



**PARIS SCHOOL OF ECONOMICS**  
ÉCOLE D'ÉCONOMIE DE PARIS

**WORKING PAPER N° 2013 – 10**

**Learning in a Black Box**

**Heinrich H. Nax  
Maxwell N. Burton-Chellew  
Stuart A. West  
H. Peyton Young**

**JEL Codes : C70, C73, C91, D83, H41**

**Keywords: Learning; Information; Public goods games**



**PARIS-JOURDAN SCIENCES ÉCONOMIQUES**

48, Bd JOURDAN – E.N.S. – 75014 PARIS  
TÉL. : 33(0) 1 43 13 63 00 – FAX : 33 (0) 1 43 13 63 10  
[www.pse.ens.fr](http://www.pse.ens.fr)

# Learning in a Black Box\*

H.H. Nax<sup>†</sup>, M.N. Burton-Chellew<sup>‡</sup>, S.A. West<sup>§</sup>, H.P. Young<sup>¶</sup>

First version: April 18, 2013

This version: September 30, 2013

---

\*We thank Margherita Comola for help at an early stage of the project and Johannes Abeler, Richard Bradley, Fabrice Etilé, Philippe Jehiel, Thierry Verdier, Vince Crawford, Colin Camerer, Guillaume Hollard, Muriel Niederle, Edoardo Gallo, Amnon Rapoport, Guillaume Fréchette and participants of the Economics and Psychology seminar at MSE Paris, the Choice Group at LSE and the 24th Game Theory Workshop at Stony Brook for helpful discussions, suggestions and comments. Correspondence to: H. H. Nax, 48 Boulevard Jourdan, Bureau A201, Paris School of Economics, 75014 Paris, France.

<sup>†</sup>Heinrich H. Nax, Paris School of Economics, École Normale Supérieure. Nax's research was in part supported by the Office of Naval Research grant N00014-09-1-0751 and under project NET of the Agence Nationale de Recherche.

<sup>‡</sup>Maxwell N. Burton-Chellew, Nuffield College, Oxford & Department of Zoology, University of Oxford (The Tinbergen Building S Parks Rd, Oxford OX1 3PS, United Kingdom).

<sup>§</sup>Stuart A. West, Department of Zoology, University of Oxford (The Tinbergen Building S Parks Rd, Oxford OX1 3PS, United Kingdom). Burton-Chellew and West's research was supported by the European Research Council.

<sup>¶</sup>H. Peyton Young, Department of Economics, University of Oxford (Nuffield College, New Rd, Oxford OX1 1NF, United Kingdom). Young's research was supported in part by the Office of Naval Research grant N00014-09-1-0751.

## Abstract

Many interactive environments can be represented as games, but they are so large and complex that individual players are in the dark about others' actions and the payoff structure. This paper analyzes learning behavior in such 'black box' environments, where players' only source of information is their own history of actions taken and payoffs received. Specifically we study voluntary contributions games. We identify two robust features of the players' learning dynamics: search volatility and trend-following. These features are clearly present when players have no information about the game; but also when players have full information. Convergence to Nash equilibrium occurs at about the same rate in both situations.

*JEL classifications:* C70, C73, C91, D83, H41

*Keywords:* learning, information, public goods games

## 1 Introduction

Many interactive environments can be represented as games, but they are so complex and involve so many individuals that for all practical purposes the game itself is unknown to the players themselves. Examples include bidding in on-line markets, threading one's way through urban traffic, or participating in a group effort where the actions of the other members of the group are difficult to observe (guerrilla warfare, neighborhood watch programs, tax evasion). In each of these cases a given individual will have only the haziest idea of what the other players are doing but their payoffs are strongly influenced by the actions of the others. How do individuals learn in such environments and under what circumstances does their learning behavior lead to equilibrium?

In this paper we investigate these questions in a laboratory environment where the underlying game has the structure of a public goods game. Players take actions and receive payoffs that depend on others' actions, but (in the baseline case) they have no information about the others and they do not know what the overall structure of the game is. This distinguishes our experimental setup from other "low-information" setups such as Rapoport et al. (2002) and Weber (2003) who provide information about the structure of the game but withhold information about the distribution of players' types and the outcome of play, or Friedman et al. (2012) who withhold information about the payoff structure of the game but provide information about other players' actions and payoffs. The closest "low-information" set-up to ours is that of Bayer et al. (2013) who also implement experiments without information about the structure of the game and about

other players. However, their set-up purposefully induces a “confusion” element that the game structure itself may also change. In our setup, the structure of the game, although it is unknown to the agents, remains fixed and agents know that. Perhaps the closest antecedent of our ‘black box’ treatments are repeated two-by-two zero sum games used to test the early Markov learning models by Suppes & Atkinson (1959).<sup>1</sup>

As the game proceeds, all agents know is the amount that they themselves contributed and the payoff they received as a result. Learning in such an explicitly repeated game environment is said to be completely uncoupled (Foster & Young, 2006; Young, 2009; see also Hart & Mas-Colell, 2003, 2006;). In this setting many of the standard learning rules in the empirical literature, such as experience-weighted attraction, k-level reasoning, and imitation, do not apply because these rules use the actions of the other players as inputs (Björnerstedt & Weibull, 1993; Stahl & Wilson, 1995; Nagel, 1995; Ho et al., 1998; Camerer & Ho, 1999; Costa-Gomes et al., 2001; Costa-Gomes & Crawford, 2006).<sup>2</sup>

Nevertheless our experiments show that learning ‘inside the box’ can and does take place.<sup>3</sup> The learning process exhibits certain distinctive patterns that have been identified in other settings. Two especially noteworthy features are the following:

*Search volatility.* A decrease in a player’s realized payoff triggers greater variance in his choice of action next period, whereas an increase in his realized payoff results in lower variance next period.

*Trend-following.* If a higher contribution leads to a higher payoff this period, the player will tend to raise his contribution next period, whereas if a higher contribution leads to a lower payoff this period, he will tend to reduce his contribution next period. Similarly, if a lower contribution leads to a higher payoff the player will tend to lower his contribution, whereas if a lower contribution leads to a lower payoff he will tend to increase his contribution next period.

Learning rules with high and low search volatility have been proposed in a variety of settings, including biology (Thuijsman et al., 1995), organizational learning (March, 1991),

---

<sup>1</sup>Note that subjects in that setting have only two strategies and were explicitly ‘fooled’ not to think that they were playing a game. The observed learning in this setting follows reinforcement.

<sup>2</sup>Indeed most impulse-based learning models (see Crawford, 2013) except for reinforcement learning (Erev & Roth, 1998) are thus excluded. We also note that social preferences cannot enter into subjects’ behavior because they cannot compare their own payoffs with others’ payoffs. Thus the black box environment eliminates some of the explanations that have been advanced for behavior in public goods games (Fehr & Schmidt, 1999; Fehr & Gächter, 2000, 2002; Fehr & Camerer, 2007).

<sup>3</sup>Oechssler & Schipper (2003) show that, in the context of 2x2 games, players can learn to identify the other player’s payoff structure on the basis of their own payoffs and information about other’s actions. In our set-up players may or may not learn the game, but even without information about the relationship between their own and others’ payoffs they do learn to play Nash equilibrium. See Erev & Haruvy (2013) for a recent survey of equilibrium learning.

computer science (Eiben & Schippers, 1998), and psychology (Coates, 2012). (In computer science and organization theory this type of learning is known as ‘exploration versus exploitation’, ‘win-stay-lose-shift’ or ‘fast and slow’ learning.) Search volatility is also a crucial feature of rules that have recently been proposed in the theoretical literature; in particular there is a family of such rules that lead to Nash equilibrium in generic  $n$ -person games (Foster & Young, 2006; Germano & Lugosi, 2007; Marden et al., 2009; Young, 2009; Pradelski & Young, 2012). To our knowledge, despite systematic differences in behavior following positive and negative feedback being noted (e.g. Ding & Nicklisch, 2013), our model is the first experimental text of an explicit learning model with such this feature.

Trend-following also has antecedents in the experimental literature, where it is sometimes called directional learning (Selten & Stoecker, 1986; Selten & Buchta, 1998).<sup>4</sup> Directional learning describes behavioral dynamics based on local adjustments of strategy and aspiration (Sauermann & Selten, 1962; Cross, 1983). The underlying behavioral components are illustrated in a bargaining game: if a higher demand results in a higher realized payoff the agent increases his demand, whereas a decrease in realized payoff leads to a demand decrease (Tietz & Weber, 1972; Roth & Erev, 1995). Other experiments such as Erev & Roth (1998) and Erev & Rapoport (1998) corroborate these bargaining findings and Nash equilibrium is learned based on reinforcement dynamics in various settings. Recent theoretical generalization’s of directional learning when strategies are unidimensional have been proposed by Laslier & Walliser (2011). Bayer et al. (2013) recover a directional learning model in their learning condition even if the structure of the game potentially changes.

In this paper we show that these two patterns help explain the behavior of subjects who are learning in a ‘black box’ environment. Moreover, despite the lack of information and the simplicity of the adjustment rules, the behavior of the group approaches Nash equilibrium and does so at approximately the same rate as in the situation where subjects have full information (Burton-Chellew & West, 2012). As we shall see, this is true both when free riding is the dominant strategy, and also when full contribution is the dominant strategy. We conclude that even when subjects have virtually no information about the game, learning follows certain distinctive patterns that have been documented in a wide variety of learning environments. In public goods games this behavior leads to Nash equilibrium at a rate that is very similar to learning under full information. Moreover, search volatility turns out to be a robust behavioral component even when more information becomes available and when players gain experience.

---

<sup>4</sup>See Harstad & Selten (2013) for a recent survey of these models.

