

TALKING OURSELVES TO EFFICIENCY: COORDINATION IN INTER-GENERATIONAL MINIMUM EFFORT GAMES WITH PRIVATE, ALMOST COMMON AND COMMON KNOWLEDGE OF ADVICE*

Ananish Chaudhuri, Andrew Schotter and Barry Sopher

We use experiments to investigate the use of advice as a coordinating device in the ‘Minimum Effort Game’ which is a coordination game with weak strategic complementarities and Pareto-ranked equilibria. The game is played by non-overlapping generations of players who, after they are done, pass on advice to their successors who take their place in the game. We conjectured that this inter-generational design might enable subjects to converge to the payoff-dominant outcome. We find that coordination is most likely to result when the advice is made public and also distributed in a manner that makes it common knowledge.

Coordinating the actions of multiple agents is a central issue in many economic problems and the successful resolution of coordination failures is crucial to achieving optimal outcomes in such cases. For example, one of the most important problems in macroeconomics involves the issue of under-employment where the economy can be trapped in a Pareto-inferior equilibrium (Cooper, 1999). In such instances no firm wishes to expand production unless it can be assured that others will do so, yet not doing so leads to an outcome that is worse for everyone concerned. Kremer (1993) in his O-Ring theory of development proposes extensive coordination failures as the cause of under-development in many countries. Here countries may be caught in a low-level equilibrium ‘trap’ when development requires the simultaneous industrialisation of many sectors of the economy but no sector can break even industrialising alone; see also Murphy *et al.* (1989). Similar considerations arise in models of currency crises or speculative attacks (Morris and Shin, 1998), models of bank-runs (Diamond and Dybvig, 1983) and models of political revolution (Kuran, 1987, 1995) where it is politically risky to attempt to overturn a government unless it is commonly known that others will rise up together.¹ Such coordination problems are ubiquitous and typically characterised by both strategic complementarities and spill-over. In most instances these phenomena give rise to multiple Pareto-ranked equilibria.

The central policy question here is how to break out of unsatisfactory equilibria when they occur. Obviously this can only be achieved if some event occurs that convinces people that others have interpreted this event to mean that everyone will coordinate

* Resources for the research were provided by the National Science Foundation through grant number SES 0111640, by both the Center for Experimental Social Science and the C.V. Starr Center for Applied Economics at New York University, by the Rutgers University Research Council, and by the University of Auckland Research Council. We thank Shachar Kariv for both his comments and research assistance and Bogachan Celen for his comments. In addition, we thank John Van Huyck for providing us with a copy of his instructions as well as the Indian Statistical Institute, Kolkata for allowing us to conduct some pilot experiments there. Finally, we thank Leonardo Felli and two anonymous referees for their comments and suggestions which have led to improvements in the article’s exposition. We are responsible for all errors in the article.

¹ Knez and Simester (2001) provide an insightful account about how successful resolution of coordination failures in various operations led to the remarkable turn-around at Continental Airlines after 1995. Ichniowski *et al.* (1997) describe how successful steel-mills adopt innovative human resource management practices that foster coordination along their production lines in an attempt to boost productivity.

(i.e., increase output or revolt). In a macro context such an event might be an announcement of the central bank to lower interest rates or even some jawboning (public announcement) by a politician. In a revolution it might be a public demonstration (Chwe, 2001). In other words, it must be common knowledge that such an event is interpreted in an identical manner by all agents. The question then becomes: how easy is it to create such an expectation? In his book, *Rational Ritual: Culture, Coordination and Common Knowledge*, Chwe (2001) discusses the fine points of arriving at common knowledge and investigates various institutional devices that societies commonly use in order to do so. In this article we show that

- (1) indeed such common knowledge regarding the expected actions to be taken by others can lead to successful coordination but
- (2) creating such common knowledge is often challenging and
- (3) that small deviations from such common knowledge can lead to seriously sub-optimal outcomes.

