

Article

The Black Box as a Control for Payoff-Based Learning in Economic Games

Maxwell N. Burton-Chellew ^{1,*}  and Stuart A. West ²¹ Department of Economics, University of Lausanne, CH-1015 Lausanne, Switzerland² Department of Biology, University of Oxford, Oxford OX1 3RB, UK

* Correspondence: maxwell.burton@unil.ch

Abstract: The black box method was developed as an “asocial control” to allow for payoff-based learning while eliminating social responses in repeated public goods games. Players are told they must decide how many virtual coins they want to input into a virtual black box that will provide uncertain returns. However, in truth, they are playing with each other in a repeated social game. By “black boxing” the game’s social aspects and payoff structure, the method creates a population of self-interested but ignorant or confused individuals that must learn the game’s payoffs. This low-information environment, stripped of social concerns, provides an alternative, empirically derived null hypothesis for testing social behaviours, as opposed to the theoretical predictions of rational self-interested agents (*Homo economicus*). However, a potential problem is that participants can unwittingly affect the learning of other participants. Here, we test a solution to this problem in a range of public goods games by making participants interact, unknowingly, with simulated players (“computerised black box”). We find no significant differences in rates of learning between the original and the computerised black box, therefore either method can be used to investigate learning in games. These results, along with the fact that simulated agents can be programmed to behave in different ways, mean that the computerised black box has great potential for complementing studies of how individuals and groups learn under different environments in social dilemmas.

Keywords: altruism; asocial control; behavioural economics; conditional cooperation; confusion; directional learning; reinforcement learning; social preferences



Citation: Burton-Chellew, M.N.; West, S.A. The Black Box as a Control for Payoff-Based Learning in Economic Games. *Games* **2022**, *13*, 76. <https://doi.org/10.3390/g13060076>

Academic Editors: Kjell Hausken and Ulrich Berger

Received: 29 September 2022

Accepted: 14 November 2022

Published: 16 November 2022

Publisher’s Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2022 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Understanding human behaviour in social dilemmas is of crucial importance to solving many global issues [1–7]. Experiments using economic games provide a useful tool for investigating social behaviours [8]. By making participants pay for their decisions, experimenters hope to measure social preferences on the assumption that participants pay for preferred outcomes [9]. By using games with repeated decisions (repeated games), experimenters hope to measure how individuals respond to the behaviours of others (social responses) [10–18]. However, in repeated games, social responses can be confounded by individuals responding to their payoffs and learning how to play the game (payoff-based learning) [19–23]. This problem is particularly acute if many participants start the experiment without fully understanding the game’s payoffs [24–28]. Consequently, experimental control treatments are required to control for potentially confounding factors such as payoff-based learning.

One solution to this problem of confounding is to make individuals face the same decision but in a low-information environment stripped of all social concerns (“asocial controls”) [19,25,27,29–33]. For example, Burton-Chellew et al. introduced the black box method as an “asocial control” to decouple social responses from payoff-based learning in repeated public-goods games [19–21,34]. Specifically, individuals interacted with a virtual black box, with which they could make voluntary inputs of “virtual coins” to

obtain uncertain returns over several rounds¹. However, in reality, the experiment actually involved groups of real participants playing a typical public goods game, with all the usual payoffs and social connections, just unknowingly². By repeating the black box game for multiple rounds, one could measure how inputs evolved in populations of ignorant individuals with no social concerns [19–21]. The black box thus aimed to capture the psychology of self-interested but ignorant/confused individuals that use trial and error learning to improve their earnings. In this way, it was consistent with a rich history of prior studies that investigated how individuals learn in low-information environments and how reinforcement learning can affect cooperation [35–42].

The black box as an asocial control provided an alternative null hypothesis, empirically derived from behavioural observations, to the usual theoretical null hypotheses of a population of perfectly rational and selfish agents (*Homo economicus*). This “baseline” measure could then be compared to behaviour in versions of the normal, “revealed” public goods game to test if the addition of social information affected aggregate behaviour. Burton-Chellaw and West’s original results showed that aggregate contributions in the black box treatment were largely indistinguishable from those in the standard “revealed” public goods game, where individuals can observe their groupmates’ decisions, consistent with models of payoff-based learning [19]. In both cases, despite the income maximising decision in the one-shot version of the game being to contribute 0%, initial levels of contributions averaged around 40–50%, before gradually declining to approximately 15% by round 16, the final round. In both cases, most individuals contributed 0% in the final round, but approximately four percent of individuals still contributed fully. While these similarities did not confirm that individuals were using payoff-based learning, they did mean that one could not reject the null hypothesis of self-interest unless one assumed the participants perfectly understood the revealed game and that the similar levels of cooperation were mere coincidence.

Although there were large similarities between the black box results and the typical results, it is important to keep in mind that the black box was providing a simplified model of behaviour based on the extreme assumption that all players are ignorant/confused and respond only to their own payoffs [19]. However, the black box also allows for examinations of how individuals learn and for estimating parameters within explicit hypothesised learning rules [20,21]. For example, subsequent collaborations with H. Nax and H. Peyton Young analysed individual-level data to estimate how much individuals value the earnings of their groupmates [20] and how individuals use payoff-based learning in the non-social and two social settings [21]. Not surprisingly, some differences in behaviour were found across the three treatments (the black box is, after all, like *Homo economicus*, a rather extreme hypothesis/model). When individuals could observe their groupmates’ decisions, they showed some conditional responses, but only if they could not also observe their groupmates’ payoffs (which is technically redundant information if individuals fully understand the game). However, payoff-based learning was significant in all three treatments [20], including even the manner of such learning [21], suggesting participants were motivated to try and increase their own income in all game forms. Together, these results suggested that conditional cooperation was more a function of social learning rather than a social preference for equal outcomes, although learning and social preferences may also interact [43].

The black box can also be easily modified and adapted to test different hypotheses. For example, Burton-Chellaw and West [32] also subsequently used the black box to show that payoff-based learning is impeded when either group size (N) or the marginal per capita return from contributing is large (MPCR). By testing behaviour in three different black boxes that varied in either group size ($N = 3$ or 12) or the marginal per capita return from contributing ($MPCR = 0.4$ or 0.8), they showed that a large group size and/or a high MPCR and thus reduces the correlation between personal contributions and personal payoffs, thereby impeding payoff-based learning and potentially explaining why the rate of decline in cooperation varies across studies. They confirmed this hypothesis with a

comparative analysis that compared the rates of decline in 237 published public goods games. They found that rates of decline in contributions were slower when either group size or MPCR was large, and more specifically, when the estimated correlation between personal contributions and personal payoffs was weaker, a principle proved in their black box experiment [32].

However, one potential issue with the black box is that because participants are interacting, albeit unknowingly, the learning of one participant changes the learning environment for other participants. While this is also true for revealed social games, it may complicate efforts to discern individual learning from collective learning [44]. Another possible issue is that individuals do not know they can provide benefits to other participants, which may raise ethical concerns for some reviewers (however we do not think this omission of externalities constitutes deception).

Here, we present a modified black box method that solves these two potential issues. Our solution is to make individuals still interact with a black box but change the set-up so that individuals are grouped not with each other but with computerised players (computerised black box) (Figure 1). Otherwise, the set-up remains the same for the participants. There are several advantages to this approach: (1) individuals do not affect other participants, and thus can be treated as independent data points, providing more statistical power for given costs; (2) the learning environment can be maintained constant; (3) individuals are not affecting each other's payoffs, thereby removing any potential ethical concerns; (4) computerized players receive no earnings making the study of behaviour in large groups more affordable; and (5) computerised players can be programmed to play in different, interesting ways, allowing one to test various hypotheses that would otherwise be unfeasible without using deception.