The paper is structured as follows. In section 2 we describe the experimental set-up in detail. In section 3 we present the empirical findings together with statistical tests of significance. We conclude in section 4. An appendix contains experimental instructions and supplementary regression tables.

## 2 Experimental set-up

Participants were recruited from a subject pool that had not previously been involved in public goods experiments.<sup>5</sup> The subjects were not limited to university students, but included different age groups with diverse educational backgrounds.<sup>6</sup> Each subject played four separate contribution games, where each game was repeated twenty times with randomly allocated subgroups. The twenty-fold repetition of a given game will be called a ‘phase’ of the experiment. This yields a total of 236 games and 18,880 observations. During sixteen separate sessions in groups of sixteen or twelve, subjects play the four phases in a row.

Games differ with respect to two *rates of return* (‘low’ and ‘high’) such that either ‘free-riding’ or ‘fully contributing’ is the dominant strategy. Treatments differ with respect to the information that is revealed to the agents (‘black box’, ‘standard’ and ‘enhanced’). Every player plays both black box games, and either the two standard or enhanced information games. Either two black box games are played first, or two games of one of the other two treatments.<sup>7</sup> The two first phases constitute stage one, the second two constitute stage two. Each possible information order is played in four separate sessions, each of which with a different order of the two rates of return.<sup>8</sup>

In this section, we shall describe the underlying stage game, the structure of each repeated game, and the different information treatments.

---

<sup>5</sup>The subjects were recruited through ORSEE (Online Recruiting System for Economic Experiments; Greiner, 2004) with this request specifically made.

<sup>6</sup>The experiment was programmed and conducted with the software z-Tree (Fischbacher, 2007). All experiments were conducted at the Centre for Experimental Social Sciences (CESS) at Nuffield College, University of Oxford.

<sup>7</sup>Of the 18,880 observations, 9,440 are black box, 4,640 are standard, and 4,800 are enhanced.

<sup>8</sup>This yields sixteen different scenarios each played in one of the sessions given by the possible orders and permutations; four (combinations) times four (orders) equals sixteen (sessions).

## 2.1 Linear public goods game

Consider the following linear public goods game. Every player  $i$  in population  $N = \{1, 2, \dots, n\}$  makes a nonnegative, real-numbered contribution,  $a_i$ , from a finite budget  $B \in \mathbb{N}^+$ . Given a resulting vector of all players' contributions,  $\mathbf{a} = \{a_1, a_2, \dots, a_n\}$ , for some *rate of return*,  $e \geq 1$ , the public good is provided,  $R(\mathbf{a}) = e \sum_{i \in N} a_i$ , and split equally amongst the players.<sup>9</sup> Given others' contributions  $\mathbf{a}_{-i}$ , player  $i$ 's contribution  $a_i$  therefore results in a payoff of

$$\phi_i = \frac{e}{n} \sum_{i \in N} a_i + (B - a_i),$$

where  $\frac{e}{n}$  is the marginal rate of return. Write  $\phi$  for the payoff vector  $\{\phi_1, \dots, \phi_n\}$ .

**Nash equilibria.** Depending on whether the rate of return is low ( $e < n$ ) or high ( $e > n$ ), individual contribution of either zero (free-riding) or  $B$  (fully contributing) is the strictly dominant strategy for all players; the respective Nash equilibrium results in either nonprovision or full provision of the public good.<sup>10</sup>

## 2.2 Repeated game

In any session, the same population  $S$  (with  $|S| = 12$  or  $16$ ) plays four phases. Each phase is a twenty-times repeated, symmetric public goods game. In each period  $t$  of each game, players in  $S$  are randomly matched to groups of four to play the stage game.<sup>11</sup> The budget every period is a new  $B = 40$  to every player of which he can invest any amount, but he cannot invest money carried over from previous rounds. The rate of return is either *low* ( $e = 1.6$ ) or *high* ( $e = 6.4$ ) throughout the game.<sup>12</sup> Write  $N_4^t$  for any of the four-player groups matched at time  $t$ , and  $\rho_4^t$  for the partition of  $S$  into such groups. Given  $\mathbf{a}^t$ ,  $\rho_4^t$ ,  $N_4^t$ , in any period  $t$  such that  $i \in N_4^t \in \rho_4^t$ , each  $i$  receives

$$\phi_i = \sum_{t=1, \dots, 20} \phi_i^t = \sum_{t=1, \dots, 20} [e/4 \sum_{j \in N_4^t} a_j^t + (40 - a_i^t)].$$

For every  $i$ ,  $\phi_i$  represents a real monetary value that is paid after the game.<sup>13</sup>

<sup>9</sup>If  $e < 1$ ,  $R(\mathbf{a})$  is a public bad.

<sup>10</sup>When  $e = n$ , any contribution is a best reply to any level of contribution by the others, hence any level of provision is supported by at least one pure-strategy Nash equilibrium.

<sup>11</sup>This follows the random 'stranger' rematching of Andreoni (1988).

<sup>12</sup>The Nash equilibrium payoffs of the stage games are 40 when  $e = 1.6$ , and 256 when  $e = 6.4$ .

<sup>13</sup>Hundred coins are worth £0.15 to pay £19.2 maximal individual earnings from the whole experiment.

## 2.3 Information treatments

Before the first game begins, each player is told that two separate experiments will be played. (From the analyst's point of view the experiment consists of both of these taken together.) At no point before, during, or in between the separate phases of the experiment are players allowed to communicate. Depending on which treatment is played, the following information is revealed at the start of each stage when a new treatment begins.

**All treatments.** Two separate twenty-times repeated stage games are played. In each, the underlying stage game is unchanged for 20 periods. Each player receives 40 monetary units each period of which he can invest any amount. Moreover, after investments are made, each player gets a nonnegative return each period which, at the end of each game, he receives together with his uninvested money according to a known exchange rate.

**Black box.** No information about the structure of the game, or about other players' actions or payoffs is revealed. Throughout the game players only know their own contribution and payoffs.

**Standard.** The rules of the game are revealed including production of the public good, rate of return, and how groups in the session form to play each stage game. In addition, at the end of each period, players receive a summary of the relevant contributions in their previous-period group; without knowing who these players are and what payoffs they receive.<sup>14</sup>

**Enhanced.** In addition to the information available in the standard treatment, players learn about the other group members' previous-period payoffs.

Our main treatment is black box. Indeed, all our results concerning the learning rules are based on black box data from sessions when black box is played first, in which case subjects really have no knowledge of the structure of the game. For this reason, we also required the recruiting system (ORSEE) to select only 'first timers' in public-goods experiments. (See Appendix A for full black box instructions and Appendix B for the output screens displayed during the experiment for each treatment.)

---

<sup>14</sup>This follows the standard information treatment design as in Fehr & Gächter (2002) and random 'stranger' rematching of subjects as introduced by Andreoni (1988).



## 3 Findings

First, we analyze the black box data when players have no prior experience playing public goods games in our experiments (or in other prior experiments). In this ‘black box’ treatment, subjects have no knowledge of the structure of the game and receive no information about others’ actions or payoffs.

### 3.1 Black box

In this section, we analyze the black box treatments when played before the standard or enhanced treatments. Subjects were chosen only if they had not previously participated in any public goods related experiment prior to ours.<sup>15</sup> There were eight such sessions and 4,960 observations. By choosing these sessions, we ensure that players not only have no explicit knowledge of the game structure and no information about others, but also have no “experience” of related laboratory experiments that they could use to make structural inferences. Unless stated otherwise, “significance” refers to a 95% confidence interval (5% significance) when rejecting the null and to a 90% confidence interval (10% significance) when the null cannot be rejected.

#### 3.1.1 Contributions

For both rates of return independent of whether played first or second (phase one or two), the mean initial contribution lies between one third and half the budget. The mean is 15.3 in the games with  $e = 1.6$ , and 13.9 in the games with  $e = 6.4$ . Using period-specific Mann-Whitney tests, we cannot reject the null of equal contributions for games with different rates of return and for phase one versus phase two, that is, similar observations are made in all games (with different rates of return) played in all initial black box treatment games. This is evidence of “restart effects” (observed throughout the literature).<sup>16</sup> The period-two adjustments to the initial contribution have means close to zero.<sup>17</sup> Notably, a large part of these contributions lie at 0, 10, 20, 30, 40, with several contributions made also in the range from 1 to 9; contributions at other numerical values are rare.<sup>18</sup> Using period-specific Mann-Whitney tests, we cannot reject the null of equal contributions for different rates of return for periods 1 through to period 3.

