nature COMMUNICATIONS

ARTICLE

Received 14 Feb 2014 | Accepted 25 Apr 2014 | Published 29 May 2014

DOI: 10.1038/ncomms4976

OPEN

Extortion subdues human players but is finally punished in the prisoner's dilemma

Christian Hilbe[1,2], Torsten Röhl[1] & Manfred Milinski[3]

Extortion is the practice of obtaining advantages through explicit forces and threats. Recently, it was demonstrated that even the repeated prisoner's dilemma, one of the key models to explain mutual cooperation, allows for implicit forms of extortion. According to the theory, extortioners demand and receive an excessive share of any surplus, which allows them to outperform any adapting co-player. To explore the performance of such strategies against humans, we have designed an economic experiment in which participants were matched either with an extortioner or with a generous co-player. Although extortioners succeeded against each of their human opponents, extortion resulted in lower payoffs than generosity. Human subjects showed a strong concern for fairness: they punished extortion by refusing to fully cooperate, thereby reducing their own, and even more so, the extortioner's gains. Thus, the prospects of extorting others in social relationships seem limited; in the long run, generosity is more profitable.

[1] Evolutionary Theory Group, Max-Planck-Institute for Evolutionary Biology, August-Thienemann-Strasse 2, 24306 Plön, Germany. [2] Program for Evolutionary Dynamics, Harvard University, One Brattle Square, Cambridge, Massachusetts 02138, USA. [3] Department of Evolutionary Ecology, Max-Planck-Institute for Evolutionary Biology, August-Thienemann-Strasse 2, 24306 Plön, Germany. Correspondence and requests for materials should be addressed to C.H. (email: hilbe@fas.harvard.edu).

The repeated prisoner's dilemma has a long tradition of serving as a key model to explore the evolution of cooperation[1–6]. The rules of this stylized game are simple: in each round, two subjects simultaneously decide whether to cooperate or to defect. When both subjects cooperate they each receive a payoff R, which exceeds the payoff P for mutual defection. However, when a cooperating subject encounters a defector, the defector gets the highest possible payoff T, whereas the cooperator obtains the lowest payoff S. Although mutual defection is inefficient, it is the unique equilibrium if the prisoner's dilemma is only played for a single round. However, if subjects have the option to reciprocate past actions in future encounters, a considerable body of evidence suggests that mutual cooperation becomes feasible[7–12], and that it is in fact favoured by evolutionary forces[13–18].

Recently, the conclusion that repetition naturally promotes mutual cooperation has been challenged. With an elegant mathematical proof, Press and Dyson[19] have demonstrated that the repeated prisoner's dilemma also contains sophisticated strategies that aim to dominate the co-player. Such extortionate strategies have three remarkable properties: (i) they enforce a linear relationship between the player's own payoff and the opponent's payoff (strategies with this property were called zero-determinant strategies or ZD strategies); (ii) they prescribe to cooperate sufficiently often, such that the opponent's best response is to be fully cooperative; (iii) at the same time, extortioners aim to cooperate less often than their opponent, to gain higher payoffs. As a result, extortioners are unbeatable: in a pairwise encounter, they cannot be outperformed by any opponent. These surprising findings have attracted considerable attention[20], as they suggest that sophisticated players aware of such strategies are able to manipulate and exploit their partners, which should result in an evolutionary advantage.

Despite this relative strength, extortioners have problems to succeed in evolving populations[21–23]. Extortion is unstable: as a homogeneous population of extortioners ends up with the mutual defection payoff P, more cooperative strategies can easily invade and take over the population. Eventually, this dynamics may even promote the emergence of generous ZD strategies, which may be considered as the more benevolent counterpart to extortioners[24]. Generous ZD strategies share the first two properties of extortioners: they enforce a linear relationship between the payoffs of the two players, and they provide incentives for the opponent to cooperate. However, as opppsed to extortioners who aim to outcompete their opponents, the payoff of generous players never exceeds the payoff of the co-player. Although generous strategies seem to be too modest to succeed, they evolve under a wide range of conditions[25–27]. Extortionate strategies, on the other hand, require specific assumptions to be successful: extortioners either need to be stubborn and to stick to their strategy[19], or they need to adopt new strategies at a slower rate than their co-players[21,28–30].

Although these previous theoretical studies offer a fascinating new perspective on direct reciprocity and repeated games, they raise great expectations for studying how the two strategy classes, extortion and generosity, perform against real subjects. To this end, we have designed an economic experiment with four different treatments (see Table 1 and Methods). In each treatment, human subjects played 60 rounds of the prisoner's dilemma against a predefined computer programme (subjects did not receive any information about the length of the game or the nature of their opponent). The four treatments differed in the implemented ZD strategy of the computer programme, which was either strongly extortionate (ES), mildly extortionate (EM), mildly generous (GM) or strongly generous (GS).

For all treatments, theory predicts that humans maximize their expected payoff by cooperating in every round. In that case, extortioners do not only outperform their human opponents, but they are also expected to receive higher average payoffs than the generous ZD strategies. In the experiment, however, we find that although extortionate strategies indeed dominate their human co-players, this success comes at a cost. Humans are significantly less cooperative against extortioners. As a result, generosity is the more profitable strategy.

Results

Performance of ZD strategies against humans. Figure 1 shows the resulting average payoffs over all 60 rounds of the game, across the 4 treatments. These results confirm that the two extortionate ZD strategies indeed gain higher payoffs than their human co-players. For example, in the strong extortion treatment, the computer programme obtained an average payoff of $\pi_{ES} = €0.192$ per round, whereas the human subjects earned on average $\tilde{\pi}_{ES} = €0.128$ (Wilcoxon matched-pairs signed-rank test, $n_{ES} = 16$ human co-players, $Z = 3.523$, $P < 0.001$). Similarly, the mildly extortionate ZD strategy received a payoff of $\pi_{EM} = €0.208$, which clearly exceeds the mean payoff of the human opponents, $\tilde{\pi}_{EM} = €0.165$ (Wilcoxon matched-pairs signed-rank test, $n_{EM} = 14$, $Z = 3.181$, $P = 0.001$).

Conversely, in the two generosity treatments human subjects had the upper hand, as expected. In the mild generosity treatment, the ZD strategy earned $\pi_{GM} = €0.235$, as compared with the human subjects' mean payoff $\tilde{\pi}_{GM} = €0.260$ (Wilcoxon matched-pairs signed-rank test, $n_{GM} = 14$, $Z = -2.527$, $P = 0.012$). Lastly, the strong generosity treatment resulted in an average payoff of $\pi_{GS} = €0.237$ for the ZD strategy and $\tilde{\pi}_{GS} = €0.280$ for the human co-players (Wilcoxon matched-pairs

Table 1 | Overview of the experimental design.

Treatment	Number of human co-players	Cooperation probabilities					Slope
		p_0	p_R	p_S	p_T	p_P	s
ES	16	0.000	0.692	0.000	0.538	0.000	1/3
EM	14	0.000	0.857	0.000	0.786	0.000	2/3
GM	14	1.000	1.000	0.077	1.000	0.154	2/3
GS	16	1.000	1.000	0.182	1.000	0.364	1/3

ES, strong extortion; EM, mild extortion; GM, mild generosity; GS, strong generosity; ZD, zero determinant.
In each of the four treatments, the computer played according to a different ZD strategy. ZD strategies are defined by five probabilities: p_0 is the probability to cooperate in round $m = 1$, and for $i \in \{R, S, T, P\}$ the value of p_i is the probability to cooperate in round $m > 1$ after receiving the payoff i in round $m - 1$, see refs 6,31. Extortionate strategies do not cooperate in the first round, and they never cooperate after mutual defection. Generous strategies, on the other hand, cooperate in the first round and they always cooperate after mutual cooperation. For a derivation of the implemented cooperation probabilities, we refer to the Supplementary Methods. The parameter s determines the slope of the predicted payoff relation: for example, a slope of $s = 2/3$ implies that for each Cent that the ZD strategist earns additionally, the human co-player's additional payoff is 2/3 Cents. In general, a smaller slope increases the payoff inequality between players: decreasing the value of s makes extortionate ZD strategies even more extortionate, whereas it makes generous ZD strategies even more generous. For this experiment, we followed the parameters of ref. 3, that is, the payoffs were set to $T = €0.50$, $R = €0.30$, $P = €0.10$ and $S = €0.00$.

signed-rank test, $n_{GS}=16$, $Z=-2.521$, $P=0.012$). Thus, extortionate strategies dominated their respective co-players, whereas generous strategies let their co-players succeed. These results are in line with the theory of ZD strategies, which in fact makes virtually no assumptions about human play[19]. In addition, the relationship between the payoffs of the ZD strategist and the human co-player fits reasonably to the linear prediction, as illustrated by Fig. 2, despite the fact that the experimental game is only played for a finite number of rounds (see Methods).