We replicate the experimental design from Burton-Chellew and West, 2021, which used three different black boxes that varied in either group size or the cost of contributing to create one “easy” learning condition with a small group size and low MPCR ($N = 3$ and $MPCR = 0.4$) and two “difficult” learning conditions with either a large MPCR ($N = 3$ and $MPCR = 0.8$) or a large group size ($N = 12$ and $MPCR = 0.4$) [32]. Individuals could input 0–20 virtual coins in each round. However, instead of connecting human participants together, here we use computerised groupmates (programmed to input a random integer drawn from a uniform distribution of 0–20 coins). The payoff formula remained identical for all rounds and was the same for the human and computerised black boxes. This allowed us to compare rates of learning in the two methods, depending on both group size and MPCR. If behaviour with computerised black boxes qualitatively replicates behaviour with human black boxes, then the computerised black box method can be used as a complementary method to test hypotheses without any concerns about participants affecting each other's behaviour and/or earnings.

We also address a related research question on payoff-based learning. As mentioned above, Burton-Chellew and West, 2021, previously showed that payoff-based learning is impeded when groups are large or MPCR is high [32]. In such conditions, participants still contributed around 50% at the end of 16 rounds, despite the Nash equilibrium being 0%, indicative of zero learning. Nevertheless, it may be that individuals just need more time to learn in these challenging conditions and will eventually learn not to contribute. To test this, we repeated the black boxes with the difficult learning conditions, but under two conditions, a short game and a long game (16 versus 40 rounds).

In all cases, we measured learning in two ways; (1) how quickly incentivised contributions converged towards the Nash equilibrium of 0 contributions, and (2) by asking participants at the end of the experiment to report their belief about what was the best number to contribute (“input”) into the black box (this was unincentivised).

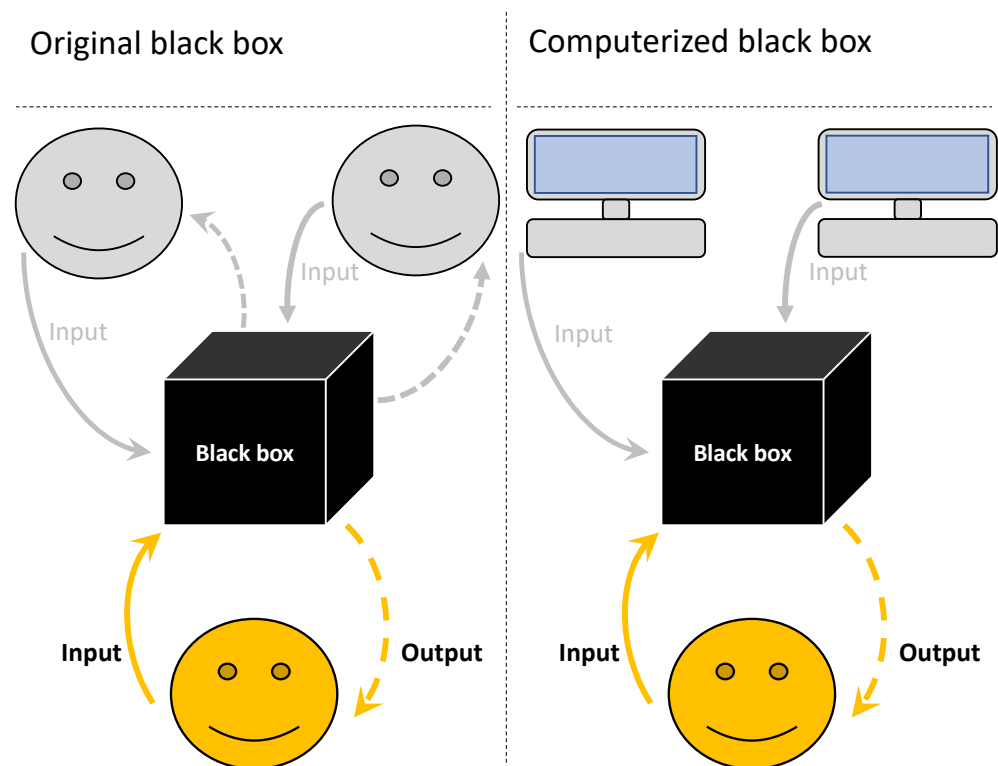


Figure 1. Black box methodologies. In the original black box, participants are connected online and interact in the usual experimental manner for economic games. However, by “black-boxing” the social aspect of the game or the game’s rules and payoffs, one can investigate how participants learn under certain conditions. If one is concerned about individuals affecting either the learning or the payoffs of other participants, one can replace the focal player’s interaction partners with programmed computerised/virtual players (computerised black box). This also allows for more control over the learning environment, as partners can be programmed to be more/less cooperative, etc.

2. Results

2.1. Learning with Hidden Humans or Hidden Computers

We found that there was no significant difference in contributions (“inputs”) depending upon whether individuals were grouped with humans or computers. The rate of decline in contributions across all 16 rounds of the short games did not significantly differ between the original black box with humans and the computerised black box in any of the three black boxes (Figure 2; Table 1). Specifically, the game round \times groupmates interaction was non-significant in all three black boxes (generalised linear mixed models controlling for autocorrelation among groups/individuals: when $N = 3$ and $MPCR = 0.4$, $Z = 0.2$, $p = 0.821$; when $N = 3$ and $MPCR = 0.8$, $Z = -1.2$, $p = 0.222$; and when $N = 12$ and $MPCR = 0.4$, $Z = -0.3$, $p = 0.760$, Table 1).

As an additional check, we also compared final round contributions (“inputs”), which could be argued to be the best measure of learning. Again, we found no significant differences between playing with humans or with computers in any of the three black boxes (Table 2). Specifically, when learning was easy ($N = 3$ and $MPCR = 0.4$), mean \pm SE final round inputs (0–20 virtual coins) were 3.0 ± 0.57 coins with humans and 4.7 ± 0.90 coins with computers (Wilcoxon rank-sum test: $W = 537$, $p = 0.797$). When learning was difficult, mean final round inputs were typically around 50% (10 virtual coins) with both humans and computerised groupmates ($N = 3$ and $MPCR = 0.8$, with humans = 10.1 ± 0.91 coins, with computers = 9.4 ± 1.12 coins, $W = 573.5$, $p = 0.563$; $N = 12$ and $MPCR = 0.4$, with humans = 10.3 ± 0.89 coins, with computers = 9.7 ± 1.12 coins, $W = 128$, $p = 0.806$).

