A mechanism for optimal enforcement of coordination: Sidestepping theory of mind

Alexander Funcke (funcke@sas.upenn.edu)[1] and
Daniel Cownden (dc89@st-andrews.ac.uk)[2]

[1]Philosophy, Politics & Economics, University of Pennsylvania
[2]School of Biology, University of St Andrews

February 13, 2015

Abstract

Mechanisms employing fines and rewards may be introduced in multi-equilibrium situations to enforce a certain equilibrium. The mechanism does two things. First, it produces a signal disrupting the normal dynamics of repeated play; potentially encouraging agents to reconsider their expectations. Secondly, it changes the payoffs. In deciding what behavior to engage in after introduction of a new mechanism a rational agent need to consider to what degree the signal has been received and convinced others, which in turn depends on what those others believe about the reception of the signal, and so on. The epistemic mess is a challenge for both the agents and for a policy maker interested in facilitating a re-coordination. The latter begs the question: How large must fines and rewards be to ensure re-coordination? We show that the result we get from classic game theory is more "heavy handed" than necessary; Far less intervention is actually required. Specifically we will outline a mechanism that ensures re-coordination, regardless of the idiosyncratic belief formation processes of the population, while at the same time making minimal interventions.

Keywords: coordination, fine and rewards, mechanism design

JEL Classifications: C79, D80, Z13

1 Introduction

Social life is rife with strategic situations that can be characterized as coordination problems; mere conventions, such as whether to drive on the left or right side of the road, how to greet one another or in determining an appropriate tip. Coordination problems are not restricted to the trivial. Some of humanity's greatest challenges have aspects that can be viewed as coordination problems. In the fight against corruption (Barr and Serra, 2010; Dong et al., 2012), upholding sanitation (Curtis et al., 2011), stifling environmental degradation (Shove, 2010), curbing violence (Fabiano et al., 2003) and facilitating growth (Whiteley, 2000). In game theory, coordination problems are characterized as games with multiple Nash equilibria. I.e. strategic situations where there are more than one way for everyone to choose a behavior such that no agent has an incentive to single-handedly deviate. In multiple equilibria situations, an agents' preferred action is conditional on the actions of the others, hence the need to coordinate. In these cases, classic game theory can not make precise predictions about strategy selection as a rational agent's choice is conditional on her expectation on the other agents' choices, and the other agents' choices on the yet other agents' choices, and so on. This infinite regress of reasoning gives no clear behavioral prediction from our usual rationality assumptions. Empirically, outcomes in coordination games tend not to be toss-ups, rather, some classes of equilibria are selected more often than other (Sugden, 1989). There is thus an ongoing search for a predictive theory of equilibrium selection in coordinative situations.

Within the framework of rationality, various refinements of the Nash equilibrium have been proposed as solution concepts singling out one of a game's multiple equilibria as unique. E.g. the payoff dominant equilibrium (Harsanyi and Selten, 1988), the risk-dominant equilibrium (Kandori et al., 1993a), the sub-game perfect equilibrium (Selten, 1965), or the Bayesian perfect equilibrium (Harsanyi, 1967). Each of the refinements is predictive in some situations, but no refinement is universally applicable (Fudenberg et al., 1988). Bounded rationality has also been used to design equilibrium concepts. The type of infinite depth reasoning we saw earlier is of course nothing like human reasoning (Nagel, 1995), the mismatch has inspired an alternative method, known as cognitive hierarchy or level-k reasoning, where agents only reason k levels into the regress, where k typically is between 0 and 3. While this method does single out a specific equilibrium as salient, its predictions are far from consistently accurate (Hedden and Zhang, 2002; Kawagoe and Takizawa, 2009). Evolutionary methods, where agents change their strategies based on learning rules, or where a population of strategies reproduces themselves, also provide alternatives to strictly rational behavior, see Hofbauer and Sigmund (2003) for an excellent review. Inspired by the empirical findings of behavioral economics, yet another approach has been to formalize salience as an important aspect of equilibrium selection (Binmore and Samuelson, 2006; Sugden, 2011).

