on average.

Table 3: Summary statistics

	Mean	Max	Min
Group payoff (ECU, all treatments)	1,077	1,224	878
Individual payoff (ECU, all treatments)	180	336	63
Share of meritocratic decisions	0.52	0.96	0.12
Unique alliances/group	1.3	5	0
Share of rounds in alliance	0.32	1	0
Alliance length <sup>†</sup>	6.3	25	2
Group loss	-173	-372	-6
Round of 1st alliance (if any)	10	23	1
Share of groups with any alliance	0.93		
Age	21.7	37	17
Group male share	0.54	0.83	0.33
Happiness (1= V. Unhappy ...5=V. Happy)	3.81	5	1
Political (1= Left ...10= Right)	5.53	10	1
Family wealth (1= Wealthy ...3= Poor)	1.95	3	1
People v skills (1=People...5=Skills)	2.49	5	1
Marital status (1= Partnered)	0.26		
International student (1=Yes)	0.38		
Business student (1=Yes)	0.21		
Club members (1= Yes)	0.56		
Private School (1=Yes)	0.40		

<sup>†</sup> Of groups with any alliance

# Friendships, and racial character statistics are for groups who played the identification treatment. The data was unable to be collected in the completely anonymous treatment groups.

\* Mean rating for each subject by other group members.

For socio-demographic characteristics, survey questions and coding are in Table A.2.

Some other features are worth highlighting. First is the wide range of outcomes in terms of experimental earnings for individuals and groups. The highest group earnings were 39% higher than the lowest group, with such variation the result of some alliances lasting a whole 25 round treatment. For individuals, the highest earning subject made a payoff five times higher than the lowest earner. Given the repeated nature of the game, such divergence is to be expected, and is in accordance with the idea that a social norms of meritocracy *or* back-scratching can emerge in response to early decisions, generating expectations about the play of others in later rounds.

<sup>4</sup>An alliance is observable in these graphs as an alternating step-change in payoffs of two players in a group.



Additionally, we see that players often renege on an alliance, only to later form an alliance with a different player, with groups having up to 5 unique alliance pairs formed in a treatment (with an average of 1.3). This possibility is a new feature in our experimental setup.

In terms of socio-demographics our participants cover a broad sample of university students. There is a large share of international students, 38%, and a high representation of business students, 21%, and a roughly even gender split. In order to look at social cues in the transparency treatment racial appearance is of interest, as it offers one social signal upon which to base decisions about with whom form an alliance. Subjects are classified by the experimenter<sup>5</sup> in categories of Caucasian, Black, Indian and Asian, and are predominantly Asian (46%) and Caucasian (44%). Subjects also self-report their religion, with most subjects reporting being Atheist (42%) or Christian (28%). For groups who played the transparency treatment (20 out of 30) the attractiveness of subjects was rated by others in their group, and they also recorded whether they were friends with other subjects, with only 7% of subjects having a friend in their own group.

## 4.2 Treatment effects

To answer our main research question, we compare indicators of alliance behaviour between the baseline and transparency treatment. Table 4 summarises the mean outcomes of a variety of indicators of interest, including the earlier-defined alliance choice types, and compares the distributions of outcomes in the baseline and treatment. The frequency of these alliance choice types by treatment and round is plotted in Figure A.2 of the Appendix in order for a visual comparison. We are also able to observe the average learning effect by comparing the first and second treatments played by groups in this experiment, which are in the right two columns of Table 4.

Our main indicator of interest, the number of meritocratic rounds, is unchanged between the baseline and treatment, being 0.51 of the total rounds, which is not statistically different from the 0.53 share in the baseline. However, the number of AI choices, which represent attempts to form new alliances, is higher in the transparency treatment. The round of first alliance reduced significantly in the treatment, meaning that despite the more frequent alliance initiation attempts, alliances were in fact established more quickly. This is consistent with the higher number of unique alliance pairs per group in the treatment. Together with the observed increase in AI rounds and unique alliance per group, these results suggest that the effect of identification via photographs in this game is to both facilitate more rapid alliance formation, and also to increase the social pressure to renege on ‘unfair’ long alliances. This evidence does not meet our prior expectation of overall pro-social effects for the whole group from photographs, but does corroborate earlier results on the effect of transparency in ? and ?. It also reveals that the lack of net effect on meritocratic play may not be evidence of the lack of underlying effects on behaviour.

In terms of alliances between particular players of recognised groups, such as race and gender, we find no significant change in the same-group alliances in the treatment, however we see that players who identified each other in the experimental survey as being friends outside of the lab were able to find each other and form alliances, though this occurred in only one group, and a useful statistical test of this effect is therefore lacking. At this level of analysis there is little

---

<sup>5</sup>Classifications were made based on inspection of player photos immediately following the experiment. Players classified as Indian appear to have some South-Asian ancestry, while players classified as Black may have some African ancestry.

