

How History and Convention Create Norms: An Experimental Study

Francesco Guala and Luigi Mittone

Final version published in *Journal of Economic Psychology* 31 (2010), pp. 349-56.

Abstract

According to a tradition that goes back to David Hume, social conventions have a natural tendency to turn into norms. Normativity increases compliance and stabilizes individual behaviour in spite of changes in incentives. In this paper we report experimental data that confirm this insight and encourage mildly optimistic conclusions regarding human sociality: habits provide extra glue that keeps individuals together, and prevents them from succumbing to anti-social temptation even when punishment is unlikely.

JEL Classification: H11, D02, C73

Key words: coordination games, experimental economics, conventions norms

Two men, who pull the oars of a boat, do it by an agreement or convention, tho' they have never given promises to each other. Nor is the rule concerning the stability of possession the less deriv'd from human conventions, that it arises gradually, and acquires force by a slow progression, and by our repeated experience of the inconveniences of transgressing it. On the contrary, this experience assures us still more, that the sense of interest has become common to all our fellows, and gives us a confidence of the future regularity of their conduct: And 'tis only on the expectation of this, that our moderation and abstinence are founded. (Hume 1740: Part II, Section II)

1. Introduction

David Hume is considered one of the great precursors of the game-theoretic analysis of social institutions. In a famous paragraph of the *Treatise of Human Nature*, he compares social conventions with the action of two rowers who must work in perfect synchrony to move their boat ahead. In modern terms, Hume models society as a sequence of coordination games, and

highlights the role of habits and customs to select among the many possible equilibria of the game of life.

Hume's insights were revived in the twentieth century, when rational choice theorists were finally able to model and analyze rigorously social institutions (e.g. Ullmann Margalit 1977, Schotter 1981, Sugden 1986, Binmore 1994, 1998, Skyrms 1996, 2004). Theorists in the Humean tradition have traditionally modelled *social conventions* as solutions to coordination games with multiple equilibria (Lewis 1969), and *social norms* as solutions to dilemmas of cooperation (Bicchieri 2006). The same theorists, however, have also acknowledged that the distinction between norms and conventions may be in practice less sharp than suggested by theory, for conventions have a tendency to acquire normative status, or to turn themselves into norms.

This story, as we shall see, sounds plausible enough from a psychological and evolutionary point of view. However, no attempt has been made so far to investigate empirically the normative status of social conventions. In this paper we study the emergence of an endogenous convention during repeated play of a symmetric coordination game. To investigate its normative power, after a few rounds we offer players the opportunity to earn extra money by breaching the convention. Many experimental subjects prefer not to take this opportunity and stick to the convention that has evolved during the early phase of the game. Since this pattern of behaviour can be explained in various ways, we introduce a control condition where subjects are faced with the same payoff structure but without being previously exposed to repeated coordination play. Self-interested choices are significantly more frequent in this condition, which indicates that conventions are endowed with normative power as suggested by theorists working in the Humean tradition.

We shall proceed as follows: in the next section, we set the stage by summarizing the theory of coordination games and of the emergence of social conventions. In section 3 we illustrate how habits and customs may help solve the problem of temptation, and we present experimental data that corroborate this idea. Section 4 concludes with general comments and directions for future research.

2. Conventions and norms

A simple symmetric coordination game is represented in Table 1. Each player has two possible strategies, Left or Right, and the numbers in each cell represent their payoffs (the first number for the row player, the second one for column). Standard game theory, to be sure, does not give particularly helpful advice in situations of this kind: a perfectly rational calculator cannot do better than flipping a coin and choosing a strategy at random. In a seminal study of coordination games Thomas Schelling (1960) however noticed that in many real-life interactions we are much more successful than purely rational calculators. He argued that seemingly irrelevant features of the environment, such as the position of the objects of choice or the way they are labelled, function as cues that help us converge on a common solution. A strategy that is made salient by such features is called a “focal point”, and in the course of repeated interaction is likely to become a point of attraction for the individuals in a given population (see also Sugden 1986; Bacharach and Bernasconi 1997).

