

Leadership under the shadow of the future: Intelligence and strategy choice in infinitely repeated games

Drouvelis, Michalis; Pearce, Graeme

European Economic Review

DOI:

<https://doi.org/10.1016/j.euroecorev.2023.104372>

Published: 01/02/2023

Peer reviewed version

[Cyswllt i'r cyhoeddiad / Link to publication](#)

Dyfyniad o'r fersiwn a gyhoeddwyd / Citation for published version (APA):

Drouvelis, M., & Pearce, G. (2023). Leadership under the shadow of the future: Intelligence and strategy choice in infinitely repeated games. *European Economic Review*, 152, Article 104372. <https://doi.org/10.1016/j.euroecorev.2023.104372>

Hawliau Cyffredinol / General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal ?

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Journal Pre-proof

Leadership under the shadow of the future: Intelligence and strategy choice in infinitely repeated games

Michalis Drouvelis, Graeme Pearce

PII: S0014-2921(23)00001-6
DOI: <https://doi.org/10.1016/j.eurocorev.2023.104372>
Reference: EER 104372

To appear in: *European Economic Review*

Received date: 6 June 2022
Revised date: 28 November 2022
Accepted date: 2 December 2022



Please cite this article as: M. Drouvelis and G. Pearce, Leadership under the shadow of the future: Intelligence and strategy choice in infinitely repeated games. *European Economic Review* (2023), doi: <https://doi.org/10.1016/j.eurocorev.2023.104372>.

This is a PDF file of an article that has undergone enhancements after acceptance, such as the addition of a cover page and metadata, and formatting for readability, but it is not yet the definitive version of record. This version will undergo additional copyediting, typesetting and review before it is published in its final form, but we are providing this version to give early visibility of the article. Please note that, during the production process, errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

© 2023 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

Leadership under the shadow of the future: intelligence and strategy choice in infinitely repeated games*

Michalis Drouvelis^{*,*} and Graeme Pearce[◇]

^{*}*University of Birmingham*

^{*}*CESifo, Munich*

[◇]*Bangor University*

Abstract

We examine the impact of intelligence on decision making in an infinitely repeated sequential public goods game. Using a two-part experiment, we collect data on subjects' intelligence and a wide range of preference characteristics, and match these to their full contingent strategy profiles. We find that leaders are less likely to play a free-riding strategy as their intelligence increases. Followers are less likely to play a grim-trigger strategy as intelligence increases. Performing simulations using players' strategies, we find that groups contribute more and are more profitable as intelligence increases. Our results have implications for the design of policies promoting group success.

Keywords: intelligence, IQ, leadership, infinitely repeated games, strategy elicitation, experiments

JEL codes: H41, C72, C92

1 Introduction

Infinitely repeated games characterise key aspects of many of our everyday relations. In a range of these interactions individuals face incentives in which personal interest and group benefits are in conflict. Existing evidence from the infinitely repeated games literature suggests that such interactions may be beneficial to curb free-riding behaviours (Dal Bó & Fréchet 2019), allowing players to respond to past actions in the future, and therefore enable the enforcement of efficient or “reasonable” outcomes (Wen 2002). Although strategies in infinitely repeated simultaneous-move games have been well studied (Aoyagi & Fréchet 2009, Breitmoser 2015, Dal Bó & Fréchet 2011, Fudenberg et al. 2012, Romero & Rosokha 2018), the strategies that individuals select in games in which one player assumes a leadership role, or takes the first

*We thank Guillaume Fréchet, Louis Putterman, Eugenio Proto and Emanuel Vespa for helpful comments and suggestions. We acknowledge financial support from the University of Birmingham, and thank Hamideh Mohtashami for helping conduct the experiments. This experiment was pre-registered at the AEA RCT registry, trial number AEARCTR-0004988.

action, surprisingly have received limited attention. In situations where personal and collective interest are in conflict, but where there exists the opportunity for profitable cooperation, the success of a group often depends upon a leader's ability to solve a complex social problem, or find the solution to a complex social dilemma (De Cremer 2006, Choi & Mai-Dalton 1998). As highlighted by Ghidoni & Suetens (2020), as well as Kartal & Müller (2021), sequentiality reduces the strategic risk for a player who moves second. A follower can therefore potentially reap the benefits of cooperation, and avoid being betrayed by a leader, if and only if the leader initiates cooperative behaviour. If the leader is cognitively sophisticated enough to be able to understand this, they too face reduced strategic risk in comparison to simultaneous move games. It then seems plausible that intelligence is a key determinant of cooperation in sequential games. In this paper, we fill the gap in the literature by examining if leaders' and followers' intelligence impacts behaviour in infinitely repeated sequential public good games.¹

The social dilemmas that leaders need to solve are cognitively demanding, and require them to think strategically in order to try and predict the behaviour of their followers (Rustichini 2015) and act appropriately. They must take their own and their followers' incentives and beliefs into account when choosing what action to take, or which strategy to implement (Costa-Gomes et al. 2001). Kosfeld (2020) highlights three criteria a leader's strategy must satisfy in order to be successful. First, they have to place trust in the motivated in order to initiate motivation; second, they have to incentivise cooperation with rewards, and punish those who are not motivated to cooperate in order to encourage it; and finally, try and attract those followers who respond to these incentives. In support of Kosfeld (2020), Gächter & Renner (2018) highlight the importance that leaders play in managing followers' beliefs in order to keep followers motivated, studying how initial actions, and beliefs about actions, are crucial for cooperation in the future. In interactions that are repeated infinitely, or indefinitely, this is particularly relevant, as leaders must consider how their actions affect the decisions of followers in the current interaction, but also in all future interactions. Given the previous literature examining the role of intelligence in decision making (Costa-Gomes et al. 2001, Frederick 2005, Proto et al. 2019), it seems reasonable to predict that those leaders that are the most cognitively sophisticated, and therefore the most able to understand the benefits to be accrued in the future from their actions today, should be best placed to choose the most successful strategies.

Although a hypothesised link between leadership intelligence, cooperation and efficiency seems sensible, it is important to acknowledge that the evidence on the importance of leadership in social dilemmas is mixed. For example, Rivas & Sutter (2011) find that leadership has a strong positive effect on cooperation in groups when endogenized, although the effect is more muted when imposed exogenously. Haigner & Wakolbinger (2010) and Cappelen et al. (2016) corroborate this finding, and show that the endogenous selection of leaders has positive effects. In contrast Sahin et al. (2015) report evidence from the laboratory that leaders who set an example, and those who use messages in order to try and promote cooperation among group members, are both highly ineffective at increasing cooperation and efficiency. They suspect this may be a consequence of the parameters of their experiment. This suggests that the incentive structure of the dilemma and leadership effectiveness interact. In a similar laboratory setting, Figuières et al. (2012) report evidence that any positive effects of leadership vanish when the dilemma is played repeatedly, and the leadership role is randomised in each repetition. The

¹Following Dal Bó & Fréchette (2018), we use the terms 'infinitely' and 'indefinitely' repeated games interchangeably throughout the paper.

mixed evidence on the role and importance of leadership for cooperation highlights the need to understand more clearly the dimensions and characteristics that make leaders successful.

The purpose of this paper is to examine the role that intelligence plays in the strategy choices of leaders and followers in social dilemmas. Specifically, we examine the extent to which a leader's intelligence links to the criteria outlined by Kosfeld (2020). This is done using a novel two part experiment designed in the spirit of work conducted by Selten et al. (1997), Dal Bó (2005) and Proto et al. (2019). In one part subjects complete a number of incentivised tasks from which we elicit a comprehensive range of economically significant preference characteristics, and obtain a variety of demographic measures and personality scores using questionnaires. We elicit social preferences using the equality equivalence test (Kerschbamer 2015), risk attitudes following Holt & Laury (2002) and provide a measure of intelligence by using the Raven Test (see Foulds & Raven (1950)). Although what constitutes intelligence and how to measure it is hotly debated (see Sternberg & Kaufman (2011) for a comprehensive recent overview and Burke (1958) for an older discussion), the Raven Test has been used extensively in research in psychology, as a tool in hiring, the military (Burke 1958, Sundet et al. 2004) and education to examine an individual's problem solving ability, abstract reasoning, or what educational psychologists call fluid intelligence (Cattell 1963).² Descriptively, we categorise subjects using this measure into what we call *Low Raven*, if they score below or equal to the average Raven score and *High Raven* if they score above. However, our analyses uses Raven score *per se*, ruling out our results being driven by what might be seen as an arbitrary categorisation. Subjects are also never told their own or other subjects' Raven scores, and are unaware that the experimental focus is on their score in this test. This rules out any sort of status effects driving our results (Kumru & Vesterlund 2010, Jack & Recalde 2015).

In another part, subjects play a two player infinitely repeated sequential public goods game, in which they are randomly assigned to be either a first mover (*Leader*) or a second mover (*Follower*). *Leaders* first decide how much of their endowment to contribute to the public good, *Followers* observe this, and then decide how much of their endowment to contribute. Following Roth & Murnighan (1978) we induce an infinitely repeated game by repeating the game indefinitely, randomly continuing the game at the end of each period of play. This game has features that closely resemble the dilemmas faced by leaders in organizational contexts and a wide range of applications (e.g. employer employee relations, borrower-lender relations, trade) (Ghidoni & Suetens 2020). In addition, the sequential nature of the game provides us with the opportunity to observe how *Leaders'* strategies influence the actions of *Followers* and how outcomes evolve as the interaction is repeated. The indefinitely repeated nature of the game is a distinguishing feature of our design, and the study of leadership in this setting has so far been neglected, making this a unique contribution to the literature.

Another feature of our design is that we elicit subjects' full strategy profiles for the indefinitely repeated game following the approach of Axelrod (1980), Selten et al. (1997) and Dal Bó & Fréchette (2019). This approach provides a number of advantages, which in turn provide contributions to the growing literature examining leadership, and that which seeks to examine the effect of intelligence on economic decision making. First, we examine strategy choices before any type of learning has taken place, and as such can rule out learning dynamics within the experiment as an explanation for the differences conditional on intelligence. Learn-

²Cattell (1963) provides a discussion of the definitions and differences between what are widely regarded as the two different types of intelligence, fluid and crystallised intelligence.

ing dynamics have been shown to be important for equilibrium selection in range of games, including cognitively demanding repeated social dilemmas (Macy & Flache 2002) and market games (Bosch-Rosa et al. 2018), with evidence suggesting there exists a strong link between intelligence and learning dynamics (Chen 2015). For example, Gill & Prowse (2016) examine the micro-processes that drive differences in behaviour between those of high and low intelligence. They find that those with high cognitive ability converge more frequently to equilibrium play and earn substantially more than those with low cognitive ability in p -beauty contest games. However, we still know little about the pure impact of intelligence on strategy choices when learning dynamics are muted.

Second, as we observe full strategy profiles, we can rule out differences in beliefs as driving the differences in most of the observed behaviour conditional on intelligence; although *Leaders* must take an initial action, all strategies are then defined for all possible contingencies, so subjects do not have to choose strategies based on their expectations about what the other player will do. Previous work has shown this to be important, especially for leaders, who are regarded as ‘belief managers’ in sequential social dilemmas (Gächter & Renner 2018). The perfect monitoring environment we choose further simplifies the identification of strategies being chosen, and enables us to use cluster analysis to group strategies together that share common characteristics. Third, our design accommodates us to conduct simulations of interactions between the strategies of all *Leaders* and *Followers*, and thus we can consider how the intelligence composition of groups influences how cooperative and successful they are. This is akin to the computerized tournaments conducted by Axelrod (1980). This approach means we can consider how successful *Leaders* are with respect to the population of subjects, not just in a single interaction in a single experimental session. Fourth, the strategy elicitation method we use eliminates the issue of identifying strategies econometrically and dealing with the challenge of some histories of play having only few observations, or a small number of realised periods of play. Finally, learning which strategies *Leaders* and *Followers* actually use is of interest for a variety of reasons, such as informing future theoretical work, understanding the characteristics of individuals who play strategies regarded as best responses, and identifying the environments in which we might expect cooperation to emerge.