Table 4: Treatment effects

	Baseline	Transparency	First Treatment	Second Treatment
Share M rounds	0.53	0.51	0.55	0.50
Share AI rounds	0.23	0.30*	0.30	0.21***
Share AR rounds	0.24	0.20	0.16	0.29**
Share IA rounds	0.34	0.30	0.24	0.41**
Share of groups (any IA)	0.83	0.80	0.80	0.83
Mean alliance length <sup>#</sup>	7.3	4.9*	4.3	8.9***
Round of first alliance <sup>#</sup>	11.3	7.2**	10.7	9.2
Unique alliances/group	1.2	1.6	1.4	1.1
Mean group earnings	1,078	1,076	1,086	1,068
Mean individual earnings	180	179	181	178
Equality (group Gini)	0.15	0.13	0.12	0.16*
Group losses (ECU)	-172	-174	-164	-182
Share same gender alliance	0.48	0.48		
Share same race alliance	0.56	0.31		
Share friends alliance	0.00	0.06		

\*10%, \*\*5%, \*\*\*1% p-value significance of two-sided t-test: baseline compared to transparency; second compared to first treatment.

<sup>#</sup> Of groups where there as any alliance.

Same gender, race and friends alliances are only for the 20 groups who played the Identity treatment where such data was able to be collected.

support for our expectation that players can use photographs to form alliances with others from recognisable social groups, such as race and gender.

Lastly, our comparison of first and second treatments show a learning effect in terms of alliance formation, whereby second treatments had significantly higher AR choice frequency (0.29 share of all rounds instead of 0.16 in the first treatment). Additionally, there are more rounds in alliances (0.41 compared to 0.24), longer alliances (8.9 rounds on average compared to 4.3), and subsequently lower group earnings and greater group inequality.

### 4.3 Who forms alliances?

Our second research question regarding who is more prone to form alliances motivates us to explore the characteristics of players forming alliances, and those who do form them, based on alliance choice types. The alliance play choices comprise the two described earlier; an alliance initiation (AI) is defined to occur if a player  $i$  chooses player  $j$  to receive the payment when player  $j$  does not have the highest productivity in that round and where those players are not already in an alliance, and an alliance reciprocation (AR), is a non-meritocratic reciprocal allocation of the payment in the round immediately following an alliance initiation, or any other non-meritocratic decision within an alliance period.

We wish to know what the characteristics of these alliance players are, for which we estimate a linear probability model of the decision variable  $AP_{igt} \in \{0, 1\}$ , which is either AI or AR, as

$$AP_{igt} = W_i X'_{it} \alpha + v_{gt}$$

in which an individual  $i$  in group  $g$  is said to initiate an alliance at time  $t$  if at that time (s)he is the allocator and makes an AI or AR choice. For most individuals and periods,  $AP_{igt}$  is a missing variable, since the only players who can make an AI choice are the allocators of a round who are not currently in an alliance, and those who can make AR choices must already have been initiated into an alliance.  $X'_{it}$  is a matrix of socio-demographic variables and treatment controls. The error term  $v_t$  is clustered by group  $g$ .

We deal with the unbalanced sample of decisions across players by weighting each decision by the inverse of the frequency of decisions by that player in the vector  $W_i$ , to provide a player-balanced estimation of the vector of coefficients,  $\alpha$ . Table A.3 of the Appendix shows the results of these estimations.

A number of socio-economic variables closely relate to individual AI choices in our sample of subjects. International students were around 13% more likely to initiate an alliance, though not to reciprocate, a result perhaps driven by greater misunderstanding of the game amongst this group.<sup>6</sup> Male subjects were around 11% more likely to initiate alliances, and 20% more likely to maintain them by reciprocating. In terms identifiable characteristics that communicate information in photographs, such as race, Indian participants were more likely to initiate alliances, though the very small number of Indian participants may be contributing to the size of that result. Participants identifying with a particular religion had not-significantly lower alliance initiation activity compared to the non-religious players. Participants from wealthier families and private schools showed no differences in their alliance behaviour. As expected, players who answered our survey question about whether forming an alliance was a fair way to play the game (after they played it) were more likely to have initiated and reciprocated alliances.

Subjects in relationships show a propensity to reciprocate alliances but not to initiate them, being 22% more likely to reciprocate, with a similar pattern found amongst business students, indicating that these groups were better at responding to alliance play than creating it themselves.

In terms of our research question about whether some types of people are more likely to form alliances, we can summarise that non-religious international student males in relationships appear the most ‘alliance-prone’ in terms of their initiation and reciprocation choices.