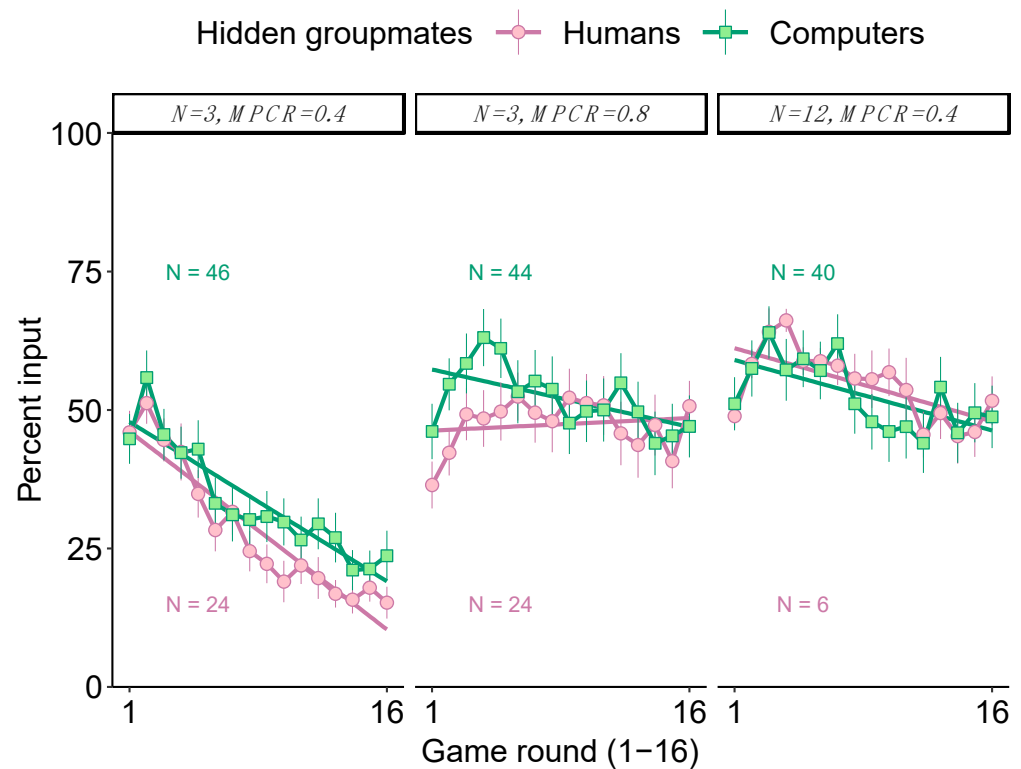


Figure 2. Learning in a black box, either with hidden human or computerised groupmates. We varied both the group size (N) and the benefit of contributing ($MPCR$) across three black boxes. Participants played all three black boxes, in counter-balanced order, but here we only show naïve behaviour (their first black box). Data show the mean contribution per round, with 95% confidence intervals based on the group means (in games with computers, the independent group is just one individual). The rate of learning was broadly similar regardless of playing with humans or computers. The linear regressions do not account for random effects/repeated measures and are therefore for illustration purposes only. The figures are annotated with the sample sizes of independent replicates (groups of humans or individuals grouped with computers).

Table 1. Contributions over time. Analysis of how contributions (“inputs”) change during the game for each black box depending on if groupmates were humans or computers. Generalised linear mixed model with a binomial logit link, and random intercepts for both groups and individuals and random slopes for individuals.

Fixed Effects	N = 3, MPCR = 0.4		N = 3, MPCR = 0.8		N = 12, MPCR = 0.4	
	Z	p	Z	p	Z	p
Intercept (humans)	0.7	0.476	−1.4	0.159	3.6	<0.001
Round	−8.0	<0.001	2.0	0.046	−1.0	0.298
Groupmates (computers)	0.3	0.795	2.2	0.030	−0.1	0.896
Round x Groupmates	0.2	0.821	−1.2	0.222	−0.3	0.760
N. obs.	1888		1856		1792	
N. individuals	118		116		112	
N. groups	70		68		46	
Random effects	Variance	St. dev.	Variance	St. dev.	Variance	St. dev.
Individual intercept	1.12	1.059	1.36	1.168	3.00	1.733
Individual slope	0.03	0.167	0.07	0.255	0.02	0.170
Group intercept	0.60	0.777	1.08	1.038	0.00	0.00

Table 2. Final contributions. Comparison of mean final round contributions (“inputs”) of virtual coins into the black box (0–20 coins). Comparisons made with Wilcoxon rank-sum test.

Black Box	Input: Humans	Input: Computers (Short)	Input: Computers (Long)	W ¹	P ¹	W ²	P ²
N = 3, MPCR = 0.4	3.0 ± 0.57	4.7 ± 0.90	/	573	0.797	/	/
N = 3, MPCR = 0.8	10.1 ± 0.91	9.4 ± 1.12	6.2 ± 0.97	573.5	0.563	1287.5	0.025
N = 12, MPCR = 0.4	10.3 ± 0.89	9.7 ± 1.12	5.4 ± 0.98	128	0.806	1238.5	0.006

¹ Comparing final inputs with human or computerised groupmates. ² Comparing final inputs in short or long games.

We also asked the participants at the end of the 16 rounds if they thought there was a best number to input and if so, what it was (methods). Again, there were no significant differences depending on whether playing with humans or with computers (Figure 3; Table 3). Specifically, the mean ± SE stated beliefs (0–20 coins) for when N = 3 and MPCR = 0.4 were 1.5 ± 0.79 coins with humans and 3.0 ± 0.80 coins with computers (Wilcoxon rank-sum test, $W = 348$, $p = 0.072$); for when N = 3 and MPCR = 0.8, they were 8.8 ± 1.32 coins with humans and 8.9 ± 1.59 coins with computers ($W = 469.5$, $p = 0.888$); and for when N = 12 and MPCR = 0.4, were 9.8 ± 1.38 coins with humans and 8.0 ± 1.37 coins with computers ($W = 387$, $p = 0.349$).

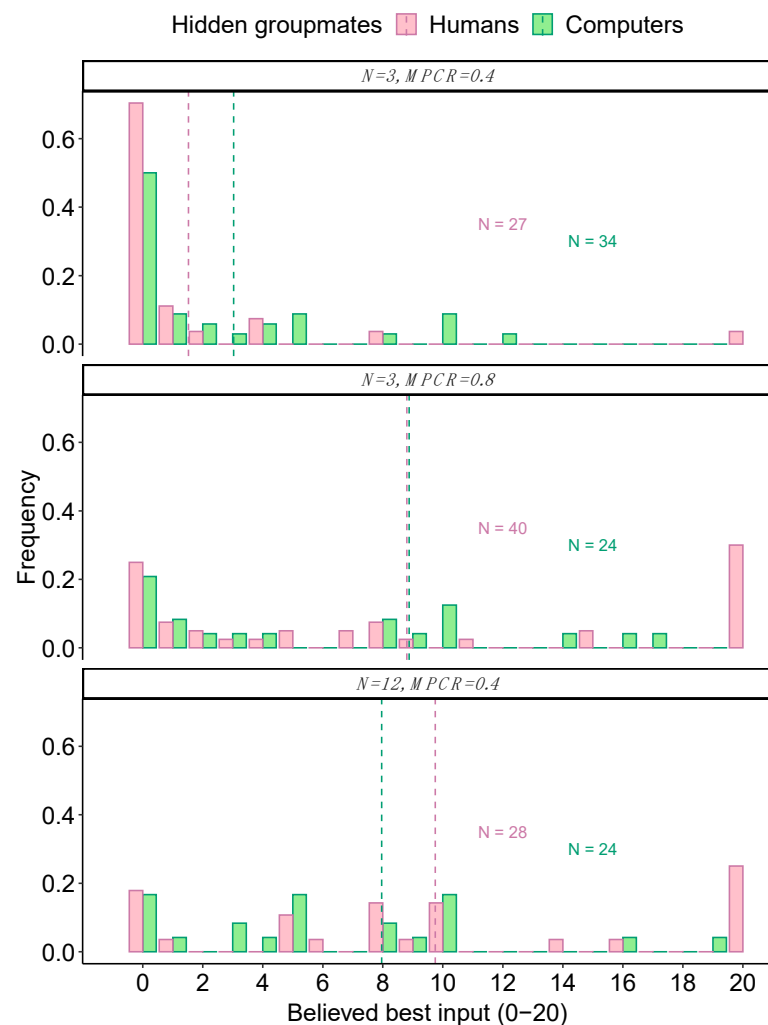


Figure 3. Groupmates and beliefs. Histograms show the frequency of each stated belief about what was the best number to input into the black box. Dashed vertical lines show the mean response. All responses are from naïve participants after finishing their first black box. The figures are annotated with the number of individuals.