We report a number of observations. First, we observe that around 50% of subjects play Tit-for-Tat type strategies regardless of Raven score: subjects are willing to contribute to the public good as long as the other player also contributes. This is similar to the proportion of subjects playing this strategy previously found in the finitely repeated games literature (Fischbacher et al. 2001). However, it is important to note that a conditionally cooperative strategy is part of an equilibrium in the current experiment, but that is not the case in finitely repeated games or in one-shot games. Kosfeld (2020) highlights these strategies as being crucial for leadership success. We also find that *Leaders* are less likely to play a free-riding strategy as their intelligence increases: a one point increase in Raven score reduces the probability that a *Leader* is a Free-Rider by around 1%. We report no significant differences in beliefs, or Period 1 play in *Leaders* conditional on Raven score, ruling out disparities in beliefs and first period play as driving any of our observations (Gächter & Renner 2018). With regard to *Followers*, we find that higher levels of intelligence are associated with higher levels of contributions in the first period of play. Intelligence also lowers the probability that the *Follower* plays a strategy that is similar to Grim Trigger, a well studied strategy identified in the prisoner’s dilemma

literature (Axelrod 1980).³

Using each subject's unique strategy profile to simulate over 9000 unique interactions, we find that groups comprised of subjects with above average intelligence (*High Raven*) *Leader* and *Follower* contribute significantly more to the public good and make significantly higher profits than groups where both players are below average intelligence (*Low Raven*). We find that *Leaders'* Raven score increases the earnings and contributions of *Followers*, and that *Followers'* Raven score increases the earnings and contributions of *Leaders*. However, the effect of the *Leader* on the *Follower* is estimated to be significantly larger than *Follower* on the *Leader*. We also find an interesting interaction effect between Raven scores, with the *Follower's* Raven score having a greater impact on the *Leader's* contributions as the *Leader's* intelligence increases.

Our paper makes a number of contributions. First, we find evidence to suggest that intelligence can play an important role for strategy choice in sequential social dilemma games, influencing the probability that free-riding and grim trigger strategies are selected. We show this in an indefinitely repeated setting, a context that closely maps to the dilemmas faced by individuals in organizations in terms of incentives, and a number of other relevant applications (e.g. employer employee relations, borrower-lender relations, trade). Finally, our results complement and extend the literature that examines the implications that intelligence has for economic behaviour and outcomes, and our results have implications for the solutions to organizational problems. Specifically, if an organization wants to minimise the number of free-riders and maximise the number of conditional cooperators, Raven tests could be employed to that end.

The remainder of this paper is organized as follows. Section 2 outlines our experimental procedure and design, Section 3 presents the results from the experiment and simulations, and Section 4 concludes.

2 Experimental procedure and design

The experiment was designed to examine how intelligence impacts the strategy choices of *Leaders* and *Followers* in sequential social dilemmas. To do this we follow Proto et al. (2019), and design a two part experiment. In Part A, subjects make a number of decisions that enable us to measure, and therefore control for, their intelligence level and a range of economically significant characteristics. In Part B, subjects then define a full contingent strategy profile for a two player sequential public goods game that is indefinitely repeated. No feedback about earnings or outcomes from either Part of the experiment was given until all decisions and questions had been made and answered. The order in which participants completed Part A and Part B was randomised in order to control for any link between decision order and behaviour.

³Throughout the paper, we refer to what the some papers in the public goods game literature might call 'conditional cooperation' as Tit-for-Tat type strategies. This is done in order to distinguish between different types of conditional cooperation, because other strategies such as grim trigger are also a type of conditional cooperation.

2.1 Procedure

2.1.1 Part A - Preferences and individual characteristics

In Part A, subjects completed three tasks: The Equality Equivalence Test (EET) (Kerschbamer 2015) to provide a categorisation of social preferences, a ten item Holt and Laury lottery choice list (LCL) (Holt & Laury 2002) to elicit risk preferences, and a 36 item Raven Test in order to provide a measure of intelligence. These were completed in a random order to control for order effects. Subjects were paid for one of the tasks, chosen at random; if the EET or LCL was chosen for payment, one decision was selected at random for payment, if the Raven Test was selected, three questions were selected at random and paid if correct.

For the EET, subjects made decisions over two sets of five binary decisions. In each decision, they chose between two allocations, one that resulted in an equal payoff to themselves, and a charity, and one that resulted in an unequal payoff to themselves and charity.⁴ We chose a charity rather than another subject to receive the payment in order to reduce any income effects, or beliefs about receiving additional earnings, impacting decisions within Part B. Following the procedure of Kerschbamer (2015), these ten decisions can be used to categorise the social preference type of each subject. We chose the EET over other tests because it provides a measure of social preferences without having to make restrictive assumptions about functional forms, the selection of specific functional forms which the researcher wishes to estimate, or other modelling variants. Table I outlines the ten binary decisions.

Decision	LEFT	RIGHT
1.	£3 to you, £8 to the charity	£5 to you, £5 to the charity
2.	£4 to you, £8 to the charity	£5 to you, £5 to the charity
3.	£5 to you, £8 to the charity	£5 to you, £5 to the charity
4.	£6 to you, £8 to the charity	£5 to you, £5 to the charity
5.	£7 to you, £8 to the charity	£5 to you, £5 to the charity
6.	£3 to you, £2 to the charity	£5 to you, £5 to the charity
7.	£4 to you, £2 to the charity	£5 to you, £5 to the charity
8.	£5 to you, £2 to the charity	£5 to you, £5 to the charity
9.	£6 to you, £2 to the charity	£5 to you, £5 to the charity
10.	£7 to you, £2 to the charity	£5 to you, £5 to the charity

Table I: Equality Equivalence Test

In the LCL, subjects made ten binary decisions over two lotteries. In each lottery, the payoffs were kept constant, but the probabilities were varied. Table II outlines the lotteries. The row number they switch from Left to Right provides an estimate of their level of risk aversion - lower switches suggest a higher level of risk aversion. We selected the LCL over other methods because it has been used widely in the literature and is suitable for the population (university students) that we are studying (Harrison et al. 2008).⁵

⁴We chose UNICEF, as this charity isn't related to any specific political party, religion, or ideology, and is known internationally.

⁵In both the EET and LCL, we interpret subjects with multiple switching points as being indifferent between the alternatives (Anderson & Mellor 2009). As a consequence, we use the first switching point in the estimate of risk aversion / social preferences, although our results are not sensitive to using alternative procedures.

Decision	LEFT	RIGHT
1.	10% chance of £2, 90% chance of £1.80	10% chance of £3.85, 90% chance of £0.10
2.	20% chance of £2, 80% chance of £1.80	20% chance of £3.85, 80% chance of £0.10
3.	30% chance of £2, 70% chance of £1.80	30% chance of £3.85, 70% chance of £0.10
4.	40% chance of £2, 60% chance of £1.80	40% chance of £3.85, 60% chance of £0.10
5.	50% chance of £2, 50% chance of £1.80	50% chance of £3.85, 50% chance of £0.10
6.	60% chance of £2, 40% chance of £1.80	60% chance of £3.85, 40% chance of £0.10
7.	70% chance of £2, 30% chance of £1.80	70% chance of £3.85, 30% chance of £0.10
8.	80% chance of £2, 20% chance of £1.80	80% chance of £3.85, 20% chance of £0.10
9.	90% chance of £2, 10% chance of £1.80	90% chance of £3.85, 10% chance of £0.10

Table II: Lottery choice lists

The Raven test (Foulds & Raven 1950) we use and the way in which we implement it is identical to that employed by Proto et al. (2019). We use a 36 item test from the Advance Progressive Matrices (APM) Set E, with subjects limited to 30 seconds for each item. For each question of the APM, subjects are shown a pattern, with one item in the sequence missing, and subjects must select the correct answer from a choice of eight in order to complete it. The matrices get more difficult as the subject progresses. We rewarded subjects with £2 per correct answer out of three randomly chosen questions if this test was chosen for payment. This was done to incentivise subjects to put effort into answering the questions, and although this is not typical for the Raven test, we do this following Proto et al. (2019).

Throughout the paper, we divide subjects into *Low* and *High* Raven depending on their score in the Raven Test. Those below and equal to the mean are identified as *Low*, and those above the mean, we describe as *High*. We use this split for presenting descriptive statistics, however due to it being an arbitrary divide, we focus on the impact of Raven score *per se* on decisions for the analyses.

2.1.2 Part B - Strategies for the infinitely repeated game

In Part B, subjects are randomly assigned a role, either as a First Mover, herein *Leader* or as a Second Mover, herein *Follower*, and then matched into pairs. They then play a single indefinitely repeated sequential public goods game that has the following structure. In each period of the game, subjects have 20 tokens and have five actions: contribute 0, 5, 10, 15 or 20 tokens to a ‘Group Project’. The *Leader* takes their action first, the *Follower* observes it, and then takes their action. We implemented a marginal per capita return of 0.75; both subjects received 0.75 tokens for each token contributed to the Group Project. Following Proto et al. (2019), who find empirically that a high continuation probability is most likely to induce gains from intelligence, we implement the indefinitely repeated game with a continuation probability of $\delta = 0.75$. Rather than elicit subjects’ actions in each period of the game, we instead follow a similar procedure to Dal Bó & Fréchet (2019), whereby subjects had to define a complete contingent strategy profile before playing a game and any feedback had been received. Once all subjects had defined this strategy, the computer would then match each *Leader* to one *Follower*, implement the strategy and that of the other player, and determine the outcome of the game.

Strategies were elicited as follows. *Leaders* first decided how many tokens to contribute in Period $t = 1$, specifying either 0, 5, 10, 15 or 20. The remainder of their strategy profile, which we call their Plan of Action is defined by answering twenty five questions. The subjects’

answers to these questions determines what their action would be in all periods after Period 1 once their Plan of Action is matched to a *Follower*. The exact questions that *Leaders* had to answer after specifying Period 1 contributions were as follows:

- After Period 1, if I last contributed 0 tokens and the second mover contributed 0 tokens, then contribute _____ tokens.
- After Period 1, if I last contributed 0 tokens and the second mover contributed 5 tokens, then contribute _____ tokens.
- ...
- After Period 1, if I last contributed 20 tokens and the second mover contributed 15 tokens, then contribute _____ tokens.
- After Period 1, if I last contributed 20 tokens and the second mover contributed 20 tokens, then contribute _____ tokens.

In doing so we observe the *Leaders*' full contingent strategy profile.

Followers strategies were elicited in almost the same way, however due to the sequential nature of the game, *Followers* are able to condition Period 1 contributions on the *Leader's* contribution. Thus, *Followers* first define what we call a Period 1 Plan, and then a Plan of Action. We elicit the Period 1 Plan using the following five questions:

- If the first mover contributes 0 tokens in Period 1, then contribute _____ tokens.
- If the first mover contributes 5 tokens in Period 1, then contribute _____ tokens.
- If the first mover contributes 10 tokens in Period 1, then contribute _____ tokens.
- If the first mover contributes 15 tokens in Period 1, then contribute _____ tokens.
- If the first mover contributes 20 tokens in Period 1, then contribute _____ tokens.