[Table 1 about here]

Schelling’s hypothesis was mainly based on anecdotal examples and rudimentary experiments, but recent more systematic studies have confirmed its validity (Mehta et al. 1994). Among the “irrelevant” details that can make a strategy salient, Schelling argued, *history* plays an important role: “Precedence seems to exercise an influence that greatly exceeds its logical importance or legal force”, and “there is [...] a strong attraction to the *status quo ante*” (1960: 67-8). This idea was further articulated by David Lewis (1969) in his classic book on *Convention*. Lewis, who was mainly interested in the conventional nature of language, argued that a given strategy can become salient simply in virtue of the fact that it was played by a sufficiently large number of people in previous rounds of the game. When this happens, we shall say that a “convention” has emerged in a given population.¹

In one important respect the Hume-Schelling-Lewis theory sits firmly in the rational choice tradition. Our main motivation to follow a convention is selfish: we drive on the left because we want to avoid accidents; we say “cat” rather than “tac” because we want to be understood by our interlocutors; we wear black at funerals because we want to communicate our grief.

¹ Inductive learning in repeated coordination games has been studied theoretically and experimentally by e.g. Crawford and Haller (1990), Blume and Gneezy (2000), Haruvy and Stahl (2004).

Even robust conventions, however, face threats of various kinds.² There are always “trembling hand” accidents, to begin with, when individuals deviate by mistake from the established equilibrium. These deviations may have only minor effects in payoff terms, but can nevertheless generate confusion in the other players, insinuating doubts regarding the kind of game that is actually being played. And sometimes the payoff structure *does* change, because of some change in players’ preferences or in their material incentives. In all such cases rational self-interest alone may be too fragile to sustain social coordination in the long run.

How can social coordination survive in spite of such hindrances? It is likely that more than one mechanism is in place to sustain social conventions, and some plausible factors have already been studied experimentally. Chaudhuri et al. (2009) for example have discovered that *advice* from a previous generation of players enhances coordination in repeated games, over and above the coordination rates achieved by simple induction from the past. Another mechanism which has been widely investigated is *punishment* by offended individuals or bystanders. Experimenters have focused in particular on retaliation triggered by violations of norms of equality and reciprocity (e.g. Fehr and Gächter 2002, Boyd et al. 2003) but it is possible that punishment reduces deviance from social conventions too.

In this paper we explore yet another mechanism that has been occasionally discussed in the theoretical literature but, as far as we know, has never been investigated empirically before. The basic idea is that players may have *normative* reasons to conform to conventions. Hume’s followers, in particular, have argued that conventions are a species of norms, or that they tend to *turn* into social norms (Lewis 1969: 97; Ullmann Margalit 1977: 88-89; Gilbert 1989: 349-355; Young 1998: 144-145). As Robert Sugden has put it,

We expect that our dealing with other people will be regulated by convention, but this expectation is more than a judgement of fact: we feel entitled to expect others to follow conventions when they deal with us, and we recognize that they are entitled to expect the same of us (Sugden 1986: 154).

Normative expectations have many virtues. Absent-minded players will be more careful not to breach a convention, if they know that such a breach will be treated as an offense by the

² See e.g. Hume (1740, Part II, 7; 1777: Part I, 4), Sugden (1986: 151-154).

other members of their group. Moreover, a normative convention may become so strong to exert its influence beyond the narrow boundaries of the original coordination game that gave rise to it. That you *ought* to act in a certain way means, typically, that we expect you to do so in spite of incentives to take a different course of action. So it is possible that normative expectations enforce conformity when self-interest alone would encourage deviation from a given rule.