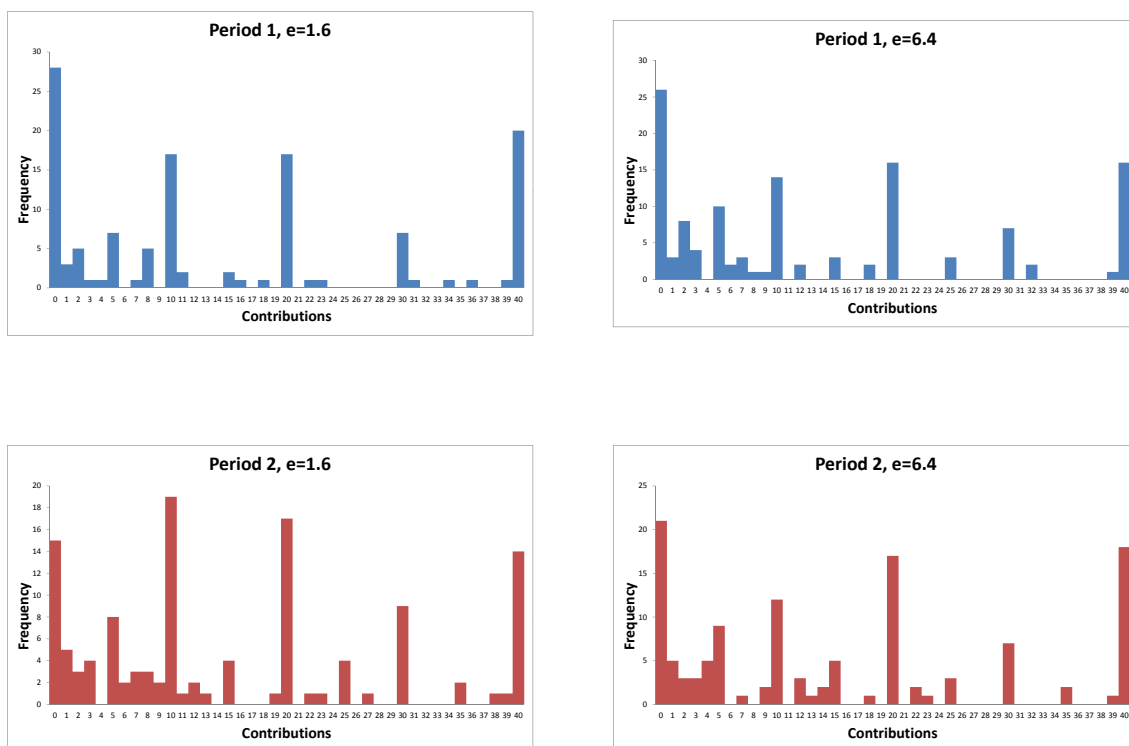
<sup>15</sup>We used ORSEE to recruit our subjects with this ‘first timer’ characteristic. In Section 4, we will test for experience and information effects by comparison with the rest of the data.

<sup>16</sup>See Andreoni, 1988; Croson, 1996; Cookson, 2000; Camerer, 2003.

<sup>17</sup>Means are 0.4 when  $e = 1.6$  and 1.6 when  $e = 6.4$ .

<sup>18</sup>As is usually the case; see Camerer, 2003.

Figure 1: Initial contributions (phases 1,2).



*Initial contributions often lie at 0, 10, 20, 30, 40, with several contributions made also in the range from 1 to 9; contributions at other numerical values are rare.*


(Figure 1 illustrates. See Appendix C, tables 4 and 5 for details.)

In the early periods of black box treatments, this is the behavior one would expect and it is consistent with hundreds of previous experiments on public goods.<sup>19</sup> Since players have no knowledge of the differences in the game structures, they initially make random contributions and play different games in much the same way. Note that players initially contribute less than half of their budget, another observation consistent with previous experimental results. A possible explanation for this is ambiguity aversion (Epstein, 1999).<sup>20</sup> At the beginning of the game, players cannot judge their initial payoffs to be particularly “good” or “bad”, and therefore have no comparison concerning the performance (“success” vs. “failure”) of different strategies. One would consequently not expect an obvious direction for the initial-period adjustments, and indeed they are similarly random until period 3 for the different rates of return.

<sup>19</sup>As surveyed, for example, in Chaudhuri, 2011.

<sup>20</sup>In Fehr & Schmidt, 1999, for example, mean contributions are ca. 40 percent of the budget; this would mean ca. 16 in our experiment.

Figure 2: Mean black box play (phases 1,2).



*Contributions deteriorate in the free-riding game. This trend is accentuated compared to the full-contribution games which display a weak positive trend.*

(Figure 2 illustrates.)


If players were not to learn at all, one would expect similar behavior throughout the games. Using period-specific Mann-Whitney tests, however, we reject the null of equal contributions for different rates of return in all periods after period 3. Moreover, play evolves following distinctly different patterns: the contributions towards the end of the game are closer to the respective stage game Nash equilibrium.<sup>21</sup> Simple linear regressions of contributions on a time trend (controlling for phase and group) reveal a trend of -0.6 when  $e = 1.6$  (significant at 1%), and of +0.1 when  $e = 6.4$  (significant at 10%). Note the trend is accentuated in the games with  $e = 1.6$ .<sup>22</sup> When  $e = 6.4$ , the mean contribution never exceeds twenty and thus remains more than “halfway” from the Nash contribution. Again, players’ ambiguity aversion may explain this phenomenon.<sup>23</sup>

<sup>21</sup>Means are 5.7 when  $e = 1.6$  and 18.5 when  $e = 6.4$ .

<sup>22</sup>This pattern is consistent with previous experiments as, for example, noted in Ledyard, 1995. Note our contributions decline more sharply than in Bayer et al. (2013), possibly indicating a “within-game restart effect” in their experiments since players are lead to believe that underlying stage games are changing.

<sup>23</sup>A contribution decrease represents less ambiguity and is, ceteris paribus, preferred by ambiguity-averse agents, a contribution increase results in more ambiguity (Schmeidler, 1989).

Figure 3: Final contributions.



*Final contributions amass close to zero in the free-riding game, and split between contributions close to zero and 40 for the full-contribution game.*

(Figure 3 illustrates. See Appendix C, table 6 for details.)

To summarize our observations regarding aggregate contributions in free-riding games, we reconfirm the standard findings from the literature regarding initial contributions, restart effects and contribution deterioration. Moreover, we note that in games where full contribution is a Nash equilibrium, the observed trends to Nash equilibrium are substantially weaker. Higher contributions seem harder to learn.

### 3.1.2 Difference in variation

**Adjustment.** The *adjustment* of a player  $i$  in period  $t$  is  $a_i^{t+1} - a_i^t$ .

**Success versus failure.** A player  $i$  experiences *success* in period  $t + 1$  if his realized payoff does not go down ( $\phi_i^{t+1} \geq \phi_i^t$ ); otherwise he experiences *failure*.

Based on a Levene's test (robust variance), the hypothesis of equal variances of adjustments following success versus failure is rejected with 99% confidence.<sup>24</sup> Furthermore, we regress absolute adjustments on success versus failure controlling for phase, period, individual, group, and two lags of contributions and payoffs. We obtain a significantly (with 99% confidence) smaller coefficient of adjustment following success (by -1.2) than after failure. The phase dummies are not significant. None of the period dummies is significant after period 3. The two previous contributions are significant. The previous-period payoffs are significant, those lagged two periods are not significant. We conclude that players' adjustments after successful contributions have a smaller variation than adjustments after unsuccessful contributions. As discussed in sections 1 and 2 of this paper,

<sup>24</sup>The robust variance test is not rejected dependent on whether  $e = 1.6$  (mean -0.5) or 6.4 (0.2).

Table 1: Adjustments after success and failure (phases 1,2).

statistic	success	failure
variance	123.9	188.5
mean	+0.8	-1.2

*SEARCH: the variance after success is lower than after failure; the average adjustments are close to zero.*

the phenomenon of “different adjustments following success versus failure” is a feature of “search” in several recent learning models.<sup>25</sup> Subsequently, we shall refer to this feature as SEARCH.

(Table 1 summarizes. See also Appendix C, table 7 for details.)

### 3.1.3 Directional learning

Our games have the following structure, of which subjects in the black box are unaware. First, despite knowing that they can make any contribution in the range from 0 to 40, subjects do not know that the Nash equilibrium lies at either end. Second, an individual’s payoff rises with the contributions of the other players but may rise or fall with his own adjustment depending on the game’s underlying rate of return. Importantly, in varying his own contribution, the individual agent is unable to distinguish between a self-induced success versus failure, and one that is caused by others: in the interplay with others, incremental adjustments in either direction may lead to higher or lower payoffs.<sup>26</sup>

The structure of the action space calls for a directional learning model (Selten & Stoecker, 1986; Selten & Buchta, 1998) based on success-versus-failure stimuli. Given period  $t$ , a player  $i$  may increase or decrease his contribution relative to his contribution in the previous period. As a result, he experiences either success or failure. Table 2 shows the model of directional adjustments that we propose and shall subsequently refer to as TREND.

To test the TREND hypothesis, given two subsequent payoffs,  $\phi_i^{t+1}$  and  $\phi_i^t$ , we regress adjustments relative to two lags of contributions ( $c_i^{t+2} - c_i^t$ ) and adjustments relative to

<sup>25</sup>See Young, 2009; Marden et al., 2011; Pradelski & Young, 2012.

<sup>26</sup>The effects that a player’s own actions and those of others have on one’s own payoff are important features in Young, 2009; Marden et al., 2011; Pradelski & Young, 2012.

Table 2: TREND hypothesis.

Player $i$ in period $t + 1$	given success in period $t$ ( $\phi_i^t \geq \phi_i^{t-1}$ )	given failure in period $t$ ( $\phi_i^t < \phi_i^{t-1}$ )
given an increase in period $t$ ( $a_i^t > a_i^{t-1}$ )	<i>increases contribution with respect to period <math>t - 1</math></i>	<i>decreases contribution with respect to period <math>t</math></i>
given a decrease in period $t$ ( $a_i^t < a_i^{t-1}$ )	<i>decreases contribution with respect to period <math>t - 1</math></i>	<i>increases contribution with respect to period <math>t</math></i>

previous contributions ( $c_i^{t+2} - c_i^{t-1}$ ) separately.


