


THE COOPERATION PUZZLE: WHAT REALLY DRIVES PEOPLE TO COOPERATE?

Ivanna Piven and Shrey Patel

Supervised by Dr Konstantinos Ioannidis



Laidlaw Undergraduate Research and Leadership Programme 2025



Introduction

Human cooperation, broadly categorised by the collaboration and investment of individuals, is a costly act consuming time or – in this case – money. In the context of public goods games, ‘cooperation’ (or ‘contribution’) is exhibited when individuals give their own money¹ towards the communal ‘public good’. The sum of the money in this public good can then be multiplied and the money distributed between all participants – rather crucially – regardless of whether they ‘cooperated’ in the first instance.

However, in the context of economic games, being selfish or a ‘free rider’ (meaning one does not cooperate) is a profit-maximising strategy where one can rely on the contributions of others for their own gain. We even observe this evolutionarily: being selfish is essentially a survival instinct designed for us to retain our resources. Despite this cooperative and altruistic action has nevertheless persisted and seems to be deeply engrained within human behaviour. This paradoxical phenomenon has become a central question in behavioural economics, begging the question of why cooperation is still so widespread in economic games such as the public goods games and what factors influence these cooperative behaviours.

Methodology

Our starting point was recognising that “Why do humans cooperate?” is far too broad to tackle in six weeks. To refine the scope, we used Zelmer’s (2003) meta-analysis as a baseline. Zelmer synthesised 27 studies up to 2001, highlighting factors such as marginal per capita return (MPCR), communication, group composition, framing, subject experience, and belief elicitation as significant determinants of cooperation.

We asked: *Since 2003, have new experiments changed our understanding of which variables affect contribution rates in linear public goods games?* Specifically:

- Do design features such as communication, punishment, framing, and group structure consistently increase cooperation?
- Do subject characteristics like prior experience, cultural background, or online vs. lab environments moderate behaviour?

¹ ‘money’ referring to either genuine money or tokens distributed for the purpose of the PGG.

- Are Zelmer’s “stylised facts” still robust in the modern experimental landscape?

By narrowing to *linear voluntary contribution mechanism (VCM) experiments with MPCR < 1*, we ensured comparability across studies.

Search Strategy and Identification of Studies

We applied the PRISMA (Preferred Reporting Items for Systematic Reviews and Meta-Analyses) framework to systematically search, identify, and screen relevant studies.

Databases and sources:

- Google Scholar (primary database).
- EconLit, IDEAS

We experimented with different search terms such as combinations of the following

- “Public goods game” AND “experiment”
- “Voluntary contribution mechanism” AND “experiment”
- “Cooperation” AND “public good”

These were combined with publication year filters (2003–2025) and English language.

Yield:

- First 20 pages of Google Scholar results produced 200 records.
- After removing duplicates: 159 records remained.
- Screening by title/abstract excluded 98 (e.g. theoretical papers, field studies, or non-linear designs).
- 61 full texts were retrieved and assessed.
- We included the first 20 that met all criteria and these were included in the final dataset.

This process is summarised in the PRISMA flowchart included in our poster.

Inclusion and Exclusion Criteria

We designed inclusion/exclusion rules to balance rigour with feasibility.

Inclusion:

- Laboratory experiments using linear VCMs with a single public good.

- Reported group-level or treatment-level contribution data (mean % of endowment contributed).
- MPCR strictly less than 1.
- Published after 2003, peer-reviewed or working papers.
- Full text accessible.

Exclusion:

- Non-linear public goods games or threshold/step-return designs.
- Multi-public-good settings or hybrid games.
- Papers without extractable outcome data.
- Purely theoretical or commentary articles.

Where reporting was ambiguous, we followed Zelmer's principle of *imputation rules*: for example, assuming subjects were inexperienced unless otherwise stated. This ensured consistency across studies.

Screening and Data Extraction

For the first three studies (Gächter & Herrmann 2011, Arechar 2018, and Fehr & Gächter 2000), we both independently screened abstracts and coded variables to minimise bias. We were given full freedom to code everything we felt was important and we both coded with over 90% similarity, giving us confidence our results would be reproducible and consistent. For any disputes, our supervisor kindly also coded the same data and we referred to his judgement. The takeaways we got from these meetings was:

- When multiple treatments were reported (e.g. baseline, communication, punishment), we entered them separately as treatment-level observations.
- When variance measures were missing, we relied on reported means only (acknowledging reduced statistical precision).
- If contributions were reported over time, we coded average contributions across all periods, while noting end-game decay separately.