Comparison of the performance of different ZD strategies. Surprisingly, however, both extortionate ZD strategies yielded a lower payoff than their two generous counterparts. Indeed, when we pool the two extortionate treatments and the two generous treatments, we find that generosity resulted in a >18% increase in payoffs (Mann–Whitney U-test, $n_E=n_G=30$, $Z=-2.544$, $P=0.011$). Against an extortionate ZD strategy, the mean cooperation rate of the human co-players was 34.2%, which is only half of the cooperation rate against generous ZD strategies, 67.7% (Mann–Whitney U-test, $n_E=n_G=30$, $Z=-3.625$, $P<0.001$). This gap comes unexpected, as the different ZD strategies provide similar incentives for their human co-players to cooperate (as indicated by the matching slope values in Table 1). However, a comparison of the human decisions over the course of the game suggests that the treatments followed a different dynamical pattern (Fig. 3). Generous ZD strategies were more successful in motivating their human co-players towards more cooperation: in the two generous treatments, humans had a cooperation rate of 53.0% during the first ten rounds, as compared with 76.0% during the last ten rounds (Wilcoxon matched-pairs signed-rank test, $n_G=30$, $Z=3.161$, $P=0.002$). In contrast, when paired with an extortionate ZD strategy, the cooperation rate of human subjects only slightly increased from 30.3%

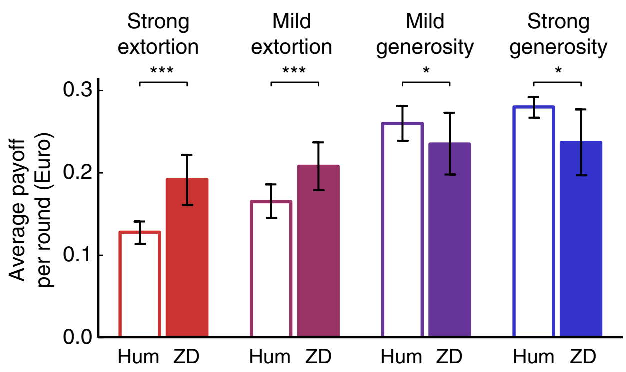

Figure 1 | Average payoffs across the four treatments for humans (empty bars) and the ZD strategies implemented by the computer programme (filled bars). In line with the theory, extortioners succeed against their human co-players, whereas generous ZD strategies lag behind their human opponents. Throughout the paper, we use two-tailed non-parametric tests for our statistical analysis, with each iterated game between a human co-player and the computer as our statistical unit (thus we have 16 independent observations for each of the 2 strong treatments, and 14 independent observations for each of the 2 weak treatments). In the above graph, three stars indicate significance at the level $\alpha=0.001$, and one star means significance for $\alpha=0.05$ (using Wilcoxon matched-pairs signed-rank tests with $n_{ES}=n_{GS}=16$, $n_{EM}=n_{GM}=14$). As an auxiliary information, we also provide error bars indicating the 95% confidence interval. Individual results for all 60 individuals are presented in the Supplementary Table 1.

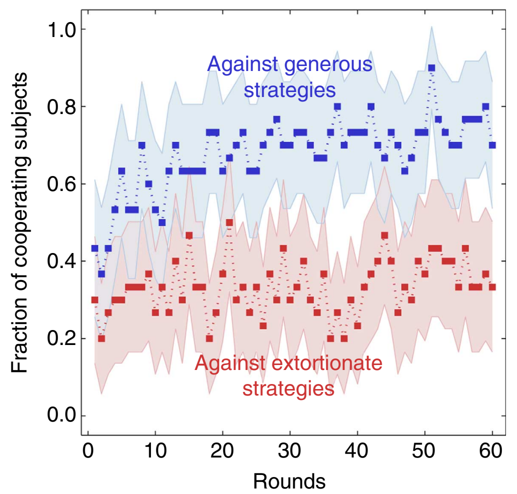

Figure 3 | Human cooperation rates over the course of the game. The graph shows the fraction of cooperating human subjects for each round for the two generosity treatments and the two extortion treatments. Dots represent the outcome of the experiment, with the shaded areas depicting the 95% confidence interval. Both curves start with cooperation rates around 30–40%. However, for the generous strategies we find a significant trend towards more cooperation, whereas for the extortionate strategies the average cooperation rates remain stable.

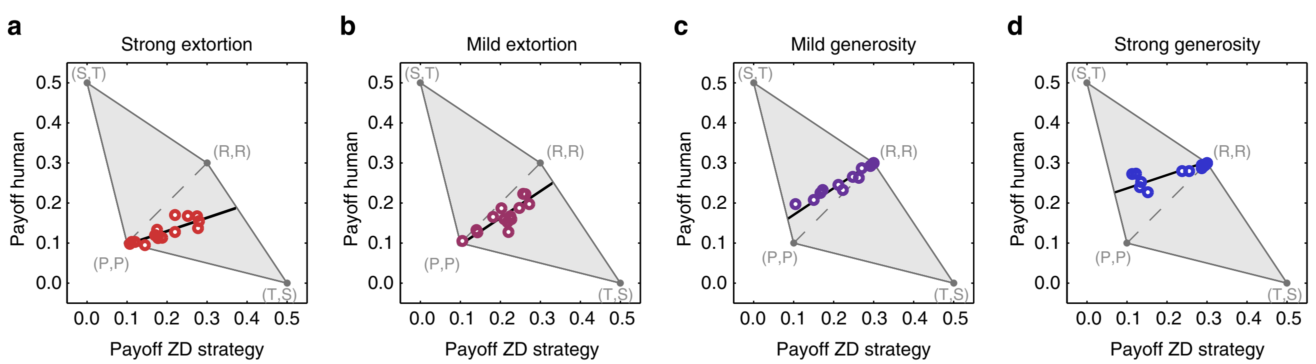

Figure 2 | Comparison of experimental results to the theoretical prediction. The grey-shaded area depicts the space of possible payoffs for the two players, that is, the ZD strategy implemented by the computer programme (x axis) and the human co-player (y axis). The black line corresponds to the theoretical prediction for the expected payoffs (as explained in the Methods) and the open circles indicate the outcome of the experiment. For the extortion treatments (**a**,**b**), these circles are below the diagonal (that is, extortioners outcompete their human co-players), whereas for the generosity treatments (**c**,**d**) these circles are above the diagonal (that is, generous players let their co-players succeed).

during the first ten rounds to 39.7% during the last ten rounds (this increase was not significant, Wilcoxon matched-pairs signed-rank test, $n_E = 30$, $Z = 1.131$, $P = 0.258$).

These results suggest that humans were somewhat reluctant to cooperate against extortioners. In fact, in the extortion treatments $<14\%$ of the human co-players were fully cooperative during the last ten rounds of the game, as compared with $>63\%$ in the generosity treatments (see Supplementary Fig. 1). On the other extreme, a third of the human subjects refused to cooperate against an extortionate co-player during the last ten rounds of the game, whereas only 1 out of 30 subjects did so in the generosity treatments. Thus, although the different treatments provided similar monetary incentives for cooperation, subjects were more hesitant to cooperate against an extortionate co-player. Withholding cooperation against these ZD strategies can be considered as a form of costly punishment (Fig. 4 and Supplementary Fig. 2). For example, reducing one's cooperation rate by 10% against strong extortioners decreased the opponent's mean payoff per round by € 0.029, but it also diminished the own payoff by € 0.011. The resulting fine-to-cost ratio for punishment, $0.029/0.011 \approx 2.6$, is close to typical values used in experiments on costly punishment[32]. Being less cooperative thus led to a strong reduction in the co-player's payoff, but it also turned out to be costly for the punishing individual itself.