Despite the ubiquity of coordination games, and the many efforts to understand them, the dynamics of coordination remain an enigma. Empirically, in the case of repeated interactions, history is key (Schmidt et al., 2003). The importance of history should be expected as the past choices of others is the only direct clue individuals have about how the others will act in the future. Coordinative situations tend, perhaps, for this reason, to display strong path-dependency. This dependency can keep a population coordinated on an inferior equilibrium for historical reasons despite the presence of better options.

One of the primary motivation for studying coordination games is that they form the backbone of the analysis of social norms. (Bicchieri, 2006). When policymakers try to alter social norms, i.e. to re-coordinate a population's empirical expectations about how the others will act, they will typically impose sanctions via laws. The theoretical law literature suggests that this might be efficient, as social norms and law may work as complementaries in regulating behavior (Funk, 2005; Zasu, 2007). It remains to determine the right level of sanctions to achieve coordination. If we take a simple incentives-based approach to crime, the conclusion is reached that the optimal fine is the maximal fine (Becker, 1968). This first simplistic conclusion has been severely critiqued (Garoupa, 1997). E.g., sanctions may be socially costly (Kaplow, 1990), especially if the extent of harm caused by the sanction is unknown (Polinsky and Shavell, 1992) or if the sanctioned agent has limited wealth (Polinsky and Shavell, 1991).

In this paper, we lay out a set of criteria for an optimal enforcement mechanism and show that such an optimal enforcement mechanism does exist. "Optimal" will be taken to mean the least intrusive mechanism that still suffices to re-coordinate all rational agents.

For clarity and simplicity, we will take as our setting the repeated play of a 2-player stag hunt where players are randomly drawn from a finite population. In the stag hunt game, there are three equilibria of interest: a payoff dominant equilibrium, a risk-dominant equilibrium, and a mixed-Nash equilibrium. The first is an equilibrium that is Pareto superior to all others, and the second is the equilibrium with the biggest basin of attraction under most equilibrium selection methods, such as replicator dynamics (Harsanyi and Selten, 1988). Coordination failure is here taken to mean the failure to coordinate on the payoff dominant equilibrium. Unfortunately, empirical work tells us that avoiding this kind of failure is "an extremely unlikely outcome either initially or in repeated play" (Van Huyck et al., 1990). Additionally, the theoretical results Young (1993); Kandori et al. (1993b) concur with this empirical finding, showing that for certain stochastic strategy dynamics the limit is a unique fixed point that corresponds identically to the risk-dominant equilibrium. The third equilibrium is a mixed strategy equilibrium. It can be interpreted as a situation where the population is divided into their choice of strategies in such a way that the expected payoff of each strategy is the same. This third equilibrium is unstable under essentially any learning or evolutionary dynamic that might be used to model the population. Never the less, this equilibrium point is important as it corresponds to the critical threshold or tipping point at which all rational payoff maximizers will change their behavior from one strategy to the other.

The introduction of a system of fines and rewards to this game serves two distinct but related purposes. The first, and most straightforward, is to alter the payoff structure of the game. It is important from an equilibrium selection standpoint as it alters the value of the mixed strategy equilibrium mentioned above. That is; it changes the proportion of the population a rational payoff maximizer needs to believe will engage in a given strategy before that rational agent chooses that same strategy. The second purpose is to serve as a re-coordination signal. Empirically, in small group situations the *magnitude* of fines and rewards are often relatively unimportant in re-coordinating behavior (Brandts and Cooper, 2006). Similarly, in small groups, pre-game communication (Blume and Ortmann, 2007) can suffice to re-coordinate. The standard interpretation of these findings is that all that is required for re-coordination is a signal, and a system of fines and rewards can readily serve as this signal. Specifically, a signal causes a break in the normal dynamics of repeated play, by encouraging agents to reconsider their expectations of what other players might do in the future, in light of this shared signal.[1]

How an agent responds to the introduction of a new institution thus depends on the interaction between both signaling and the shifting of the tipping point. In deciding which behavior to engage in after the introduction of a new institution an agent must consider to what extent the introduction induced signal has been received throughout the population. Further, she need to consider how other members of the population will respond to the signal if it reaches them, which in turn depends on what she believe about what the others' beliefs about the signal reception and impact, and so on. This process of belief formation is epistemically challenging for both the decision makers within the population and for a policymaker wishing to facilitate re-coordination.