#### 4.4 Who are in an alliance?

To look in more detail at whether players use social cues from photos to coordinate back-scratching, we first look at the effect of the combined characteristics of players in an alliance, even in the absence of visible signals about the other player. Even though mostly anonymous, the joint characteristics of successful alliances can matter merely as a result of interaction of heterogeneous players in the repeated experimental setting. To this end estimate a Probit model of the odds that person  $i$  is in an alliance with  $j$  at round  $t$ , denoted as  $A_{ijt} \in \{0, 1\}$ , with errors clustered by individual,  $v_{ijt} = e_i + e_{ijt}$ , in the following form:

---

<sup>6</sup>As a robustness check we test the same models using only rounds 10 to 24 of a treatment and compare with results using only rounds 1 to 15. If the result holds more in later rounds compared to earlier rounds we have a clue that either there are more persistent mistakes occurring. We find that in later rounds the coefficient for alliance initiation is higher and significant, while for reciprocation it is lower and remains insignificant, suggesting the possibility that international students both formed alliances, and reneged on, them faster than other players.

$$A_{ijt} = \alpha + X_i' \beta + X_{ij}' \gamma + v_{ijt}$$

where the variables of particular interest are now in  $X_{ij}'$ , denoting joint variables, including whether or not both players are of the same gender, same ethnicity, same age, students in the same degree course, and so forth. The high number of observations comes from the many potential alliance combinations in each round, which for a group of  $n$  players is  $(n^2 - n)/2$ . For each individual in a group we know in each of the 25 rounds whether or not they have formed an alliance with each of the others. Table A.4 of the Appendix shows model results.<sup>7</sup>

We see again that males and subjects in relationships are more likely to be in alliances, as well as those that self-report taking on leadership positions in their social group activities (though this could also be the result of their feelings of leadership that arise from being in an alliance in the experiment). Considering the joint characteristics of potential alliance partners, we see negative coefficients in political views, wealth, age, happiness and marital status indicating a mixing of these players in alliances pairs. When using the data from groups who played the transparency treatment only (20 of the 30 groups), the change in magnitude and direction of the fitted coefficients for some joint variables, such as happiness (-0.23 to 0.11), clubs (0.05 to 0.27), and business students (-0.09 to 0.04), suggests that the treatment may have had some influence of the joint characteristics of alliance partners.

## 4.5 The effect of transparency

The previously estimated effects of joint characteristics of alliance partners may be due to an accident of group composition. We now take the additional step of conditioning these joint characteristics on the transparency treatment to answer our question of whether social cues from identification of players facilitates the formation of certain alliances. To that end we use the data from the 20 groups who played the transparency treatment to estimate the Probit model with errors clustered by individual:

$$A_{ijt} = \alpha + X_{it}' \beta + X_{ijt}' (\gamma + Z_{it} \delta) + v_{ijt}$$

where  $Z_{it}$  now includes indicators for whether we are looking at the baseline or identification treatments. Results are in Table A.5 of the Appendix. The coefficients on particular joint player characteristics interacted with the transparency treatment show that identification might increase the ability of some ‘alliance-prone’ individuals to find each other, while deterring others.

Subjects in relationships are more likely to be in an alliance (marital coeff. 0.27), though in terms of joint characteristics they are more likely to match with similar marital types only in the transparency treatment. Wealthy (and non-wealthy) subjects follow a similar pattern of being more likely to be in alliances with each other only in transparency. The same is true with business students (coefficient of joint interaction 0.37, and coefficient of joint business -0.09). Transparency did not facilitate common gender or age alliances, against our expectation, nor

---

<sup>7</sup>The low Pseudo  $R^2$  and low marginal average effects (for example the marginal average effect is 1% for the marital status dummy variable which has a coefficient of 0.23 in the Probit model) are due to the large counterfactual sample size. Of the 45,000 observations, alliances can at most occur in 3,000 observations in the case that one out of the 15 potential alliance pairs in a group occurred in every round of the full experiment. The significance and direction of the effect is the primary interest in this analysis.

were alliance pairs with common race or religion more frequent in the transparency treatment.<sup>8</sup> Overall however, these results are rather inconclusive, as only some common player features responded to transparency, which could be by chance.

## 4.6 Robustness checks

A number of checks are made to examine the robustness of these results, and our interpretations. First, we look at the problem of interpreting the finite experiment as a proxy of an infinitely repeated one. We check the frequency of alliance breakdown near the end of each treatment with a Wilcoxon signed rank test of the changes to alliance frequency in the final three rounds of the second treatment played by each group. We find a significant increase in alliance breakdown in only the final three rounds or less of a group treatment played (p value = 0.03 for third last round). To control for these ‘end-of-game’ effects we reproduced the above analysis stripping out the choices made in the final three, four and five rounds, and also the first five rounds to account for a degree of learning. In all cases the direction and magnitude of regression coefficients are similar, and our result thus appear relatively robust to such effects.