Discussion

Repeated games, and in particular the repeated prisoner's dilemma, are model cases to explore the tension between cooperation and conflict in long-term social relations[33]. Although repetition was previously thought to promote cooperation, it has recently been suggested that iterated games may open the door for the systematic manipulation of opponents[19]. The newly discovered ZD strategies are surprisingly simple: they do not require to take the whole history of the game into account—it is sufficient to consider the last round only. Although previous literature on ZD strategies has focused on infinitely repeated games, social relationships in reality (and also our experiment) have a finite though fuzzy horizon. However, as we show in the Methods, this does not notably diminish the power of ZD strategies; if there is a sufficient number of rounds, ZD strategists have a similar amount of control as in the infinitely repeated game.

Two subclasses of ZD strategies have received particular attention: extortioners, as they are able to outcompete their direct opponents[19], and generous ZD strategies, as they allow for stable mutual cooperation[25,26]. Herein, we have investigated the performance of these two strategy classes against human subjects. Our results confirm that extortioners dominated their direct opponents, but unexpectedly generosity turned out to be the more profitable strategy. In a way, extortion meant to 'win each battle, but at the expense of losing the war'. These findings are superficially in line with previous evolutionary studies, which suggested that natural selection in well-mixed populations favours generous ZD strategies[26,27]. However, in these theoretical studies the success of generosity was based on a different argument; in an evolving population extortion does not prevail because mutual extortion is unstable, which leads extortioners to change their strategy[21]. In our experiment, the strategy of the extortioners was fixed, but extortioners were unable to motivate their co-players to cooperate fully, despite setting up appropriate incentives.

There are two possible explanations why humans were reluctant to cooperate against extortioners. On the one hand, subjects may have strived for high payoffs, but they did not have enough time to learn that they need to fully cooperate to reach this aim. This seems to be especially relevant as their opponents' strategies were stochastic and thus not straightforward to predict. However, this argument does not explain why generous ZD players were more successful to catalyse cooperation than extortioners—after all, the implemented ZD strategies were equally complex and they provided comparable monetary incentives to promote cooperation. Instead, our results suggest that the subjects were not only driven by monetary considerations, but that they were willing to apply reciprocal strategies to oppose extortionate behaviours. In fact, several behavioural studies have reported that a large fraction of humans can be described as conditional cooperators[34–36]. In line with this hypothesis, we find that humans were almost four times more likely to cooperate in a given round if their co-player did so in the previous round (human cooperation rates were 81.1% if the co-player cooperated in the previous round, and 22.0% otherwise, see Supplementary Table 2). Reciprocal behaviours in turn can have various behavioural roots, such as conformism, or the wish to enforce fair outcomes[37–39]. In our generosity treatments the two possible objectives, payoff-maximization and fairness, were perfectly aligned; by maximizing their expected payoffs humans also ensured equal outcomes. In contrast, in the extortion treatments there was a trade-off; humans that aimed to maximize their payoffs had to accept the most unfair outcome. As more than half of the participants declared in the post-experiment questionnaire that equality motives affected their decisions, the wish to ensure fair outcomes may have been an important reason for the downfall of extortion.

However, unlike in other strategic situations as in the ultimatum game[40], unfairness was not straightforward to detect in our behavioural experiment. It is not a single selfish decision

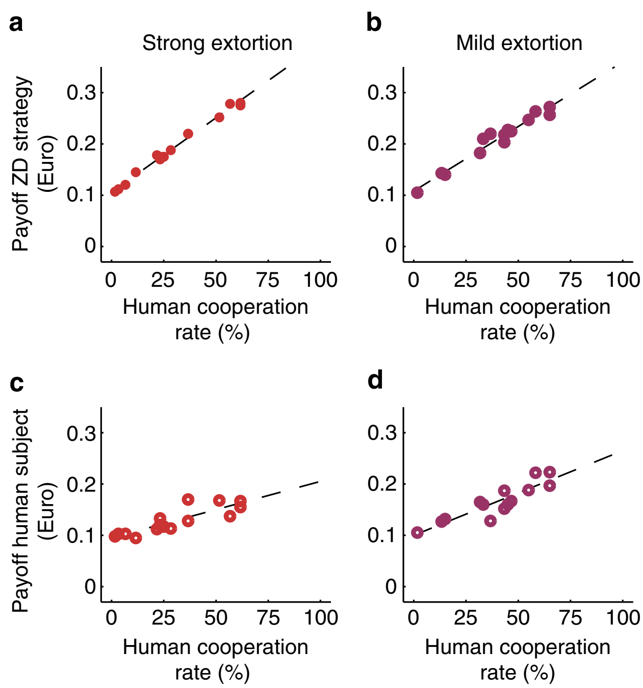

Figure 4 | Withholding cooperation as a form of costly punishment. The graph shows the effects of of human cooperation on the payoffs of ZD strategies (**a**,**b**) and on the human subjects' payoffs (**c**,**d**). The horizontal axis shows the fraction of rounds in which the human players cooperated. Coloured dots represent the outcome of the experiment, whereas the dashed line depicts the linear regression curve based on a least squares analysis. Human cooperation had a strongly positive impact on the co-player's payoff, and a weakly positive impact on the own payoff. Thus withholding cooperation punishes extortion.

that makes an opponent behaving extortionate. Rather, it is the systematic interplay of selfishness and cooperation, which only unfolds itself over the course of the game. At first sight, the extortionate strategies described by Press and Dyson[19] look rather inconspicuous (which may be one of the reasons why these strategies were discovered only recently). Extortioners apply a simple, conditionally cooperative strategy—with a slight bias to their own advantage. Although this more implicit form of selfishness seems to be more difficult to detect, humans have evolved mechanisms such as conditional cooperation that prevent them from being exploited.

Although our experiment did not entail an explicit punishment option, we found that by withholding contributions, subjects applied an implicit form of costly punishment. Such an effect has not been reported previously. In fact, it seems difficult to show such an effect with a conventional experiment, in which two human subjects play against each other. One would have to demonstrate that withholding cooperation is indeed individually costly. However, this seems almost impossible, as long as the co-player's strategy is unknown (for example, against an unconditional defector, withholding cooperation is the best response and hence no instance of costly punishment). Overall, our results thus suggest that sufficient monetary incentives alone are not enough to induce cooperation in long-term social relationships. Instead, humans take additional motives such as individual intentions and fairness considerations into account, and they are ready to fight back when they feel exploited.

Methods

Experimental design. Experiments were conducted in November and December 2013 at the universities of Kiel and Hamburg, Germany, with subjects recruited from a first-year course in biology. All participants gave their informed consent to participate. For each of ten experimental sessions, we invited six volunteers to participate in a game. To ensure the subjects' anonymity, participants were separated by opaque partitions, they were playing under a neutral pseudonym and they were not allowed to talk to each other during or after the experiment. All experimental decisions were made on a computer screen using the experimental software Z-Tree[41]. As we were interested in the relative performance of extortionate and generous strategies, participants were not playing against each other, but against a randomly determined computer strategy (out of the four alternatives ES, EM, GM or GS, as outlined in Table 1). The subjects' instructions were kept in a neutral way, that is, subjects were neither told that they would interact with a computer opponent nor that they would play against each other (see Supplementary Methods for a translation of the experiment's instructions). The game consisted of 60 rounds of the prisoner's dilemma (subjects were not informed about the exact duration of the game, but rather that they would play over many rounds). The experiment took $\sim$1 h. Including the show-up fee of € 10, individual earnings were on average between € 17.65 (in the strong extortion treatment) and € 26.78 (in the strong generosity treatment).