The mechanism begs the question, how large must the fines and rewards be to ensure re-coordination on a desired equilibrium. Because the answer to this question hinges on the essentially unknowable, and quite possibly irrational, belief formation processes of the population, it might seem at first glance that the best one could hope for is to introduce fines and rewards that are so substantial that they alter the strategic structure of the game entirely; completely removing the undesired equilibrium. Here we will show that such an approach is far more "heavy handed" than necessary. Far less intervention is required. Specifically we will outline a mechanism that ensures re-coordination, regardless of the idiosyncratic belief formation processes of the population, while at the same time making minimal interventions.

Our mechanism arises from three observations:

1. An agent's subjective expected utility of a given strategy will depend on their idiosyncratic beliefs concerning the ratio of other agents choosing it.

2. The act of enforcement requires inspection, from which a policymaker can derive an unbiased estimate of the frequencies of each strategy in play.

3. The fines and rewards imposed upon agents can be made dynamic and in particular contingent on the observed frequencies of behaviors.

[1] These results do provide effective mechanisms for re-coordination on a small interpersonal scale, but as groups become larger the efficacy of these mechanisms drops off. Suggesting that signaling alone will not reliably affect re-coordination on a societal level.

	C	B
C	1	0
B	T	P

Figure 1: All stag hunt games can be normalized to this form, with P and T satisfying the following conditions, $0 < P \leq T < 1$ and $1 < T + P$.

Taken together, this means that the rewards and fines can be tuned such that the expected utility of an agent's strategic choices (the game payoff plus the expected fine/reward) follow the policymaker's preference ordering, regardless of agents idiosyncratic belief formation processes. Reducing each agent's epistemically challenging problem to a trivial decision.

The remainder of the paper will formalize the re-coordination problem and our proposed solution to this problem. In the conclusion, we will briefly discuss the general implications of this work.

2 Formalizing the Recoordination Problem

As mentioned earlier, for simplicity and clarity we formalize the problem of re-coordination in terms of repeated play of a stag hunt game. In particular the stag hunt, like the weakest link game is a worst case for efficient coordination, as the payoff dominant equilibrium is different from the risk-dominant equilibrium.

Pairs of agents are repeatedly drawn from a population of economic agents. These agents play a symmetric stag hunt game. The payoffs for such a game are shown in figure 1. The two strategies of the game are C, for "[C]omply", and B, for "[B]reach". The strategies are in relation to the policymaker's desire for the payoff dominant equilibrium to be selected over the risk-dominant equilibrium.

Each agent holds an idiosyncratic belief about the ratio of C-choosers in the rest of the population. We use x to denote the actual ratio of C-choosers in the population, and $x_i := E_i[x]$ to denote agent i's belief about this ratio. We make no assumption about how agents form their beliefs, or how these beliefs relate to the actual ratio. The agents are assumed to be indistinguishable to each other for the purposes of this game, thus an agent i's best estimate of the probability that her randomly drawn opponent chooses C is x_i.

Agents are assumed to be expected payoff maximizers. The decision problem for each agent is to choose the strategy that maximizes expected payoff. The expected payoff for agent i choosing strategy C is

$$E_i[U_G(C)] = E_i[x] = x_i \tag{1}$$

and for the same agent choosing strategy B,

$$\begin{aligned} E_i[u_G(B)] &= E_i[x]T + E_i[1-x]P \\ &= Tx_i + P(1-x_i). \end{aligned} \tag{2}$$

Accordingly, agent i will choose C if and only if

$$x_i - Tx_i - P(1 - x_i) \geq \Delta, \tag{3}$$

where Δ is the minimum difference between the two options that are salient to all agents of the population. Note that whether this inequality holds depends not only on the details of the payoff function, but also upon agent i's belief about the ratio of C-choosers in the rest of the population. For this reason determining the dynamics of the ratio of actual C-choosers requires an explicit belief formation process for x_i.

Suppose now that the population is in a state where every agent in the population reliably chooses strategy B. Given that this is the case, and that the experience of the agents confirms that this is the case, payoff maximizing agents will persist in choosing strategy B. We term this persistent selection of the payoff inefficient equilibrium *coordination failure.*

We take as our problem the design an intervention that will effectively re-coordinate the population on the payoff efficient equilibrium. In particular we suppose that some authority is instituted with the right to inspect the agents' play of this game, and further to levy fines and rewards on the basis of these inspections. Specifically, we seek a system of inspections, fines, and rewards, to ensure the population re-coordinates. We assume that this intervention has the following form:

Timing 1.