After success ( $\phi_i^{t+1} \geq \phi_i^t$ ), we obtain significant coefficients in line with the TREND hypothesis for adjustments relative to two lags of contributions ( $c_i^{t+2} - c_i^t$ ): we observe a significant increase with respect to contributions two periods ago (+7.1; significant at 1%) following an increase; and we observe a significant decrease with respect to contributions two periods ago (-12.5; significant at 1%) following a decrease. The previous contribution has a negative effect (-0.2; significant at 1%). All other coefficients are insignificant as in the failure case. Trend following after success therefore summarizes as follows:

- *If an increase is a success in the current period, the average next-period adjustment relative to the previous period –net of negative level effects– is another increase. Similarly, if a decrease is a success in the current period, the average next-period adjustment relative to the previous period –net of negative level effects– is another decrease.*

After failure ( $\phi_i^{t+1} < \phi_i^t$ ), we obtain significant coefficients in line with the TREND hypothesis for adjustments relative to previous contributions ( $c_i^{t+1} - c_i^t$ ): we observe significant coefficients for decrease with respect to the previous-period contribution (-3.4; significant at 1%) following an increase; and we observe significant increase with respect to the previous-period contribution (+2.2; significant at 1%) following a decrease. None of the phase or group dummies is significant. None of the period dummies after period 4 is significant. Past payoffs have significant but marginal (less than 0.05) effects. The previous contribution has a negative effect (-0.6; significant at 1%). Trend reversal after failure therefore summarizes as follows:

- *If an increase is a failure in the current period, the average next-period adjustment relative to the current period –net of negative level effects– is a decrease instead.*

Figure 4: TREND without level effect.



*TREND without level effects: success versus failure leaves an intermediate range over which predictions are ambiguous.*


*Similarly, if a decrease is a failure in the current period, the average next-period adjustment relative to the current period –net of negative level effects– is an increase instead.*

Based on TREND, given two previous contributions and a success versus failure stimulus, the location of the expected next-period contribution is known relative to the relevant previous contribution, but, given the same two previous contributions, the relative location of the contribution depending on success versus failure is unknown. This is because TREND implies adjustment tendencies only with respect to the current period after failure, and with respects to the previous period after success. It is unclear where, given both current and previous contributions, the next-period contribution will lie after either success or failure is experienced.

(Figure 4 illustrates. See also Appendix C, table 8 for details.)

We shall illustrate this point with an example as illustrated in Figure 4, if an agent contributes 10 in one period and 20 in the next, on average he contributes at least 10 if 20 was a success, and less than 20 if 20 was a failure. The following (stronger)

Figure 5: Average adjustment hypothesis (Bayer et al., 2013).



*The average adjustment hypothesis ignores level effects and places further structure on TREND, making predictions over the intermediate range where TREND remained ambiguous.*

hypothesis has been considered (Bayer et al., 2013): given two previous contributions, the next contribution is, in expectation, closer to the contribution which resulted in a higher payoff. If the stimulus is trend-neutral, the expectation lies exactly at the average. As illustrated in Figure 5, if an agent contributes 10 in one period and 20 in the next, on average he contributes at least 15 if 20 was a success, and less than 15 if 20 was a failure.

(Figure 5 illustrates. See Appendix C, table 9 for details.)

We regress contribution adjustments on the four trends controlling for phase, period, individual, group, two payoff lags, previous contribution. This yields the following effects:

given  $a_i^t$  and  $a_i^{t+1}$  of some individual  $i$ , the net average adjustments, controlling for level effects, are



$\mathbf{E}(a_i^{t+2}; a_i^{t+1}, a_i^t) - \frac{a_i^t + a_i^{t-1}}{2} =$	$\phi_i^t \geq \phi_i^{t-1}$	$\phi_i^t < \phi_i^{t-1}$
$a_i^t > a_i^{t-1}$	+3.6**	+0.3 <sup>o</sup>
$a_i^t < a_i^{t-1}$	-6.5**	-1.9* <sup>o</sup>

- \*\* : significant at 1%, consistent with the average hypothesis.
- <sup>o</sup> : insignificant, inconsistent with the average hypothesis.
- \*<sup>o</sup> : significant at 1%, inconsistent with the average hypothesis.

Note that the adjustments following success point in the “right” direction (following success) and are significant. The adjustments following failure point in the “wrong” direction (following failure) and are significant in case of decrease. As before, the previous contribution has a significant negative effect of -0.4 (with 99% confidence). Previous profits have a marginal but significant effect. If we drop the two payoff lags and the previous contribution from the regression (since they are not included in the average adjustment model), we improve in terms of the model’s predictions but not regarding all aspects:

given  $a_i^t$  and  $a_i^{t+1}$  of some individual  $i$ , the gross average adjustments, not controlling for level effects, are

$\mathbf{E}(a_i^{t+2}; a_i^{t+1}, a_i^t) - \frac{a_i^t + a_i^{t-1}}{2} =$	$\phi_i^t \geq \phi_i^{t-1}$	$\phi_i^t < \phi_i^{t-1}$
$a_i^t > a_i^{t-1}$	+1.4*	-3.3**
$a_i^t < a_i^{t-1}$	-5.4**	-0.9 <sup>o</sup>

- \* : significant at 5%, consistent with the average hypothesis.
- \*\* : significant at 1%, consistent with the average hypothesis.
- <sup>o</sup> : insignificant, inconsistent with the average hypothesis.

We conclude that there is strong evidence for TREND including negative previous-period

contribution level effects. *Ceteris paribus*, a higher previous-period contribution leads to a negative adjustment. Depending on trend stimulus (success or failure), success trends are followed with respect to the contribution two periods ago, failure trends are reversed with respect to the previous contribution. The size of the adjustments relative to the previous two-period average is ambiguous and depends on the level of the previous-period contribution as well as the spread of the two previous contributions. A learning model assuming average adjustments without level effects gives consistent predictions in the success and increase directions; given a decrease-failure, however, it may lead to wrong predictions. It is worth noting that in the context where the average adjustment rule was originally formulated (Bayer et al., 2013), the problematic decrease-failure trends are relatively rare because only games where free-riding is Nash are played. The most common failure in those games comes from contributing more, not less.<sup>27</sup>

### 3.2 Experience effects

In this section, we analyze the whole black box data including sixteen sessions and 9,440 observations. In 4,480 observations, individuals have “experience” in the sense that they play the black box games in the second stage having previously played a non-black box treatment in the first stage where they received information about the game structure in the form of detailed game instructions and about other players’ past actions (and payoffs). Even though subjects are explicitly told that a separate experiment is started after the first stage of the experiment, there is evidence that experience matters. This can be explained by subjects linking the black box game to the previous voluntary contributions games which seems reasonable given that payoffs will be similar and the action space is the same.

(Figure 6 illustrates. See Appendix C, table 10 for details.)


First, we test for distributional differences in contributions using period-specific Mann-Whitney tests for the two rates of return separately.<sup>28</sup> The following observations are made. In the initial period, restart effects are noted as previously. For both rates of return separately, the null hypothesis of equal contributions dependent on experience versus no experience cannot be rejected. Furthermore, there is no significant difference in between the two rates of return which is natural since players have no basis to distinguish between

---

<sup>27</sup>Conversely, in the games where fully contributing is Nash, decreases are the more common failures.

<sup>28</sup>Appendix C also contains regressions that cluster on individuals for the initial period and all other periods separately. The basic message being that there are significant differences if  $e = 6.4$  but no significant differences in the initial periods or if  $e = 1.6$

Figure 6: Mean black box play (phases 1,2 versus 3,4).



*Experience plays no visible role in the free-riding games. Experience has a positive effect in the full-contribution games.*

them.<sup>29</sup> If  $e = 6.4$ , the null hypothesis of equal contributions dependent on experience versus no experience can be rejected with 90% confidence after period 5 except for periods 16 and 17. If  $e = 1.6$ , it cannot be rejected in any period.

Second, we compare the linear time trends of contributions (Appendix C, table 11 contains the details). Simple linear regressions controlling for phase and group reveal a similar trend as before if  $e = 1.6$  (-0.5; significant at 1%) but a stronger drift  $e = 6.4$  (+0.3; significant at 1%). Previously the respective trends were -0.6 (99%) and +0.1 (90%). Indeed, the time trends are not significant from each other if  $e = 1.6$ , but the time trend is significantly larger with experience than without experience if  $e = 6.4$ .<sup>30</sup> Hence the contributions in games with high rates of return improve faster with experience, but there is no significant difference in play of games with low rates or return.

Finally, we perform the linear regressions with respect to SEARCH (absolute size of adjustments) and TREND (directional adjustments) from the previous section controlling for the same effects and adding experience dummies for both rates of return. We obtain

<sup>29</sup>The initial contributions for both rates of return lie between 14 and 15 in all four phases of the experiment.

<sup>30</sup>See Appendix C for the confidence intervals.

the following results. The absolute size of adjustments (SEARCH) is smaller after the initial two periods (all period dummies after period 3 are negative and significant), and decreasing with experience. As before, the previous contribution is a positive and significant factor. Payoff level effects are significant but marginal. All other effects (including phase dummies) are not significant. With respect to directional adjustments (TREND), experience leads to larger contributions but the direction of adjustments is unchanged. As before, the level of previous contributions is a negative and significant factor. All period dummies after period 4 are negative. All other effects are as before.

We conclude that experience has two main effects. On the one hand, experience decreases the SEARCH effect but preserves its sign. On the other hand, experience leads to larger contributions over time when  $e = 6.4$ , but has no significant effect when  $e = 1.6$ . In much a different setting, our findings complement the findings concerning experience in games with low rates of return from Marwell & Ames (1981), Isaac et al. (1984), Isaac et al. (1988) by analysis of games with high rates of return and by identifying a common SEARCH feature when standard information is withheld. In fact, it turns out that experience has a more significant effect in games where free-riding is not Nash (when  $e = 6.4$ ) in which case experienced players contribute more and learn Nash equilibrium at a faster rate. A possible explanation for this is that experience reduces the agents' perceived ambiguity of the game.