Theoretical predictions. The extortionate and generous strategies used for the experiment are instances of a more general strategy class, the class of ZD strategies. In an infinitely repeated prisoner's dilemma, a ZD strategist can unilaterally enforce a linear relation between his own payoff π and the co-player's payoff $\tilde{\pi}$. That is, payoffs obey a linear relation of the form[19,21,25,26]

$$-\tilde{\pi}+s\pi+(1-s)l=0, \qquad (1)$$

where l and s are characteristic properties of the applied ZD strategy[27]. The baseline payoff l can be interpreted as the payoff of a ZD strategy against itself (for the two extortionate strategies $l=P$, and for the two generous strategies $l=R$). The slope s determines how strongly the payoffs of the two players are correlated (for the two mild treatments, we have used $s=2/3$, corresponding to a rather high correlation; for the two strong treatments, we have used $s=1/3$). If the prisoner's dilemma is only repeated for a finite number of rounds M, equation (1) does not need to be satisfied any longer. Nevertheless, one can derive the following estimate for the players' expected payoffs (see Supplementary Methods),

$$-\frac{p_0}{\phi M}\leq -\tilde{\pi}+s\pi+(1-s)l\leq\frac{1-p_0}{\phi M}, \qquad (2)$$

where p_0 is the probability that the ZD-strategist cooperates in the first round and ϕ is a constant. In Fig. 2, the expected payoff range according to equation (2) is depicted as a thin black area. See Supplementary Methods for further details.

References

1. Rapoport, A. & Chammah, A. M. *Prisoner's Dilemma* (University of Michigan Press, 1965).
2. Trivers, R. L. The evolution of reciprocal altruism. *Q. Rev. Biol.* **46,** 35–57 (1971).
3. Axelrod, R. & Hamilton, W. D. The evolution of cooperation. *Science* **211,** 1390–1396 (1981).
4. Boyd, R. & Lorberbaum, J. No pure strategy is evolutionary stable in the iterated prisoner's dilemma game. *Nature* **327,** 58–59 (1987).
5. Nowak, M. A. Five rules for the evolution of cooperation. *Science* **314,** 1560–1563 (2006).
6. Sigmund, K. *The Calculus of Selfishness* (Princeton University Press, 2010).
7. Friedman, J. A non-cooperative equilibrium for supergames. *Rev. Econ. Stud.* **38,** 1–12 (1971).
8. Aumann, R. J. *Survey of Repeated Games* (Wissenschaftsverlag, Mannheim, 1981).
9. Fudenberg, D. & Maskin, E. Evolution and cooperation in noisy repeated games. *Am. Econ. Rev.* **80,** 274–279 (1990).
10. Wedekind, C. & Milinski, M. Human cooperation in the simultaneous and the alternating prisoner's dilemma: Pavlov versus generous tit-for-tat. *Proc. Natl Acad. Sci. USA* **93,** 2686–2689 (1996).
11. Milinski, M. & Wedekind, C. Working memory constrains human cooperation in the prisoner's dilemma. *Proc. Natl Acad. Sci. USA* **95,** 13755–13758 (1998).
12. Dal Bó, P. & Fréchette, G. R. The evolution of cooperation in infinitely repeated games: experimental evidence. *Am. Econ. Rev.* **101,** 412–429 (2011).
13. Molander, P. The optimal level of generosity in a selfish, uncertain environment. *J. Conflict. Resol.* **29,** 611–618 (1985).
14. Nowak, M. A. & Sigmund, K. Tit for tat in heterogeneous populations. *Nature* **355,** 250–253 (1992).
15. Stephens, D. W., McLinn, C. M. & Stevens, J. R. Discounting and reciprocity in an iterated prisoner's dilemma. *Science* **298,** 2216–2218 (2002).
16. Nowak, M. A., Sasaki, A., Taylor, C. & Fudenberg, D. Emergence of cooperation and evolutionary stability in finite populations. *Nature* **428,** 646–650 (2004).
17. St. Pierre, A., Larose, K. & Dubois, F. Long-term social bonds promote cooperation in the iterated prisoner's dilemma. *Proc. R. Soc. B Biol. Sci.* **276,** 4223–4228 (2009).
18. Fischer, I. *et al.* Fusing enacted and expected mimicry generates a winning strategy that promotes the evolution of cooperation. *Proc. Natl Acad. Sci. USA* **110,** 10229–10233 (2013).
19. Press, W. H. & Dyson, F. D. Iterated prisoner's dilemma contains strategies that dominate any evolutionary opponent. *Proc. Natl Acad. Sci. USA* **109,** 10409–10413 (2012).
20. Ball, P. Physicists suggest selfishness can pay. *Nature* doi:10.1038/nature.2012.11254 (2012).
21. Hilbe, C., Nowak, M. A. & Sigmund, K. The evolution of extortion in iterated prisoner's dilemma games. *Proc. Natl Acad. Sci. USA* **110,** 6913–6918 (2013).
22. Adami, C. & Hintze, A. Evolutionary instability of zero-determinant strategies demonstrates that winning is not everything. *Nat. Commun.* **4,** 2193 (2013).
23. Szolnoki, A. & Perc, M. Evolution of extortion in structured populations. *Phys. Rev. E* **89,** 022804 (2014).
24. Stewart, A. J. & Plotkin, J. B. Extortion and cooperation in the prisoner's dilemma. *Proc. Natl Acad. Sci. USA* **109,** 10134–10135 (2012).
25. Akin, E. Stable cooperative solutions for the iterated prisoner's dilemma. Preprint at http://arxiv.org/abs/1211.0969 (2013).
26. Stewart, A. J. & Plotkin, J. B. From extortion to generosity, evolution in the iterated prisoner's dilemma. *Proc. Natl Acad. Sci. USA* **110,** 15348–15353 (2013).
27. Hilbe, C., Nowak, M. A. & Traulsen, A. Adaptive dynamics of exortion and compliance. *PLoS One* **8,** e77886 (2013).
28. Bergstrom, C. T. & Lachmann, M. The Red King Effect: when the slowest runner wins the coevolutionary race. *Proc. Natl Acad. Sci. USA* **100,** 593–598 (2003).
29. Frean, M. R. & Abraham, E. R. Adaptation and enslavement in endosymbiont-host associations. *Phys. Rev. E* **69,** 051913 (2004).
30. Damore, J. A. & Gore, J. A slowly evolving host moves first in symbiotic interactions. *Evolution* **65,** 2391–2398 (2011).
31. Nowak, M. A. & Sigmund, K. A strategy of win-stay, lose-shift that outperforms tit-for-tat in the Prisoner's Dilemma game. *Nature* **364,** 56–58 (1993).
32. Fehr, E. & Gächter, S. Cooperation and punishment in public goods experiments. *Am. Econ. Rev.* **90,** 980–994 (2000).
33. Colman, A. M. *Game Theory and its Applications in the Social and Biological Sciences* (Butterworth-Heinemann, 1995).
34. Keser, C. & van Winden, F. Conditional cooperation and voluntary contributions to public goods. *Scand. J. Econ.* **102,** 23–39 (2000).

35. Fischbacher, U., Gächter, S. & Fehr, E. Are people conditionally cooperative? Evidence from a public goods experiment. *Econ. Lett.* **71,** 397–404 (2001).
36. Fischbacher, U. & Gächter, S. Social preferences, beliefs, and the dynamics of free riding in public goods experiments. *Am. Econ. Rev.* **100,** 541–556 (2010).
37. Fehr, E. & Schmidt, K. A theory of fairness, competition, and cooperation. *Q. J. Econ.* **114,** 817–868 (1999).
38. Brosnan, S. F. & de Waal, F. B. M. Monkeys reject unequal pay. *Nature* **425,** 294–297 (2003).
39. Oechssler, J. Finitely repeated games with social preferences. *Exper. Econ.* **16,** 222–231 (2013).
40. Güth, W., Schmittberger, R. & Schwarze, B. An experimental analysis of ultimatum bargaining. *J. Econ. Behav. Organ.* **3,** 376–388 (1982).
41. Fischbacher, U. z-tree: Zurich toolbox for ready-made economic experiments. *Exper. Econ.* **10,** 171–178 (2007).