1. *The policymaker chooses the proportion, k/N, of games played which will be inspected.*
2. *The policy maker sets out rules according to which fines, F, and rewards, R, are assigned. The magnitude of the fines and rewards are allowed to vary with the outcome of each batch of inspections. In each N-game batch, k of the interacting pairs will be inspected, for a total of $2k$ inspected agents per batch. We use F_s and R_s, with $s \in \{0, 1, \ldots, 2k\}$, to denote the fines and rewards applied to inspected agents when s of the $2k$ agents inspected are C-choosers.*
3. *The policy maker informs the population about the chosen system of inspection, fines, and rewards.*
4. *The policy maker divides the interactions into N-game batches, with k games inspected in each batch. For each interaction in a batch*
   (a) *Nature randomly draws two agents i and j.*
   (b) *Agents i and j choose their strategies.*
   (c) *Nature determines whether this interaction is one of the k of N interactions that is inspected.*
   (d) *Player receive payoffs from the interaction.*

5. *The policy determines the magnitude of the fines and rewards to be assigned, potentially on the basis of, s, the number of compliant agents in the sampled population.*

6. *The policy maker distributes fines and rewards to the inspected agents according to the agents' strategic choices, i.e. a reward if C is chosen and a fine if B is chosen.*

7. *Continue to the next batch of choices and proceed to Step 4*

Under an intervention of this form the expected utility of engaging in behavior C is

$$E_i[U_{\bar{G}}(C)] = E_i[x + (k/N)R(x_s)] = x_i + (k/N)E_i[R(x_s)] \tag{4}$$

and for the same agent choosing strategy B,

$$\begin{aligned} E_i[u_{\bar{G}}(B)] &= E_i[xT + (1-x)P - (k/N)F(x_s)] \\ &= Tx_i + P(1-x_i) - (k/N)E_i[F(x_s)]. \end{aligned} \tag{5}$$

Thus, agent i will choose C if and only if

$$x_i - Tx_i - P(1-x_i) + (k/N)E_i[R(x_s) + F(x_s)] \geq \Delta. \tag{6}$$

Recall that Δ is the minimum salient difference between the two options. Again, whether this inequality holds depends not only on the details of the payoff function and reward/fine functions, but also upon agent i's belief about the ratio of C-choosers in the rest of the population. What makes the re-coordination problem difficult is that the agents' idiosyncratic beliefs and belief formation processes are difficult to predict or model, both for the policymaker, and for the agents within the population. Our solution hinges on a mechanism that renders these idiosyncratic and essentially unknowable agent beliefs irrelevant in determining the expected utility of a given action. The essence of the mechanism is a carefully chosen dynamic system of rewards and fines.

Naively, this problem of belief formation could be avoided by ignoring rewards and only introducing a constant fine, F, so that

$$\begin{aligned} & x_i - [Tx_i + P(1-x_i)] + (k/N)F \geq \Delta \forall x_i \in [0,1] \\ \iff & F \geq (\Delta + P)N/k \end{aligned}$$

Such a fine will maintain the desired inequality for all possible beliefs an agent might have, i.e. even in the worst case scenario where an agent believes that every other agent in the population will choose B. It is observations like this that are responsible for the suggestion that maximal fines are in some sense optimal. However, worries about corruption and further harm introduced by fines suggests that transfers due to fines and rewards should be minimized. Formally then, we would like the net fines and rewards levied for any given sampling outcome to be as small as possible, with increasing penalty given to larger transfers, that is we seek to minimize

$$|F_0|^2 + |R_{2k}|^2 + \sum_{s=1}^{2k-1} (|F_s| + |R_s|)^2 \tag{7}$$

while still ensuring re-coordination.

Finally, we assume that there is a budget constraint upon the fines and rewards requiring that

$$R_s s - F_s(2k - s) \leq Y \tag{8}$$

for all possible sampling outcomes.