After further discussion with our supervisor, we created preliminary variables to code based off the first few studies, and we were advised to add extra variables if they were reported frequently. We also included extra details for each study, in case we added extra variables – to avoid having to revisit each previous study. We also took inspiration from Zelmer's variables for comparison. The final variables included:

- **Study characteristics:** Author, year, publication status, location, venue (lab/online).
- **Design factors:** group size, MPCR, number of rounds, stranger vs. partner matching, framing, communication, punishment, belief elicitation, subject experience.
- **Sample characteristics:** subject pool (students, children, general population), age, gender share, prior economics training.
- **Incentives:** endowment, currency conversion, show-up fees, average earnings.
- **Outcomes:** mean contribution (% of endowment), sometimes reported by treatment (e.g. communication vs. no-communication).

We then spent a few weeks creating the final dataset, which included 33 variables coded for 76 treatment-level observations, across 20 studies.

Quality Assessment and Bias Considerations

We assessed each study along three dimensions:

1. **Completeness of reporting** – did the paper clearly describe group size, MPCR, and outcome measures?
2. **Internal validity** – were randomisation and incentives robust? Were payments salient?
3. **Risk of bias** – was selective reporting likely (e.g. missing treatment arms)?

Although we did not apply a formal risk-of-bias tool (like Cochrane’s ROB-2, used in medicine), we flagged studies with incomplete reporting.

Publication bias was considered through funnel plots and Egger’s regression (though with only 20 studies, the power of the test was limited).

Statistical Synthesis and Comparison with Zelmer

Following Zelmer (2003), we structured our analysis as a **meta-regression** rather than a pure effect-size aggregation.

- **Dependent variable:** Mean Contribution (% of endowment).
- **Independent variables:** MPCR, group size, number of periods, communication, punishment, framing, subject experience, belief elicitation, online vs. lab environment, cultural region.
- **Weights:** observations weighted by number of groups/subjects to account for sample size differences.

We estimated pooled coefficients using weighted least squares. We used this approach as it allowed us to retain comparability with Zelmer while incorporating new moderators (e.g. online vs. lab, cultural variation).

We also conducted sensitivity analyses: excluding low-quality studies, comparing European vs. non-European samples, and examining first vs. last period contributions.

To collaborate effectively and ensure transparency and reproducibility, we documented all steps:

- **Google Drive** was used for file sharing and version control.
- Annotated notes were kept for each study
- Results were synthesised in the shared Excel sheet, which served as the input for statistical analysis.

This collaborative workflow mirrored best practices in systematic reviews, while remaining feasible within the six-week timeframe.

Compared to Zelmer (2003):

- We restricted to studies after 2003, ensuring we built on rather than duplicated her work.
- We expanded the scope to include online experiments and cross-cultural samples, which were rare two decades ago.
- We placed greater emphasis on coding institutional features like punishment regimes, belief elicitation, and framing.

Limitations remain:

- Small sample size (20 studies) limited statistical power.
- Reporting heterogeneity (e.g. missing variances) constrained the depth of effect-size analysis.
- Potential selection bias from focusing primarily on Google Scholar rather than exhaustively searching EconLit / IDEAS.

Results

Firstly, we found that cooperation between participants in public goods games increased when competitive treatments were applied. Augenblick (2014) explored competitiveness within public goods games, using the format of posting postcards that

urged monetary donations towards political parties. These postcards either contained no motivational message (control group) or information about the previous contributions of those in the same or the opposing political group. Augenblick discovered that postcards with competitive messages induced significantly higher levels of contribution compared to the control group postcards.

Cooperation rates were also highest when individuals within public goods games were aware that people around them were contributing, an indicator of ‘prosocial’ behaviour or conditional cooperation. Arechar’s (2018) and Frey’s (2004) public goods games experiments exhibited pro-social behaviour with evidence of people being more willing to contribute when others contribute. During Frey’s experiment, in one treatment, participants were told that a high percentage of students had contributed to this public fund in the past whereas other participants were told that a low percentage of students had contributed to this same fund in the past. Results show a 2.3% increase in contribution in the first treatment compared to the second treatment. These experiments represent a recurring theme in public goods games with conditional cooperation (people cooperating under the condition that others are too) remaining a primary driver for cooperation in public goods games.

Cooperation was also high with individuals who were allowed to communicate before making their public goods decisions either face-to-face or via an online chatroom. Bochet (2003) ran a public goods game experiment, adjusting degrees of communication with the players before the game began. People either did not communicate, could communicate via chat room with their anonymity preserved or communicated for 5 minutes face-to-face. Bochet discovered that individuals who communicated face-to-face had the highest levels of cooperation (with online chat room communication being a close second) while individuals who did not communicate at all had the lowest levels of cooperation. So, direct communication before rounds of a public goods game – whether face-to-face or online – seems to significantly impact cooperation in some way.