Let us define a *norm* as a rule of conduct (or strategy) S in a game G such that: (a) all players prefer to play S provided everybody else does the same; (b) this preference is common knowledge among players; (c) the players who do not play S typically incur some costs that reduce their utility payoffs (Bicchieri 2006). The latter condition (c) is particularly important: in laboratory experiments we can detect the effect of norms on individual behaviour by observing whether the subjects are willing to forego individual earnings in order to conform to a given rule of conduct. When a norm is in place, each player expects that the others will bear the cost of conformity, because she is planning her choices based on the assumption that they will conform. The existence of a norm is manifested in the decision to “leave some money on the table” and privilege conformity with the rule with respect to one’s own private gain.³

This hypothesis can be easily backed up by speculative evolutionary stories. As social animals, human beings have likely acquired through evolution a “basic desire” to keep the good will of the members of their group (Sugden 1986: 156). This basic desire makes us feel uneasy when we become the object of resentment, as for instance when we breach a convention. Add to this fact that we resent it when other people frustrate our expectations, and we have the beginning of an explanation of the normativity of convention.

³ There is abundant evidence of this type of behaviour in the experimental literature. The most thoroughly investigated case is people’s willingness to pay a material cost to punish behaviour that violates norms of fairness (e.g. Fehr and Gächter 2002). In our experiments we do not provide the other players with a chance to punish deviants because we are interested in studying precisely those free-riding opportunities that are unlikely to be punished (either by a central authority or by fellow citizens). Of course we do not claim that all free-riding opportunities are of this kind. But they certainly exist, and they provide the strongest possible case for the creation of a centralized policing authority. Our point is that in spite of the absence of material punishment, we observe individuals forgoing material gains (paying opportunity costs) in order to abide to conventions and norms.

3. An experiment on conventions and norms

We study the emergence of an endogenous convention in a three-player repeated coordination game (the “Baseline” condition) similar to the one in Table 1. The choice of symmetric payoffs is not meant to imply that this situation is true of social coordination *in general*. Rather, we use this setting for its simplicity, and to avoid contamination with other considerations that may arise from the emergence of inequitable conventions.⁴ Because we are interested here in a specific research question (“Are conventions norms?”), we aim at isolating the “pure” effect of precedence by eliminating other confounding factors.

To test the hypothesis that conventions are endowed with normative power, we provide one player with an incentive to deviate from the convention that has emerged during repeated play of the coordination game (section 3.1). To eliminate the “shadow of the future”, this incentive is offered only in the very last round of the game, and in a somewhat surprising manner. As we shall see, several players decide to ignore the incentive and comply with the convention that has evolved in the early rounds of the game. Since this behaviour can be explained in more than one way, in section 3.2 we introduce an experimental design (the “One-shot” condition) aimed at disentangling the factors that inhibit deviance from the convention. The experimental data are discussed in section 3.3.

3.1. Baseline condition

The main task in the Baseline condition is a repeated coordination game with two options labelled “Red” and “Blue”. At each round each subject chooses an option by clicking a button on the screen and then receives feedback regarding the choices of the other two group members. If all members choose the same option, they earn 10 tokens each; otherwise, they receive nothing. The payoffs of this stage-game are represented in Table 2; the first number represents the payoff of Player 1, the second number of Player 2, and the third one of Player 3. When this game is implemented in the laboratory for real money, most groups converge on a common strategy after a couple of rounds, and continue to coordinate for the rest of the game.

⁴ Previous experiments suggest that the salience of an equilibrium may be partly offset by payoff asymmetries (see Schotter and Sopher 2003; Crawford et al. 2008).

[Table 2 about here]

Subjects are told in the instructions that the game will last for ten rounds only. They are also told that the payoffs might change during the course of the game, but that all players will be informed in advance when this will happen. In the instructions no specific details are provided regarding the payoff structure of these “special rounds”. As a matter of fact, in the tenth and last round all groups face a “temptation” game with the payoff structure represented in Table 3. Within each group, one player (the “potential deviant”) is given an incentive to deviate from the convention that has emerged during earlier rounds. In Table 3, by stipulation, the potential deviant is Player 3. Each row represents a possible strategy for the potential deviant, with respective payoffs.

[Table 3 about here]

Notice that if all players were allowed to choose simultaneously (as in the earlier coordination rounds) the temptation stage-game would have no pure-strategy equilibrium. Suppose the convention is Red: Players 1 and 2 could try to anticipate Player 3’s deviation and abandon the convention to try to converge with Player 3 on the other colour (Blue). Player 3 should take this into account, and keep playing Red to outwit them. But Players 1 and 2 should anticipate this too ... and so on indefinitely. Whatever behaviour we observed, we would be unable to infer Player 3’s motivations from the data.