Here we assume that inspection does not interrupt the play of the game, so that when a game is inspected the resultant payoffs are the sum of usual payoffs and the fines and rewards.

3 A minimal, effective re-coordination solution

As noted above, a system of rewards and fines gives a policymaker the power to arbitrarily reorder the agents' relevant preference orderings. However, a naive system of rewards and fines may in order to ensure that all agents play C, as the policymaker desires, need to make significant interventions, and thus be politically untenable. We will show that it is possible to introduce a very particular sets of rewards and fines such that the payoff maximizing agents will re-coordinate. Further, they will do this regardless of their beliefs about the proportion of C-choosers, while at the same time keeping fines and rewards to a minimum, and respecting a budget constraint. The key is to make the size of the fines and rewards dependent upon the proportion of the population playing C.

For any targeted fine and reward system, there must also be an inspection mechanism to identify which agents should receive fines and rewards. We make the simple observation that such an inspection mechanism may also provide the policymaker with an unbiased estimate of the ratio of compliance within the population. Thus, inspection may be used to produce a posthoc, unbiased estimate of the true proportion of the population playing C, and the fines and rewards applied can be made contingent upon this estimate.

Given the form of the intervention outlined above, we are tasked with finding fine and reward values F_s and R_s, $s \in \{0, 1, \ldots, 2k\}$ such that re-coordination is guaranteed to occur, the intervention is within the budgetary constraint, and net transfers are as small as possible given that these first two constraints are met.

Proposition 1. *For every pair of positive integers N and k with $k \leq N$, there exists a unique set of real numbers $\{F_0, F_1, \ldots F_{2k-1}\}$ and $\{R_1, \ldots R_{2k}\}$ which minimize $|F_0|^2 + |R_{2k}|^2 + \sum_{s=1}^{2k-1}(|F_s| + |R_s|)^2$ subject to the constraints:*

1.
$$x - Tx - P(1-x) + (k/N)\left[x^{2k}R_{2k} + (1-x)^{2k}F_0 \right.$$
$$\left. + \sum_{s=1}^{2k-1}\binom{2k}{s}(x)^s(1-x)^{(2k-s)}(R_s+F_s)\right] \geq \Delta \forall x \in [0,1]$$

and

2. $R_s s = F_s(2k-s) \forall s \in \{1,2,\ldots,2k-1\}$ *and* $R_{2k} = 0$

Proof. For convenience we introduce $g_s = |R_s| + |F_s|$, with $\mathbf{g} = (g_0, g_1, \ldots, g_{2k})$, and $b_s(x) = \binom{2k}{s}x^s(1-x)^{(2k-s)}$ with $\mathbf{b}(x) = (b_0(x), b_1(x), \ldots, b_{2k}(x))$. Letting l be the linear function $l(x) = N/k(\Delta + P - (1-T+P)x)$ and "$\cdot$" denote the standard scalar product in R^{2k+1}, we may rewrite the first constraint as

$$\mathbf{g} \cdot \mathbf{b}(x) \geq l(x) \forall x \in [0,1], \tag{9}$$

and, rearranging the second constraint we have that

$$g_s = \frac{Y + 2kF_s}{s}. \tag{10}$$

The first set of constraints bounds the possible values of g_s from below. The second set of constraints does not bound g_s but only places constraints on the ratio between R_s and F_s for a given value of g_s. Thus, our problem is reduced to finding the smallest $\mathbf{g}$, in the sense of the standard Euclidean norm, which satisfies the constraints in equation 9.

For each $x \in [0,1]$, equation 9 says that $\mathbf{g}$ must lie in the halfspace on or above the hyperplane defined by all $\mathbf{g}$ such that $\mathbf{g} \cdot \mathbf{b}(x) = l(x)$. Taken together then, 9 says that $\mathbf{g}$ must lie in the intersection of the infinite collection of these halfspaces. The intersection of halfspaces forms a convex set $\mathcal{C}$, and the Euclidean norm of a vector is a convex function. Thus, the basic results of convex optimization theory guarantee that there is a unique vector $\mathbf{g}^*$ which will minimize $|\mathbf{g}|$ while satisfying the given constraints. □

In the discussion, we look at three particular examples where the unique minimizer, shown to exist here, has been computed numerically.