We then elicit the *Followers* contributions for all other periods in the same way as the *Leader*, asking them the following twenty five questions to elicit their Plan of Action:

- After Period 1, if I contributed 0 tokens in the last period and the first mover contributed 0 tokens in this period, then contribute _____ tokens.
- After Period 1, if I contributed 5 tokens in the last period and the first mover contributed 0 tokens in this period, then contribute _____ tokens.
- After Period 1, if I contributed 10 tokens in the last period and the first mover contributed 0 tokens in this period, then contribute _____ tokens.
- ...
- After Period 1, if I contributed 20 tokens in the last period and the first mover contributed 20 tokens in this period, then contribute _____ tokens.

We limit the *Followers* plan so that it can only be conditioned on the previous two actions, i.e. their own action in Period $t - 1$ and the *Leader's* action from Period t . This is done in order to keep strategies between players comparable. We also elicited all subjects' beliefs about what the other player would contribute in Period 1, which was done prior to their Plan of Action being defined. This was incentivised, with correct beliefs rewarded with £2.

Once both players' full strategy profiles are defined, the computer plays out the indefinitely repeated game using the strategies exactly as defined. The game was played once, with no feedback until the game had ended and all questionnaires and responses had been elicited. Although we place some restrictions on the players' strategies, the majority of strategies used and studied in the public goods game literature can still be played, such as Tit-for-Tat and Free-Ride, as well as strategies identified in the prisoners' dilemma literature, such as Grim Trigger, Punish/Reward and Always Cooperate. This is despite the limitations we place on the history of play subjects can condition their strategy on.

The main consideration associated with increasing the history of play that subjects are able to condition their strategy on is that the number of questions they need to answer to define their strategy increases exponentially. We made a conscious design choice to keep the procedure as simple as possible, as otherwise subjects' ability to understand the procedure may otherwise be driving the differences in strategies between intelligence levels. This also motivated our decision to incorporate perfect monitoring, rather than public or private monitoring, in the experimental environment. As [Aoyagi et al. \(2019\)](#) show, this should reduce complexity in strategy choice, making it easier to identify the strategies subjects play. Considerations associated with the 'strategy method' are that it may force subjects to think differently about the game in comparison to the 'direct response method', which in turn may produce differences in behaviour. However, previous work has found the two methods to both be behaviourally valid, with the majority of studies finding no differences in behaviour between the two methods (see [Brandts & Charness \(2011\)](#) for a review of the literature). There are, however, a number of additional advantages of eliciting strategies rather than observing actions. First, we observe the exact strategy each subject is playing, rather than having to estimate it. This reduces the possibility of error or issues due to the estimation procedure. Given the continuation probability we use of 0.75, we would expect to observe each interaction lasting only four Periods - which is far less than the number of possible actions. This contrasts with a prisoner's dilemma where there are only two actions, and where such an approach may be more advantageous. This would then make the estimation of strategies difficult, and may rely on a large number of assumptions that may weaken the analysis.

An alternative design choice might be to repeat the number of interactions, and enable subjects to modify their strategies after learning had taken place. Previous work has shown that leading figures can teach others to play optimally, and that this emerges with experience ([Camerer et al. 2002](#), [Hyndman et al. 2012](#), [Vostroknutov et al. 2018](#)). We made a conscious decision to avoid this, as we wanted to avoid the possibility of learning, and belief updating that may influence behaviour. Both learning and belief updating may vary ambiguously with levels of intelligence and may have made the interpretation of our data more difficult. Our focus in this paper is on the pure effect of intelligence on strategy choice, absent learning within the experiment and holding beliefs constant. ⁶

⁶We acknowledge that providing subjects with the ability to play the game by specifying actions, rather than strategies, prior to committing to a single strategy would have given them the ability to learn. Had we instead implemented this design choice, we may have observed different behaviour.


Panel A		
<i>Raven Score</i>	<i>Leader</i>	<i>Follower</i>
<i>Low</i>	42	48
<i>High</i>	53	47
Panel B		
<i>Treatment</i>	<i>Simulated Interactions</i>	
<i>Low Raven</i>	$42 \times 48 = 2016$	
<i>High Raven</i>	$53 \times 47 = 2491$	
<i>Mixed Raven</i>	$42 \times 47 + 53 \times 48 = 4518$	
Total	9025	

Table III: Observations

As we elicit subjects' strategies in Part B rather than their actions, we are able to examine how each *Leader's* strategy would perform against every other *Follower's* strategy, and vice versa. This is achieved by simulating how the game would have been played for all the possible strategy combinations. For each simulated interaction, we simulate a game consisting of four periods, because this is the expected number of periods the game should run for, given the continuation probability of 0.75. We do this in the spirit of the seminal work by Axelrod (1980) who examined which strategies were most effective in prisoner's dilemma games, examining strategies submitted by different academics. Axelrod (1980) found that tit-for-tat was the most effective. Similarly Selten et al. (1997) examined strategy performance in indefinitely repeated Cournot Oligopoly games, and used tournaments to examine which was most effective. In this paper, we match each *Leader's* strategy we observe to each *Follower's* strategy exactly once and examine how each subject's intelligence level impacts their own, and the other players, earnings and contributions.

To describe the simulated games, we refer to *Low* and *High* Raven categorisations of the subjects: in the *High Raven* treatment, both the *Leader* and the *Follower* have a Raven score above the mean; in the *Low Raven* treatment, both players have Raven scores below the mean; in the *Mixed Raven* treatment one player has a Raven score above, and one a Raven score below, the mean. The experiment was conducted in FEELE at the University of Exeter, and in BEEL at the University of Birmingham in February 2020. 190 undergraduate students were recruited through ORSEE (Greiner 2015) (Exeter) and SONA (Birmingham). The experiments were conducted using zTree (Fischbacher 2007). Once Part A and Part B of the experiment were completed, we then obtained individual demographics and asked participants to complete the BIG 5 personality questionnaire.⁷ Table III Panel A presents the number of observations we obtain for both *Leaders* and *Followers*, disaggregated by their Raven score, and Panel B displays the number of simulated interactions.

⁷All experimental materials are available in the Appendix.



Note: Mean score is 20.4, standard deviation is 4.74. Median score is 21. $N = 190$.

Figure 1: Distribution of Raven Scores

3 Results

This section outlines the experimental results. A number of common features are present throughout. Where non-parametric tests are used, we present the test used and p -value in parentheses. Unless otherwise stated, all tests are two-sided. As described in Section 2, we divide subjects into *Low* and *High* Raven groups for the descriptive statistics, but focus on the impact of Raven score for the analyses.

3.1 Data summary

Table IV presents the range of characteristics we elicit, with those classified as *Low Raven* presented in the left column, those as *High Raven* in the centre column, and the results of Robust Rank Order Tests comparing the averages for each of the variables in the right hand column.

As can be seen, across all observable variables only Raven scores are significantly different between *High* and *Low* Raven subjects ($p < 0.001$, Robust Rank Order test). We take this as initial suggestive evidence that any behavioural differences observed between subjects of *Low* and *High* Raven score in the infinitely repeated game are unlikely to be driven by differences in other individual characteristics. Figure 1 presents the distribution of Raven scores.

In order to shed light on how subjects are actually playing the infinitely game, and how this might be influenced by their cognitive ability, we group subjects together who play strategies that have similar characteristics. We categorise *Leaders* and *Followers* into types based on (1) their Period 1 contribution or Period 1 Plan and (2) their Plan of Action. As we observe subjects' strategies precisely, this is done using a popular and unsupervised machine learning algorithm, the k -means clustering algorithm. Describing the algorithm applied to the Plans of Action, first k reference plans are selected at random from the 190 plans provided by subjects. Each subjects' Plan of Action is a twenty five element vector, $P_k = (p_1, p_2, p_3, \dots, p_{25})$, where each element corresponds to one of the questions answered by the subject during the experiment. We then compare each subject's Plan of Action, S_i to each of the k reference plans, P_k , by calculating the Manhattan Distance between them. For example, the distance between subject j 's twenty five element plan of action, $S_j = (s_1, s_2, s_3, \dots, s_{25})$ and the twenty five element

<i>Measure</i>	<i>Low Raven</i>	<i>High Raven</i>	<i>H₀ : High=Low</i>
<i>Raven Score</i>	16.467 (3.373)	23.93 (2.463)	$p < 0.001$
<i>Risk Score</i>	6.578 (2.077)	7.08 (1.68)	$p > 0.1$
<i>EET x score</i>	0.678 (0.815)	0.87 (0.812)	$p > 0.1$
<i>EET y score</i>	1.222 (0.9)	1.25 (0.914)	$p > 0.1$
<i>Proportion of males</i>	0.544 (0.501)	0.52 (0.502)	$p > 0.1$
<i>Proportion of Birmingham</i>	0.444 (0.5)	0.4 (0.492)	$p > 0.1$
<i>Age</i>	20.589 (2.481)	20.78 (3.498)	$p > 0.1$
<i>Political Score</i>	4.533 (1.743)	4.42 (1.665)	$p > 0.1$
<i>Agreeableness</i>	117.7 (19.195)	116.64 (16.663)	$p > 0.1$
<i>Mother's Education level</i>	4.156 (1.595)	4.4 (1.498)	$p > 0.1$
<i>Father's Education level</i>	3.589 (1.336)	3.55 (1.253)	$p > 0.1$
<i>Subjects</i>	90	100	

Notes: The mean *Raven Score* is 20. All measures compared using two sided Robust Rank Order Tests, using individual level observations. Standard deviation in parentheses. Risk score is the average switching point in the Holt and Laury lottery choice list. EET *x* and *y* scores are calculated following [Kerschbamer \(2015\)](#). Political Score is a measure of how 'Right Wing' an individual regards themselves, with higher scores being more 'Right Wing' (1–7). Proportion of males/Birmingham outlines the proportion of male subjects, and those from the University of Birmingham (others from Exeter). Agreeableness is calculated from the Big 5 personality test, with higher scores being more agreeable subjects. Mother's and Father's education level is an ordinal categorical variable, with higher numbers meaning a higher education level (1–6).

Table IV: Summary Statistics - Observable Characteristics

reference plan k , P_k , is calculated as follows

$$d_{j,k}(S_j, P_k) = \sum_{i=1}^{25} |s_i - p_i| \quad (1)$$

Whichever of the k distances is shortest, the subject's i Plan of Action is assigned to Cluster k . Once all subjects' plans are categorised, we then re-calculate each cluster as being the average of all the Plans of action assigned to that cluster. This is done as follows,

$$V_k = \left(\frac{1}{n_k} \right) \sum_{i=1}^{n_k} S_{i,k} \quad (2)$$

where n_k is the number of Plans of Action assigned to each cluster k . This entire procedure is then repeated until each Plan of Action remains in the same cluster, and no Plans are assigned to new clusters.⁸ As outlined, we use the k -means algorithm to individually cluster all subjects' Plans of Action and the Period 1 Plans of *Followers*. The only difference is that *Followers* Period 1 plans are 5 element vectors, rather than 25.