On the contrary, certain factors make individuals refrain from cooperating in public goods games and even enable people to exhibit anti-social behaviours. For instance, some public goods games reveal a pattern between individuals with prior experience of participating in such games or with economic backgrounds cooperating significantly less than individuals with no prior experience can be observed in some public goods game experiments. Conte (2018) selected both experienced and inexperienced individuals for his public goods game experiment, hypothesising that the more experienced participants will contribute less. His results confirmed his hypothesis: he found that the experienced participants systematically contribute smaller amounts than the inexperienced participants and that the beliefs of the experienced participants tended to be more accurate than those of the inexperienced participants. Therefore, having previous experience with participation in public goods games seems to be a factor for less frequent cooperation, likely because these individuals are more familiar with profit-maximising tactics (which include the ‘free rider’ action of not contributing but receiving benefits regardless) and can better predict the decisions of the other participants.

Furthermore, one study showed that soliciting information about how cooperation benefits others can lead to *lower* levels of cooperation, thus providing evidence for anti-social behaviour within the public goods game domain and defying otherwise pro-social expectations. This is the case in Burton-Chellew’s 2012 experiment: in this experiment, some participants were not told about others’ contributions, some were told how much others contributed, and some received an even more detailed breakdown of how much everyone contributed and earned in that round. The participants with this last treatment contributed the least out of everyone. This confirms that having an increased knowledge of the consequences of you cooperating on other’s peoples earnings leads to lower overall levels of cooperation. These observations go against the ‘prosocial preferences hypothesis’ which would suggest that increased information would either not influence or would increase the level of cooperation (Burton-Chellew, 2012) since increased information about other participants – and in particular how their contributions would benefit others – in fact *decreases* contributions.

Additionally, a common pattern among public goods games is that cooperation decreases as the game is being played. Burton-Chellew (2012) noticed the notably fewer contributions in public goods games as the game continues to be played. Similarly, Page (2005) noticed the decay in cooperation with repeated rounds and even proposed regrouping strategies to mitigate this cooperation decrease and sustain substantial levels of cooperation.

Another treatment commonly explored in public goods games is punishment, meaning that participants have the opportunity to punish others within the game, even if it comes at a cost to them². Punishment in public goods games can take several forms but in the literature we have reviewed, this most often means that ‘punished’ individuals are fined, forced to skip a round or are expelled from the game.

Some might expect that the looming possibility of being punished for not cooperating would overall increase the cooperation rate in the game with free riders aware of the consequences of failing to cooperate. This can indeed be observed in several public goods games, such as Rand’s 2011 experiment. When the ‘restricted punishment’ treatment was applied in this experiment (meaning participants could punish others, but only non-cooperators) cooperation rates were the highest at 87%. This might imply that the threat of punishment deters free riders thus overall increasing cooperation.

Although punishment does evidently have cooperation-enhancing effects, sometimes it works conversely and can negatively affect cooperation rates. One observation we noticed was anti-social punishment; this takes the form of punishment being directed at cooperators rather than free riders meaning that cooperation instead of free riding was punishable. Gächter and Herrmann noticed this in their 2011 experiment: in all four subject pools, punishment was primarily directed at *contributors*. This resulted in cooperation being lower in the presence of punishment than in its absence. Gächter and Herrman conclude that in their experiment that punishment had no cooperation-enhancing effect. Similarly, Rand (2011) found that cooperation rates only increased when the ability to punish was restricted to cooperators punishing defectors. However, when the full set of punishment strategies (ability to punish cooperators *as well as*

² Punishing others in a PGG can sometimes cost money (for example, in Gächter’s and Herrman’s 2011 experiment punishing others costs 1 token)

defectors) was available, the positive effect of punishment almost entirely disappeared.

Conclusion

To conclude, the results from our analysis of 20 public goods game experiments represent the factors that influence cooperation in public goods games the most. Across treatments three findings stood out most consistently. First, conditional cooperation remains a central driver; participants contribute more when they believe others are also contributing, underscoring the role of social norms and expectations. Secondly, communication – whether face-to-face or online – substantially increases contributions, demonstrating that opportunities to connect with others before the game begins plays a vital role in maintaining high cooperation levels. Lastly, punishment has dual effect: restricted punishment can sustain cooperation unlike unrestricted punishment which often leads to antisocial dynamics leading to suppressed contributions.