4 Discussion

In this paper a simple idea was introduced: a mechanism using dynamic fines and rewards calibrated to offset the players' expected utilities such that they no longer depend on the expected choices of others. The mechanism was introduced using a 2-player stag hunt, it may, however, enforce a strategy in any coordinative situation. As discussed in the introduction, the dynamics of coordinative situations are hard to model, they require an explicit *theory of mind*, i.e. a way for players to form beliefs about others' beliefs. We lack a useful theory about such a theory. The mechanism introduced here lets us sidestep this problem, by transforming the complex social decision to a trivial individual decision for each player.

	C	B
C	1	0
B	.75	.5

G_1

	C	B
C	1	0
B	.95	.9

G_2

	C	B
C	.5	0
B	1	.25

G_3

Figure 2: Payoff matrices for the exemplary games. G_1 and G_2 are stag hunts, G_3 is a prisoner's dilemma.

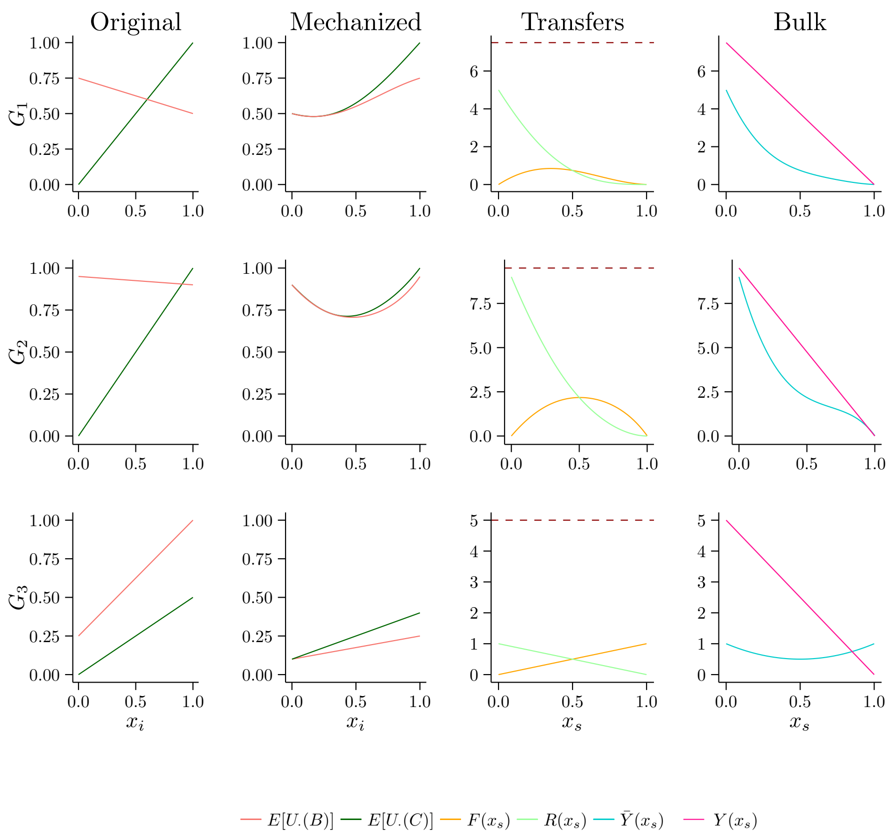

Figure 3: The first column show the expected payoffs for the two strategies as a function of agent belief in the *Original* version of G_1, G_2 and G_3 as seen in Figure 2. The second column depicts the expected payoffs for the mechanized version of the same games. The third column shows the fines and rewards as a function of the proportion of compliant agents in the sampled population. The minimal static fine required to guarantee re-coordination is depicted as a dashed line for reference. In the fourth column the *Bulk* transfers under the dynamic regime, $\bar{Y}$ is compared to the bulk transfers under a minimal static fine regime, Y.

In this section we consider two stag hunts, G_1 and G_2 and one prisoner's dilemma game, G_3. The games are depicted in Figure 2.