One survey question asked “In your opinion, forming an alliance in this game is...” with multiple choice answers from never justifiable to always justifiable. We find a large and significant correlation between a subject’s actual alliance formation and their reported justifiability. Table 5 shows in the first two columns this relationship, and notably, the effect of transparency is to further increase a subject’s justification of their alliance choices. This suggests either a) a degree of *ex post* rationalisation of alliance behaviour, which is amplified by transparency, or b) that subjects who see alliance play as fair are more likely to form alliances under transparency.

Table 5: Subject rationalisation of alliances

	Justified		Conceal	
Total alliance rounds	0.08***	0.07***	0.01*	0.01
Total alliance rounds $\times$ Transparency		0.03**		0.01**
R <sup>2</sup>	0.15	0.17	0.02	0.02

Number of subjects in the data is N=180.

It is also of interest to see whether subjects observed to be in an alliance attempted to conceal their alliances under transparency. Our survey at the end of the experiment asked “Were you part of an alliance that excluded the player with highest productivity number at any point throughout the whole game?” The multiple choice answer allowed subjects to reveal in which treatment they formed an alliance, if any. We match these answers with our own alliance specification to determine the degree of honesty in self-reported alliance formation, with results in Figure 1. Assuming errors in answering are the same for those who did and did not have alliances, we see a clear discrepancy, with more dishonesty (i.e. concealment) by subjects who had alliances. We also test the effect of transparency on how much subjects concealed their alliances by lying on the survey, reporting the results of linear regressions of dishonesty on alliance play in the last two columns of Table 5. It is clear that the concealment observed is mostly dependent on the

<sup>8</sup>Players who reported being friends outside of the experiment were more likely to find each other in the transparency treatment. Though our statistical analysis is limited due to the very low prevalence of friends being assigned into the same experimental group, and the zero number of cases of friends forming alliances outside of the transparency treatment.

number of rounds played under transparency, rather than when anonymous, suggesting that our transparency treatment prompted some level of self-deception, or guilt, from participants who acted ‘unfairly’.




Figure 1: *Frequency of subject self-reported alliances according to observed alliances*

A high degree of path-dependence has been observed in this experimental setup (?), and we test for whether subject outcomes are sensitive to early round play. We find a slight negative correlation between a subject’s first round being allocator and their total payoff, though first round allocators are significantly more likely to be the highest earning subject in a group.<sup>9</sup> Groups themselves appear to learn social norms over time, with the number of M choices in the first 10 rounds of each treatment strongly predicting the number of M choices in the last 10 rounds.<sup>10</sup>

In terms of the alliance behaviour, we should expect a degree of rationality to be adhered to, in that initiating an alliance with the second-highest productivity subject would be preferred to initiating an alliance with players of lower productivity. Only 32% of initiation attempts were to the second highest productivity player (AI to Prod. 1 = 24%, Prod. 2 = 19%, Prod. 3 = 26%, Prod. 4 = 32%). The greater initiation towards players with a productivity of 1 suggests that payoff-maximisation behaviour may be dominated by the process of establishing alliances, and that initiating alliances with lower productivity players may provide a more costly, but reliable, signal of the intention to form an alliance.

We also look for more general evidence that the observed effect of transparency leading to faster alliances, seen in Table 4, is due to overall social cues from photographs. To that end, for each observed alliance, we generate a metric of alliance pair similarity based on whether the alliance partners share the same features based on the following characteristics: international student, gender, age,<sup>11</sup> business student, and religion. From this we count the number of similar features of each alliance pair to get similarity measure (from zero to five), and count how many rounds from the start of the treatment to the alliance (or from a previous alliance pair) to get a measure of alliance formation time. We then fit a linear model to see whether the joint interaction of transparency and similarity, shows the expected relationship with alliance formation time, and could partly be driving our results.<sup>12</sup> Table 6 shows the results of this modelling exercise, with

<sup>9</sup>p value = 0.01 of Mann-Whitney-Wilcoxon test that the distribution of top-ranked first allocators differs from the expected one-sixth result.

<sup>10</sup>In a linear estimation of group meritocratic decisions,  $M_{16-25} = \alpha + \beta M_{1-10} + \epsilon$ , we find that  $\alpha = 1.43$  ( $p = 0.01$ ),  $\beta = 0.62$  ( $p = 0.00$ ).