We investigate a well-known coordination problem called the ‘minimum effort game’ (which is an n-person variant of the stag-hunt game and a strategic form representation of Bryant’s (1983) Keynesian coordination game) with a long history in the experimental literature (Van Huyck *et al.*, 1990), hereafter referred to as VBB. The game appears in Table 1:

Consider this game as being played by a set of N players (in our experiments we use 8 subjects). Each player can choose an integer, c_i from the set $\{1, 2, 3, 4, 5, 6, 7\}$. Individual payoffs are determined for each subject by the payoff function: $\pi_i = k + a[\min(c_i \dots c_n)] - bc_i$. In other words, the payoff for any player i is equal to k plus a constant, a , times the minimum choice of any subject minus another constant, b , times i ’s choice. Choosing $k = \$0.60$, $a = \$0.20$ and $b = \$0.10$ defines the game depicted in Table 1 which is the identical payoff table used by VBB. In this payoff matrix the Nash equilibria are displayed along the diagonal and are Pareto-ranked. The best payoff occurs when all subjects choose 7 but since the cost of one’s choice is subtracted from the common payoff to all, higher choices are more risky. In fact, the mini-max or secure strategy choice is to choose $c_i = 1$.

When this game is played in the laboratory subjects routinely select the Pareto-worst outcome where everyone chooses $c_i = 1$. While this result seems odd, it is consistent with a large number of game theoretic papers all of which demonstrate that

Table 1
Payoff Table in Van Huyck et al.’s (1990) Minimum Effort Game

		Smallest Value of X Chosen						
Your Choice of X		7	6	5	4	3	2	1
	7	1.30	1.10	0.90	0.70	0.50	0.30	0.10
	6	-	1.20	1.00	0.80	0.60	0.40	0.20
	5	-	-	1.10	0.90	0.70	0.50	0.30
	4	-	-	-	1.00	0.80	0.60	0.40
	3	-	-	-	-	0.90	0.70	0.50
	2	-	-	-	-	-	0.80	0.60
	1	-	-	-	-	-	-	0.70

Pareto-inferior outcomes are likely in a setting where the Pareto-dominant outcome is also the most risky. For example, Crawford (1991, 1995) demonstrates that the VBB results (and similar results from other studies) are consistent with results in evolutionary game theory and also shows how such results could be the outcome of a learning process which converges to the Pareto-worst equilibrium. Carlsson and Van Damme (1993) study 2×2 global games in which the risk dominant equilibrium is the unique equilibrium expected to be selected despite the existence of another, Pareto-dominant, equilibrium. Morris *et al.* (1995) generalise this result and provide conditions under which we would expect that to occur. Finally, Rubinstein (1989) shows that such results are possible when the players have ‘almost’ but not ‘complete’ common knowledge of the payoffs of the game. Rubinstein adds, however, that while such an outcome is logically consistent, it is not intuitively appealing and hence we should search for other heuristics that decision makers might use in such contexts that lead them to the Pareto-dominant outcome.

As argued above, the path to successful coordination involves being able to reach a state where it is common knowledge that all players intend to cooperate (choose 7 in the VBB minimum effort game). But is common knowledge a necessary condition or might it be possible to have some notion of ‘almost common knowledge’ which is good enough for successful coordination? Monderer and Samet (1989) develop the notion of a common p -belief which makes common knowledge something less than an all or nothing proposition and which we will use to explain the results of our experiment. The question then is: what devices do we have at our disposal that can generate such common knowledge and can extensive coordination failures be resolved if we only have ‘almost’ common knowledge rather than common knowledge?

In this article we investigate the use of advice left by players who have experience participating in the game as a coordinating device or event. In the experiments reported here we present an inter-generational version of the VBB experiment. The experimental approach is similar to that of Schotter and Sopher (2003, 2006, 2007) and also to the theoretical work of Anderlini and Lagunoff (2005) and Anderlini *et al.* (2007) on dynastic games where economic agents have finite lives but play games of infinite duration with inter-generational inter-dependent utility functions.