The public goods game format acts like a framework for many real-life scenarios including voting, donating to charity or completing community projects. This is one of the reasons why it is so vital to continue studying the public goods game, due to how much it reveals about the human psyche and how much we can learn about human interaction. Thus, our transparent protocol, coding dictionary, and dataset provide a replicable foundation for future extensions.

References

Anderson, L.R., Mellor, J.M. and Milyo, J. (2004) 'Social capital and contributions in a public-goods experiment', *American Economic Review*, 94(2), pp. 373–376.

Anwar, C.M.S. and Georgalos, K. (2024) 'Position uncertainty in a sequential public goods game: An experiment', Unpublished working paper, University of Nottingham.

Arechar, A.A., Gächter, S. and Molleman, L. (2018) 'Conducting interactive experiments online', *Experimental Economics*, 21(1), pp. 99–131.

Auerswald, H., Schmidt, C., Thum, M. and Torsvik, G. (2018) 'Teams in a public goods experiment with punishment', *Journal of Behavioral and Experimental Economics*, 76, pp. 21–34.

Augenblick, N. and Cunha, J.M. (2014) 'Competition and cooperation in a public goods game: A field experiment', *The Economic Journal*, 124(574), pp. 105–144.

Barron, K. and Nurminen, T. (2020) 'Nudging cooperation in public goods provision', *Journal of Behavioral and Experimental Economics*, 85, 101521.
doi:10.1016/j.socec.2020.101521.

Burton-Chellew, M.N. and West, S.A. (2012) 'Prosocial preferences do not explain human cooperation in public-goods games', *Proceedings of the National Academy of Sciences*, 109(2), pp. 759–764.

Chaudhuri, A. and Paichayontvijit, T. (2017) 'On the long-run efficacy of punishments and recommendations in a laboratory public goods game', *Scientific Reports*, 7, 12286.
doi:10.1038/s41598-017-12490-5.

Conte, A., Levati, M.V. and Montinari, N. (2019) 'Experience in public goods experiments', *Theory and Decision*, 86(1), pp. 65–93.

Dong, Y., Zhang, B. and Tao, Y. (2016) 'The dynamics of human behavior in the public goods game with institutional incentives', *Scientific Reports*, 6, 28809.
doi:10.1038/srep28809.

Eckel, C.C., Fatas, E., Godoy, S. and Wilson, R.K. (2016) 'Group-level selection increases cooperation in the public goods game', *PLoS ONE*, 11(8), e0157670.
doi:10.1371/journal.pone.0157670.

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

Frey, B.S. and Meier, S. (2004) 'Social comparisons and pro-social behavior: Testing "conditional cooperation" in a field experiment', *American Economic Review*, 94(5), pp. 1717–1722.

Gächter, S. and Herrmann, B. (2011) 'The limits of self-governance when cooperators get punished: Experimental evidence from urban and rural Russia', *European Economic Review*, 55(2), pp. 193–210.

Gereke, J., Schaub, M. and Baldassarri, D. (2021) 'Immigration, integration and cooperation: Experimental evidence from a public goods game in Italy', *Journal of Ethnic and Migration Studies*, 47(1), pp. 1–21.

Lotito, G., Migheli, M. and Ortona, G. (2011) 'Is cooperation instinctive? Evidence from the response times in a public goods game', *POLIS Working Papers*, 190, Università del Piemonte Orientale.

Nakagawa, Y., Onuma, Y. and Yamagishi, T. (2022) 'Cultural differences in sustaining cooperation in public goods games: Evidence from Japan', *Asian Journal of Social Psychology*, 25(3), pp. 291–304.

Nguyen, T., Lê, T. and Phan, V. (2022) 'Exploring cooperation and punishment in public goods games in Vietnam', *Frontiers in Psychology*, 13, 800701. doi:10.3389/fpsyg.2022.800701.

Ozono, H., Jin, N., Watabe, M. and Shimizu, K. (2016) 'Solving the second-order free rider problem in a public goods game: An experiment using a leader support system', *Scientific Reports*, 6, 38349. doi:10.1038/srep38349.

Page, T., Putterman, L. and Unel, B. (2005) 'Voluntary association in public goods experiments: Reciprocity, mimicry and efficiency', *The Economic Journal*, 115(506), pp. 1032–1053.

Rand, D.G., Greene, J.D. and Nowak, M.A. (2011) 'Spontaneous giving and calculated greed', *Nature Communications*, 2, 434. doi:10.1038/ncomms1442.