### 3.3 Comparison with information treatments

In this section, we shall investigate whether SEARCH and TREND are only features of black box behavior or whether they persist in the standard and enhanced treatments, that is, if players have explicit information about the structure of the game and others' actions (and payoffs). We shall find that SEARCH is very robust and persists in all treatments, while TREND is fairly robust but more sensitive to the treatment setting and rate of return.

#### 3.3.1 Distributional differences

Using period-specific Mann-Whitney tests we test for differences in contributions in the three information treatments for both rates of return.<sup>31</sup> It turns out that initial black box contributions are significantly lower than non-black box contributions for both rates of

---

<sup>31</sup>This yields three tests for each period for both high (when  $e = 6.4$ ) and low (when  $e = 1.6$ ) rates of return as summarized in Table 3.

Table 3: Mann-Whitney tests of equal contributions.

comparison		period						
		1	2	3	4	5	. . .	20
$e =$	treatments							
1.6	enhanced vs. standard	=	**	**	**	**	. . .	**
	black box vs. enhanced	**	=	***	***	***	. . .	***
	black box vs. standard	**	*	=	=	=	. . .	=
6.4	enhanced vs. standard	=	=	**	**	***	. . .	***
	black box vs. enhanced	***	***	***	***	***	. . .	***
	black box vs. standard	***	***	***	***	***	. . .	***

Significance: \*: 10%; \*\*: 5%; \*\*\*: 1%; =: not significant.

return. There are no differences in initial contributions between standard and enhanced treatments but significant differences with the black box treatment.<sup>32</sup> After period two, all treatments are significantly different from each other in all periods except standard and black box play when  $e = 1.6$  which are not significantly different in any period. When  $e = 6.4$ , standard treatment contributions significantly exceed enhanced treatment contributions which exceed black box contributions. When  $e = 1.6$ , standard and black box contributions are not significantly different from each other but exceed enhanced treatment contributions. In the final period, of all three treatments, enhanced information is closest to Nash equilibrium when  $e = 1.6$ , while standard and black box are closer to Nash when  $e = 6.4$ .

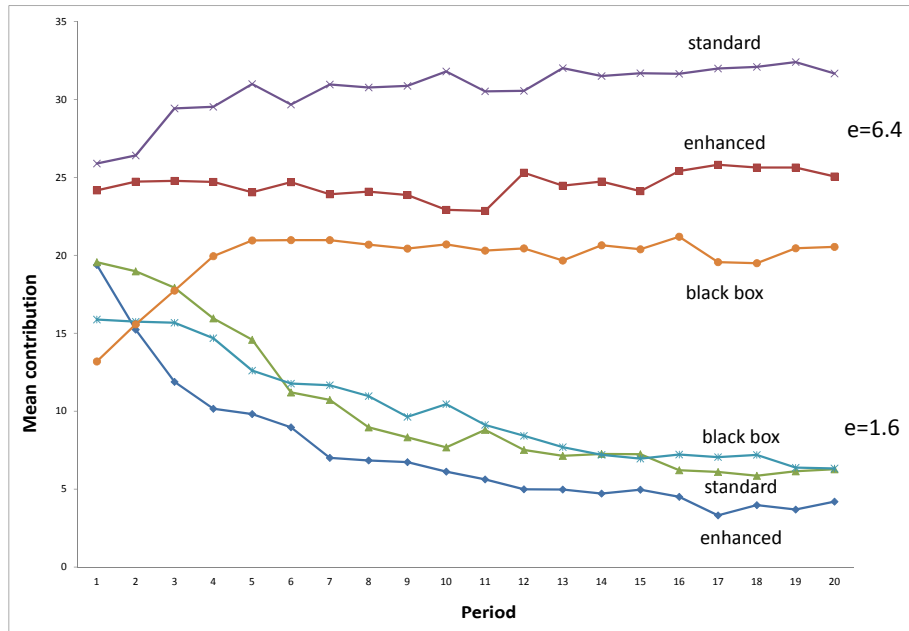
(Table 3 and Figures 7 and 8 illustrate.)

### 3.3.2 SEARCH

Testing for SEARCH, we perform tests for each rate of return separately, clustering on individuals and regressing absolute adjustments on black box, standard and enhanced information effects. We make the following observations. First, (more) information leads to less variation in adjustments in between time periods: for both rates of return, the standard and enhanced treatment dummies are negative. If  $e = 1.6$ , the enhanced treatment dummy is significantly more negative than the standard dummy. If  $e = 6.4$ , there


<sup>32</sup>Recall that initial black box contributions lie below 20 and show no differences for the two rates of return (see section 3.1.1). Initial contributions are higher for both standard and enhanced treatments; close to 20 when  $e = 1.6$ , and above 20 when  $e = 6.4$ .

Figure 7: Mean play all treatments (black box, standard, enhanced).



*Black box: initial contributions are low; trend when  $e = 1.6$ , weak trend when  $e = 6.4$ .  
 Standard and enhanced: initial contributions are higher; enhanced trend steeper than standard trends when  $e = 1.6$ , weak trends if  $e = 6.4$ .*

Figure 8: SEARCH and mean absolute adjustments (black box, standard, enhanced).



*SEARCH volatility is observed in all treatments; more information leads to less variation in adjustments in between time periods for both rates of return.*

is no significant difference. In all treatments and for both rates of return, success in the previous period leads to less variation in adjustments between time periods than failure in the previous period.

### 3.3.3 TREND

Testing for TREND, we perform tests for each rate of return separately (with individual clustering). We regress one-period and two-period adjustments treatment fixed effects and on TREND, controlling also for fixed effects from experience and of two lags of contributions and payoffs. We know that, overall, enhanced information leads to play over time that gets closer to the Nash equilibrium than black box and standard treatment play in the games with  $e = 1.6$ , while standard information play gets closer to the Nash equilibrium than enhanced and black box treatment play in the games with  $e = 6.4$ .

At the TREND level, we observe a violation of TREND in two cases: *i.* when a contribution decrease leads to a lower payoff (down and failure) in games with  $e = 1.6$ , the expected next-period contribution is nevertheless smaller in both standard and enhanced treatments (down is not reversed), *ii.* when a contribution decrease leads to a higher payoff (down and success) in games with  $e = 1.6$ , the expected next-period contribution is nevertheless larger in the standard treatment (down is not followed). The remaining TREND components are preserved in a weak sense (pointing in the right way) in both the standard and enhanced treatments even if some TREND components are no longer statistically significant.

The remaining differences concern the size of the fixed effects and effects of earlier period contribution lags. First, contributions two periods ago now are also significant (previously only contributions in the previous period but not two periods ago were significant). With 99% confidence, the level of contribution two periods ago has a negative effect on the contribution adjustments relative to two periods ago and a positive effect on the contribution adjustments relative to the previous period. The constant for both adjustments relative to one and two periods ago is negative in games when  $e = 1.6$  and positive when  $e = 6.4$ . This suggests longer “memory” effects in the non-black box treatments. The standard and enhanced information fixed effects in terms of adjustments of contributions relative to two periods ago are negative and significant (at different levels of significance), and positive but insignificant for adjustments relative to the previous period. Experience has no significant effects.

### 3.3.4 Summary

Comparing black box with standard and enhanced treatments at a purely distributional level, we observe the following differences. Initial contributions are higher in both standard and enhanced treatments (and not different from each other), and this finding includes the games with  $e = 1.6$  where such contributions are worse replies in terms of the game's Nash equilibrium. After period two, there are significant differences between all three treatments when  $e = 6.4$ . When  $e = 1.6$ , there are no differences between black box and standard treatments, and enhanced treatment contributions are significantly lower. In fact, enhancing the information relative to standard treatments leads to lower contributions in all games; these are higher than black box contributions when  $e = 6.4$  and lower than black box contributions when  $e = 1.6$ .<sup>33</sup>

Let us also briefly summarize our results regarding SEARCH and TREND. SEARCH is robust in all treatments. More information reduces the size of this effect but preserves its sign. TREND including level effects is robust in black box treatments with or without experience. Inside the box, adjustments are made dependent on whether the previous adjustment was successful or not. There is only partial evidence for TREND adjustments in the non-black box data and further contribution lags are significant. TREND fails to explain non-black box adjustments, and subjects' behavior is likely to interact and depend on other factors such as social preferences and social learning.

## 4 Conclusion

Much of the prior empirical work on learning in games has focussed on situations where players have a substantial amount of information about the structure of the game, and they can observe the behavior of others as learning proceeds. In this paper by contrast, we have examined situations in which players have no information about the strategic environment, and they must feel their way based solely on the pattern of realized payoffs. We identified two key features of such completely uncoupled learning dynamics – search volatility and directional adjustments, both of which have antecedents in the psychology literature. Search volatility in particular has not been examined previously in the context of experimental game theory and turns out to be very robust even when players gain

---

<sup>33</sup>Decreasing contributions due to more information have previously been observed (Huck et al., 2011). Indeed, this may be explained by considerations of relative payoffs. In both games, independent of the Nash equilibrium, those who contribute less are relatively better off. This generally aids learning in the games with  $e = 1.6$  but may lead to an adverse learning effect in the games with  $e = 6.4$ .



experience and information. Whether this remains true for other classes of games is an open question for future research.

When no information regarding other players' actions and payoffs is available and individuals are limited in their computational capacity and memory, individual learning amounts to trial-and-error type behavior. Outside game theory, these rules have a long tradition dating back to Thorndike (1898) and constitute the heart of experimental psychology of animal behavior processes. In our black box treatment, we enforce completely uncoupled behavior in the context of public goods games played in the economics laboratory. The key components of individual adjustments that we identify depart from traditional economic theory in a way not captured by existing models of bounded rationality linking reinforcement learning, directional learning and search volatility models.