<sup>11</sup>Same age is within 4 years

<sup>12</sup>We also checked whether the observed alliances between self-reported friends lasted longer in the transparency treatment. There were only two friendship alliances observed, both in transparency, which lasted only 6 and 2 rounds, which was lower than average.

a significant interaction term in Model 3 suggestive that knowledge of closer similarity between players speeds up their alliance formation.

Table 6: Results of alliance formation time models

Model	1	2	3	4
Transparency	-2.24	1.38		1.69
Similarity	-0.78	-0.12	-0.35	
Similarity $\times$ Transparency		-1.45	-0.98**	-1.57*
R <sup>2</sup>	0.05	0.07	0.06	0.07

Number of alliances in data is N=77. Dependent variable is number of rounds to form alliance.

## 4.7 Summary of results

**Question 1:** Does identifying players increase meritocratic play and increase group payoffs?

The transparency treatment in our setup, which provided photographs of players, resulted in no observable increase in meritocratic play or group payoffs. Moreover, alliance initiation choices increased, as did the number of alliance pairs formed per group. However, we also observed a decrease in the average length alliance, which could indicate that a fairness norm was present but was insufficient to shift groups towards the meritocratic equilibrium. Our choice models reinforce these results.

**Question 2:** Who are more likely to form alliances, and do social cues from identification facilitate alliance formation between these players?

On this question we find some suggestive results. In terms of the overall tendency to form alliances in our setup, men are more likely to both initiate and reciprocate alliances, which is consistent with the results of ?, who found that men are more confident that bribes will be reciprocated. Business students and subjects in relationships showed no tendency towards higher rates of alliance initiation, though they did engage in much higher levels of alliance reciprocation.

Our results suggest that transparency through photo identification did change the composition of alliances, though exactly how is not fully clear. Transparency facilitated faster alliance formation in general, and between overall similar subjects as we have defined them. However, some player characteristics one would expect to facilitate alliance formation had no effect, such as age, gender and race, or had the opposite effect, such as having the same religion, being social club members.

Transparency also affected the way subject’s rationalised their choices, leading to them to both justify alliance play, and conceal it from the experimenter in survey responses.

## 5 Discussion and conclusion

In this paper we have studied the effect of transparency as an anti-corruption policy in an experimental game designed to capture the costly back-scratching process underlying much corruption. In our setup, such a process entails a violation of meritocracy in order to establish an alliance

based on reciprocation of discretionary favours, and in which the costs of this process are directly borne by those players not chosen to be in an alliance. Our main results are that introducing a transparency policy, by identifying with photographs those who are back-scratching at the expense of others, generated no overall improvement in group payoffs compared to the baseline treatment where decisions were anonymous. Alliances in the treatments formed more quickly, but were shorter, with more different alliance pairs per group. To understand why that is the case, we looked at the characteristics of players forming alliances with each other, to see if the story that there are social cues in photographs that might support costly back-scratching alliances is a plausible one. Our data is consistent with this story, though the specific elements by which social cues work to change behaviour, in which direction, and why, remains unclear.

The unique part of our setup is that there are conflicting pro-social choices; either cooperate as an alliance pair, or cooperate meritocratically with the whole group. It may be the case that identification does promote cooperation in general, but in our setup improved both types of cooperation simultaneously. This is consistent with our result of faster alliance formation, but short alliances, under transparency.

Understanding our results in terms of preferences is a challenge. While the lack of effect of transparency conforms to the expected utility result based solely on payoffs, such an approach cannot predict the observed decrease in alliance length and increase in the number of alliance pairs per group was observed. Nor does it predict the frequent use of the most costly alliance initiation attempts (allocating to the lowest productivity player, rather than the second highest). Together with the path dependence of group behaviour observed, this suggests that in general the effect of transparency is likely to be sensitive to initial social conditions which provide the common knowledge necessary to reason about strategy choice.

A key remaining question is how our results might inform the design of institutions, and how transparency policies in public institutions need to balance these competing coordination effects. ? proposed to mandate the anonymity of political donors in order to disrupt the trade in political favours. Just as we do, they argued that if donors are anonymous, politicians have one less signal about who to trade favours with, though of course this signal disruption will only work if donors have no credible means of showing what they have donated.

Another example where anonymity, rather than transparency, might reduce corruption is where highly discretionary decisions will be made by a public committee, such as when a secret public jury decides on allocating government contracts. Unable to see who decides on their fate, firms seeking favourable treatment will be unable to focus on exploiting relationships and establishing expectations of reciprocation with committee members, and instead may be forced to compete on the quality of their tenders.

Given our experimental results, we are skeptical that transparency will reduce corruption in cases where discretionary favours can be made without fear of punishment, and where transparency only helps in the coordination of favours: transparency might expose corruption, but it does not by itself stop it. Transparency might be a stepping-stone to reducing corruption though if it leads the majority to impose additional measures, such as punishment.