Table 3. Beliefs about the best number. The mean \pm SE value participants stated as the best number to input at the end of the game. Comparisons made with Wilcoxon rank-sum test.

Black Box	Humans (N) *	Computers— Short (N)	Computers— Long (N)	W ¹	P ¹	W ²	P ²
N = 3, MPCR = 0.4	1.5 \pm 0.79 (27) *	3.0 \pm 0.80 (34)	/	348	0.072	/	/
N = 3, MPCR = 0.8	8.8 \pm 1.32 (40)	8.9 \pm 1.59 (24)	3.5 \pm 1.10 (24)	469.5	0.888	410.5	0.010
N = 12, MPCR = 0.4	9.8 \pm 1.38 (28) *	8.0 \pm 1.37 (24)	4.0 \pm 0.93 (30)	387	0.349	496.5	0.016

* An error prevented data collection from some participants in the first seven sessions with human groupmates.

¹ Comparing beliefs after play with human or computerised groupmates. ² Comparing beliefs after short or long games.

Overall, we found no significant differences between either inputs or beliefs, depending on if individuals were grouped with humans or computers. Our experiments with computerised groupmates replicated the results from the prior study with human groupmates [32]. Rates of learning were qualitatively similar regardless of groupmates being humans or computers in all three black box settings (Figure 2). These results mean that the original black box with humans can be used without having to worry too much about collective learning, or alternatively that the new, computerized, black box method can be used in certain contexts to obtain qualitatively similar results. However, we caution that for the “easy” black box (N = 3 and MPCR = 0.4), the final round inputs and the post-game beliefs about the value of the best input were lower, but not significantly, in the human black box. Looking at Figure 2, it may be that the learning rates for the “easy” black box (N = 3 and MPCR = 0.4) would have diverged if the experiment had continued for longer than 16 rounds, but we find no statistical support for this prediction within our data.

2.2. Learning in Longer Games

We found clear evidence of payoff-based learning in the long-run games (Figure 4). Overall, the estimated rate of decline was significantly negative in both black boxes (Table 4, generalised linear mixed model controlling for individual: when N = 3 and MPCR = 0.8, $Z = -2.8$, $p = 0.005$; when N = 12 and MPCR = 0.4, $Z = -3.3$, $p < 0.001$, depending on black box). However, for both black boxes, the rate of decline was not significantly different between the short and long games, suggesting that the rate of learning is relatively constant within these time frames despite being undetectable in the short games (Table 4, round \times game length interaction: N = 3 and MPCR = 0.8, $Z = 1.0$, $p = 0.308$; N = 12 and MPCR = 0.4, $Z = 0.5$, $p = 0.643$).

Table 4. The effect of game length. Analysis of how inputs change during the game for each black box depending on game length (16 or 40 rounds). Generalised linear mixed model with a binomial logit link and random intercepts and slopes for individuals.

Fixed Effects	N = 3, MPCR = 0.8		N = 12, MPCR = 0.4	
	Z	p	Z	p
Intercept	1.1	0.257	1.7	0.080
Round	−2.8	0.005	−3.3	<0.001
Game length (short)	0.6	0.575	0.4	0.666
Round \times Game length	1.0	0.308	0.5	0.643
N. obs.	2544		2480	
N. individuals	90		86	
N. groups	90		86	
Random effects	Variance	St. dev.	Variance	St. dev.
Individual intercept	3.00	1.731	5.67	2.381
Individual slope	0.01	0.116	0.01	0.122

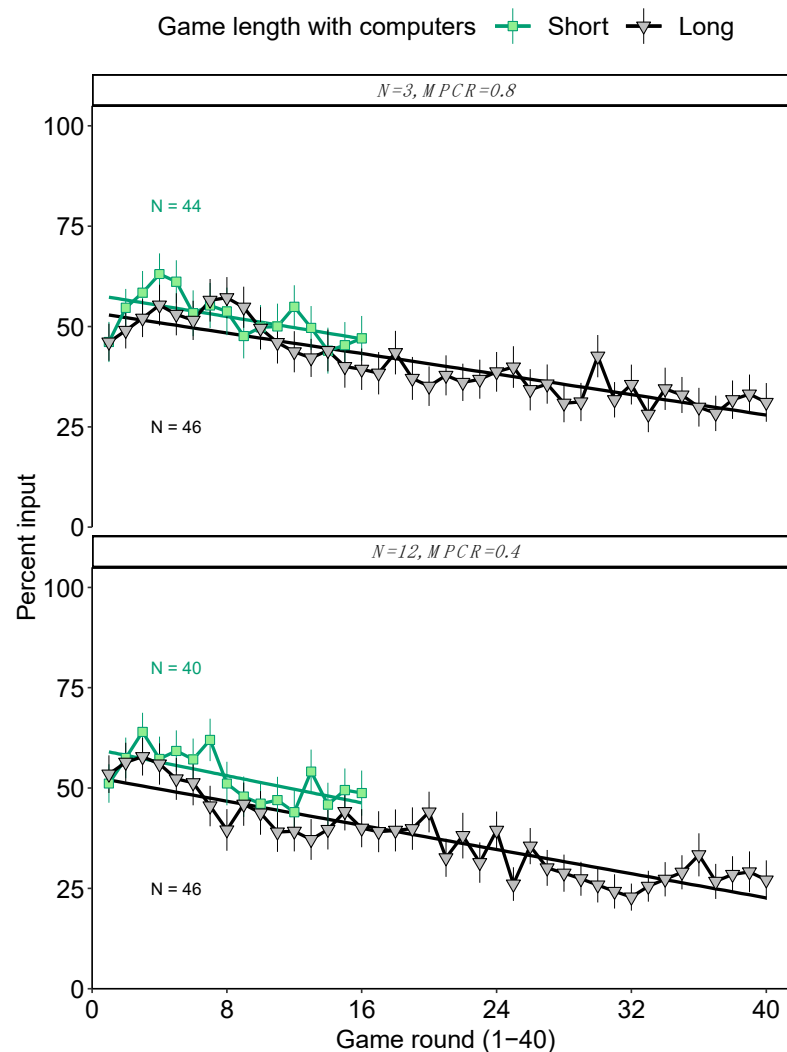


Figure 4. Learning and game length. The green data are the same as in Figure 2. Data show mean contributions with 95% confidence intervals, depending on game length, for two different black box parameter settings. The rate of learning was broadly similar in both black boxes regardless of game length, but because learning is slow in these parameter settings (large groups or high MPCR), the learning is only evident in long games. The linear regressions do not account for random effects/repeated measures and are therefore for illustration purposes only. The figures are annotated with the number of independent replicates (individuals grouped with computers).

Again, we compared the mean final contributions (“inputs”). These were significantly smaller, and thus closer to the income-maximising input of 0 coins, after the long game than after the short game in both black boxes (Table 2). Specifically, for the black box, where $N = 3$ and $MPCR = 0.8$, mean \pm SE final inputs were 9.4 ± 1.12 coins in the short game and 6.2 ± 0.97 coins in the long game (Wilcoxon rank-sum test, $W = 1287.5$, $p = 0.025$). For the black box where $N = 12$ and $MPCR = 0.4$, final inputs were 9.7 ± 1.12 coins in the short game and 5.4 ± 0.98 coins in the long game ($W = 1238.5$, $p = 0.006$).