In Figure 3 we plot how the mechanism apply to the games in Figure 2. All games are mechanized using batches of 100 interaction out of which 10 are inspected. In the left most column of panels, the *Original* column, a player i's expected utility of the original games, $E_i[U_\cdot(C)]$ and $E_i[U_\cdot(B)]$, are plotted for her expectation of the proportion of compliers, x_i. Similarly, in the second column of panels from the left, the *Mechanized* column, depicts her expected utility of the mechanized game, $E[U_{\bar{\cdot}}(C)]$ and $E[U_{\bar{\cdot}}(B)]$. For G_2 and G_3, we can see how in the *Mechanized* column how the expected utility difference is such that "comply" dominates.

In the third column of panels from the left, the *Transfers* column, we plot the expected fine and reward function, $E[F(x_s)|x_i]$ and $E[R(x_s)|x_i]$, of the mechanism, as well as the static fine level that would suffice for re-coordination, $\hat{F}(x_s)$, all with respect to the proportion of compliers in a sample, x_s. The rewards are high with few compliers in an inspected sample and decreases to zero as the proportion of compliers increases. The fines, which finance the rewards, do the opposite, as long as some level of enforcement is necessary. In the case of the stag hunts, if the expected proportion of compliers overshoot the game's tipping point the payoffs of the original game will start enforcing the "comply" strategy. This allows fines to decrease. As fines reaches zero, enforcement cease. For the prisoner's dilemma there is no tipping point, hence to finance the rewards and uphold the enforcement, fines keep increasing throughout. Note that for all the games the smallest sufficient static fine level is higher or equal to any dynamic fine level.

In the fourth column of panels, the *Bulk* column, we plot the expected bulk of transfers per individual under a minimal static fine regime, $(1 - x_i)\hat{F}(x_i)$, and for our mechanism, $x_i E[R(x_s)|x_i] + (1 - x_i)E[F(x_s)[x_i]$. For the stag hunts, we see how the mechanism's bulk transfers undershoot the minimal static fine regime. In the case of the prisoner's dilemma, the condition of a balanced budget forces us to spend the income of the few huge fines needed in order to break the temptation to breach. Making the mechanism less minimal than the other regime for samples where the proportion of compliers is high.

In order to keep the presentation simple, we chose to analyze the mechanism in an idealized settings. We assume that there is a policymaker that has the power, the will, and the know-how to implement a dynamic fine and reward system. We also assume that the inspection is incorruptible and perfect, that every agent has the same utility function and that any two players are equally likely to play each other. The mechanism does therefore not lend itself to direct application for most coordinative situation though it can inform a more complex analysis.

5 Conclusion

In order to model the dynamics of coordinative situations one needs to make assumptions about the agents' empathetic skills: how they perceive each others' inductive standard, preferences, background information and the extent of the common knowledge thereof.

This paper develops a simple idea. Faced with a coordinative problem a policymaker may introduce an institution that balances the uncertainty about how others will act. It achieves this by introducing a fine and reward system that depends on an approximation of the same uncertain variable as the original game: the ratio of the inspected population acting in accordance with the policymaker's will. Reward and fine functions are determined such that they make a given strategy dominant. We thereby show that we can change norms and enforce conventions without unbalancing the budget, miss out on utility or having expected fine and reward levels different from zero.

References

Abigail Barr and Danila Serra. Corruption and culture: An experimental analysis. *Journal of Public Economics*, 94(11-12):862–869, 2010.

Gary S Becker. Crime and punishment: An economic approach. *The Journal of Political Economy*, 76(2):169–217, 1968.

Cristina Bicchieri. *The grammar of society: The nature and dynamics of social norms*. Cambridge University Press, 2006.

Ken Binmore and Larry Samuelson. The evolution of focal points. *Games and Economic Behavior*, 55(1):21–42, 2006.

Andreas Blume and Andreas Ortmann. The effects of costless pre-play communication: Experimental evidence from games with pareto-ranked equilibria. *Journal of Economic theory*, 132(1):274–290, 2007.

Jordi Brandts and David J Cooper. A change would do you good.... an experimental study on how to overcome coordination failure in organizations. *The American Economic Review*, pages 669–693, 2006.

Val Curtis, Wolf Schmidt, Stephen Luby, Rocio Florez, Ousmane Touré, and Adam Biran. Hygiene: new hopes, new horizons. *The Lancet infectious diseases*, 11(4):312–321, 2011.