## References



Appendix

# Baseline

## YOU ARE THE ALLOCATOR

Please choose who you would like to allocate the \$25 payment to this round by selecting their shape, then click Allocate.

You are

Phase 1

Round 4

Productivity number	2	1	3	5	4
Player					

ALLOCATE

## YOU ARE AN OBSERVER

Your position in the game for this round is shown below

You are

Phase

1

Round

1

The allocator is

Productivity number	2				
Player					

Click continue when you have understood the game for this round

Continue

# Transparency

## YOU ARE THE ALLOCATOR

Please choose who you would like to allocate the \$25 payment to this round by selecting their shape, then click Allocate.

Phase 1

Round 3

Productivity number	3	2	3	1	4
Player					

Allocate

## YOU ARE AN OBSERVER

Your position in the game for this round is shown below

Phase 1

Round 2

The allocator is

Productivity number	4				
Player					

Click continue when you have understood the game for this round

Continue

Table A.1: Experiment screenshots (top row is Allocator, bottom row is other players)

Table A.2: Survey questions

Variable	Survey question	Response type
People v Skills	Reflecting on your own life experience, how important do you think it is to know the right people versus having the rights skills in order to succeed?	1= Mostly People ... 5 = Mostly Skills
Club	Are you a member of a student club, society, or sports club?	0 = No, 1 = Yes
Leader	Did your parents encourage you to be in leadership positions at school?	0 = No, 1 = Yes
Wealth	Would you say that your family is ...	1 = Wealthy, 2 = Average, 3 = Poorer than average
Religion	I identify as a...	Christian, Muslim, Buddhist, Hindu , Atheist, Other
Politics	In political matters, people talk of ' the left ' and ' the right '. How would you place your views on this scale, generally speaking?	1 = Left ... 10 = Right
Fair	In your opinion, forming an alliance in this game is...	1 = Never Justifiable ... 5 = Always justifiable
Friend	Before today, when did you last communicate with this person?	1 = Yesterday, 2 = Last week, 3 =In the last month 4 =In the last year, 5= Over a year ago, 6= Never
Beauty	How attractive do you think the average Australian would rate this player?	1 = Very unattractive... 7 = Very attractive
Happy	All things considered in your life, how happy would you say you are usually?	1 = Very unhappy ... 5 = Very happy
Age	Please enter your age in years	Numerical whole year variable
Gender	I am	0 = Female, 1 = Male
Inter. stud.	Are you an international student?	0 = No, 1 = Yes
Marital	I am	0 = Single, 1 = Partnered, 2 = Married
Bus. stud.	Please enter the name of you current degree	Free form text input. Experimenter coded.

For the analysis partnered and married responses are groups together to create a binary single or not-single variable.




Figure A.1: Accumulated player payoffs, by group and round, with transparency treatment shaded