To deal with this problem, in our experiment we implement a *sequential* version of the temptation stage-game. In this version, Player 1 and Player 2 move first, followed by Player 3. Player 3 thus can see what her two partners have chosen, and the problem of strategic reasoning is solved. This solution, however, raises another subtle concern: if Players 1 and 2 were given the option of sequentially choosing their preferred colour, Player 2 would have opportunity to breach the convention (which would result in an equal payoff of zero tokens for all players), and make Player 3’s choice irrelevant. This would indeed be a (weakly) dominant strategy for a pessimistic Player 2 who anticipates Player 3’s defection and believes that she has nothing to lose by deviating first. By *not* derailing the game (that is, by successfully coordinating on either Red or Blue), the first two players thus send an implicit message to the potential deviant that may be read as an offer of cooperation. Since there is extensive evidence from the experimental literature that “nice” moves of this kind can trigger

norms of reciprocity (see e.g. Charness and Rabin 2002, Falk et al. 2003, McCabe et al. 2003, Cox 2004), it is possible that Player 3 would respond “nice” with “nice” to comply with the norm of reciprocity, rather than with the convention per se.

To control for this confounding factor we eliminate the opportunity for Players 1 and 2 to breach the convention: instead of making their choices autonomously, a computer programme assigns the same colour (i.e. the convention) that they have played in the previous round of the game (round 9). At this point, Player 3 is offered the opportunity to breach the convention. Round 10 then takes the form of a dictator game for Player 3, with two options labelled in the same way as the options of earlier coordination rounds. In Round 10 the third player cannot read the moves of Players 1 and 2 as “nice” intentional offers, and the confounding effect of reciprocity is eliminated.⁵

3.2. *One-shot condition*

Let us suppose that the potential deviant chooses the same colour as her partners in Round 10 of the Baseline condition. The prospect of breaching a convention generates a psychological cost, which the potential deviant is not willing to pay. But what exactly is the source of this cost? Intuitively, three sets of considerations may prompt subjects to seek the maximization of group rather than individual earnings: (1) *altruistic* considerations prescribing to maximize the earnings of the other players; (2) *inequality aversion* prescribing to avoid inequitable outcomes; and/or (3) *utilitarian* considerations prescribing to maximize the total payoff of the group. In either case, these considerations would encourage conformist behaviour with respect to the established convention: altruist, egalitarian, and utilitarian players should all choose Red-Red or Blue-Blue in the temptation round of the Baseline condition.

To disentangle the normative power of conventions from these “extrinsic” normative considerations, we compare behaviour in Round 10 of the Baseline condition with behaviour in a structurally similar one-shot game, where subjects face a payoff structure identical to the

⁵ We experimented with various alternative designs before settling on this solution. All these experiments, which we now treat as pilots, generated similar results to those reported here. See Guala (2007) and Guala and Mittone (2008) (Experiment 2 in particular).

one in Table 3, but *without being exposed to a history of successful coordination in previous rounds of the game*.⁶

In the “One-shot” condition subjects again play in groups of three, anonymously and without communication. The payoff structure is transparent for all players, and the game is played in two stages as follows: in the *first stage*, the computer assigns a colour (Red or Blue) to the first two players. In the *second stage* of the game, the third player (the potential deviant) faces a choice between: (a) playing the same colour assigned to the other players in the previous round, or (b) playing a different colour.⁷ The monetary payoffs are identical to those of Table 3.

The One-shot condition and the final round of the Baseline condition thus are perfectly isomorphic from a game-theoretic point of view. In either condition the potential deviant takes a colour as given, and faces a straightforward dilemma between an individualistic and a cooperative choice. The key point for our purposes is that altruistic, utilitarian, and inequality-aversion considerations apply exactly in the One-shot game as in the Baseline condition. The only difference is that in the Baseline the potential deviant may be influenced by *precedence* – a history of repeated convergence on a common strategy – whereas in the One-shot condition this element has been removed by design. If individualistic behaviour was more common in the One-shot than in the Baseline condition, then, we would obtain direct evidence that repeated group play generates an intrinsic normative pressure to conform to the group’s conventions.