Moreover, when asked about their beliefs about a possible best number, stated beliefs were significantly lower on average at the end of the long games compared to the short games (Figure 5, Table 3). Specifically, the mean \pm SE stated beliefs (0–20 coins) for when $N = 3$ and $MPCR = 0.8$, were 8.9 ± 1.59 coins in the short game and 3.5 ± 1.10 coins in the long game (Wilcoxon rank-sum test, $W = 410.5$, $p = 0.010$); for when $N = 12$ and $MPCR = 0.4$, were 8.0 ± 1.37 coins in the short game and 4.0 ± 0.93 coins in the long game ($W = 496.5$, $p = 0.016$).

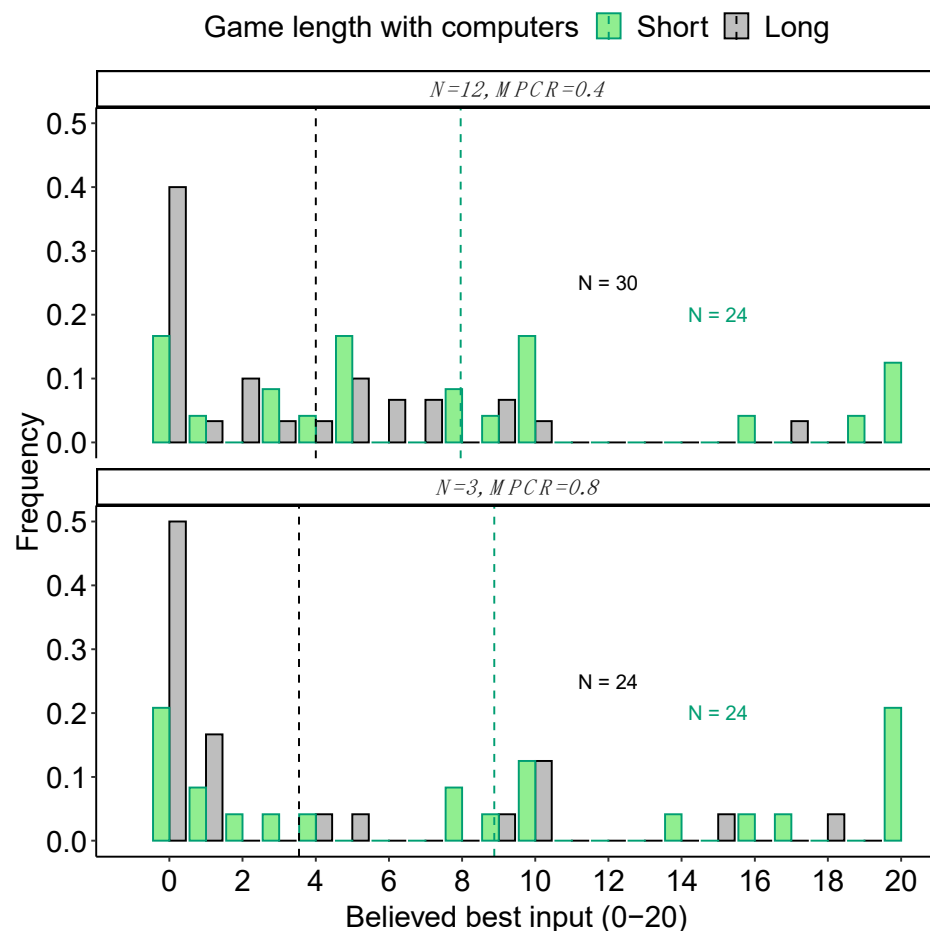


Figure 5. Game length and beliefs. Histograms show the frequency of each stated belief about what was the best number to input into the black box. Dashed vertical lines show the mean response. Beliefs were more accurate after a longer game. All responses are from naïve participants after finishing their first black box. The figures are annotated with the number of individuals.

3. Discussion

3.1. Payoff-Based Learning in Public-Goods Games

Our results show clear evidence that participants are capable of payoff-based learning in public goods games (Figures 2 and 4). Such learning does not require being “collective”, because we found strong evidence of learning even when we grouped individuals with computerised players that just played randomly and did not learn (Figure 2). Overall, we found clear evidence for payoff-based learning towards the income-maximising decision in both our short and long form games, suggesting that even when the learning environment is difficult, e.g., when group sizes are large ($N = 12$), or when the costs of “wrong” decisions are relatively small, e.g., when $MPCR = 0.8$, individuals can still learn to improve their income (Figure 4).

It has long been appreciated that participants will likely need time to learn in economic experiments [18,45–49]. While the role of learning has been clear in non-social dilemma studies as an explanation for initial non-income-maximising behaviour, it has often been disputed in social games [50]. Instead, many researchers have assumed that their participants understood the game from the beginning and were instead responding to other participants rather than their payoffs [16,51]. However, it is interesting to note that the standard procedure in repeated games is to show participants their payoffs after each round, which begs the question of why this was assumed necessary.

3.2. The Value of Control Treatments in Economic Experiments

One reason for this difference in approach between non-social and social experiments could be because deviations from income-maximising behaviour can always be rationalised in social games as behaviour motivated by the social consequences [9,52]. However, an approach of simply “measuring” social behaviours and preferences is problematic because it does not control for other behavioural processes, such as payoff-based learning.

Instead, when attempting to measure social behaviours, especially in artificial settings, one needs adequate behavioural controls to test whether the social factors are motivating behaviour. For instance, one can enhance how information is presented or framed to see if this affects social behaviours. This can be a particularly powerful approach when the information is technically redundant and thus provides no new information to participants yet still affects social behaviours [19,23].

Structurally, one can change the game’s payoffs so that failures to maximise income harm rather than benefit social partners. Such a reversal of social consequences should largely eliminate failures to maximise income among truly prosocial participants, a hypothesis falsified by studies that converted public-goods games into “public-delight” games [19,26,53,54]. In public-delight games, the return from contributing is set to be greater than 1, meaning that both selfish and prosocial individuals should contribute fully [19,26,53,54]. However, results show that many individuals still initially fail to maximise income before learning to moderately increase (rather than decrease) their contributions as the game is repeated, a result not easily explained by any rational social preferences [19,26,53,54]. Two reasons the increase may only be moderate when the MPCR > 1 are that in the public-delight game (1) the personal benefit of contributing is easily dwarfed by the benefits obtained from groupmates’ contributions, meaning that individuals have relatively less influence over their own payoff, impeding payoff-based learning [32]; and (2) contributions always produce profits, meaning individuals never suffer losses, which perhaps are more salient mistakes than failures to maximise profits.

Alternatively, one can remove social factors to create “asocial controls”, either by presenting the game differently, as is done by the black box, or structurally, by using revealed games played with computerised partners, which eliminate social concerns. There is a long history of using games with computerised partners, and such experiments have been useful to show that apparently social behaviours are not unique to games with human partners [25,27,29,30,33]. However, the interpretation of games with computers can be disputed, because although they clearly show that individuals often fail to maximise income even when there are no social consequences, they cannot rule out that individuals are psychologically motivated to help even computers.