As k is arbitrarily chosen, we conduct the procedure with $k = 4$, $k = 6$ and $k = 8$, then select k for the analysis conditional on which one 'best fits' the data. In order to assess which k provides the best fit, once subjects' Plans have been clustered we use silhouette analysis (Rousseeuw 1987), and calculate the following silhouette statistic for the Plan of Action of each subject i , for each k ,

$$h_{i,k} = \frac{b_{i,k} - a_{i,k}}{\max(a_{i,k}, b_{i,k})} \quad (3)$$

where $a_{i,k}$ is the mean Manhattan distance of subject i 's Plan of Action to the other subjects' plans of action that are in the same category, and $b_{i,k}$ is the mean Manhattan distance to subjects of the next closest category. A silhouette statistic of 1 implies that the subjects' plan of action falls perfectly into one distinct cluster; whereas a silhouette statistic of -1 implies a subject has been perfectly mis-assigned and their plan of action is most similar to subjects in another category, rather than the one in which our procedure has placed them. We select the k that produces the highest average silhouette statistic.


We use Manhattan Distance over Euclidean Distance due to the latter's sensitivity to high dimensions and the former's simple interpretation: Manhattan distance tells us how many discrete changes the subject needs to make to their plan of action in order for them to be playing the exact reference plan being considered. For example, as subjects can only make contribution choices in multiples of five from 0 to 20, a Manhattan distance of five means our subject needs to make a single change to their plan in order for it to be identical to the reference plan. A distance of 20 implies four changes must be made to their plan in order for it to be identical to the reference plan.⁹

⁸In some cases this may not be achievable. We implement this procedure using a maximum of 1000 iterations.

⁹As an example: a subject with the Plan $S = (0, 0, 0, \dots, 0, 5)$ has a Manhattan distance of five to the plan $P = (0, 0, 0, \dots, 0, 0)$. If the subject made a single change to their plan, contributing 0 instead of 5 (a 'single' change as subjects can only contribute in multiples of five) in the final element of their plan, S and P would then be identical. Similarly, a subject with the plan S has a distance of ten to the plan $F = (5, 0, 0, \dots, 0, 0)$, and would need to make two changes for them to be identical.

3.1.1 Period 1 contributions and Period 1 Plans

We begin by examining the Period 1 contributions of the *Leader*, with Figure 2 presenting average Period 1 contributions graphically, disaggregated by Raven score. As can be seen, *Leaders* contribute on average around 15 tokens in Period 1, with *High* Raven subjects contributing more than *Low*, although this difference is not significant at conventional levels ($p = 0.11$, Robust Rank Order Test).



Note: Vertical bars represent 95% confidence intervals.

Figure 2: Leaders' contributions in Period 1


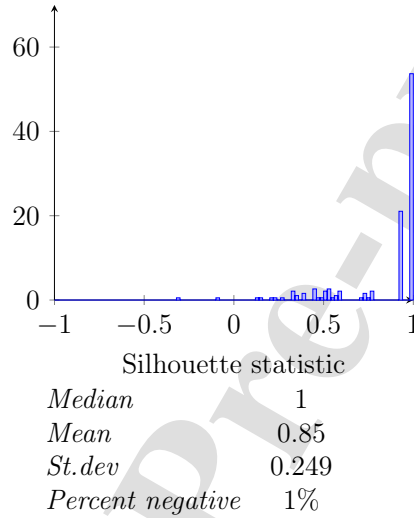


Figure 3: Followers' Period 1 plans

We now examine *Followers'* Period 1 Plans. Figure 3 displays two matrix plots. In each plot, a darker coloured cell represents a contribution closer to zero, whilst a lighter one represents a contribution closer to 20 - the colour key is given on the right. On the y axis is the contribution of the *Leader* in Period 1. The cell at the top of each diagram represents the *Follower's* contribution in Period 1 if the *Leader* contributes 20 tokens in Period 1; the cell at the bottom of each diagram represents the *Follower's* contribution in Period 1 if the *Leader* contributes 0 tokens in Period 1. As can be seen, lower (higher) contributions by the *Leader* means the *Follower*, regardless of Raven score, would contribute less (more). A simple initial comparison reveals that there are no discernible differences in the average Period 1 Plans between *Low* and *High* Raven groups ($p > 0.1$ in all cases, Fisher's Exact tests).

In order to examine the rich individual heterogeneity within the Period 1 Plans of *Followers*, we implement the k -means procedure outlined above, clustering the plans using $k = 4$, $k = 6$ and $k = 8$ reference plans. To determine which k best fits the data, we run the procedure and present the distributions of silhouette statistics for each of the different number of clusters, calculated following Equation 3. These distributions, as well as the mean, median and standard deviation of the silhouette statistics are presented in the appendix, along with the percent of subjects with negative silhouette scores - those observations mis-assigned by the algorithm.



Note: Percent of observations on the y -axis. Silhouette statistic calculated from *Followers'* Period 1 Plans, following Equation 3.

Figure 4: Distribution of silhouette statistics, Period 1 Plans $k = 6$ clusters

We find that $k = 6$ clusters produces a significantly higher silhouette score than both $k = 4$ and $k = 8$ ($p < 0.001$ in both cases, Signed-Rank tests), and not a single Period 1 Plan is misassigned. We therefore focus our attention on using six clusters. Figure 4 presents the silhouette statistic distribution for $k = 6$ clusters. We present the average Period 1 Plans for each of the six clusters as matrix plots in Figure 5; darker colours represent contributions closer to 0, and lighter closer to 20; the y axis is the contribution of the *Leader* in Period 1. In order to avoid any confusion regarding the Period 1 Plans, we do not name them or attempt to link them to strategies in the literature, and instead name them Cluster 1 - 6.

To more clearly explain how the diagrams can be interpreted, as an example consider Cluster 1. This shows a light colour (white) at 20, and progressively darker colours towards 0 (black). This shows that, if the *Leader* were to contribute 20 tokens in Period 1, the *Follower* would also contribute 20; instead, if the *Leader* contributed 0 in Period 1, the *Follower* would contribute 0. Cluster 1 is therefore a Period 1 Plan in which the *Follower* makes contributions conditional on the *Leader's* contributions. In contrast, Cluster 6 is coloured entirely black: this means that, regardless of the contributions of the *Leader*, the *Follower* will contribute 0.

Table V presents the total number of subjects in each of the six clusters and the average silhouette statistic for that cluster, with the information disaggregated by *Low* and *High* Raven scores.

Cluster	N	Total	N	Low Raven	N	High Raven
		Silhouette		Silhouette		Silhouette
Cluster 1	58	0.81	31	0.78	27	0.85
Cluster 2	11	0.45	4	0.47	7	0.45
Cluster 3	8	0.22	6	0.19	2	0.28
Cluster 4	9	0.50	2	0.4	7	0.52
Cluster 5	2	0.4	1	0.58	1	0.23
Cluster 6	7	1	4	1	3	1

Table V: Followers assigned to each cluster

As shown in Table V, Cluster 1 is clearly the most played, regardless of Raven score. Cluster 1 also has a high average silhouette statistic, which suggests that plans in this cluster are very similar to each other. Cluster 2 is similar to Cluster 1, however will contribute more than or equal to Cluster 1 for every contribution the *Leader* makes. Interestingly, Cluster 6 has a silhouette statistic of 1, suggesting the plans are identical to each other. Cluster 6 is a plan that always contribute 0 tokens, regardless of what the *Leader* contributes.

To formally examine if Raven score impacts Period 1 Plan choice parametrically, we estimate the marginal effects from a number of Multinomial Logit regressions. In each regression, the dependent variable is a categorical variable that takes a different value for each cluster. In all regressions, the variable of interest is the subject's Raven score. In model 1, we only include Raven score. In model 2, we control for subjects' risk preferences, social preferences, their gender, their beliefs about the other player's contributions in period 1. In model 3, we add additional controls for subjects' political attitudes and where the experiment took place. The marginal effect of Raven score on the probability of each plan is presented in Table VI.¹⁰

	Cluster 1	Cluster 2	Cluster 3	Cluster 4	Cluster 5	Cluster 6
Model 1	-0.015 (0.011)	0.012 (0.008)	-0.007 (0.006)	0.012* (0.007)	-0.002 (0.003)	0.00 (0.006)
Model 2	-0.018 (0.012)	0.015* (0.009)	-0.009 (0.006)	0.011 (0.007)	-0.003 (0.004)	0.003 (0.006)
Model 3	-0.025** (0.012)	0.022** (0.009)	-0.01* (0.006)	0.009 (0.007)	-0.002 (0.004)	0.005 (0.005)

Notes: Standard errors in parentheses. *, ** and *** denote significance at the 10%, 5% and 1% level. Coefficients are the marginal effect of Raven score on the probability of the subject being assigned to that cluster, estimated from Multinomial Logit models. All models estimated using 95 observations. In model 1, we include Raven score. In model 2, we add controls for subjects' risk preferences, social preferences, their gender, and their beliefs about the other player's contributions in Period 1. In model 3, we add additional controls for subjects' political attitudes and where the experiment took place.

Table VI: Marginal effect of Raven score on Followers' Period 1 Plans

As can be seen in Table VI across all three models, Raven score has a positive impact on the probability that a *Follower* plays a Cluster 2 plan ($p = 0.08$, in Model 2, $p < 0.05$ in Model 3, T-tests). Raven score also has a negative impact on the probability of playing a Cluster 1

¹⁰Once we correct the p -values for potential multiplicity using the Holm-Bonferroni correction procedure, only the marginal effects coefficients in Model 3, Cluster 1 and 2, remain significant at the 10% level.

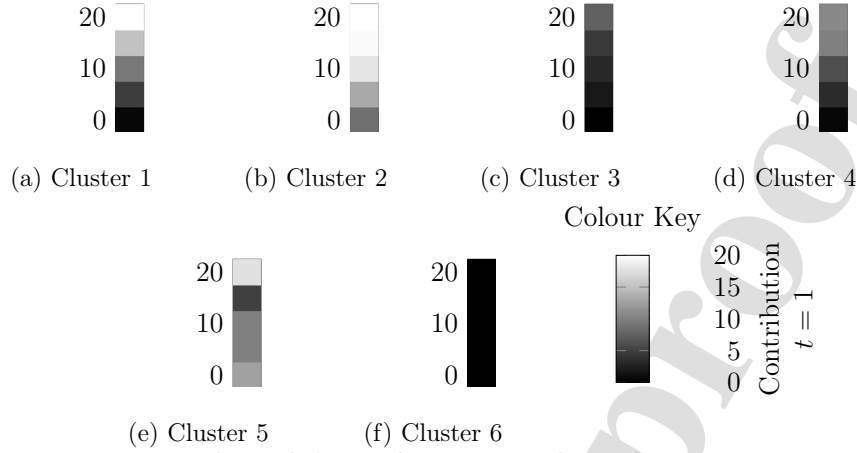



Figure 5: Average Period 1 Plans for each cluster

plan, however this is only significant in Model 3 ($p < 0.05$, Model 3, T-Tests). Importantly, the sign on the coefficients for Cluster 1 and Cluster 2 are robust across the three models: we take this as evidence that an increased Raven score reduces the probability of a subject playing a Cluster 1 Period 1 Plan, but increases the probability of playing Cluster 2. In all models, no other coefficients are estimated to be significant at the 5% level ($p > 0.05$ in all other cases, T-Tests).

Finally, we examine the subjects' beliefs about the other player's contribution in Period 1. Figure 6 presents the average beliefs of *Leaders* and *Followers*. We find no significant differences in average beliefs between *Low* and *High* Raven groups for *Leaders* or *Followers* ($p > 0.1$, in both cases, Robust Rank Order Tests). This suggests, regardless of Raven score, that initial beliefs about the contributions of the other player are identical. This rules out any observed differences in earnings and contributions between *Low* and *High* Raven groups as being the result of path dependency stemming from initial beliefs. This leads to our first observation.