Acknowledgements

We are grateful to A. Traulsen for his insightful comments. Moreover, we thank K. Hagel, H. Brendelberger and S. Dobler for support in performing the experiment and the 60 students for their participation. C.H. acknowledges generous funding from the excellence scholarship of the University of Vienna and from the Schrödinger scholarship of the Austrian Science Fund (FWF) J3475.

Author contributions

C.H., T.R. and M.M. designed the research; C.H. and M.M. performed the experiment and wrote the paper.

Additional information

Supplementary Information accompanies this paper at http://www.nature.com/naturecommunications

Competing financial interests: The authors declare no competing financial interests.

Reprints and permission information is available online at http://npg.nature.com/reprintsandpermissions/

How to cite this article: Hilbe, C. *et al.* Extortion subdues human players but is finally punished in the prisoner's dilemma. *Nat. Commun.* 5:3976 doi: 10.1038/ncomms4976 (2014).

Supplementary Figures

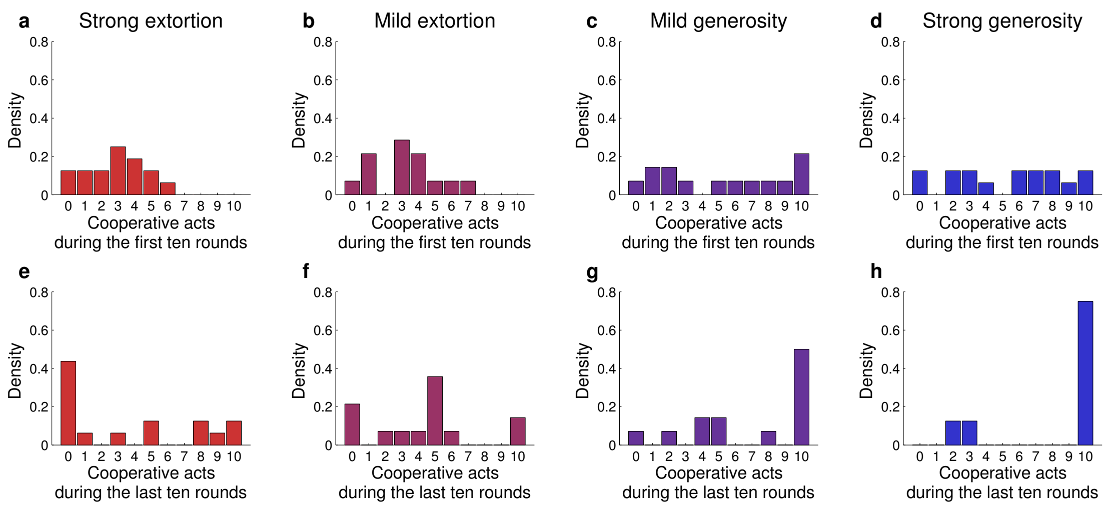

Supplementary Figure 1: Comparison between human cooperation rates in the beginning and in the end of the experiment. In each graph, the horizontal axis shows how often human subjects cooperated during the first ten rounds (**a - d**) or during the last ten rounds (**e - h**). In the beginning of the experiment, behaviors were rather evenly distributed. By the end of the experiment, however, most subjects either had a high cooperation rate (in the two generosity treatments) or they had a low or moderate cooperation rate (in the two extortion treatments). Nevertheless there were also a few subjects in the extortion treatments that were fully cooperative during the last ten rounds (2 out of 16 subjects in the strong extortion treatment, and 2 out of 14 subjects in the mild extortion treatment). In the post-experiment questionnaire, these cooperative subjects stated that they wanted to establish a regime of mutual cooperation, although most of them realized that their opponent was somewhat more selfish.

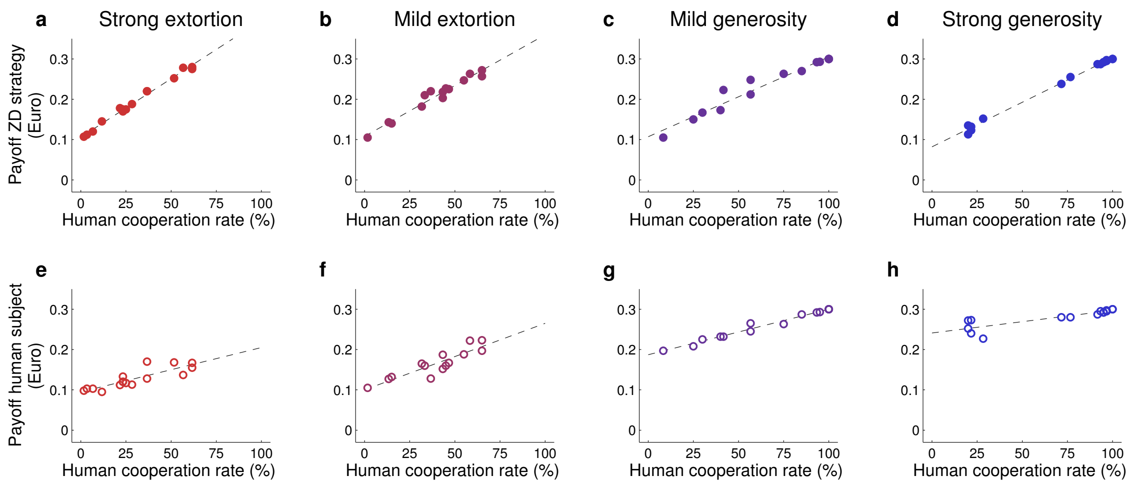

Supplementary Figure 2: Effects of human cooperation on the payoffs of ZD strategies (**a** - **d**) and on the human subjects' payoffs (**e** - **h**). In all graphs, the horizontal axis shows the fraction of rounds in which the human players cooperated. Colored dots represent the outcome of the experiment, whereas the dashed line depicts the linear regression curve based on a least squares analysis. (**a** - **d**) As expected, human cooperation had a positive impact on the payoffs of their ZD co-players: if humans increased their cooperation rate by 10 %, a linear regression analysis suggests that the payoffs per round of the ZD strategies increased by € 0.029 (ES), € 0.025 (EM), € 0.020 (GM), and € 0.022 (GS), respectively. (**e** - **h**) Also the human subjects themselves benefited from being more cooperative: an increase of their cooperation rates by 10 % resulted in an increase of their own payoffs by € 0.011 (ES), € 0.017 (EM), € 0.011 (GM), and € 0.006 (GS), respectively.

Supplementary Tables

Treatment		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	∅
Strong extortion	f	21	17	9	11	8	8	22	8	1	5	24	1	16	0	2	8	**10.1**
	$\tilde{f}$	31	34	14	22	14	15	37	17	2	13	37	7	22	1	4	14	**17.7**
	π	25.2	27.8	17.5	22.0	17.0	17.5	28.0	18.8	11.2	17.8	27.5	14.5	22.0	10.7	12.0	17.0	**19.2**
	$\tilde{\pi}$	16.8	13.7	13.3	12.8	12.0	11.7	15.5	11.3	10.3	11.2	16.7	9.5	17.0	9.8	10.3	12.0	**12.8**
Mild extortion	f	30	19	26	1	30	11	17	6	14	24	18	21	35	8	-	-	**18.6**
	$\tilde{f}$	39	27	33	1	35	22	19	8	20	26	26	28	39	9	-	-	**23.7**
	π	27.2	22.7	24.7	10.5	26.3	22.0	18.2	14.3	21.0	20.3	21.8	22.5	25.7	14.0	-	-	**20.8**
	$\tilde{\pi}$	19.7	16.0	18.8	10.5	22.2	12.8	16.5	12.7	16.0	18.7	15.2	16.7	22.3	13.2	-	-	**16.5**
Mild generosity	f	60	38	22	57	53	26	56	60	25	31	60	16	45	36	-	-	**41.8**
	$\tilde{f}$	60	34	15	57	51	25	56	60	18	24	60	5	45	34	-	-	**38.9**
	π	30.0	21.2	15.0	29.3	27.0	22.3	29.2	30.0	16.7	17.3	30.0	10.5	26.3	24.8	-	-	**23.5**
	$\tilde{\pi}$	30.0	24.5	20.8	29.3	28.7	23.2	29.2	30.0	22.5	23.2	30.0	19.7	26.3	26.5	-	-	**26.0**
Strong generosity	f	58	31	58	26	57	49	60	58	26	60	31	57	48	26	57	55	**47.3**
	$\tilde{f}$	58	13	58	13	57	46	60	58	17	60	12	56	43	12	57	55	**42.2**
	π	29.5	12.3	29.5	13.2	29.2	25.5	30.0	29.7	15.2	30.0	11.3	28.7	23.8	13.5	29.2	28.7	**23.7**
	$\tilde{\pi}$	29.5	27.3	29.5	24.0	29.2	28.0	30.0	29.7	22.7	30.0	27.2	29.5	28.0	25.2	29.2	28.7	**28.0**