## References

- J. Andreoni, "Why Free Ride? Strategies and Learning in Public Goods Experiments", *Journal of Public Economics* 37, 291-304, 1988.
- R.-C. Bayer, E. Renner & R. Sausgruber, "Confusion and learning in the voluntary contributions game", *Experimental Economics*, forthcoming, 2013.
- J. Björnerstedt, & J. Weibull, "Nash Equilibrium and Evolution by Imitation", in Arrow, K. and Colomatto, E. (Eds.), *Rationality in Economics*, New York: Macmillan, 1993.
- M. N. Burton-Chellew & S. A. West, "Pro-social preferences do not explain human cooperation in public-goods games", *Proceedings of the National Academy of Science* 110, 216-221, 2013.
- C. F. Camerer, *Behavioral Game Theory: Experiments on Strategic Interaction*, Princeton University Press, 2003.
- C. F. Camerer & T.-H. Ho, "Experience-weighted Attraction Learning in Normal Form Games", *Econometrica* 67, 827-874, 1999.
- A. Chaudhuri, "Sustaining cooperation in laboratory public goods experiments: a selective survey of the literature", *Experimental Economics* 14, 47-83, 2011.
- J. Coates, "The Hour Between Dog and Wolf. Risk Taking, Gut Feelings, and the Biology of Boom and Bust", *Penguin Press USA*, 2012.
- R. Cookson, "Framing effects in public goods experiments", *Experimental Economics* 3, 55-79, 2000.
- M. Costa-Gomes, V. Crawford & B. Broseta, "Cognition and behavior in normal-form games: An experimental study", *Econometrica* 69, 1193-1235, 2001.
- M. Costa-Gomes & V. Crawford, "Cognition and Behavior in Two-Person Guessing Games: An Experimental Study", *American Economic Review* 96, 1737-1768, 2006.
- V. P. Crawford, "Boundedly Rational versus Optimization-Based Models of Strategic Thinking and Learning in Games", *Journal of Economic Literature*, forthcoming, 2013.
- R. Croson, "Partners and Strangers Revisited", *Economics Letters* 53, 25-32, 1996.
- J. G. Cross, *A Theory of Adaptive Economic Behavior*, Cambridge University Press, 1983.

- J. Ding & A. Nicklisch, “On the Impulse in Impulse Learning”, *MPI Collective Goods Preprint 13/02*, 2013.
- A. E. Eiben & C. Schippers, “On evolutionary exploration and exploitation”, *Fundamenta Informaticae 35*, 35-50, 1998.
- I. Erev & A. Rapoport, “Coordination, ‘magic,’ and reinforcement learning in a market entry game” *Games and Economic Behavior*, 146-175, 1998.
- I. Erev & E. Haruvy, “Learning and the economics of small decisions”, J. H. Kagel & A. E. Roth (Eds.), *The Handbook of Experimental Economics*, forthcoming, Princeton University Press, 2013.
- I. Erev & A. E. Roth, “Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria”, *American Economic Review 88*, 848-881, 1998.
- L. G. Epstein, “A Definition of Uncertainty Aversion”, *The Review of Economic Studies 66*, 1999.
- E. Fehr & S. Gächter, “Cooperation and Punishment in public goods experiments”, *American Economic Review 90*, 980-994, 2000.
- E. Fehr & S. Gächter, “Altruistic punishment in humans”, *Nature 415*, 137-140, 2002.
- E. Fehr & K. M. Schmidt, “A Theory of Fairness, Competition and Cooperation”, *Quarterly Journal of Economics 114*, 817-868, 1999.
- E. Fehr & Camerer, “Social neuroeconomics: the neural circuitry of social preferences”, *Trends in Cognitive Sciences 11*, 419-427, 2007.
- U. Fischbacher, “z-Tree: Zurich Toolbox for Ready-made Economic Experiments”, *Experimental Economics 10*, 171-178, 2007.
- D. Foster & H. P. Young, “Regret testing: Learning to play Nash equilibrium without knowing you have an opponent”, *Theoretical Economics 1*, 341-367, 2006.
- D. Friedman, S. Huck, R. Oprea & S. Weidenholzer, “From imitation to collusion: Long-run learning in a low-information environment”, *Discussion Papers in Economics of Change SP II 2012-301*, WZB , 2012.
- F. Germano & G. Lugosi, “Global Nash convergence of Foster and Young’s regret testing”, *Games and Economic Behavior 60*, 135-154, 2007.
- B. Greiner, “An Online Recruitment System for Economic Experiments”, in: K. Kremer

& V. Macho (Eds.): *Forschung und wissenschaftliches Rechnen 2003*. GWDG Bericht 63, 79-93, 2004.

R. M. Harstad & R. Selten, "Bounded-Rationality Models: Tasks to Become Intellectually Competitive", *Journal of Economic Literature*, forthcoming, 2013.

S. Hart & A. Mas-Colell, "Uncoupled Dynamics Do Not Lead to Nash Equilibrium", *American Economic Review* 93, 1830-1836, 2003.

S. Hart & A. Mas-Colell, "Stochastic Uncoupled Dynamics and Nash Equilibrium," *Games and Economic Behavior* 57, 286-303, 2006.

T.-H. Ho, C. Camerer & K. Weigelt, "Iterated Dominance and Iterated Best-response in p-Beauty Contests", *American Economic Review* 88, 947-969, 1998.

S. Huck, P. Jehiel & T. Rutter, "Feedback spillover and analogy-based expectations: A multi-game experiment", *Games and Economic Behavior* 71, 351-365, 2011.

M. Isaac, J. Walker & S. Thomas, "Divergent Expectations on Free Riding: An Experimental Examination of Possible Explanations", *Public Choice* 43, 113-149, 1984.

M. Isaac & J. Walker, "Group Size Effects in Public Goods Provision: The Voluntary Contributions Mechanism", *Quarterly Journal of Economics* 103, 179-199, 1988.

J.-F. Laslier & B. Walliser, "Stubborn Learning", Ecole Polytechnique WP-2011-12, 2011.

J. O. Ledyard, "Public Goods: A Survey of Experimental Research", in J. H. Kagel & A. E. Roth (Eds.), *Handbook of experimental economics*, Princeton University Press, 111-194, 1995.

J. G. March, "Exploration and Exploitation in Organizational Learning", *Organization Science* 2, 71-87, 1991.

J. R. Marden, H. P. Young, G. Arslan, J. Shamma, "Payoff-based dynamics for multi-player weakly acyclic games", *SIAM Journal on Control and Optimization* 48, special issue on "Control and Optimization in Cooperative Networks", 373-396, 2009.

J. R. Marden, H. P. Young, L. Pao, "Achieving Pareto Optimality Through Distributed Learning", *discussion paper*, University of Colorado and University of Oxford, 2011.

G. Marwell & R. Ames, "Economists Free Ride, Does Anyone Else?", *Journal of Public Economics* 15, 295-310, 1981.

R. Nagel, "Unraveling in Guessing Games: An Experimental Study", *American Economic*

*Review 85*, 1313-1326, 1995.

H. H. Nax, “Completely uncoupled bargaining”, *mimeo.*, 2013.

H. H. Nax, B. S. R. Pradelski & H. P. Young, “The Evolution of Core Stability in Decentralized Matching Markets”, *Department of Economics Discussion Paper 607*, University of Oxford, 2012.

J. Oechssler & B. Schipper, “Can you guess the game you are playing?”, *Games and Economic Behavior 43*, 137-152, 2003.

B. S. R. Pradelski & H. P. Young, “Learning Efficient Nash Equilibria in Distributed Systems”, *Games and Economic Behavior 75*, 882-897, 2012.

A. Rapoport, D. A. Seale, & J. E. Parco, “Coordination in the aggregate without common knowledge or outcome information”, in: R. Zwick & A. Rapoport (Eds.): *Experimental Business Research*, 69-99, 2002.

A. E. Roth & I. Erev, “Learning in extensive form games: Experimental data and simple dynamic models in the intermediate term”, *Games and Economic Behavior 8*, 164-212, 1995.

H. Sauermann & R. Selten, “Anspruchsanpassungstheorie der Unternehmung”, *Zeitschrift für die Gesamte Staatswissenschaft 118*, 577-597, 1962.

D. Schmeidler, “Subjective Probability and Expected Utility without Additivity”, *Econometrica 57*, 571-587, 1989.

R. Selten & J. Buchta, “Experimental Sealed Bid First Price Auction with Directly Observed Bid Functions”, in: Budescu, D., Zwick, I.E. & R. (Eds.), *Games and Human Behavior*, Essays in Honor of Amnon Rapoport, 1998.

R. Selten & R. Stoecker, “End Behavior in Sequences of Finite Prisoner’s Dilemma Supergames: A Learning Theory Approach”, *Journal of Economic Behavior and Organization 7*, 47-70, 1986.

D. Stahl & P. Wilson, “On Players’ Models of Other Players: Theory and Experimental Evidence”, *Games and Economic Behavior 10*, 218-254, 1995.

P. Suppes & A. R. Atkinson, *Markov Learning Models for Multiperson Situations*, Stanford University Press, 1959.

F. Thuijsman, B. Peleg, M. Amitai & A. Shmida, “Automata, matching and foraging behavior of bees”, *Journal of Theoretical Biology 175*, 305-316, 1995.

E. Thorndike, "Animal Intelligence: An Experimental Study of the Associative Processes in Animals", *Psychological Review* 8, 1898.

R. Tietz & H. Weber, "On the nature of the bargaining process in the Kresko-game", in: Sauermann, H. (Ed.), *Contributions to experimental economics*, Vol. 3, pp. 305-334, 1972.