Observation 1. *Followers are less likely to play a Cluster 1 Period 1 Plan, and more likely to play a Cluster 2 Period 1 Plan, as their Raven score increases. However, there is no link between Leaders' contributions in Period 1, or between subjects' beliefs about contributions, and Raven score.*

Observation 1 suggests that some *Followers* are more likely to contribute to the public good in Period 1 as their intelligence increases. This is because a Cluster 2 Period 1 Plan will always contribute more than, or the same amount as, a Cluster 1 Period 1 Plan. Observation 1 therefore supports the notion that intelligence works to reduce opportunistic behaviour in social dilemmas (Proto et al. 2019).



Note: Vertical bars represent 95% confidence intervals.

Figure 6: Beliefs in Period 1

3.1.2 Plans of Action

We now consider Plans of Action. Figure 7 summarizes the plans of action defined by *Leaders*, disaggregated by Raven score. The figure presents a three-way matrix plot that displays the *Leader's* contribution in Period t , conditional on their own contribution in Period $t - 1$ and the *Follower's* contribution in Period $t - 1$. Figure 8 presents the same for *Followers*, plotting the *Follower's* contribution in Period t , conditional on their own contribution in Period $t - 1$ and the *Leader's* contribution in Period t . In both figures, a darker coloured cell represents a contribution closer to zero, whilst a lighter one represents a contribution closer to 20.

Some clear patterns in the average Plans of Action emerge, regardless of Raven score. First, both *Leaders* and *Followers* define Plans of Action that contribute more the higher the contribution of the other player. This is revealed by lighter colours in the top rows of the matrices, and darker cells in the bottom rows. This is evidence of Tit-for-Tat type strategies. There are also clear divides along the diagonal, which suggests players contribute more when the other player contributes more than them, and less in the opposite case. This is shown by lighter cells in the top left of the matrices, and darker cells in the bottom right. This is similar to previous

findings in the finitely repeated games literature where such strategies are not an equilibrium, as reported by Fischbacher et al. (2001) and Chaudhuri (2011), and indicates that the average Plan of Action is at least partially conditionally cooperative. This is true for both *Leaders* and *Followers*.




Figure 7: Aggregated plan of action, Leaders


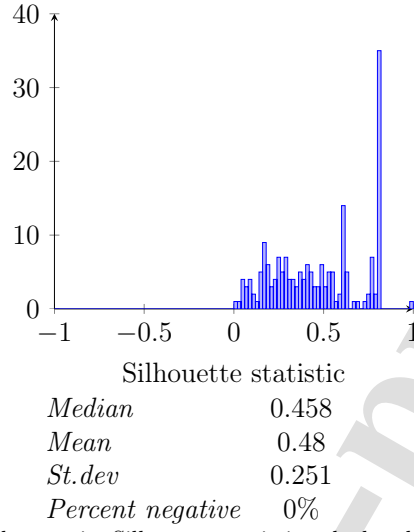


Figure 8: Aggregated plan of action, Followers

We now turn to cluster analysis in order to draw out any patterns from the data, and shed light on the individual Plans of Action being played by subjects. This is done following the procedure described in Section 3.1. The figures given in the appendix present the distribution of silhouette statistics for $k = 4$, $k = 6$ and $k = 8$ clusters, which we use to determine how many clusters to use in our analysis. It's clear from the presented averages that $k = 6$ provides the best fit. The difference in silhouette scores between $k = 6$ and $k = 8$ is significantly different ($p < 0.001$, Signed Rank Test). When $k = 6$, the average is also significantly different to when $k = 4$ ($p < 0.001$, Signed Rank Test). Figure 9 presents the silhouette plot for $k = 6$ clusters.

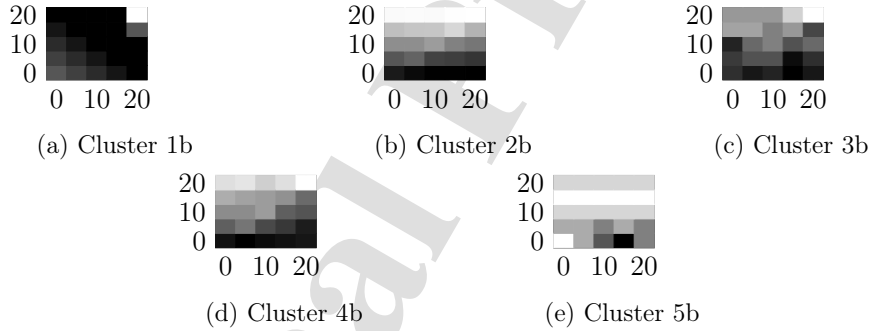
We present the average Plan of Action for each of the six clusters in Figure 10 for *Leaders* and Figure 11 for *Followers*, which we suffix with a 'b' to distinguish them from the Period 1 Plans in Section 3.1.1.

Although we do not explicitly name any of the clusters given in Figures 10 and 11, some of them have close analogues to strategies studied in the literature. For example, Cluster 1b appears to represent a 'Grim Trigger' type strategy. This is because the Plan of Action will only contribute twenty tokens (and approximately 0 otherwise) unless the other player last contributed 20 tokens, and they themselves also last contributed 20. Cluster 2b is very close



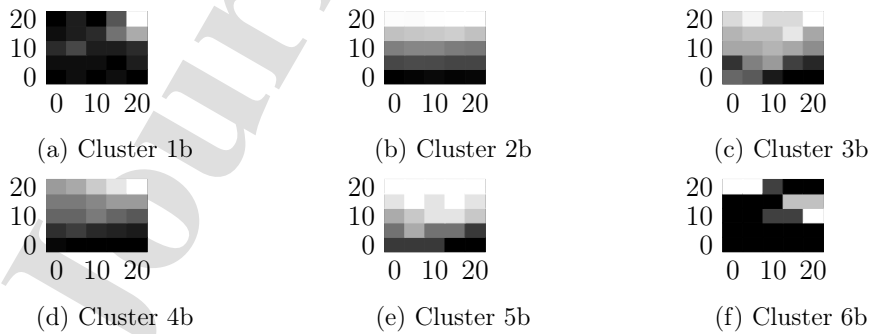
Note: Percent of observations on the y -axis. Silhouette statistic calculated from subjects' plans of action, following Equation 3.

Figure 9: Distribution of silhouette statistics, $k = 6$ clusters



Note: The x -axis shows the *Leaders'* contribution in $t - 1$, and y -axis the *Follower's* contribution in $t - 1$.

Figure 10: Leaders' plans of action by cluster



Note: The x axis shows the *Follower's* contribution in $t - 1$ and the *Leader's* contribution in t on the y -axis.

Figure 11: Followers' plans of action by cluster

to being a Tit-for-Tat type strategy, with the horizontal line patterns implying the subject will contribute the same number of tokens as those contributed by the other player. Cluster 3b and 4b look similar to a ‘Punish/Reward’ type strategy, with a pattern that suggests subjects contribute more when the other player contributes more than them, but zero otherwise. Cluster 5b and 6b, when considering *Leaders* and *Followers* together, are more difficult to link to the literature, although it’s clear that Cluster 6b for *Leaders* is always contribute 0. Table VII presents the total number of subjects in each of the clusters and the average silhouette statistic, as well as the information disaggregated by Raven scores.

	Leaders						Followers					
	Total		Low Raven		High Raven		Total		Low Raven		High Raven	
Cluster	N	Silhouette	N	Silhouette	N	Silhouette	N	Silhouette	N	Silhouette	N	Silhouette
Cluster 1b	13	0.378 (0.212)	7	0.297 (0.213)	6	0.471 (0.184)	18	0.519 (0.145)	14	0.516 (0.15)	4	0.531 (0.147)
Cluster 2b	47	0.579 (0.228)	21	0.595 (0.239)	26	0.565 (0.221)	47	0.589 (0.233)	22	0.624 (0.23)	25	0.56 (0.236)
Cluster 3b	12	0.185 (0.11)	5	0.161 (0.05)	7	0.202 (0.147)	10	0.162 (0.1)	4	0.233 (0.087)	6	0.114 (0.091)
Cluster 4b	19	0.286 (0.112)	7	0.245 (0.128)	12	0.309 (0.101)	13	0.292 (0.131)	5	0.203 (0.078)	8	0.349 (0.129)
Cluster 5b	4	0.645 (0.245)	2	0.761 (0.00)	2	0.529 (0.354)	6	0.491 (0.285)	2	0.41 (0.498)	4	0.532 (0.276)
Cluster 6b	0	-	0	-	0	-	1	1 (0.00)	1	1 (0.00)	0	-

Note: Standard deviations in parentheses.

Table VII: Plans of action assigned to each cluster

Table VII shows that Cluster 2b is the most played, with around 50% of subjects being assigned to this cluster. This is very similar to the percentage of subjects playing Conditional Cooperate as reported by Fischbacher et al. (2001). Cluster 4b, the Punish/Reward type is second, and then Cluster 1b, Grim Trigger.

Despite our observation that a Cluster 1b Plan of Action would form part of a Grim Trigger type strategy, subjects clustered into Cluster 1b cannot be distinguished from those playing a Free-Rider strategy (always contribute zero) without taking into account their Period 1 contributions. This is because if the player ensured that they contributed less than 20 tokens in Period 1, and then played a Cluster 1b Plan of Action, this would mean they would then always contribute 0. To examine this more closely, and determine what strategy subjects are playing, Table VIII distinguishes between subjects by their Period 1 behaviour. Panel A presents *Leaders*, and Panel B *Followers*. In each case, we disaggregate by Raven score.¹¹

Panel A				
Period 1 Contribution	<i>Leaders playing Cluster 1b</i>			<i>Implied strategy</i>
	<i>Total</i>	<i>Low Raven</i>	<i>High Raven</i>	
0	3	2	1	Free-Ride
5	6	4	2	Free-Ride
20	4	1	3	Grim Trigger

Panel B				
Period 1 Plan	<i>Followers playing Cluster 1b</i>			<i>Implied strategy</i>
	<i>Total</i>	<i>Low Raven</i>	<i>High Raven</i>	
Cluster 1	5	5	0	Grim Trigger
Cluster 3	5	4	1	Free-Ride
Cluster 4	1	1	0	Free-Ride
Cluster 6	7	4	3	Free-Ride

Note: The table presents the number of subjects playing Cluster 1b Plans of Action, disaggregated by their Period 1 decisions. Panel A presents the number of *Leaders* and Panel B the number of *Followers*. The final column presents the implied strategy, conditional on Period 1 and Plan of Action clusters.

Table VIII: Distinguishing between Grim Trigger and Free-Riding strategies

As can be seen in Table VIII, there are slight differences between *High* and *Low* Raven subjects. As seen in Panel A, *Low* Raven *Leaders* appear more inclined to play a Free-Riding strategy, whereas Panel B suggests *Low* *Followers* are more likely to play Grim Trigger. In contrast, *High* Raven *Followers* do not play Grim Trigger, choosing instead to contribute zero in Period 1, meaning those *High* *Followers* playing Cluster 1b are Free-Riders.