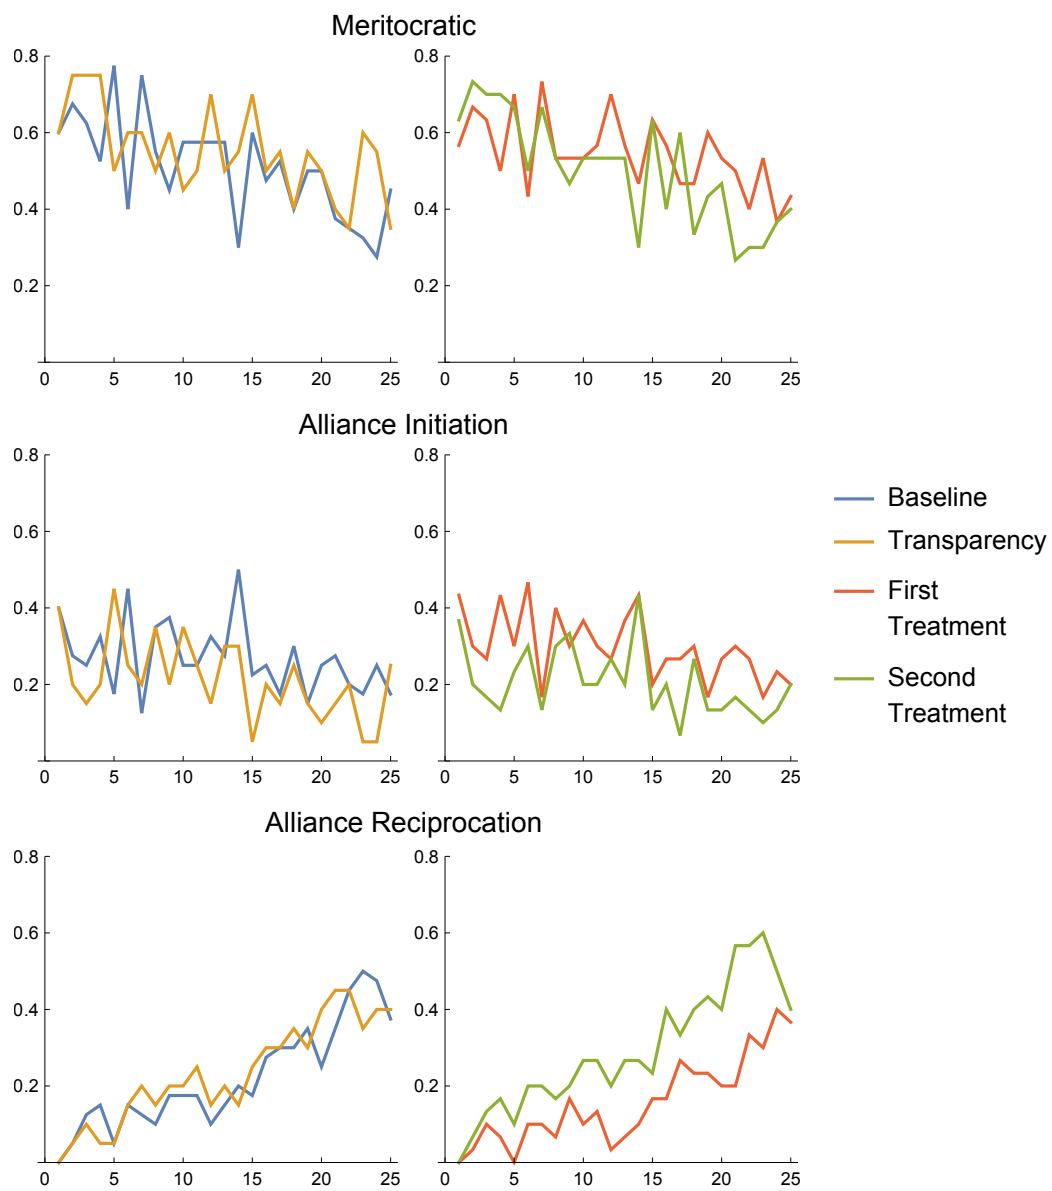


Figure A.2: Choice frequency by treatment and order

Table A.3: Results of alliance choice models

Alliance initiation						Alliance reciprocation				
Age	-0.01				-0.01	-0.00				-0.00
Gender	0.06				0.11**	0.17*				0.21*
Beauty	-0.04				0.03	0.06				0.06
Black	-0.09				-0.07	-0.05				-0.12
Indian	0.23**				0.32**	0.01				-0.03
Asian	0.08				0.07	0.06				0.09
Buddhist		-0.10			-0.12		-0.15			-0.16
Christian		-0.02			-0.02		-0.07			-0.02
Hindu		-0.00			-0.19		0.01			0.09
Inter. stud.			0.11**		0.13**			-0.04		-0.00
Bus. stud.			0.00		0.02			0.23*		0.29**
Marital			0.10		0.10*			0.22*		0.23**
Happy			0.00		0.02			-0.04		-0.02
Politics			-0.01		-0.02			-0.00		-0.00
Wealth			-0.02		-0.00			0.08		0.08
Private sch.				0.02	0.05				-0.01	0.02
Fairness				0.07***					0.14***	
Clubs				-0.03	-0.05				-0.01	-0.03
Skills				-0.00					-0.05	
Leader				-0.01					0.10	
Photo	0.02	0.05	0.04	0.04	0.02	-0.05	-0.02	0.01	-0.07	-0.02
Order	-0.04	-0.02	-0.02	-0.03	-0.03	0.15	0.18*	0.15	0.14	0.12
Obs.	1,090	1,090	1,090	1,090	1,090	544	544	544	544	544
N (AP= 1)	377	377	377	377	377	338	338	338	338	338
R <sup>2</sup>	0.03	0.01	0.03	0.03	0.07	0.07	0.05	0.10	0.20	0.17

# Models weighted by inverse of individual player decision frequency.  $p$  values from Wilcoxon signed-rank test \* < 0.10, \*\* < 0.05, \*\*\* < 0.01. Order is 1 for second treatment played by a group, 0 for first treatment.