The black box method provides a complementary form of asocial control to games with computers because, instead of removing social interactions, it hides them in a low information environment [19,31]. Consequently, the game is so devoid of social factors and framing that it is implausible to argue participants are behaving according to social psychology. Instead, the black box provides a clean measure of how a population of self-interested participants will collectively learn with experience, which can serve as a useful baseline measure of what behaviour to expect: if the addition of social information does not amplify social responses, then it may be more parsimonious to assume payoff-based learning is responsible. Specific learning hypotheses can then be tested with individual level data [20,21]. If one is concerned about participants in the original version of the black box affecting the learning and earnings of other participants, one can use the computerised black box to obtain similar results. One can even test various hypotheses by programming how the computerised partners will behave. One could then test if and how such learning spills over into other games or games with real humans, perhaps black box learning could be used to improve cooperation in games where cooperation is favoured [55].

4. Materials and Methods

4.1. Participants and Location

We analysed previously published data from Burton-Chellew and West, 2021 [32], and compared these data with those from a new experiment with computerised black boxes. The new experiment was conducted in February 2020 using 222 participants across 14 sessions over 3 days at the Centre for Experimental Social Sciences, Oxford (CESS). CESS recruited the participants from their entire database with the sole restriction that they could not have participated in a prior study by Burton-Chellew and West, 2021 [32]. CESS deployed assistants to manage participant reception, consent, and payments. MNBC conducted the experiment in a laboratory with 25 computer stations. For each session, we made 25 spaces available, and attendance varied from 7 to 25. The experiment was coded and conducted in z-TREE and participants were recruited using the ORSEE software [56,57].

The total sample was 222 participants. Their age ranged from 18 to 81 years old. The mean \pm SD was 31.8 years \pm 15.9 ($N = 192$, this does not include seven participants that declined to answer nor all 23 participants from the first session). According to the self-reported genders, we had 139 females, 81 males, one other and one declined to answer.

4.2. Experiment Design

A copy of the instructions is available in Appendix A and online at the Open Science Framework, along with the consent form, data and analysis script. Our design replicated the three “short” treatments from Burton-Chellew, 2021, which varied both the group size (N) and the return from contributing (Marginal Per Capita Return, MPCR). We also added two “long” treatments, played for 40 rounds instead of 16 rounds. However, instead of connecting real individuals with each other, we connected each participant with $N-1$ computerised virtual players that simulated random decisions drawn from a uniform decision. This way, the payoff was entirely consistent with a public goods game, but players had no knowledge of the game, and there were no ethical concerns about possible deception or a lack of informed consent.

Participants either played all three short treatments or one long treatment. When playing all three short treatments, we counterbalanced the order across sessions (although not perfectly, as the six permutations required a multiple of six sessions and we had eight). However, here we only analyse the “naïve” data from participants playing their first game. The reason we made participants play all three short black boxes was to standardise payments across sessions (a higher MPCR leads to increased mean payoffs).

The five treatments and naïve sample sizes were: black box with $N = 3$ and MPCR = 0.4 first = 46 (over three sessions, 23, 11, 12); black box with $N = 3$ and MPCR = 0.8 first = 44 (over three sessions, 22, 12, 10); black box with $N = 12$ and MPCR = 0.4 first = 40 (over two sessions, 18 and 22); black box with $N = 3$ and MPCR = 0.8 long version = 46 (over three sessions, 22, 7, 17); black box with $N = 12$ and MPCR = 0.4 long version = 46 (over three sessions, 25, 9, 12).

The exchange rate was 0.65 pence to 1 virtual coin, except in the first session, where it was 0.75. The total endowment was 960 coins (GBP 6.24) in the short treatment sessions and 800 coins (GBP 5.20) in the long treatments. Payments were rounded up to the nearest 10 pence. Earnings ranged from GBP 8.20 to GBP 16.90 (mean = GBP 12.38), to which a GBP 5 show up fee was added.

After each black box, we asked participants two questions: (1) “Do you think there was a best number to put in this black box? Please enter 1 for ‘Yes’ and 0 for ‘No’.”; and then (2) “If yes, what do you think was the best number (0–20)? If no, then please enter 99.” Then after the experiment, we asked the participants to complete a brief survey on their self-reported gender, age, and personality traits, before asking them, “In a few words, please tell us what, if anything, you think the experiment was about?”. We did this to check that they did not perceive the experiment as a social dilemma. As responses to this same question have already been analysed by Burton-Chellew and West, 2021, who found

that only 2% ($N = 5/216$) of participants ($N = 5/216$) “mentioned anything that could be construed as social”, we do not analyse the responses here [32].

4.3. Analyses

All tests are two-tailed and we conducted all analyses in RStudio [58]. All our data and analysis files are freely available online at the Open Science Framework, <https://osf.io/9uxdv/> [59].

Author Contributions: Conceptualization, M.N.B.-C. and S.A.W.; Methodology, M.N.B.-C. and S.A.W.; Software, M.N.B.-C.; Validation, M.N.B.-C. and S.A.W.; Formal analysis, M.N.B.-C.; Investigation, M.N.B.-C. and S.A.W.; Resources, M.N.B.-C. and S.A.W.; Data curation, M.N.B.-C.; Writing—original draft, M.N.B.-C. and S.A.W.; Writing—review & editing, M.N.B.-C. and S.A.W.; Visualization, M.N.B.-C. and S.A.W.; Supervision, S.A.W.; Project administration, M.N.B.-C.; Funding acquisition, S.A.W. All authors have read and agreed to the published version of the manuscript.

Funding: Calleva Research Centre for Evolution and Human Sciences, Magdalen College.

Data Availability Statement: Data and analysis files are freely available online at the Open Science Framework, <https://osf.io/9uxdv/> [59].

Acknowledgments: The CESS laboratory staff for support; Heinrich Nax and referees for comments; the ERC for funding.

Conflicts of Interest: The authors declare no conflict of interest.

Appendix A

A copy of the instructions for the sessions with three short black boxes. For the long games, the details were changed accordingly. Instructions were available on-screen and on a paper handout.

Welcome to the Experiment!

A copy of these instructions is also available on-screen.

We are going to give you some virtual coins. Each ‘coin’ is worth real money.

You are going to make a decision regarding the investment of these ‘coins’.

This decision may increase or decrease the number of coins you have.

The more coins you have at the end of the experiment, the more money you will receive at the end.

At the end of the experiment the total amount of ‘coins’ you have earned will be converted to pennies at the following rate: 100 coins = 65 pennies, or = GBP 0.65.

In total, you will be given 960 coins GBP 6.24 with which to make decisions and your final total, which may be more or less than 960 coins, will depend on these decisions.

The Decision

You will face the same decision many times. Each time we will give you 20 virtual ‘coins’. Then you must decide on how many of your 20 coins to input into a virtual ‘black box’.

This ‘black box’ performs a mathematical function that converts the number of ‘coins’ inputted into a number of ‘coins’ to be outputted.

The mathematical function contains two components, one constant, deterministic, component which acts upon your input, and one ‘chance’ component.

You will play with this ‘black box’ for many rounds (more on this later), and the mathematical function will not change, but the chance component means that if you put the same amount of coins into the ‘black box’ over successive rounds, you will not necessarily get the same output each time.

The number outputted may be more or less than you put in, but it will never be a negative number, so the lowest outcome possible is to get 0 (zero) back.

If you chose to input 0 (zero) coins, you may still get some back from the black box.

All coins not inputted into the black box will be automatically ‘banked’ into your private account.

All coins outputted from the black box will also be ‘banked’ and go into your private account.

You will be paid all the coins from your private account at the end of the experiment.