3.3. Results

We ran this study at the Computable and Experimental Economics laboratory of the University of Trento. Subjects were recruited using email lists from the population of graduate and undergraduate students. Volunteers registered for one of the experimental

⁶ To control for income effects, the One-shot condition was preceded by an incentivized questionnaire where subjects were given the chance of earning an amount of tokens roughly equivalent to that earned by the participants in the Baseline condition during the first nine (“normal”) rounds of that game. The same questionnaire (aimed at facilitating understanding of the game) was also administered before the Baseline condition, but without monetary incentives.

⁷ The computer device is introduced to avoid reciprocity effects and to achieve perfect symmetry with the Baseline condition.

sessions and, as they arrived at the lab, were seated randomly at 18 computer terminals separated by partitions. After signing a consent form, they were asked to read the experimental instructions illustrating the main features of the task. The task varied depending on the experimental condition. Each subject participated in one condition only, and all comparisons took place across subjects. In both conditions subjects played in groups of three players, with random selection of group membership, anonymity, and without communication. They received a show-up fee of 3 Euro, and on top of that received whatever they earned in the experimental task. Individual payoffs were calculated in terms of “experimental tokens” which were converted into real money at the exchange rate of 3 cents per token (or 1 Euro = 33 tokens).⁸ In each condition we obtained one observation regarding potential deviants per each group of three players. In total 174 subjects participated (90 in the Baseline and 84 in the One-shot condition) resulting in 28 usable data-points per treatment.⁹

[Table 4 about here]

Data from the Baseline condition (summarized in Table 4) give us a few interesting insights into the mechanics of coordination. The dominant convention is Red, played by 78% of the groups. The explanation may be cultural (we read and write from left to right, and the Red button was placed on the left-hand side of the screen) or physiological: recent neuroeconomic studies suggest that when faced with a choice between two indifferent options, we tend to choose the one upon which our sight fixated first, and it’s twice as likely that we fixate on the left-hand side than on the right-hand side object (Armel et al. 2008).

Deviation in the tenth (temptation) round of the game is uncorrelated with the type of convention that has evolved: although Red-type conventions are breached slightly less frequently (36%) than Blue ones (50%), the difference is not statistically significant ($\chi^2 = .368$, $df = 1$, $p = .54$). Coordination is reached on average in two rounds, but there are a surprising number of “trembling hand” deviations: in 40% of the groups at least one subject experiments with breaching the convention.

⁸ Average earnings ranged between nine and eleven Euro, depending on the condition, for about half an hour of experimentation.

⁹ The data of two subjects in the Baseline condition were not considered for the statistical analysis of data because by round nine their respective groups had not achieved a proper convention of coordination.

Only in one case the members of a group failed to coordinate in the ninth round, and this data-point was excluded for the purpose of data-analysis. In another case, a group managed to coordinate on a bizarre conventional rule, prescribing to play Blue and Red in alternating rounds. This strange convention prevents a straightforward interpretation of behaviour in the tenth round, so we also excluded this group at the stage of data-analysis.

Intuitively, the history of coordination may influence the decision to conform in the tenth round. A subject who has experienced prolonged coordination with the members of his group, for example, may feel a stronger “bond” than a subject who only managed to coordinate two or three times before the tenth period. We include in Table 4 a variable that may be used to test this hypothesis: the number of consecutive rounds of successful coordination that took place immediately before the tenth round. Surprisingly, the decision to deviate in the tenth round is not influenced by coordination history ($t = 0.206$, $df = 26$, $p = 0.84$).