We now formally examine if there exist a relationship between intelligence and strategy type. To do this, we estimate the marginal effects from a number of Multinomial Logit regressions. In each regression, the dependent variable is a categorical variable that takes a different value for each cluster. We use the cluster each Plan of Action was assigned to as the strategy type, separating Cluster 1b plans into Free-Riding and Grim-Trigger as shown in Table VIII. In all

¹¹Similar arguments can be made for Cluster 3b and Cluster 4b. For example, if a *Follower* were to play Cluster 6 in Period 1, along with Cluster 4b Plan of Action, this would correspond to a Free-Riding strategy. However, not a single subject played this combination, or one similar, that could be classified as a Free-Riding strategy.

regressions, the variable of interest is the subject's Raven score. In model 1, we only include Raven score. In model 2, we control for subjects' risk preferences, social preferences, their gender, and their beliefs about the other player's contributions in Period 1. In model 3, we add additional controls for subjects' political attitudes and where the experiment took place. The marginal effect of Raven score on the probability of each plan of action being assigned to each Cluster is presented in Table IX. We present the estimates for *Leaders* in Panel A, and *Followers* in Panel B.

Table VII, Panel A, outlines how the marginal effect of Raven score on the probability of being a Free-Rider is negative and significant for *Leaders* in all three models ($p < 0.05$, in all models, T-Tests). The Table suggests a one point increase in Raven score reduces the probability that a *Leader* is a Free-Rider by around 1%. This estimate is robust across models. This means that the average *High* Raven *Leader*, with a Raven score of 24, is around 8% less likely to play a Free-Riding strategy in comparison to the average *Low* *Leader*, who has a Raven score of 16. Similarly, Panel B outlines how Raven score has a negative marginal effect on the probability that a *Follower* plays Grim Trigger ($p < 0.05$ model 1 and 3, $p < 0.1$ in model 2, T-Tests): a one point increase in Raven score reduces the probability that a *Follower* plays Grim Trigger by around 1%.¹² This leads to our second observation.

Observation 2. *A one point increase in Raven score decreases the probability of the Leader being a Free-Rider by 1%, but decreases the probability of the Follower playing Grim Trigger by 1%.*

Although we do not know for certain, we are able to speculate as to why we observe more intelligent *Followers* being less likely to play a Grim Trigger strategy. Grim trigger is a less 'forgiving' conditionally cooperative strategy than Tit-for-Tat, as just one uncooperative action is punished forever. It may be that a more intelligent *Follower* understands more accurately that such a strategy may have negative consequences for her own payoffs in the future. Therefore, the more intelligent that *Followers* are the less likely they are to play Grim Trigger.

We can now examine what the estimated marginal effects mean for the behaviour of the average *Leader* and *Follower* conditional on their Raven score. Consider the estimated impact of Raven score on the probability that a *Leader* plays a Free-Riding strategy: a 1 point increase in Raven score lowers the probability that they play this strategy by approximately 1.3%. As there is an 8 point difference in Raven score between *Low* and *High* Raven *Leaders*, *High* *Leaders* are 10.4% less likely to be a Free-Rider than a *Low* *Leader*. A similar exercise can be done for *Followers*: a 1 point increase in Raven score lowers the probability that they play a Grim Trigger strategy by 1.1%. An 8 point difference in Raven score between *Low* and *High* Raven *Followers* means that *High* *Followers* are 8.8% less likely to play Grim-Trigger than a *Low* *Follower*. We examine the significance that these differences have on earnings and contributions in the following section.

¹²If we conduct the analysis without distinguishing between Grim Trigger and Free-Riding strategies, including just Cluster 1b instead, we find similar results: a one point increase Raven score decreases the probability that subjects play Cluster 1b by around 1%. This is significant at the 5% level in all cases ($p < 0.05$, in all cases, T-Tests).

Panel A: Leaders		Cluster					
	<i>Grim Trigger</i>	<i>Cluster 2b</i>	<i>Cluster 3b</i>	<i>Cluster 4b</i>	<i>Cluster 5b</i>		<i>Free-Rider</i>
<i>Model 1</i>	-0.001 (0.004)	0.009 (0.011)	-0.002 (0.007)	0.008 (0.009)	0.00 (0.004)		-0.013** (0.006)
<i>Model 2</i>	-0.002 (0.004)	0.001 (0.011)	0.004 (0.008)	0.008 (0.009)	-0.001 (0.005)		-0.01** (0.005)
<i>Model 3</i>	0.001 (0.003)	-0.002 (0.011)	0.004 (0.008)	0.006 (0.009)	0.002 (0.005)		-0.01** (0.005)
Panel B: Followers		Cluster					
	<i>Grim Trigger</i>	<i>Cluster 2b</i>	<i>Cluster 3b</i>	<i>Cluster 4b</i>	<i>Cluster 5b</i>	<i>Cluster 6b</i>	<i>Free-Rider</i>
<i>Model 1</i>	-0.018** (0.007)	0.012 (0.011)	0.002 (0.007)	0.011 (0.008)	-0.004 (0.005)	-0.002 (0.003)	-0.001 (0.006)
<i>Model 2</i>	-0.015** (0.006)	0.006 (0.011)	0.003 (0.007)	0.009 (0.008)	-0.006 (0.006)	0.00 (0.00)	0.003 (0.006)
<i>Model 3</i>	-0.015** (0.006)	0.004 (0.011)	0.002 (0.007)	0.01 (0.008)	-0.006 (0.006)	0.00 (0.00)	0.005 (0.006)

Notes: Standard errors in parentheses. *, ** and *** denote significance at the 10%, 5% and 1% level. Coefficients are marginal effects estimated from Multinomial Logit models. All models estimated using 95 observations. Panel A presents the estimated marginal effect of Raven score on cluster assignment for *Leaders*; Panel B for *Followers*. Model 1 includes only Raven score as an explanatory variable. Model 2 adds additional controls for subjects' risk preferences, social preferences, their gender, and the subject's belief about the other player's contributions in period 1. Model 3 additionally controls for political attitudes and where the experiment took place. Cluster6b estimates are empty in Panel A due to there being no observations.

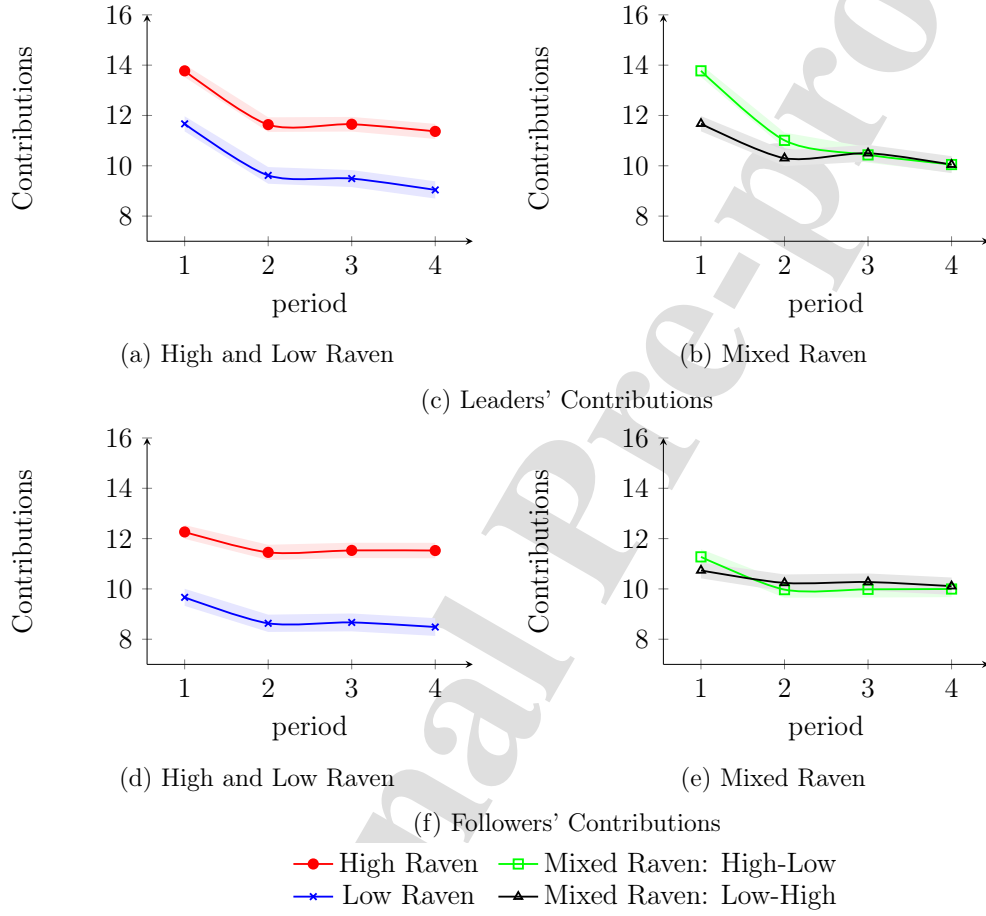
Table IX: The marginal effect of Raven score on cluster assignment

3.2 Simulated games

Although Observation 1 and 2 highlight a link between intelligence and strategy choices, it's not clear how these differences might impact outcomes, specifically contributions and earnings, in the infinitely repeated game. In order to determine how strategy choice and the Raven score of *Leaders* and *Followers* impact contributions and earnings, we simulate a game between each of our 95 *Leaders* and each of our 95 *Followers*. Given our continuation probability of 0.75, a game is expected to last four periods. Figure 12 plots the contributions made by *Leaders* and *Followers* in each period the simulated games. Figure 13 presents earnings. In both diagrams the information is disaggregated by Raven scores.

As can be seen, *High* Raven groups contribute and earn more than *Low* Raven groups in all periods. *Mixed* groups also appear to do better than *Low* Raven groups. This finding closely replicates those of Proto et al. (2021), who show how strategic interactions and cooperation are affected by the heterogeneity of cognitive skills of groups of players.

In order to formally examine if there exist difference between groups and to establish what is driving any differences, we conduct a number of Tobit regressions where the dependent variable is either average contributions or earnings in each interaction. This gives us 95 observations per subject. To control for dependence between observations we cluster standard errors at the subject level, and we conduct regressions using *Leader* and *Follower* observations separately. In each regression model, we include the *Leader's* and the *Follower's* Raven scores as our



Note: Shaded area represents 95% confidence intervals. *Low Raven* refers to simulated play where both subjects are *Low*, *High Raven* where both subjects are *High* and *Mixed- Raven* where there is one *Low* and one *High*: *High-Low* means the *Leader* is *High* and the *Follower* is *Low*.