R. A. Weber, "Learning' with no feedback in a competitive guessing game", *Games and Economic Behavior* 44, 134-144, 2003.

H. P. Young, "Learning by trial and error", *Games and Economic Behavior* 65, 626-643, 2009.

# For online publication: appendices

## Appendix A: black box instructions

Participants received the following on-screen instructions (in z-Tree) at the start of the Black Box game and had to click an on-screen button saying, “I confirm I understand the instructions” before the game would begin:

### *Instructions*

Welcome to the experiment. You have been given 40 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.

During the experiment we shall not speak of £Pounds or Pence but rather of “Coins”. During the experiment your entire earnings will be calculated in Coins. At the end of the experiment the total amount of Coins you have earned will be converted to Pence at the following rate: 100 Coins = 15 Pence. In total, each person today will be given 3,200 coins (£4.80) with which to make decisions over 2 economic experiments and their final totals, which may go up or down, will depend on these decisions.

### *The Decision*

You can choose to keep your coins (in which case they will be ‘banked’ into your private account, which you will receive at the end of the experiment), or you can choose to put some or all of them into a ‘*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 function contains a random component, so if two people were to put the same amount of coins into the ‘*black box*’, they would not necessarily get the same output. The number outputted may be more or less than the number 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 box.

Any coins outputted will also be ‘banked’ and go into your private account. So, your final income will be the initial 40 coins, minus any you put into the ‘*black box*’, plus all the coins you get back from the ‘*black box*’.

You will play this game 20 times. Each time you will be given a new set of 40 coins to use. Each game is separate but the ‘*black box*’ remains the same. This means you

cannot play with money gained from previous turns, and the maximum you can ever put into the ‘*black box*’ will be 40 coins. And you will never run out of money to play with as you get a new set of coins for each go. The mathematical function will not change over time, so it is the same for all 20 turns. However as the function contains a random component, the output is not guaranteed to stay the same if you put the same amount in each time.

After you have finished your 20 turns, you will play one further series of 20 turns but with a new, and potentially different ‘*black box*’. The two boxes may or may not have the same mathematical function as each other, but the functions will always contain a random component, and the functions will always remain the same for the 20 turns. You will be told when the 20 turns 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.

*I confirm I understand the instructions*

## **Appendix B: on-screen output in each treatment**

Supplementary Figures: the post-decision feedback information that participants received: (a) in the black-box treatment; (b) the first feedback screen in both the standard and the enhanced-information treatments; and (c) the second feedback screen, which differed between the two treatments (not shown in black-box treatment). Dashed lines border the information that was only shown in the enhanced-information treatment.



**(a)**

GAME SUMMARY	Number of Coins
Initial Coins	40
Minus (-) your input	15
Plus (+) the output returned	24
<hr/>	
Your final number of coins	49

This screen lists your decisions and the results, along with your income (for this turn).

**(b)**

Player Name	Contribution (0-40)
You	15
player A	10
player B	30
player C	5
<hr/>	
Total contributions	60
Total after growth	96

This screen lists the decisions of you and the other players (in random order) in your group (for this round). Remember the terms Player A, B, & C **are meaningless** as you play with randomly selected people in each round. The group's total contributions and the new total after the 'growth' stage are also shown.

(c)

Player Name	Contribution (0-40)	Player Income =	Credits retained	+ Credits returned	= Total credits
<b>You</b>	15	For you =	25	24	49
player A	10	player A	30	24	54
player B	30	player B	10	24	34
player C	5	player C	35	24	59

Your income from your group's contributions and subsequent 'growth' is shown.

This screen lists the decisions and earnings of you and the other players (in random order) in your group (for this round).

Remember the terms Player A, B, & C **are meaningless** as you play with randomly selected people in each round.

## Appendix C: regression outputs

Table 4: Section 3.1.1. Mann-Whitney tests for difference in contributions: period 1

Comparison	# observations	Test statistic	(P-value)
1 if $e = 1.6$	124+124	-0.614	(0.5394)
1 if "phase"=1	124+124	-1.220	(0.2223)

We use black box data from phases 1 and 2. We test for differences in first-period contributions under inexperienced black box play. We find no evidence for significant differences depending on phase and rate of return.

Table 5: **Section 3.1.1.** Mann-Whitney tests for difference in contributions: 1 if  $e = 1.6$

Period	# observations	P-value	(Significance)
1	124+124	-0.614	(0.5394)
2	124+124	-0.366	(0.7146)
3	124+124	0.099	(0.9213)
4	124+124	2.305**	(0.0212)
5	124+124	3.300***	(0.0010)
6	124+124	3.566***	(0.0004)
7	124+124	3.369***	(0.0008)
8	124+124	4.539***	(<0.0001)
...	...	...	...
20	124+124	6.232***	(<0.0001)

Significance: \*: 10%; \*\*: 5%; \*\*\*: 1%.

We use black box data from phases 1 and 2. We test for differences in contributions under inexperienced black box play depending on rate of return. We find evidence for significant differences for all periods after three.

Table 6: **Section 3.1.1.** Time trends of contributions

	$e = 6.4$ coefficient (std. error)	$e = 1.6$ coefficient (std. error)
Period	<b>0.0831*</b> (0.0514)	<b>-0.5561***</b> (0.0404)
1 if "phase"=1	4.6486*** (0.5937)	-0.2628 (0.4675)
1 if "group"=2	-1.3031 (0.8345)	0.2734 (0.6492)
1 if "group"=3	-1.4063* (0.8345)	0.0016 (0.6492)
1 if "group"=4	-1.6472* (0.8545)	-0.4944 (0.6728)
Constant	15.8716*** (0.8479)	16.1658*** (0.6676)
Observations	2,480	0.0247
Adjusted $R^2$	2,480	0.0697

Significance: \*: 10%; \*\*: 5%; \*\*\*: 1%.

We use black box data from phases 1 and 2. We perform separate OLS regressions for each rate of return of contributions on a constant and a time trend controlling for phase and group effects.

Table 7: Section 3.1.2. SEARCH: absolute adjustments following success and failure

	Coefficient (std. error)
1 if “failure”	<b>1.1914***</b> (0.3768)
L.contribution	0.1948*** (0.0129)
L2.contribution	-0.0447*** (0.0127)
L.payoff	0.0275*** (0.0050)
L2.payoff	0.0094** (0.0050)
Individual fixed effects	not listed
Periods	3, 13, 15, 16, 19 significant below 10% others not significant
1 if “group”=1	-0.8012** (0.3908)
1 if “group”=2	-0.3644 (0.3919)
1 if “group”=3	-0.7581* (0.3922)
1 if “phase”=1	-0.0397 (0.2689)
Constant	3.3822** (1.6366)
Observations	4,464
Adjusted $R^2$	0.2246

Significance: \*: 10%; \*\*: 5%; \*\*\*: 1%.

We use black box data from phases 1 and 2. We perform an OLS regression of absolute adjustments on a dummy separating success and failure including a constant and controls for two lags of contributions and payoffs, as well as phase, group, period and individual fixed effects.

Table 8: **Section 3.1.3.** TREND: directional adjustments relative to previous periods

	<b>One-period adjustment: coefficient (std. error)</b>	<b>Two-period adjustment: coefficient (std. error)</b>
1 if “up” & “success”	0.1447 (0.5989)	<b>7.0971***</b> (0.7012)
1 if “down” & “success”	-0.4467 (0.5431)	<b>-12.5189***</b> (0.6359)
1 if “up” & “failure”	<b>-3.3727***</b> (0.5428)	3.9860*** (0.6356)
1 if “down” & “failure”	<b>2.2309***</b> (0.5901)	-6.0167*** (0.6909)
L.contribution	-0.5823*** (0.0150)	-0.2114*** (0.0176)
L.payoff	0.0204*** (0.0053)	0.0189*** (0.0062)
L2.payoff	0.0284** (0.0053)	-0.0028 (0.0062)
Individual fixed effects	not listed	not listed
Periods	1-6 not significant 7-20 significant at >90%	3, 14, 16, 17 not significant others significant below 10%
1 if “group” = 1	0.5948 (0.4578)	0.5320 (0.5360)
1 if “group” = 2	0.6235 (0.4590)	0.2241 (0.5374)
1 if “group” = 3	0.4216 (0.4595)	0.3352 (0.5381)
1 if “phase” = 1	0.5553 (0.3150)	0.2149 (0.3686)
Constant	1.3711* (1.9914)	2.1540 (2.2943)
Observations	4,464	4,464
Adjusted $R^2$	0.3175	0.2071

Significance: \*: 10%; \*\*: 5%; \*\*\*: 1%.

We use black box data from phases 1 and 2. We perform an OLS regression of one-period and two-period adjustments on dummies for each directional success-failure impulse including a constant and controls for two lags of payoffs, one contribution lag, as well as phase, group, period and individual fixed effects.

Table 9: **Section 3.1.3.** Bayer et al. (2013): average adjustment hypothesis

	Coefficient (std. error)
1 if “up” & “success”	3.6210*** (0.6086)
1 if “down” & “success”	-6.4828*** (0.5520)
1 if “up” & “failure”	0.3067 (0.5517)
1 if “down” & “failure”	-1.8929*** (0.5997)
L.contribution	-0.3969*** (0.0152)
L.payoff	0.0197*** (0.0054)
L2.payoff	0.0128** (0.0054)
Individual fixed effects	not listed
Periods	1-4 not significant 5-20 significant at 5%
1 if “group”=1	0.5635 (0.4652)
1 if “group”=2	0.4238 (0.4665)
1 if “group”=3	0.3784 (0.4670)
1 if “phase”=1	0.3851 (0.3201)
Constant	1.3711 (1.9914)
Observations	4,464
Adjusted $R^2$	0.1522

Significance: \*: 10%; \*\*: 5%; \*\*\*: 1%.