Supplementary Table 1: Results of the experiment for each of the human-computer interactions of the experiment (16 interactions in each of the two strong treatments, and 14 interactions in each of the two mild treatments). The rows show the number of cooperative decisions for the computer (f) and for humans ($\tilde{f}$) over the sixty rounds of the experiments, as well as the resulting average payoffs per round in cents (π for the computer and $\tilde{\pi}$ for the human co-player, respectively).

Treatment	Decision of computer in previous period	Number of observations	Likelihood of human cooperation next round
Strong extortion	C	157	81.5%
	D	787	19.1%
Mild extortion	C	246	77.6%
	D	580	20.7%
Mild generosity	C	575	83.1%
	D	251	23.5%
Strong generosity	C	742	80.6%
	D	202	35.1%
Average over all treatments	C	1,720	81.1%
	D	1,820	22.0%

Supplementary Table 2: Evidence for conditional cooperation across treatments. Previous experiments on public good games suggest that humans are often conditionally cooperative, i.e. they wish to reciprocate the co-players' actions, see Refs. (1–3). Because generous ZD strategies are more cooperative than extortionate ZD strategies, this may explain why humans were significantly more cooperative in the generous treatments. In order to explore this explanation in more detail, we report how human subjects reacted to the outcome of the previous round. The third column shows the number of instances in which the computer cooperated during the first 59 rounds of the game (with C referring to cooperation and D referring to defection). The fourth row shows how often the human player cooperated in the following round. In all treatments, human cooperation rates were between 77.6% and 83.1% if the computer cooperated in the previous round. Otherwise, human cooperation rates were only between 19.1% and 35.1%.

Supplementary Methods

Theoretical Methods

Zero-determinant strategies in finitely repeated games

Previous studies on ZD strategies have considered infinitely repeated games (4–10), whereas the experiment had a finite number of rounds. We will demonstrate here that the definition of ZD strategies can be appropriately extended. Moreover, we will show that if the game is repeated sufficiently often, then the ZD strategies for the finitely repeated game allow a similar degree of control as in the original setup.

Let us consider a repeated prisoner's dilemma with possible payoffs R (if both players cooperate), S (if the focal player cooperates, and the co-player defects), T (if the focal player defects, and the co-player cooperates), and P (if both players defect). We assume that the typical relation $T > R > P > S$ holds. ZD strategies belong to the class of memory-one strategies: the decision whether to cooperate in a given round only depends on the outcome of the previous round. Such strategies can be written as a 5-tuple $(p_0, p_R, p_S, p_T, p_P)$. In this representation, p_0 is the player's probability to cooperate in round $m = 1$, and p_i for $i \in \{R, S, T, P\}$ is the probability to cooperate in round $m \geq 2$ after obtaining the payoff i in round $m - 1$. ZD strategies are defined as those memory-one strategies for which there are constants l, s, and $\phi > 0$ such that

$$\begin{array}{rcll} p_R &=& 1- & \phi(1-s)(R-l) \\ p_S &=& 1- & \phi\big[(1-s)(S-l)+T-S\big] \\ p_T &=& & \phi\big[(1-s)(l-T)+T-S\big] \\ p_P &=& & \phi(1-s)(l-P) \end{array} \qquad (1)$$

(this definition follows from the definition given by Press and Dyson, Ref. (4), by using the transformation $\alpha = \phi s$, $\beta = -\phi$, and $\gamma = \phi(1-s)l$). The following property of ZD strategies is central to our experiment:

Proposition 1 (An estimate for ZD strategies)
Consider a finitely repeated prisoner's dilemma with M rounds. Suppose one player applies a ZD strategy $(p_0, p_R, p_S, p_T, p_P)$ with parameters l, s, and $\phi > 0$. Let π denote the resulting average payoff per round for the ZD strategist, and let $\tilde{\pi}$ denote the respective payoff of the

co-player. Then independent of the co-player's strategy, payoffs satisfy the relation

$$-\frac{p_0}{\phi M} \;\leq\; (1-s)l + s\pi - \tilde{\pi} \;\leq\; \frac{1-p_0}{\phi M}. \tag{2}$$

In particular, it follows that

$$\big|(1-s)l + s\pi - \tilde{\pi}\big| \leq \frac{1}{\phi M}. \tag{3}$$

Proof. The proof is merely a slight variation of the proofs presented in Ref. (5) and Ref. (10). For $i \in \{R,S,T,P\}$ let $v_i(m)$ denote the probability that the payoff of the ZD strategist in round m is i. Let us introduce the following vector notation:

$$\begin{aligned}
\mathbf{v}(m) &= \big(v_R(m), v_S(m), v_T(m), v_P(m)\big)^{\mathsf{T}}\\
\mathbf{p} &= \big(p_R, p_S, p_T, p_P\big)\\
\mathbf{g} &= \big(R, S, T, P\big)\\
\tilde{\mathbf{g}} &= \big(R, T, S, P\big)\\
\mathbf{1} &= \big(1, 1, 1, 1\big)\\
\mathbf{e} &= \big(1, 1, 0, 0\big)
\end{aligned}$$

Using this notation, we can write the ZD strategist's expected payoff in round m as $\pi(m) = \mathbf{g} \cdot \mathbf{v}(m)$, and the co-player's expected payoff in that round as $\tilde{\pi}(m) = \tilde{\mathbf{g}} \cdot \mathbf{v}(m)$. Additionally, the definition of ZD strategies, Eq. (1), implies the identity

$$\mathbf{p} = \mathbf{e} + \phi\big[(1-s)(l\mathbf{1} - \mathbf{g}) + \mathbf{g} - \tilde{\mathbf{g}}\big]. \tag{4}$$

Let $q(m)$ denote the ZD strategist's probability to cooperate in round m. Using the previous notation, we can write $q(m) = \mathbf{e} \cdot \mathbf{v}(m)$ and $q(m+1) = \mathbf{p} \cdot \mathbf{v}(m)$. Thus, the quantity $w(m) := q(m+1) - q(m)$ satisfies the relation

$$w(m) = (\mathbf{p} - \mathbf{e}) \cdot \mathbf{v}(m) = \phi\big[(1-s)(l\mathbf{1} - \mathbf{g}) + \mathbf{g} - \tilde{\mathbf{g}}\big] \cdot \mathbf{v}(m) = \phi\big[(1-s)l + s\pi(m) - \tilde{\pi}(m)\big] \tag{5}$$

Taking the definition of $w(m)$, it follows that

$$\frac{1}{M}\sum_{m=1}^{M} w(m) = \frac{q(M+1) - q(1)}{M} = \frac{q(M+1) - p_0}{M} \tag{6}$$

On the other hand, by Eq. (5) we have

$$\frac{1}{M}\sum_{m=1}^{M} w(m) = \frac{\phi}{M}\sum_{m=1}^{M}\big[(1-s)l + s\pi(m) - \tilde{\pi}(m)\big] = \phi\big[(1-s)l + s\pi - \tilde{\pi}\big]. \tag{7}$$

As the two expressions (6) and (7) need to coincide, and as the definition of $q(m)$ requires $0 \le q(m) \le 1$, it follows that

$$-\frac{p_0}{\phi M} \;\le\; (1-s)l + s\pi - \tilde{\pi} \;\le\; \frac{1-p_0}{\phi M}.$$

□

We note that in the limit of infinitely repeated games, $M \to \infty$, we recover the result that ZD strategies enforce a linear relationship between payoffs,

$$\tilde{\pi} = s\pi + (1-s)l. \tag{8}$$

By choosing appropriate parameters l, s, and ϕ, the ZD strategist has a direct influence on this functional relationship.