Figure 12: Contributions in the simulated games

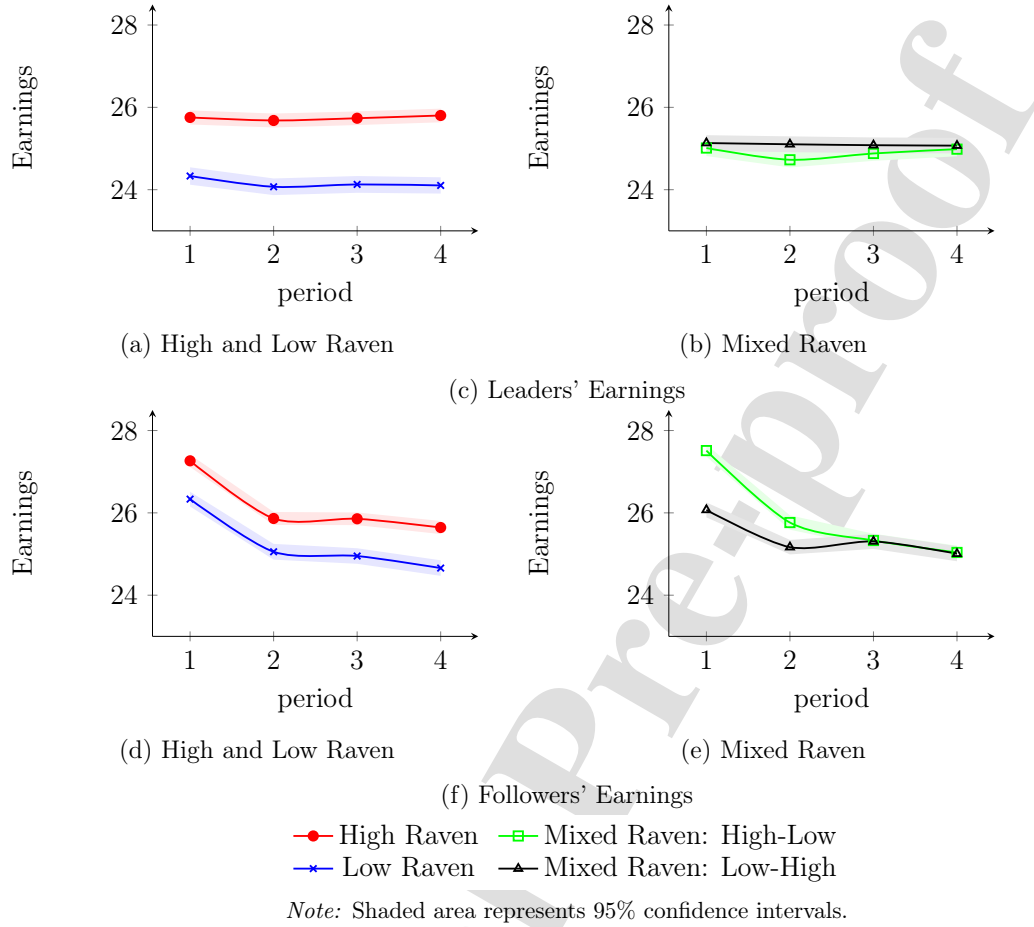


Figure 13: Earnings

explanatory variables of interest, along with the interaction of these two variables.

For the *Leader*, we control for their contribution in Period 1 as well as their beliefs about the *Follower's* contribution. In model 2, we add additional controls for risk aversion, social preferences, gender and their level of agreeableness. In model 3, we add political attitudes and where the experiment was conducted. For the *Follower*, we use the same sets of controls, however we control for the cluster their Period 1 Plan is categorised as, instead of their contribution in Period 1, as well as their beliefs about the *Leader's* contribution. From each Tobit regression we estimate the average marginal effect of the *Leader's* Raven score and the *Follower's* Raven score. We present the results in Table X, with estimates for the *Leaders* presented in Panel A, and the *Followers* in Panel B. Due to us examining *Leaders* and *Followers* separately, in Panel A *Leader's Raven* refers to the subject's own Raven score, and in Panel B, *Follower's Raven* refers to the subject's own Raven score.

Table X Panel A shows how the *Follower's* Raven score has a positive and significant marginal effect on the contributions and Earnings of the *Leader*, with a one point increase in the *Follower's* Raven score estimated to increase contributions by around 0.1 tokens and

earnings by around 0.05 tokens ($p < 0.01$ in all cases, T-Tests). Similarly, Panel B shows that the marginal effect of the *Leader's* Raven score increases *Follower's* contributions by 0.2 tokens and earnings by around 0.08 tokens ($p < 0.01$ in all cases, T-Tests).

To examine the marginal effect of the interaction of Raven scores on contributions and earnings, we estimate the marginal effect of the *Leader's* Raven score on *Follower's* contributions and earnings conditional on the *Follower's* Raven score. We do the same for *Leaders* using the *Follower's* Raven score. We estimate the conditional marginal effects from Table X Model 1, Panel A and Panel B, and present the estimates graphically in Figure 14.¹³

As can be seen in Figure 14a, the marginal effect of the *Follower's* Raven score on *Leader's* contributions is always positive and significant ($p > 0.001$ in all cases, T-tests), and increasing as the *Leader's* Raven score increases: the marginal effect of the *Follower's* Raven score when the *Leader's* Raven score is greater than or equal to twenty is significantly larger than when the *Leader's* Raven score is less than twenty ($p < 0.05$ in all cases, χ^2 tests). The marginal effect on the *Leader's* earnings is found to be constant ($p > 0.1$ in all cases, χ^2 tests). The marginal effect of the *Leader's* Raven score on the *Follower's* contributions and earnings, as shown in Figure 14b, is not found to differ with the *Follower's* Raven score ($p > 0.1$ in all cases, χ^2 tests) suggesting it is positive and constant.¹⁴

Observation 3. *The Leader's intelligence positively impacts the contributions and earnings of the Follower, and the Follower's intelligence positively impacts the contributions and earnings of the Leader. The effect of the Follower's intelligence on the Leader increases as the Leader's intelligence increases.*

The interaction effect between *Followers'* and *Leaders'* intelligence is likely driven by a number of factors. First, *Followers* are less likely to play Grim Trigger strategies as their intelligence increases. This strategy, although still conditionally cooperative, is less forgiving than other conditionally cooperative strategies. Second, it's likely that there are further nuanced differences between subjects conditional on their intelligence which, although not statistically significant individually, when taken together are likely to have a significant impact on behaviour. However, this is speculation.

A potential issue with our approach of using the expected length of an interaction for *all* the interactions we simulate may distort the balance of the value of behavior early versus late in an interaction. It's possible that this might skew the results incorrectly. An alternative approach is to consider simulations where the average length of the interactions is four periods, but each interaction is of *random* length, as determined by the continuation probability $p = 0.75$, exactly as in the experimental design. In the appendix we conduct simulations following this alternative approach, and then analyse the data in the same way as that used to produce Table X. We report marginal effects coefficients that are similar in magnitude and significance to those reported in Table X, and therefore conclude that our marginal effect estimates are reasonable.¹⁵

¹³The results are near identical if the conditional marginal effects are estimated from models 2 or 3.

¹⁴The lack of significance for the subjects' own Raven score is likely a consequence of there being small (and not significant) differences between the *Low* Raven and *Mixed* interactions, as well as between the *High* Raven and *Mixed* interactions. This is despite p -values being close to significant at the 5% and 10% level in some of the regressions. Further, the regressions use Raven score *per se* rather than the *Low/High* classifications.

¹⁵Although not the focus of this study, social preferences have previously been discussed as a potential driver of behaviour in prisoner's dilemma games. Although only exploratory, we can examine this by estimating the

Model	Contributions			Earnings		
	1	2	3	1	2	3
Panel A:						
<i>Leader's Raven</i>	0.053 (0.066)	0.044 (0.073)	0.034 (0.069)	0.043* (0.023)	0.041* (0.024)	0.04* (0.024)
<i>Follower's Raven'</i>	0.053*** (0.005)	0.053*** (0.005)	0.018*** (0.004)	0.064*** (0.005)	0.064*** (0.005)	0.019*** (0.003)
Panel B						
<i>Leader's Raven</i>	0.19*** (0.014)	0.189*** (0.013)	0.077*** (0.008)	0.093*** (0.009)	0.102*** (0.008)	0.108*** (0.008)
<i>Follower's Raven'</i>	0.036 (0.073)	-0.009 (0.073)	-0.006 (0.072)	0.018 (0.02)	0.018 (0.02)	0.018 (0.02)

Note: Robust standard errors in parentheses. *, ** and *** denote significance at the 10%, 5% and 1% level. Standard errors clustered at the subject level, 95 clusters in total in each regression. Coefficients are marginal effects estimated from Tobit regressions. Panel A presents the estimated marginal effect of Raven score on contributions/earnings for *Leaders*; Panel B for *Followers*. In Panel A *Leader's Raven* refers to the subject's own Raven score, and in Panel B, *Follower's Raven* refers to the subject's own Raven score. For the *Leader*, we control for their contribution in Period 1 as well as their beliefs about the *Follower's* contribution. In model 2, we add additional controls for risk aversion, social preferences, gender and their level of agreeableness. In model 3, we add political attitudes and where the experiment was conducted. We use the same controls for the *Follower* in each model, except control for their Period 1 Plan cluster instead of Period 1 contributions, and include their beliefs about the *Leader's* contributions.

Table X: Marginal effect of Leaders' and Followers' Raven score on Contributions and Earnings

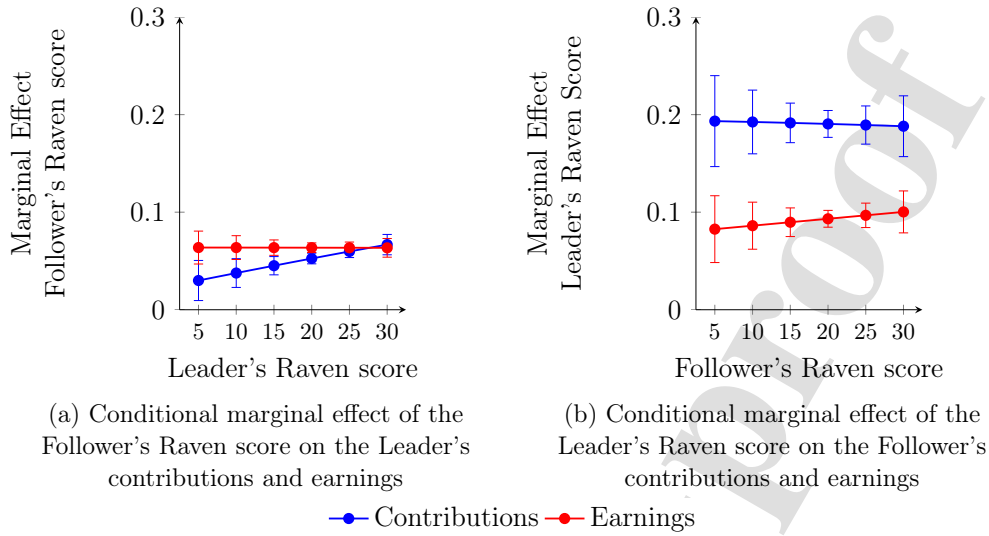
4 Conclusion

We present evidence from an experiment examining how intelligence impacts strategy choice and outcomes in an indefinitely repeated sequential public goods game. Our experiment directly elicits strategies. This approach brings a number of advantages, deepening our understanding of behaviour in social environments where individuals face strong opportunistic incentives.

We report a number of findings. With respect to strategy choices, we find that *Leaders* are less likely to play Free-Riding strategies as their intelligence increases. However, intelligence has no impact on any other types of strategies chosen by *Leaders*, for example Tit-for-Tat. Kosfeld (2020) highlights this strategy as being important for success, albeit in finitely repeated interactions where it is not an equilibrium strategy. We also report evidence that *Followers** are less likely to play Grim Trigger type strategy as their intelligence increases. We therefore find a direct link between intelligence and cooperation levels.

This result is important from an organizational and leadership perspective. In many modern corporations, human resource selection processes involve personality assessments of the employees and it has been shown that personality tests are used in firms' hiring decisions (see Autor & Scarborough (2008)). In our paper, we offer evidence that intelligence tests could

marginal effects of each social preference type on contributions and earnings. The only effect we find that is significant is that a *Leader* classified as being 'spiteful' from the EET make less contributions and profit than those classified as being 'selfish' ($p < 0.05$). We report no significant marginal effects for *Followers*.



Note: Vertical bars represent 95% confidence intervals. Conditional marginal effects estimated from Model 1 in Table X. Left hand figure shows the marginal effect of the *Follower's* Raven score estimated at different Raven scores for the *Leader*. The right hand figure shows the marginal effect of the *Leader's* Raven score estimated at different Raven scores for the *Follower*.