We use black box data from phases 1 and 2. We perform an OLS regression of adjustments to the previous two-period average on a dummy separating success and failure including a constant and controls for two lags of payoffs, one contribution lag, as well as phase, group, period and individual fixed effects.

Table 10: **Section 3.2:** Regressions on experience effects with individual-level clustering

Period		$e = 6.4$ coefficient (std. error)	$e = 1.6$ coefficient (std. error)
1	Experience dummy	-1.2964 (1.8648)	-1.5014 (1.7455)
	Constant	15.2661*** (1.2987)	13.9032*** (1.2408)
2-20	Experience dummy	<b>3.6732</b> ** (1.4738)	-0.1232 (0.8782)
	Constant	18.2920*** (0.9647)	9.8901*** (0.5651)

---

Significance: \*: 10%; \*\*: 5%; \*\*\*: 1%.

We use black box data from phases 1-4. We perform separate OLS regressions with individual subject-level clustering for each rate of return of contributions on a constant and an experience dummies for the initial period and all subsequent periods.

Table 11: **Section 3.2.** Mann-Whitney tests for difference in contributions: 1 if “experience” (“phase” > 2)

Period	# obs. (exp.+no exp.)	$e = 1.6$ test statistic (p-value)	$e = 6.4$ test statistic (p-value)
1	124+112	-0.851 (0.3949)	0.665 (0.5059)
2	124+112	0.069 (0.9448)	-0.095 (0.9241)
3	124+112	0.960 (0.3373)	-0.460 (0.6456)
4	124+112	0.041 (0.9671)	-0.883 (0.3771)
5	124+112	0.685 (0.4933)	-1.487 (0.1371)
6	124+112	0.659 (0.5098)	-1.788* (0.0738)
7	124+112	-0.386 (0.6992)	-1.957* (0.0504)
8	124+112	0.831 (0.4059)	-2.680*** (0.0074)
9	124+112	0.114 (0.9091)	-1.812* (0.0701)
10	124+112	0.143 (0.8864)	-2.049** (0.0405)
11	124+112	0.421 (0.6736)	-2.203** (0.0276)
12	124+112	0.598 (0.5499)	-3.068*** (0.0022)
13	124+112	0.118 (0.9057)	-2.696*** (0.0070)
14	124+112	0.383 (0.7018)	-2.833*** (0.0046)
15	124+112	-0.791 (0.4289)	-1.840* (0.0658)
16	124+112	0.416 (0.6774)	-1.565 (0.1175)
17	124+112	0.364 (0.7160)	-0.310 (0.7567)
18	124+112	0.730 (0.4652)	-2.582*** (0.0098)
19	124+112	0.442 (0.6585)	-2.395** (0.0166)
20	124+112	0.306 (0.7597)	-2.093** (0.0363)

Significance: \*: 10%; \*\*: 5%; \*\*\*: 1%.

We use black box data from phases 1-4. We test for differences in contributions under black box play depending on experience versus no experience for each rate of return separately. We find evidence for significant differences for most periods after five when  $e = 6.4$  but no differences when  $e = 1.6$ .



Table 12: **Section 3.2:** Time trends of contributions: experienced play

	$e = 6.4$ coefficient (std. error)	$e = 1.6$ coefficient (std. error)
Period	0.2781*** (0.0576)	-0.5287*** (0.0424)
1 if “phase” = 3	3.0366*** (0.6639)	-2.1152*** (0.4887)
1 if “group” = 1	0.3516 (0.8783)	-1.5640** (0.6465)
1 if “group” = 3	-0.9063 (0.8783)	-1.1984* (0.6465)
1 if “group” = 4	-2.3734** (1.0757)	-0.9359 (0.7918)
Constant	16.6344*** (0.9280)	17.6383*** (0.6831)
Observations	2,240	2,240
Adjusted $R^2$	0.0199	0.0749

Significance: \*: 10%; \*\*: 5%; \*\*\*: 1%.

We use black box data from phases 1-4. We perform separate OLS regressions for each rate of return of contributions on a constant and a time trend controlling for phase and group effects.

Table 13: **Section 3.3.2.** SEARCH: absolute adjustments following success and failure (cluster on individuals)

	$e = 6.4$ coefficient (std. error)	$e = 1.6$ coefficient (std. error)
1 if “failure” in black box	1.8871*** (0.2905)	3.3866*** (0.2698)
1 if “failure” in standard	1.8905*** (0.3980)	2.7538*** (0.3256)
1 if “failure” in enhanced	1.3157*** (0.4061)	1.6046*** (0.3177)
1 if standard	-2.1746*** (0.5149)	-1.0574*** (0.4130)
1 if enhanced	-1.5261*** (0.5797)	-2.1255*** (0.4337)
Constant	6.2599*** (0.3475)	5.3123*** (0.2415)
Observations	8,968	8,968
Adjusted $R^2$	0.0184	0.0437

Significance: \*: 10%; \*\*: 5%; \*\*\*: 1%.

We use all the data. We perform a regression of absolute adjustments on a dummy separating success and failure for each treatment and treatment dummies including a constant with individual clusters for each individual and first two and second two phases (accounting for experience) separately.

	One-period adj.: coeff. (std. error)	e = 1.6	Two-period adj.: coeff. (std. error)	e = 1.6
				<b>e = 1.6</b>
1 if “up & success” in black box	1.5558*** (0.5989)	1.1640* (0.5401)	1.5558*** (0.5895)	1.1639* (0.6802)
1 if “down & success” in black box	-4.0393*** (0.6101)	-2.6017*** (0.5497)	-4.0393*** (0.6101)	-2.6017*** (0.5497)
1 if “up & failure” in black box	-1.5259** (0.5431)	-1.7726*** (0.5293)	-1.5256** (0.6071)	-1.7726*** (0.5393)
1 if “down & failure” in black box	0.7230 (0.6101)	1.3349** (0.6579)	0.7231 (0.6732)	1.3349** (0.6579)
1 if “up & success” in standard	3.9850*** (0.9071)	4.3376*** (0.9407)	3.9850*** (0.9071)	4.3376*** (0.9407)
1 if “down & success” in standard	0.5618 (0.9897)	-0.5669 (0.6128)	0.5618 (0.9897)	-0.5669 (0.6128)
1 if “up & failure” in standard	-1.4162 (0.9899)	-0.9279 (0.6935)	-1.4162 (0.9899)	-0.9279 (0.6935)
1 if “down & failure” in standard	0.8114 (1.2826)	-0.1515 (1.0913)	0.8114 (1.2826)	-0.1515 (1.0621)
1 if “up & success” in enhanced	0.5114 (0.9832)	0.4568 (1.0913)	0.5114 (0.9832)	0.4568 (1.0913)
1 if “down & success” in enhanced	-2.1060** (0.9095)	-1.4127** (0.6248)	-2.1060** (0.9095)	-1.4137** (0.6248)
1 if “up & failure” in enhanced	-1.5658 (1.0649)	-0.4136 (0.7821)	-1.5658 (1.0649)	-0.4136 (0.7821)
1 if “down & failure” in enhanced	0.5936 (1.1394)	-0.4586 (0.7641)	0.5936 (1.1394)	-0.4586 (0.7641)
Constant	2.6293*** (0.6421)	-3.4252*** (0.9077)	2.6293*** (0.6421)	-3.4253*** (0.9077)
1 if standard	0.6032 (0.5115)	-1.0419* (0.5410)	0.6032 (0.5115)	-1.0419* (0.5410)
1 if enhanced	0.1016 (0.5107)	-1.2884** (0.5098)	0.1016 (0.5107)	-1.2884** (0.5098)
1 if experience	0.4160 (0.2750)	-0.3570 (0.2557)	0.4160 (0.2750)	-0.3570 (0.2557)
L.contribution	-0.5539*** (0.0211)	-0.5131*** (0.0210)	0.4461*** (0.0211)	0.4867*** (0.0210)
L2.contribution	0.3217*** (0.0211)	0.2317*** (0.0189)	-0.6783*** (0.0211)	-0.7684*** (0.0189)
L.payoff	0.0215*** (0.0033)	0.1381*** (0.0163)	0.0215*** (0.0033)	0.1381*** (0.0163)
L2.payoff	-0.0039 (0.0030)	0.0061 (0.0150)	-0.0039 (0.0030)	0.0061 (0.0150)
Observations	8,496	8,496	8,496	8,496
Adjusted $R^2$	0.2423	0.3032	0.3680	0.4471

Significance: \*: 10%; \*\*: 5%; \*\*\*: 1%.

We use all the data. We cluster on individuals and perform a regression of one-period and two-period adjustments on dummies for each directional success-failure impulse including a constant and controls for two lags of payoffs and contributions, as well as controls for treatments and experience.

Table 15: **Section 3.3:** Regressions on treatment effects with individual-level clustering

Period		$e = 6.4$ coefficient (std. error)	$e = 1.6$ coefficient (std. error)
1	Standard dummy	12.6973 (1.6095)	3.6790 (1.6285)
	Enhanced dummy	10.9760 (1.6431)	3.5020 (1.5777)
	Constant	13.1907 (0.8748)	15.8814 (0.9322)
2-20	Standard dummy	10.8300 (1.2291)	-0.2050 (0.8458)
	Enhanced dummy	4.5297 (1.3873)	-3.1088 (0.8086)
	Constant	20.0352 (0.7414)	9.8316 (0.4354)

We use all the data. We perform separate OLS regressions with individual subject-level clustering for each rate of return of contributions on a constant and dummies for standard and enhanced treatments for the initial period and all subsequent periods.