Bin Dong, Uwe Dulleck, and Benno Torgler. Conditional corruption. *Journal of Economic Psychology*, 33(3):609–627, 2012.

Patricia M Fabiano, H Wesley Perkins, Alan Berkowitz, Jeff Linkenbach, and Christopher Stark. Engaging men as social justice allies in ending violence against women: Evidence for a social norms approach. *Journal of American College Health*, 52(3):105–112, 2003.

Drew Fudenberg, David M Kreps, and David K Levine. On the robustness of equilibrium refinements. *Journal of Economic Theory*, 44(2):354–380, 1988.

Patricia Funk. Governmental action, social norms, and criminal behavior. *Journal of Institutional and Theoretical Economics (JITE)/Zeitschrift für die gesamte Staatswissenschaft*, pages 522–535, 2005.

Nuno Garoupa. The theory of optimal law enforcement. *Journal of economic surveys*, 11(3):267–295, 1997.

John C Harsanyi. Games with incomplete information played by “bayesian” players, i-iii part i. the basic model. *Management science*, 14(3):159–182, 1967.

John C Harsanyi and Reinhard Selten. A general theory of equilibrium selection in games. *MIT Press Books*, 1, 1988.

Trey Hedden and Jun Zhang. What do you think i think you think?: Strategic reasoning in matrix games. *Cognition*, 85(1):1 – 36, 2002.

Josef Hofbauer and Karl Sigmund. Evolutionary game dynamics. *Bulletin of the American Mathematical Society*, 40(4):479–519, 2003.

Michihiro Kandori, George J Mailath, and Rafael Rob. Learning, mutation, and long run equilibria in games. *Econometrica: Journal of the Econometric Society*, pages 29–56, 1993a.

Michihiro Kandori, George J Mailath, and Rafael Rob. Learning, mutation, and long run equilibria in games. *Econometrica: Journal of the Econometric Society*, pages 29–56, 1993b.

Louis Kaplow. A note on the optimal use of nonmonetary sanctions. *Journal of Public Economics*, 42(2):245–247, 1990.

Toshiji Kawagoe and Hirokazu Takizawa. Equilibrium refinement vs. level-k analysis: An experimental study of cheap-talk games with private information. *Games and Economic Behavior*, 66(1):238–255, 2009.

Rosemarie Nagel. Unraveling in guessing games: An experimental study. *The American Economic Review*, pages 1313–1326, 1995.

A Mitchell Polinsky and Steven Shavell. A note on optimal fines when wealth varies among individuals. *The American Economic Review*, pages 618–621, 1991.

A Mitchell Polinsky and Steven Shavell. Enforcement costs and the optimal magnitude and probability of fines. *Journal of Law and Economics*, pages 133–148, 1992.

David Schmidt, Robert Shupp, James M Walker, and Elinor Ostrom. Playing safe in coordination games:: the roles of risk dominance, payoff dominance, and history of play. *Games and Economic Behavior*, 42(2):281–299, 2003.

Reinhard Selten. Spieltheoretische behandlung eines oligopolmodells mit nachfrageträgheit: Teil i: Bestimmung des dynamischen preisgleichgewichts. *Zeitschrift für die gesamte Staatswissenschaft/Journal of Institutional and Theoretical Economics*, pages 301–324, 1965.

Elizabeth Shove. Beyond the abc: climate change policy and theories of social change. *Environment and planning. A*, 42(6):1273, 2010.

Robert Sugden. Spontaneous order. *The Journal of Economic Perspectives*, 3 (4):85–97, 1989.

Robert Sugden. Salience, inductive reasoning and the emergence of conventions. *Journal of Economic Behavior & Organization*, 79(1-2):35–47, 2011.

John B Van Huyck, Raymond C Battalio, and Richard O Beil. Tacit coordination games, strategic uncertainty, and coordination failure. *The American Economic Review*, pages 234–248, 1990.

Paul F Whiteley. Economic growth and social capital. *Political Studies*, 48(3): 443–466, 2000.

H Peyton Young. The evolution of conventions. *Econometrica: Journal of the Econometric Society*, pages 57–84, 1993.

Yoshinobu Zasu. Sanctions by social norms and the law: Substitutes or complements? *The Journal of Legal Studies*, 36(2):379–396, 2007.