Derivation of the strategies used for the experiment

Our experiment considers the performance of four different ZD strategies. To choose the corresponding parameters p_0, l, s and ϕ of these ZD strategies, we applied the following considerations:

- Parameter p_0: As extortioners are defined as strategies that cannot be outperformed by any opponent, they need to set $p_0 = 0$ (otherwise they would be outperformed by unconditional defectors). Analogously, a generous strategy needs to set $p_0 = 1$.

- Parameter l: Due to a similar reasoning, extortionate strategies require $l = P$, see Refs. (4, 9), whereas generous ZD strategies require $l = R$, see Refs. (7, 9).

- Parameter s: Since both strategy classes, extortioners and generous ZD strategies, require s to be in the unit interval $0 \le s \le 1$, Refs. (4, 11), we have chosen to use $s = 1/3$ (for the two "strong" treatments) and $s = 2/3$ (for the two "mild" treatments).

- Parameter ϕ: According to inequality (3), higher values of ϕ allow a ZD strategist to have a better control over the resulting payoff relations. Thus, in order to reduce the variance of our results, we have used the maximum ϕ-value (subject to the constraint that the resulting probabilities in Eq. (1) need to satisfy $0 \le p_i \le 1$, for all i).

Using the payoff values of the experiment $T = 0.5$, $R = 0.3$, $P = 0.1$ and $S = 0.0$, and Eq. (1), these parameter choices imply the following ZD strategies for the experiment (see also Table 1 in the main text):

Strong extortion ($l = P$, $s = 1/3$, $\phi = 30/13$)

$$p_0 = 0.000, \quad p_R = 0.692, \quad p_S = 0.000, \quad p_T = 0.538, \quad p_P = 0.000.$$

Mild extortion ($l = P$, $s = 2/3$, $\phi = 15/7$)

$$p_0 = 0.000, \quad p_R = 0.857, \quad p_S = 0.000, \quad p_T = 0.786, \quad p_P = 0.000.$$

Mild generosity ($l = R$, $s = 2/3$, $\phi = 30/13$)

$$p_0 = 1.000, \quad p_R = 1.000, \quad p_S = 0.077, \quad p_T = 1.000, \quad p_P = 0.154.$$

Strong generosity ($l = R$, $s = 1/3$, $\phi = 30/11$)

$$p_0 = 1.000, \quad p_R = 1.000, \quad p_S = 0.182, \quad p_T = 1.000, \quad p_P = 0.364.$$

Predictions for the experiment

In the experiment, the prisoner's dilemma was played for $M = 60$ rounds. Therefore, the inequalities in (2) predict the following relation between the expected payoff π of the ZD strategist and the expected payoff $\tilde{\pi}$ of the human co-player:

Strong extortion ($p_0 = 0$, $l = P = 0.1$, $s = 1/3$, $\phi = 30/13$)

$$\frac{1}{3} \cdot \pi + \frac{2}{3} \cdot 0.1 - \frac{13}{1800} \quad \leq \quad \tilde{\pi} \quad \leq \quad \frac{1}{3} \cdot \pi + \frac{2}{3} \cdot 0.1 \tag{9}$$

Mild extortion ($p_0 = 0$, $l = P = 0.1$, $s = 2/3$, $\phi = 15/7$)

$$\frac{2}{3} \cdot \pi + \frac{1}{3} \cdot 0.1 - \frac{7}{900} \quad \leq \quad \tilde{\pi} \quad \leq \quad \frac{2}{3} \cdot \pi + \frac{1}{3} \cdot 0.1 \tag{10}$$

Mild generosity ($p_0 = 1$, $l = R = 0.3$, $s = 2/3$, $\phi = 30/13$)

$$\frac{2}{3} \cdot \pi + \frac{1}{3} \cdot 0.3 \quad \leq \quad \tilde{\pi} \quad \leq \quad \frac{2}{3} \cdot \pi + \frac{1}{3} \cdot 0.3 + \frac{13}{1800} \tag{11}$$

Strong generosity ($p_0 = 1$, $l = R = 0.3$, $s = 1/3$, $\phi = 30/11$)

$$\frac{1}{3} \cdot \pi + \frac{2}{3} \cdot 0.3 \quad \leq \quad \tilde{\pi} \quad \leq \quad \frac{1}{3} \cdot \pi + \frac{2}{3} \cdot 0.3 + \frac{11}{1800} \tag{12}$$

In Fig. 2 of the main text, we have illustrated these inequalities: the pairs (π,$\tilde{\pi}$) that satisfy the above constraints are shown as black lines (note that the estimates (9) – (12) are statements about expected payoffs. As ZD strategies are stochastic, and as the game is only repeated for a finite number of rounds, realized payoffs do not need to satisfy the above inequalities, as can be seen in Fig. 2.) For the two extortionate treatments, these black lines are on or below the diagonal, implying that humans never yield a higher expected payoff than their extortionate co-players. Analogously, in the two generous treatments, the black lines are on or above the diagonal, implying that humans are never worse off than their generous co-players.

If human subjects aim to maximize their payoffs, their best response in all four treatments is to cooperate in all rounds. In fact, as the ZD strategist enforces a positive relation between payoffs ($s > 0$), subjects maximize their own payoff $\tilde{\pi}$ by maximizing their co-player's payoff π. When humans play their best response, expected payoffs can be calculated as

	Payoff π of ZD strategist	Payoff $\tilde{\pi}$ of human co-player
Strong extortion (ES)	0.37	0.19
Mild extortion (EM)	0.33	0.25
Mild generosity (GM)	0.30	0.30
Strong generosity (GS)	0.30	0.30

(13)

Thus, the payoffs of the ZD strategies satisfy $\pi_{ES} > \pi_{EM} > \pi_{GM} = \pi_{GS}$ – if human subjects move towards full cooperation in all four treatments, then extortionate ZD strategies should receive higher payoffs than their generous counterparts. On the other hand, for the human co-players we obtain the relation $\tilde{\pi}_{GS} = \tilde{\pi}_{GM} > \tilde{\pi}_{EM} > \tilde{\pi}_{ES}$ – not very surprisingly, humans prefer their opponents to be generous, rather than extortionate.
In summary, we have the following predictions:

1. Independent of how human subjects play, extortionate ZD strategies get at least the payoff of their co-players in each game. Generous ZD strategies obtain at most the payoff of their opponent.

2. Assuming that human subjects aim to maximize their payoffs, we would expect that their cooperation rates increase over the course of the game.

3. If there is a comparable trend towards cooperation across treatments, extortionate ZD strategies should earn higher payoffs than generous ZD strategies.