So, in summary, your income from each decision will be the initial 20 coins, minus any you put into the ‘black box’, plus all the coins you get back from the ‘black box’.

Playing the Same Box Many Times

You will play this game (make this decision) 16 times. Each time we will give you a new set of 20 coins to use.

Each Decision is Separate but the ‘Black Box’ Remains the Same

This means you cannot play with money gained from previous decisions, and the maximum you can ever put into the ‘black box’ will be 20 coins.

And you will never run out of money to play with as we will give you a new set of coins for each decision.

Please see the attached figure overleaf for a summary of the experiment.

Playing with Different Boxes

After you have finished your 16 decisions, you will play again with a new ‘black box’.

In total, you will play with 3 black boxes in the whole experiment.

All black boxes are the same in that they perform a mathematical function that converts the number of coins inputted into a number of coins to be outputted.

However Each Black Box Will Have a Different Mathematical Function

But the functions will always contain two components, one constant, deterministic, component, and one ‘chance’ component. You will play with this black box for many rounds, and the mathematical function will never change, but the chance component means that if you put the same amount of coins into the black box over successive rounds, you will not necessarily get the same output each time.

You Will Be Told When the Decisions are Finished and It Is Time to Play with a New Black Box

If you are unsure of the rules please hold up your hand and a demonstrator will help you.

Notes

- ¹ Participants were told that the black box contained a mathematical function which would remain constant for the experiment, but which contained a random component each round, meaning that a given input would not guarantee the same output, but giving the impression that the black box was in some sense solvable.
- ² Note that this is not deception as no false information is giving to participants. It is merely an omission of information about the externalities of the participant’s decisions. Crucially, participants are not going to leave the laboratory thinking that next time they play a game with humans that the humans are actually computers or actors (which is arguably the main reason for the no deception policy).

References

1. Hardin, G. The Tragedy of the Commons. *Science* **1968**, *162*, 1243–1248. [[CrossRef](#)]
2. Rustagi, D.; Engel, S.; Kosfeld, M. Conditional Cooperation and Costly Monitoring Explain Success in Forest Commons Management. *Science* **2010**, *330*, 961–965. [[CrossRef](#)] [[PubMed](#)]
3. Milinski, M.; Sommerfeld, R.D.; Krambeck, H.J.; Reed, F.A.; Marotzke, J. The collective-risk social dilemma and the prevention of simulated dangerous climate change. *Proc. Natl. Acad. Sci. USA* **2008**, *105*, 2291–2294. [[CrossRef](#)]
4. Burton-Chellew, M.N.; May, R.M.; West, S.A. Combined inequality in wealth and risk leads to disaster in the climate change game. *Clim. Chang.* **2013**, *120*, 815–830. [[CrossRef](#)]
5. Romano, A.; Balliet, D.; Yamagishi, T.; Liu, J.H. Parochial trust and cooperation across 17 societies. *Proc. Natl. Acad. Sci. USA* **2017**, *114*, 12702–12707. [[CrossRef](#)] [[PubMed](#)]

6. Bavel, J.J.V.; Baicker, K.; Boggio, P.S.; Capraro, V.; Cichocka, A.; Cikara, M.; Crockett, M.J.; Crum, A.J.; Douglas, K.M.; Druckman, J.N.; et al. Using social and behavioural science to support COVID-19 pandemic response. *Nat. Hum. Behav.* **2020**, *4*, 460–471. [\[CrossRef\]](#)
7. Ijzerman, H.; Lewis, N.A.; Przybylski, A.K.; Weinstein, N.; DeBruine, L.; Ritchie, S.J.; Vazire, S.; Forscher, P.; Morey, R.D.; Ivory, J.D.; et al. Use caution when applying behavioural science to policy. *Nat. Hum. Behav.* **2020**, *4*, 1092–1094. [\[CrossRef\]](#)
8. Thielmann, I.; Böhm, R.; Ott, M.; Hilbig, B.E. Economic Games: An Introduction and Guide for Research. *Collabra Psychol.* **2021**, *7*, 19004. [\[CrossRef\]](#)
9. Andreoni, J.; Miller, J. Giving according to garp: An experimental test of the consistency of preferences for altruism. *Econometrica* **2002**, *70*, 737–753. [\[CrossRef\]](#)
10. Fehr, E.; Gächter, S. Altruistic punishment in humans. *Nature* **2002**, *415*, 137–140. [\[CrossRef\]](#)
11. Carpenter, J.P. When in Rome: Conformity and the provision of public goods. *J. Socio-Econ.* **2004**, *33*, 395–408. [\[CrossRef\]](#)
12. Croson, R.; Fatas, E.; Neugebauer, T. Reciprocity, matching and conditional cooperation in two public goods games. *Econ. Lett.* **2005**, *87*, 95–101. [\[CrossRef\]](#)
13. Brandts, J.; Cooper, D.J.; Fatas, E. Leadership and overcoming coordination failure with asymmetric costs. *Exp. Econ.* **2007**, *10*, 269–284. [\[CrossRef\]](#)
14. Andreoni, J.; Croson, R. Partners versus strangers: Random rematching in public goods experiments. In *Handbook of Experimental Economics Results*; Plott, C.R., Smith, V.L., Eds.; North-Holland: Amsterdam, The Netherlands, 2008; p. 776.
15. Guillen, P.; Fatas, E.; Branas-Garza, P. Inducing efficient conditional cooperation patterns in public goods games, an experimental investigation. *J. Econ. Psychol.* **2010**, *31*, 872–883. [\[CrossRef\]](#)
16. Fischbacher, U.; Gächter, S. Social Preferences, Beliefs, and the Dynamics of Free Riding in Public Goods Experiments. *Am. Econ. Rev.* **2010**, *100*, 541–556. [\[CrossRef\]](#)
17. Böhm, R.; Rockenbach, B. The Inter-Group Comparison—Intra-Group Cooperation Hypothesis: Comparisons between Groups Increase Efficiency in Public Goods Provision. *PLoS ONE* **2013**, *8*, e56152. [\[CrossRef\]](#)
18. Andreozzi, L.; Ploner, M.; Saral, A.S. The stability of conditional cooperation: Beliefs alone cannot explain the decline of cooperation in social dilemmas. *Sci. Rep.-UK* **2020**, *10*, 13610. [\[CrossRef\]](#)
19. Burton-Chellew, M.N.; West, S.A. Prosocial preferences do not explain human cooperation in public-goods games. *Proc. Natl. Acad. Sci. USA* **2013**, *110*, 216–221. [\[CrossRef\]](#)
20. Burton-Chellew, M.N.; Nax, H.H.; West, S.A. Payoff-based learning explains the decline in cooperation in public goods games. *Proc. R. Soc. B-Biol. Sci.* **2015**, *282*, 20142678. [\[CrossRef\]](#)
21. Nax, H.H.; Burton-Chellew, M.N.; West, S.A.; Young, H.P. Learning in a black box. *J. Econ. Behav. Organ.* **2016**, *127*, 1–15. [\[CrossRef\]](#)
22. Nax, H.H.; Perc, M. Directional learning and the provisioning of public goods. *Sci. Rep.-UK* **2015**, *5*, srep08010. [\[CrossRef\]](#) [\[PubMed\]](#)
23. Burton-Chellew, M.N.; Guérin, C. Self-interested learning is more important than fair-minded conditional cooperation in public-goods games. *Evol. Hum. Sci.* **2022**, *4*, e46. [\[CrossRef\]](#)
24. Andreoni, J. Cooperation in public-goods experiments—Kindness or confusion. *Am. Econ. Rev.* **1995**, *85*, 891–904.
25. Ferraro, P.J.; Vossler, C.A. The Source and Significance of Confusion in Public Goods Experiments. *BE J. Econ. Anal. Policy* **2010**, *10*. [\[CrossRef\]](#)
26. Kummerli, R.; Burton-Chellew, M.N.; Ross-Gillespie, A.; West, S.A. Resistance to extreme strategies, rather than prosocial preferences, can explain human cooperation in public goods games. *Proc. Natl. Acad. Sci. USA* **2010**, *107*, 10125–10130. [\[CrossRef\]](#) [\[PubMed\]](#)
27. Burton-Chellew, M.N.; El Mouden, C.; West, S.A. Conditional cooperation and confusion in public-goods experiments. *Proc. Natl. Acad. Sci. USA* **2016**, *113*, 1291–1296. [\[CrossRef\]](#)
28. Burton-Chellew, M.N.; D’Amico, V.; Guérin, C. The Strategy Method Risks Conflating Confusion with a Social Preference for Conditional Cooperation in Public Goods Games. *Games* **2022**, *13*, 69. [\[CrossRef\]](#)
29. Houser, D.; Kurzban, R. Revisiting kindness and confusion in public goods experiments. *Am. Econ. Rev.* **2002**, *92*, 1062–1069. [\[CrossRef\]](#)
30. Shapiro, D.A. The role of utility interdependence in public good experiments. *Int. J. Game Theory* **2009**, *38*, 81–106. [\[CrossRef\]](#)
31. Bayer, R.C.; Renner, E.; Sausgruber, R. Confusion and learning in the voluntary contributions game. *Exp. Econ.* **2013**, *16*, 478–496. [\[CrossRef\]](#)
32. Burton-Chellew, M.N.; West, S.A. Payoff-based learning best explains the rate of decline in cooperation across 237 public-goods games. *Nat. Hum. Behav.* **2021**, *5*, 1330–1388. [\[CrossRef\]](#) [\[PubMed\]](#)
33. Nielsen, Y.A.; Thielmann, I.; Zettler, I.; Pfattheicher, S. Sharing Money With Humans Versus Computers: On the Role of Honesty-Humility and (Non-)Social Preferences. *Soc. Psychol. Pers. Sci.* **2021**, *13*, 1058–1068. [\[CrossRef\]](#)
34. Foster, D.P.; Peyton Young, H. Regret testing: Learning to play Nash equilibrium without knowing you have an opponent. *Theor. Econ.* **2006**, *1*, 341–367.
35. Nevo, I.; Erev, I. On Surprise, Change, and the Effect of Recent Outcomes. *Front. Psychol.* **2012**, *3*, 24. [\[CrossRef\]](#) [\[PubMed\]](#)
36. Erev, I.; Haruvy, E. Learning and the economics of small decisions. In *The Handbook of Experimental Economics*; Kagel, J.H., Roth, A.E., Eds.; Princeton University Press: Princeton, NJ, USA, 2013; Volume 2, pp. 501–512.