We collected data on subjects’ gender, age, and university degree. There is a slight prevalence of conformity in male subjects (69%) compared to females (53%), which is however non-significant ($\chi^2 = .738$, $df = 1$, $p = .39$). The variable “Degree” has no effect: economists do not behave differently from students enrolled in other programmes ($\chi^2 = .007$, $df = 1$, $p = .93$). Younger students seem to deviate more frequently from the convention than older ones, but again the difference is not statistically significant ($t = 1.289$, $df = 26$, $p = 0.20$).

[Table 5 about here]

Table 5 summarizes the data collected in the One-shot condition. We find again that on average conformists tend to be slightly older than deviants (23.55 vs. 21.85 years-old, respectively) but the effect falls just short of statistical significance ($t = 2.02$, $df = 26$, $p = 0.54$). Men and subjects studying economics conform slightly more frequently, but again there is no statistically significant effect in these variables (Gender: $\chi^2 = .197$, $df = 1$, $p = .65$; Degree: $\chi^2 = .1365$, $df = 1$, $p = .24$).

The most important comparison however concerns the decision made in the tenth round of the Baseline condition and in the One-shot condition. The deviation data to test our hypotheses can be found under the N columns in Tables 4 and 5. Comparison between the One-shot and Baseline conditions indicates that mere repetition of a collective task does enhance the conformity to a convention. The null hypothesis (H_0 : “There is no effect from repeated group play”) is rejected in a two-tailed chi-square test ($\chi^2 = 4.595$, $df = 1$, $p = .03$) and in a directional small numerosity test (p -Yates = .03).

To corroborate these results we present the outcome of two logistic regressions run by controlling for heteroskedasticity (we computed robust standard errors). The outcomes are reported in Table 6. In both regressions we consider as dependent variable a dichotomous variable assuming value 1 when a deviation occurs and value 0 when it does not. The first regression (Model 1) uses as explicatory variable only a dummy variable (ONE-SHOT) taking value 1 for the One-shot condition and value 0 for the Baseline condition. The second regression (Model 2) controls the effect produced by ONE-SHOT introducing three additional control variables. The first variable (FEMALE) is a dummy assuming value 0 if the participant is male and 1 if she is female. Another dichotomous variable (ECON) controls for the effect of education: the value of the dummy is 1 if the participant is studying economics, 0 in all the other cases. The third explanatory variable (AGE) is the participants’ age.

[Table 6 about here]

The outcome of the first regression confirms the main result. The type of treatment does significantly influence the degree of conformity. The odds of observing a deviation in the One-shot condition are about 3 times higher ($Exp(B) = 3.263$, $p = .036$) than the estimated odds in the Baseline condition. The odds ratio can be interpreted as the probability to observe a deviation divided by the probability to observe conformity. This effect is even stronger in model 2 where the odds of observing a deviation in the One-shot condition are about 7 times higher ($Exp(B) = 7.170$, $p = .006$) than in the Baseline condition. This is a further confirmation of the robustness of our results.

Similarly the second regression confirms the results of the chi-square and t-tests. None of the other explanatory variables included in Model 2, with the exception of AGE, has a significant effect on the decision to conform. AGE displays a significant negative correlation with the probability to deviate (there is a greater propensity for older participants to conform to the convention). It is worth noticing that the behavior of the older subjects reinforces the main result because the average age in the Baseline (21.64) is lower than in the One-shot condition (22.39).

The main result of the experiment thus is highly corroborated. The data support the hypothesis that repeated coordination play contributes to the emergence of a convention that is normatively binding for the members of the group. Repeated coordination play generates pressure towards conformity, and as a consequence conventions tend to turn into social norms. For several subjects violating norms is costly enough to offset the potential gains obtained by deviating. One reason why conventions are resilient is that, for many of us, they just become norms.

4. Conclusion

Coordination games are a ubiquitous form of social interaction, and understanding their functioning should be a priority in the foundations of social science. Several authors have suggested that social norms may emerge from repeated play of coordination games,¹⁰ but this insight until now was not supported by experimental data. Our Baseline design tries to capture this dynamics and can be used to study the way in which conventions acquire normative power during repeated team play. This design combines two different stage-games (a coordination and a “social dilemma” game) and thus provides insights into the way an equilibrium strategy that emerged in one context can influence behaviour in situations of a different kind. A strategy developed in one setting may have a “carry-over” effect and act as a focal point in a stage-game with a radically different payoff structure.