Table A.4: Individual and joint characteristics of alliance partners

	All groups		Photo only
Age	-0.01	-0.02*	-0.02*
Gender	0.26***	0.22**	0.15
Beauty	-0.00	0.01	-0.02
Black	0.08	0.04	-0.01
Indian	-0.02	0.10	-0.24
Asian	0.04	0.04	0.02
Buddhist	-0.17	-0.15	-0.17
Christian	-0.13	-0.11	-0.00
Hindu	-0.14	-0.12	0.40
Happy	0.02	0.04	0.03
Inter. stud.	-0.01	0.05	0.09
Bus. stud.	0.04	0.01	0.06
Marital	0.23***	0.21**	0.25**
Wealth	0.08	0.10	0.25**
Private Sch.	0.01	-0.00	0.02
Club	0.15*	0.16*	0.18*
Leader	0.23***	0.22**	0.22**
Age <sub>ij</sub>		-0.22***	-0.28***
Gender <sub>ij</sub>		0.11	-0.01
Beauty <sub>ij</sub>		-0.10	-0.10
Race <sub>ij</sub>		0.05	0.07
Religion <sub>ij</sub>		0.14*	0.13
Inter. stud. <sub>ij</sub>		0.13	0.12
Bus. stud. <sub>ij</sub>		-0.09	0.05
Marital <sub>ij</sub>		-0.06	-0.15
Happy <sub>ij</sub>		-0.24**	0.11
Politics <sub>ij</sub>		-0.14*	-0.16*
Wealth <sub>ij</sub>		-0.08	-0.19**
Private Sch. <sub>ij</sub>		0.05	0.13
Club <sub>ij</sub>		0.05	0.27***
Leader <sub>ij</sub>		0.08	0.02
Friends <sub>ij</sub>		0.11	0.19
Obs.	45,000	45,000	30,000
N (A= 1)	974	974	638
Pseudo $R^2$	0.01	0.02	0.03

$p$  values \* = 10%, \*\* = 5%, \*\*\* = 1%.

Joint common characteristics are all binary variables: Wealth<sub>ij</sub> is same response of alliance pair on 3 point family wealth scale. Age<sub>ij</sub> is age of alliance pair within 4 years. Politics<sub>ij</sub> is within 2 on a 10 point scale. Friends<sub>ij</sub> is having any identified relationship between alliance partners from the survey. Beauty<sub>ij</sub> is within 1 of each other on a 7 point scale. Happy<sub>ij</sub> is alliance pair within 1 on 5 point scale. All others joint variables are 1 is binary responses for alliance pairs are the same.

Table A.5: Identification effects on alliance partner choices

<i>In Alliance</i>			
Age	-0.02*		-0.03*
Gender	0.11		0.15
Beauty	0.01		-0.02
Black	0.19		-0.02
Indian	-0.04		-0.23
Asian	-0.00		0.07
Buddhist		-0.18	-0.15
Christian		-0.01	0.02
Hindu		0.16	0.43
Inter. stud.		0.09	0.08
Bus. stud.		0.02	0.13
Marital		0.27***	0.28***
Wealth		0.24**	0.25**
Private sch.		-0.02	-0.01
Club		0.24**	0.20*
Leader		0.22**	0.21**
Age <sub>ij</sub>	-0.32***		-0.17
Gender <sub>ij</sub>	0.09		0.02
Beauty <sub>ij</sub>	-0.21**		-0.06
Race <sub>ij</sub>	0.21**		0.19
Religion <sub>ij</sub>		0.61***	0.54***
Inter. stud. <sub>ij</sub>		0.07	0.00
Bus. stud. <sub>ij</sub>		-0.07	-0.09
Marital <sub>ij</sub>		-0.25*	-0.27**
Wealth <sub>ij</sub>		-0.52***	-0.52***
Private sch. <sub>ij</sub>		0.15	0.12
Club <sub>ij</sub>		0.66***	0.59***
Leader <sub>ij</sub>		-0.04	-0.04
Friends <sub>ij</sub>	0.06		
Age <sub>pij</sub>	0.01		-0.18
Gender <sub>pij</sub>	-0.08		-0.08
Race <sub>pij</sub>	-0.25*		-0.29
Beauty <sub>pij</sub>	0.12		-0.08
Inter. stud. <sub>pij</sub>		0.13	0.24
Private sch. <sub>pij</sub>		-0.01	0.04
Bus. stud. <sub>pij</sub>		0.28*	0.37**
Religion <sub>pij</sub>		-1.16***	-1.05***
Wealth <sub>pij</sub>		0.66***	0.73***
Marital <sub>pij</sub>		0.19	0.24
Club <sub>pij</sub>		-0.61***	-0.56***
Leader <sub>pij</sub>		0.07	0.10
Obs.	30,000	30,000	30,000
N (A= 1)	638	638	638
R <sup>2</sup>	0.005	0.057	0.070

\* = 10%, \*\* = 5%, \*\*\* = 1% significance. Dependent variable is in alliance (IA) for all possible player pairs in each round. Subscript  $p,i,j$  indicates interaction term of joint player pair characteristic and the transparency (photo) treatment.