37. Zion, U.B.; Erev, I.; Haruvy, E.; Shavit, T. Adaptive behavior leads to under-diversification. *J. Econ. Psychol.* **2010**, *31*, 985–995. [[CrossRef](#)]
38. Weber, R.A. ‘Learning’ with no feedback in a competitive guessing game. *Games Econ. Behav.* **2003**, *44*, 134–144. [[CrossRef](#)]
39. Rapoport, A.; Seale, D.A.; Parco, J.E. Coordination in the Aggregate without Common Knowledge or Outcome Information. In *Experimental Business Research*; Zwick, R., Rapoport, A., Eds.; Springer: Boston, MA, USA, 2002; pp. 69–99. [[CrossRef](#)]
40. Colman, A.M.; Pulford, B.D.; Omtzigt, D.; Al-Nowaihi, A. Learning to cooperate without awareness in multiplayer minimal social situations. *Cogn. Psychol.* **2010**, *61*, 201–227. [[CrossRef](#)]
41. Friedman, D.; Huck, S.; Oprea, R.; Weidenholzer, S. From imitation to collusion: Long-run learning in a low-information environment. *J. Econ. Theory* **2015**, *155*, 185–205. [[CrossRef](#)]
42. Bereby-Meyer, Y.; Roth, A.E. The speed of learning in noisy games: Partial reinforcement and the sustainability of cooperation. *Am. Econ. Rev.* **2006**, *96*, 1029–1042. [[CrossRef](#)]
43. Horita, Y.; Takezawa, M.; Inukai, K.; Kita, T.; Masuda, N. Reinforcement learning accounts for moody conditional cooperation behavior: Experimental results. *Sci. Rep.-UK* **2017**, *7*, 39275. [[CrossRef](#)]
44. Peyton Young, H. Learning by trial and error. *Games Econ. Behav.* **2009**, *65*, 626–643. [[CrossRef](#)]
45. Binmore, K. Why Experiment in Economics? *Econ. J.* **1999**, *109*, F16–F24. [[CrossRef](#)]
46. Binmore, K. Economic man—Or straw man? *Behav. Brain Sci.* **2005**, *28*, 817–818. [[CrossRef](#)]
47. Binmore, K. Why do people cooperate? *Politics Philos. Econ.* **2006**, *5*, 81–96. [[CrossRef](#)]
48. Smith, V.L. Theory and experiment: What are the questions? *J. Econ. Behav. Organ.* **2010**, *73*, 3–15. [[CrossRef](#)]
49. Friedman, D. Preferences, beliefs and equilibrium: What have experiments taught us? *J. Econ. Behav. Organ.* **2010**, *73*, 29–33. [[CrossRef](#)]
50. Camerer, C.F. Experimental, cultural, and neural evidence of deliberate prosociality. *Trends Cogn. Sci.* **2013**, *17*, 106–108. [[CrossRef](#)]
51. Fehr, E.; Schmidt, K.M. A theory of fairness, competition, and cooperation. *Q. J. Econ.* **1999**, *114*, 817–868. [[CrossRef](#)]
52. Sobel, J. Interdependent preferences and reciprocity. *J. Econ. Lit.* **2005**, *43*, 392–436. [[CrossRef](#)]
53. Saijo, T.; Nakamura, H. The Spite Dilemma in Voluntary Contribution Mechanism Experiments. *J. Confl. Resolut.* **1995**, *39*, 535–560. [[CrossRef](#)]
54. Brunton, D.; Hasan, R.; Mestelman, S. The ‘spite’ dilemma: Spite or no spite, is there a dilemma? *Econ. Lett.* **2001**, *71*, 405–412. [[CrossRef](#)]
55. Cherry, T.L.; Crocker, T.D.; Shogren, J.F. Rationality spillovers. *J. Environ. Econ. Manag.* **2003**, *45*, 63–84. [[CrossRef](#)]
56. Fischbacher, U. z-Tree: Zurich toolbox for ready-made economic experiments. *Exp. Econ.* **2007**, *10*, 171–178. [[CrossRef](#)]
57. Greiner, B. Subject pool recruitment procedures: Organizing experiments with ORSEE. *J. Econ. Sci. Assoc.* **2015**, *1*, 114–125. [[CrossRef](#)]
58. Team, R. *Integrated Development Environment for R*; RStudio: Boston, MA, USA, 2020.
59. Burton-Chellew, M.N.; West, S.A. Data for: The black box as a control for payoff-based learning in economic games. *Open Sci. Framew.* **2022**.