The evidence confirms the hypothesis that social conventions play an important role in sustaining cooperation even in dilemma situations, where purely self-interested motives encourage deviation from the social contract. The data indeed indicate that conventions have a tendency to *become* social norms, acquiring an intrinsic normative power that reduces

¹⁰ See Sugden (1986), Binmore (1998), Skyrms (2004).

deviation rates from an established convention. At the practical level, our results suggest that conventional modes of behaviour such as traditions, habits, and routines may be more difficult to disrupt by changing individual incentives than one would assume based on standard rational choice analysis.

References

- Armel, K.C., Beaumel, A. and Rangel, A. 2008. "Biasing Simple Choices by Manipulating Relative Visual Attention", *Judgment and Decision Making* 3: 396–403.
- Bacharach, M. and Bernasconi, M. 1997. "The Variable Frame Theory of Focal Points: an Experimental Study", *Games and Economic Behavior* 19:1-45
- Bicchieri, C. 2006. *The Grammar of Society*. New York: Cambridge University Press..
- Binmore, K. 1994. *Game Theory and the Social Contract, Vol. 1: Playing Fair*. Cambridge, Mass.: MIT Press.
- Binmore, K. 1998. *Game Theory and the Social Contract, Vol. 2: Just Playing*. Cambridge, Mass.: MIT Press.
- Blume, A. and Gneezy, U. 2000. "An Experimental Investigation of Optimal Learning in Coordination Games", *Journal of Economic Theory* 90: 161-72.
- Boyd, R., Gintis, H., Bowles, S. and Richerson, P.J. 2003. "The Evolution of Altruistic Punishment", *Proceedings of the National Academy of Sciences* 100: 3531-3535.
- Charness, G. and Rabin, M. 2002. "Understanding Social Preferences with Simple Tests", *Quarterly Journal of Economics* 117: 817-69.
- Chaudhuri, A., Schotter, A. and Sopher, B. 2009. "Taking Ourselves to Efficiency: Coordination in Intergenerational Minimal Effort Games with Private, Almost Common and Common Knowledge of Advice", *Economic Journal* 119: 91-122.
- Cox, J. 2004. "How to Identify Trust and Reciprocity", *Games and Economic Behavior* 46: 260-81.
- Crawford, V.P., Gneezy, U. and Rottenstreich, Y. 2008. "The Power of Focal Points Is Limited: Even Minute Payoff Asymmetry May Yield Large Coordination Failures", *American Economic Review* 98: 1443-58.
- Crawford, V.P. and Haller, H. 1990. "Learning How to Cooperate: Optimal Play in Repeated Coordination Games", *Econometrica* 58: 571-95.
- Falk, A., Fehr, E. and Fischbacher, U. 2003. "On the Nature of Fair Behavior", *Economic Inquiry* 41: 20-26.
- Fehr, E. and Gächter, S. 2002. "Altruistic Punishment in Humans", *Nature* 415: 137-140.