In response to our theoretical predictions, the main text reports the following findings:

1. In line with the theoretical prediction, the two extortionate strategies gained higher payoffs than their human co-players, whereas the two generous strategies obtained lower payoffs than their co-players.
2. Only in the generous treatments humans became significantly more cooperative over the course of the game (cooperation rates increased from 53 % during the first ten rounds to 76 % during the last ten rounds of the experiment, Wilcoxon matched-pairs signed-rank test, $n_G = 30$, $Z = 3.161$, $p = 0.002$). In the extortionate treatments there was only a slight trend towards more cooperation, and this trend failed to be significant (cooperation rates increased from 30.3 % to 39.7 %, Wilcoxon matched-pairs signed-rank test, $n_E = 30$, $Z = 1.131$, $p = 0.258$).
3. In contrast to our prediction, generous strategies obtained higher payoffs than extortioners ($\pi_G = €\ 0.236$ vs. $\pi_E = €\ 0.199$, $n_E = n_G = 30$, $Z = -2.544$, $p = 0.011$), because human subjects were more cooperative in the generous treatments (human cooperation rates were 67.7 % against generous strategies and only 34.2 % against extortioners, Mann-Whitney U-test, $n_E = n_G = 30$, $Z = -3.625$, $p < 0.001$).

Experimental Methods

The experiment was conducted in November and December 2013, at the University of Kiel and at the University of Hamburg, Germany. As participants, we have recruited 60 volunteers from first-year courses in Biology. These volunteers participated in ten groups of six subjects each in a computerized experiment. Before each experimental session, subjects were orally informed by one of the experimenters (M. M.) about how to operate the computers, and about the measures that were taken to ensure the subjects' anonymity. These measures included that subjects made their decisions under a neutral pseudonym, and that subjects were not allowed to talk to each other during or after the experiment. Moreover, they were informed that they would receive all their earnings in cash after the experiment. The payment procedure was organized in a way such that the anonymity of the participants was fully maintained.

After these oral instructions, participants were randomly assigned to a seat. The seats were separated by opaque partitions. Each seat came with a laptop computer, which informed participants about the rules of the game, and with which participants could communicate their decisions. The game instructions did not reveal the nature of the subjects' opponents: subjects were not told that they would play against each other, but they were also not told that they would play against a computer program (for a translation of the instructions, see Section 3). Participants were unaware of the exact length of the game; they were only informed that they would interact with their respective co-player over many rounds and that no time limit existed for their decisions to be done.

The computer opponents were randomly assigned to the six different seats. In order to avoid sequence effects, we ensured that in all experimental sessions there was at least one instance of each of the four computer strategies, ES, EM, GM, and GS. As human subjects played independently of each other, we considered each subject-computer interaction as the statistical unit of the experiment. For our analysis, we have used two-tailed tests throughout.

Instructions of the experiment

In the following, we provide the information displayed on the subjects' laptops throughout the experiment, translated from German. The instructions were the same for all four treatments.

Instructions in the beginning of the experiment

Page 1. Welcome to this experiment, in which you can earn money. At the beginning of this experiment, you will receive 10 Euros credited to your account. During the experiment you can earn more money. This may depend on your own decision and on the decision of others. Your decisions are anonymous. To ensure this, the computer assigns you a pseudonym that can be seen at the bottom left of your screen. The pseudo names are names of moons in our solar system (Ananke, Telesto, Despina, Japetus, Metis, and Kallisto). At the end of the game you will receive in cash the money in your account anonymously under your pseudo name. To render this experiment successful, it is strictly forbidden for participants to talk to each other or to communicate in any other way. After having read this text completely, please confirm by the pressing the OK-button.

Page 2. In the beginning of this experiment the computer will assign you a random co-player. Your co-player will remain the same for the whole experiment. The experiment consists of several rounds. In each round, you and your co-player face the same decision situation. An explanation of this decision situation follows on the next page. After having read this text completely, please confirm by pressing the OK-button.

Page 3. In each round, both players need to simultaneously choose a letter (either **C** or **D**). Each player needs to decide without knowing the choice of the co-player. Your payoff depends on your decision and on the decision of your co-player. The following table shows all possible payoffs. The first amount in each cell (with a blue font color) is your own payoff, and the second amount (with a green font color) is the payoff of your co-player:

		Decision of your co-player C	D
Your decision	C	€ 0.30, € 0.30	€ 0.00, € 0.50
	D	€ 0.50, € 0.00	€ 0.10, € 0.10

So there are four possible outcomes:

You: **C**	Your co-player: **C**	You get € 0.30	Your co-player gets € 0.30
You: **C**	Your co-player: **D**	You get € 0.00	Your co-player gets € 0.50
You: **D**	Your co-player: **C**	You get € 0.50	Your co-player gets € 0.00
You: **D**	Your co-player: **D**	You get € 0.10	Your co-player gets € 0.10

After having read this text completely, please confirm by pressing the OK-button.

Page 4. In each round you will have to answer the same question: "Which letter do you choose (either **C** or **D**)?" You decide by clicking on the corresponding button. Then you need to confirm your decision by clicking on the OK-button. You can only see the outcome of this round if all players have clicked on the OK-button. After having read this text completely, please confirm by pressing the OK-button.

Page 5. Examples: In each round you and your co-player are asked "Which letter do you choose (either **C** or **D**)?". After you have independently chosen a letter, the outcome of this round is displayed.

Example 1:

You: **C** Your co-player: **C** You get € 0.30 Your co-player gets € 0.30

Example 2:

You: **D** Your co-player: **C** You get € 0.50 Your co-player gets € 0.00

Example 3:

You: **D** Your co-player: **D** You get € 0.10 Your co-player gets € 0.10

Players need to confirm that they have read this page with the results by clicking on the OK-button. Then the round is over, and both players receive their payoffs credited to their account. The experiment consists of many rounds. In each round you face the same decision situation (but you can decide differently in each round). After having read this text completely, please confirm by pressing the OK-button.

Page 6. The experiment starts now! You have a credit of 10 Euros on your account. After having read this text completely, please confirm by pressing the green OK-button.

Instructions during the experiment

Page 1. Payoff table:

		Decision of your co-player	
		C	**D**
Your decision	**C**	€ 0.30, € 0.30	€ 0.00, € 0.50
	D	€ 0.50, € 0.00	€ 0.10, € 0.10

Which letter do you choose (either **C** or **D**)?

After having made your decision, please confirm by pressing the OK-button.

Page 2. Outcome of this round:

You: **X** Your co-player: **Y** You get € x Your co-player gets € y

After having read this text completely, please confirm by pressing the green OK-button.

Supplementary References

1. Claudia Keser and Frans van Winden. Conditional cooperation and voluntary contributions to public goods. *Scandinavian Journal of Economics*, 102:23–39, 2000.

2. Urs Fischbacher, Simon Gächter, and Ernst Fehr. Are people conditionally cooperative? evidence from a public goods experiment. *Economic Letters*, 71:397–404, Apr 2001.

3. Urs Fischbacher and Simon Gächter. Social preferences, beliefs, and the dynamics of free riding in public goods experiments. *American Economic Review*, 100(1):541–556, 2010.

4. W. H. Press and F. D. Dyson. Iterated prisoner's dilemma contains strategies that dominate any evolutionary opponent. *Proceedings of the National Academy of Sciences USA*, 109:10409–10413, 2012.

5. C. Hilbe, M. A. Nowak, and K. Sigmund. The evolution of extortion in iterated prisoner's dilemma games. *Proceedings of the National Academy of Sciences USA*, 110:6913–6918, 2013.

6. Ethan Akin. Stable cooperative solutions for the iterated prisoner's dilemma. *arXiv*, page 1211.0969v2, 2013.

7. A. J. Stewart and J. B. Plotkin. From extortion to generosity, evolution in the iterated prisoner's dilemma. *Proceedings of the National Academy of Sciences USA*, 110:15348–15353, 2013.

8. C. Adami and A. Hintze. Evolutionary instability of zero-determinant strategies demonstrates that winning is not everything. *Nature Communications*, 4:2193, 2013.

9. C. Hilbe, M. A. Nowak, and A. Traulsen. Adaptive dynamics of exortion and compliance. *PLoS One*, 8:e77886, 2013.

10. C. Hilbe, A. Traulsen, and K. Sigmund. Partners or rivals? Strategies for the iterated prisoner's dilemma. *Working Paper*, 2013.

11. A. J. Stewart and J. B. Plotkin. Extortion and cooperation in the prisoner's dilemma. *Proceedings of the National Academy of Sciences USA*, 109:10134–10135, 2012.