Figure 14: The interaction between Leader's and Follower's Raven scores

be an important tool for human resource managers and leaders aiming to design teamwork incentives. In particular, as Raven scores could be used to identify individuals who are most likely to cooperate, this seems to be a simple way for organizational leaders to choose the most cooperative followers (Kosfeld 2020).

Our results also draw a direct link between intelligence and earnings. Groups comprised of *High* Raven individuals cooperate more and earn significantly more than *Low* Raven groups. We also find evidence that intelligence levels interact, with *Leaders* contributing more to the public good the higher the *Follower's* level of intelligence. Our analysis shows that our findings are not a consequence of initial beliefs or first round contribution levels, and therefore rules out path dependency as an explanation. Our novel experimental design also rules out both belief learning and learning how to play the game within the experiment as potential confounds. This result is robust to controlling for relevant preference differences, including risk and social preferences, personality differences and demographics.

As with any study, this one also has its limitations. As our focus is on strategy choices in an indefinitely repeated game that is played only once, and in a setting where subjects did not experience the game under a more 'standard' implementation, subjects behaviour may differ to their behaviour situation in which they do experience a 'standard' implementation. Despite this, there exists a large literature that examines the decisions of inexperienced subjects, with limited work that examines the role of intelligence in this context. Our study therefore represents an important step in understanding how intelligence effects play without experience. Future work could examine the extent to which the results here might depend on the manner in which strategies are elicited.

Taken together, our results provide a useful guidance for policies within organisations char-

acterised by a hierarchical structure. In particular, we show that the success of cooperation and teamwork in leader-follower settings is heavily dependent on the group composition in terms of their cognitive skills. To this end, our findings give rise to a fruitful research agenda. Future work could seek to examine the robustness of our findings using varying incentives, continuation rules and intelligence measures. This would help to bolster both our findings, and those in the literature.

References

- Anderson, L. R. & Mellor, J. M. (2009), ‘Are risk preferences stable? comparing an experimental measure with a validated survey-based measure’, *Journal of Risk and Uncertainty* **39**(2), 137–160.
- Aoyagi, M., Bhaskar, V. & Fréchette, G. R. (2019), ‘The impact of monitoring in infinitely repeated games: Perfect, public, and private’, *American Economic Journal: Microeconomics* **11**(1), 1–43.
- Aoyagi, M. & Fréchette, G. (2009), ‘Collusion as public monitoring becomes noisy: Experimental evidence’, *Journal of Economic theory* **144**(3), 1135–1165.
- Autor, D. H. & Scarborough, D. (2008), ‘Does job testing harm minority workers? evidence from retail establishments’, *The Quarterly Journal of Economics* **123**(1), 219–277.
- Axelrod, R. (1980), ‘Effective choice in the prisoner’s dilemma’, *Journal of conflict resolution* **24**(1), 3–25.
- Bosch-Rosa, C., Meissner, T. & Bosch-Domènech, A. (2018), ‘Cognitive bubbles’, *Experimental Economics* **21**(1), 132–153.
- Brandts, J. & Charness, G. (2011), ‘The strategy versus the direct-response method: a first survey of experimental comparisons’, *Experimental Economics* **14**(3), 375–398.
- Breitmoser, Y. (2015), ‘Cooperation, but no reciprocity: individual strategies in the repeated prisoner’s dilemma’, *American Economic Review* **105**(9), 2882–2910.
- Burke, H. R. (1958), ‘Raven’s progressive matrices: A review and critical evaluation’, *The Journal of Genetic Psychology* **93**(2), 199–228.
- Camerer, C. F., Ho, T.-H. & Chong, J.-K. (2002), ‘Sophisticated experience-weighted attraction learning and strategic teaching in repeated games’, *Journal of Economic theory* **104**(1), 137–188.
- Cappelen, A. W., Reme, B.-A., Sørensen, E. Ø. & Tungodden, B. (2016), ‘Leadership and incentives’, *Management Science* **62**(7), 1944–1953.
- Cattell, R. B. (1963), ‘Theory of fluid and crystallized intelligence: A critical experiment.’, *Journal of educational psychology* **54**(1), 1.
- Chaudhuri, A. (2011), ‘Sustaining cooperation in laboratory public goods experiments: a selective survey of the literature’, *Experimental economics* **14**(1), 47–83.
- Chen, C. (2015), ‘Intelligence moderates reinforcement learning: a mini-review of the neural evidence’, *Journal of neurophysiology* **113**(10), 3459–3461.
- Choi, Y. & Mai-Dalton, R. R. (1998), ‘On the leadership function of self-sacrifice’, *The Leadership Quarterly* **9**(4), 475–501.

- Costa-Gomes, M., Crawford, V. P. & Broseta, B. (2001), 'Cognition and behavior in normal-form games: An experimental study', *Econometrica* **69**(5), 1193–1235.
- Dal Bó, P. (2005), 'Cooperation under the shadow of the future: experimental evidence from infinitely repeated games', *American economic review* **95**(5), 1591–1604.
- Dal Bó, P. & Fréchette, G. R. (2011), 'The evolution of cooperation in infinitely repeated games: Experimental evidence', *American Economic Review* **101**(1), 411–29.
- Dal Bó, P. & Fréchette, G. R. (2018), 'On the determinants of cooperation in infinitely repeated games: A survey', *Journal of Economic Literature* **56**(1), 60–114.
- Dal Bó, P. & Fréchette, G. R. (2019), 'Strategy choice in the infinitely repeated prisoner's dilemma', *American Economic Review* **109**(11), 3929–52.
- De Cremer, D. (2006), 'Affective and motivational consequences of leader self-sacrifice: The moderating effect of autocratic leadership', *The Leadership Quarterly* **17**(1), 79–93.
- Figuières, C., Masclet, D. & Willinger, M. (2012), 'Vanishing leadership and declining reciprocity in a sequential contribution experiment', *Economic Inquiry* **50**(3), 567–584.
- Fischbacher, U. (2007), 'z-tree: Zurich toolbox for ready-made economic experiments', *Experimental economics* **10**(2), 171–178.
- Fischbacher, U., Gächter, S. & Fehr, E. (2001), 'Are people conditionally cooperative? evidence from a public goods experiment', *Economics letters* **71**(3), 397–404.
- Foulds, G. & Raven, J. C. (1950), 'An experimental survey with progressive matrices (1947)', *British Journal of Educational Psychology* **20**(2), 104–110.
- Frederick, S. (2005), 'Cognitive reflection and decision making', *Journal of Economic Perspectives* **19**(4), 25–42.
- Fudenberg, D., Rand, D. G. & Dreber, A. (2012), 'Slow to anger and fast to forgive: Cooperation in an uncertain world', *American Economic Review* **102**(2), 720–49.
- Ghidoni, R. & Suetens, S. (2020), 'The effect of sequentiality on cooperation in repeated games', *American Economic Journal: Microeconomics*.
- Gill, D. & Prowse, V. (2016), 'Cognitive ability, character skills, and learning to play equilibrium: A level-k analysis', *Journal of Political Economy* **124**(6), 1619–1676.
- Greiner, B. (2015), 'Subject pool recruitment procedures: organizing experiments with orsee', *Journal of the Economic Science Association* **1**(1), 114–125.
- Gächter, S. & Renner, E. (2018), 'Leaders as role models and 'belief managers' in social dilemmas', *Journal of Economic Behavior & Organization* **154**, 321 – 334.
- Haigner, S. D. & Wakolbinger, F. (2010), 'To lead or not to lead: Endogenous sequencing in public goods games', *Economics Letters* **108**(1), 93–95.

- Harrison, G. W., Rutström, E. E. et al. (2008), 'Risk aversion in the laboratory', *Research in experimental economics* **12**(8), 41–196.
- Holt, C. A. & Laury, S. K. (2002), 'Risk aversion and incentive effects', *American economic review* **92**(5), 1644–1655.
- Hyndman, K., Ozbay, E. Y., Schotter, A. & Ehrblatt, W. Z. (2012), 'Convergence: an experimental study of teaching and learning in repeated games', *Journal of the European Economic Association* **10**(3), 573–604.
- Jack, B. K. & Recalde, M. P. (2015), 'Leadership and the voluntary provision of public goods: Field evidence from bolivia', *Journal of Public Economics* **122**, 80–93.
- Kartal, M. & Müller, W. (2021), 'A new approach to the analysis of cooperation under the shadow of the future: Theory and experimental evidence', *Available at SSRN 3222964*.
- Kerschbamer, R. (2015), 'The geometry of distributional preferences and a non-parametric identification approach: The equality equivalence test', *European economic review* **76**, 85–103.
- Kosfeld, M. (2020), 'The role of leaders in inducing and maintaining cooperation: The cc strategy', *The Leadership Quarterly* **31**(3), 101292. Economics and Leadership.
- Kumru, C. S. & Vesterlund, L. (2010), 'The effect of status on charitable giving', *Journal of Public Economic Theory* **12**(4), 709–735.
- Macy, M. W. & Flache, A. (2002), 'Learning dynamics in social dilemmas', *Proceedings of the National Academy of Sciences* **99**(suppl 3), 7229–7236.
- Proto, E., Rustichini, A. & Sofianos, A. (2019), 'Intelligence, personality, and gains from cooperation in repeated interactions', *Journal of Political Economy* **127**(3), 1351–1390.
- Proto, E., Rustichini, A. & Sofianos, A. (2021), 'Intelligence, Errors, and Cooperation in Repeated Interactions', *The Review of Economic Studies* **89**(5), 2723–2767.
- Rivas, M. F. & Sutter, M. (2011), 'The benefits of voluntary leadership in experimental public goods games', *Economics Letters* **112**(2), 176 – 178.
- Romero, J. & Rosokha, Y. (2018), 'Constructing strategies in the indefinitely repeated prisoner's dilemma game', *European Economic Review* **104**, 185–219.
- Roth, A. E. & Murnighan, J. K. (1978), 'Equilibrium behavior and repeated play of the prisoner's dilemma', *Journal of Mathematical psychology* **17**(2), 189–198.
- Rousseeuw, P. J. (1987), 'Silhouettes: a graphical aid to the interpretation and validation of cluster analysis', *Journal of computational and applied mathematics* **20**, 53–65.
- Rustichini, A. (2015), 'The role of intelligence in economic decision making', *Current Opinion in Behavioral Sciences* **5**, 32–36.

- Sahin, S. G., Eckel, C. & Komai, M. (2015), 'An experimental study of leadership institutions in collective action games', *Journal of the Economic Science Association* **1**(1), 100–113.
- Selten, R., Mitzkewitz, M. & Uhlich, G. R. (1997), 'Duopoly strategies programmed by experienced players', *Econometrica: Journal of the Econometric Society* pp. 517–555.
- Sternberg, R. J. & Kaufman, S. B. (2011), *The Cambridge handbook of intelligence*, Cambridge University Press.
- Sundet, J. M., Barlaug, D. G. & Torjussen, T. M. (2004), 'The end of the flynn effect?: A study of secular trends in mean intelligence test scores of norwegian conscripts during half a century', *Intelligence* **32**(4), 349–362.
- Vostroknutov, A., Polonio, L. & Coricelli, G. (2018), 'The role of intelligence in social learning', *Scientific reports* **8**(1), 1–10.
- Wen, Q. (2002), 'A folk theorem for repeated sequential games', *The Review of Economic Studies* **69**(2), 493–512.