- Guala, F. 2007. "An Experimental Study of Conventions and Norms", *ESRC Working Paper*, <http://www.esrcsocietytoday.ac.uk/>
- Guala, F. and Mittone, L. 2008. "An Experimental Study of Conventions and Norms", *CEEL Working Paper 10-08*, University of Trento, http://www-ceel.economia.unitn.it/papers/papero08_10.zip
- Haruvy, E. and Stahl, D.O. 2004. "Deductive versus Inductive Equilibrium Selection: Experimental Results", *Journal of Economic Behavior and Organization* 53: 319-31.
- Hume, D. 1740. *A Treatise of Human Nature*. London: Penguin, 2003.
- Hume, D. 1777. *Essays: Moral, Political, and Literary*. Oxford: Oxford University Press, 1963.
- Lewis, D.K. 1969. *Convention: A Philosophical Study*. Cambridge, Mass.: Harvard University Press.
- McCabe, K., Rigdon, M. and Smith, V. 2003. "Positive Reciprocity and Intentions in Trust Games", *Journal of Economic Behavior and Organization* 52: 267-275.
- Mehta, J., Starmer, C. and Sugden, R. 1994. "The Nature of Salience: An Experimental Investigation of Pure Coordination Games", *American Economic Review* 84: 658-673.
- Schelling, T. 1960. *The Strategy of Conflict*. Cambridge, Mass.: Harvard University Press.
- Schotter, A. 1981. *The Economic Theory of Social Institutions*. Cambridge University Press.
- Schotter, A. and Sopher, B. 2003. "Social Learning and Coordination Conventions in Intergenerational Games: An Experimental Study", *Journal of Political Economy* 111: 498-529.
- Skyrms, B. 1996. *Evolution of the Social Contract*. Cambridge University Press.
- Skyrms, B. 2004. *The Stag Hunt and the Evolution of Social Structure*. Cambridge University Press.
- Sugden, R. 1986. *The Economics of Rights, Cooperation and Welfare*. Oxford: Blackwell.
- Ullmann-Margalit, E. 1977. *The Emergence of Norms*. Clarendon Press.

| | | |
|-------|------|-------|
| | Left | Right |
| Left | 1, 1 | 0, 0 |
| Right | 0, 0 | 1, 1 |

Table 1: A simple coordination game.

| Outcomes | Earnings |
|---|------------|
| All players choose the same colour | 10, 10, 10 |
| One player chooses a colour different from the others | 0, 0, 0 |

Table 2: Coordination stage-game.

| Outcomes | Earnings |
|---|---------------|
| Player 3 chooses the same colour as the other players | 200, 200, 200 |
| Player 3 chooses a different colour | 0, 0, 300 |

Table 3: Temptation stage-game in the 10th round.

| Round 10 Decision | N | Convention | | Coordination History (avg) | Age (avg) | Gender | | Degree | |
|----------------------|--------------|--------------|-------------|----------------------------------|--------------|--------------|--------------|--------------|--------------|
| | | Red | Blue | | | Male | Female | Econ | Other |
| Conform | 17 (.61) | 14 (.64) | 3 (.50) | 6.5 | 22.12 | 9 (.69) | 8 (.53) | 9 (.60) | 8 (.62) |
| Deviate | 11 (.39) | 8 (.36) | 3 (.50) | 6.27 | 20.91 | 4 (.31) | 7 (.47) | 6 (.40) | 5 (.38) |
| Tot. | 28 (1.00) | 22 (1.00) | 6 (1.00) | 6.4 | 21.64 | 13 (1.00) | 15 (1.00) | 15 (1.00) | 13 (1.00) |

Table 4: Baseline condition results. All percentages (between brackets) add up vertically.

| Decision | N | Age (avg) | Gender | | Degree | |
|----------|--------------|--------------|--------------|--------------|--------------|-------------|
| | | | Male | Female | Econ | Other |
| Conform | 9 (.32) | 23.55 | 6 (.35) | 3 (.27) | 8 (.38) | 1 (.14) |
| Deviate | 19 (.68) | 21.84 | 11 (.65) | 8 (.73) | 13 (.62) | 6 (.86) |
| Tot. | 28 (1.00) | 22.39 | 17 (1.00) | 12 (1.00) | 21 (1.00) | 7 (1.00) |

Table 5: One-shot condition results. All percentages (between brackets) add up vertically.

| | Model 1 | | | Model 2 | | |
|----------|------------|------------------|---------|------------|------------------|---------|
| | Odds ratio | Robust Std. Err. | p-value | Odds ratio | Robust Std. Err. | p-value |
| ONE-SHOT | 3.263 | .565 | .036 | 7.170 | 4.792 | .006 |
| FEMALE | - | - | - | 1.411 | .879 | .574 |
| AGE | - | - | - | .743 | .102 | .037 |
| ECON | - | - | - | .673 | .412 | .532 |

Table 6: Logit regression, Model 1 and Model 2