In our study, groups of 8 subjects are recruited into the laboratory and play the same game played by subjects in the VBB paper for 10 rounds, using pen and paper. After this each subject is replaced by another subject, a laboratory ‘descendent’, or ‘successor’, who then plays the game for another 10 rounds with the new group of subjects, so the generations are non-overlapping.² Advice from a member of one generation to a

² We are primarily interested in the impact of intergenerational advice on behaviour and on social learning. Hence we have chosen a non-overlapping generational set up precisely to isolate the role of advice from that of strategy. In an overlapping-generations set up, as Kandori (1992), Cremer (1986), Salant (1991) and others have shown, one can often generate cooperation (with folk-theorem-like results) via the overlapping generations structure. We are trying to avoid this possibility and make cooperation strictly a result of advice with no strategic possibilities. However, generations in our study are tied via the payoff structure since each generation ‘cares’ about the next. Further, an overlapping generations experiment of this sort would introduce new logistical complexities as one ‘retires’ some subjects from a group and ‘hires’ new subjects. For example, if the experiment ends on one day one would have to bring back the old subjects who are remaining in the game, along with the new subjects, for a new session. This could be very complicated since we wanted to run many generations. In our experiments, when a generation was over they all went home and did not have to come back to continue with the next set of agents.

successor can be passed along via free-form messages that generation t players leave for their generation $t + 1$ successors. Finally, payoffs span generations in the sense that the payoff to a generation t player is equal to what he has earned during his lifetime plus what his successor earns. Hence, incentives exist for subjects to pass on intelligent advice. We call the method of distributing advice ‘almost common knowledge’ if players in generation $t + 1$ receive advice from *all* the generation t players in written form but the advice is not read out loud, while ‘common knowledge’ is used to denote a treatment where the advice from all the generation t players is *not only distributed on paper but actually read out loud*. Hence, for us the distinction between common and almost common knowledge is dependent on the technology of distributing information, while for Monderer and Samet (1989) it is about how such information is commonly interpreted. Still, we rely on their findings to help us interpret our results.

It was our conjecture that if we played the minimum effort game using such an inter-generational design then, over time, generations would be able to ‘talk themselves to efficiency’ in the sense that after playing the game, if any generation converged to the Pareto-worst equilibrium, they might advise the next generation to take actions consistent with an efficient outcome, knowing that following their own behaviour (as opposed to their advice) would be a first step on the path to mutually assured destruction. Wise subjects might say, ‘do as we are telling you to do, not as we did’, and such advice, if followed, might lead to a convention selecting the Pareto-superior outcome. Hence, we expected that outcomes in the inter-generational games would be more efficient than those found by VBB. (It is of course possible that, subjects would have learned the exact opposite lesson, that people cannot be trusted, and pass that on to their successors. Hence, advice could just as easily reinforce inefficiency as reduce it.)

We find that it is much harder for societies to ‘talk themselves to efficiency’ than we expected. More precisely:

- 1 Private advice between a predecessor and his successor, no matter how positive, consistently fails to result in efficiency.
- 2 With respect to public advice, i.e., advice that is publicly announced, there seems to be a quality threshold. If the advice offered is sufficiently exhortative with predecessors unanimously urging their successors to choose the strategy that is commensurate with the payoff-dominant outcome then, as long as that advice is offered in a public manner (either as common knowledge or as what we call ‘almost common knowledge’), we can expect efficient coordination to follow.
- 3 If the advice quality of the public advice is even slightly below this unanimity threshold - e.g., one subject equivocates regarding the optimal strategy and suggests a course of action that deviates from the payoff dominant equilibrium – then ‘almost common knowledge’ is no longer sufficient and the advice offered must be offered in a manner that induces complete common knowledge.

Finally, we find the interesting phenomenon that in the private advice sessions advice tends to be pessimistic in that it tends to suggest to successors that they should choose the action commensurate with the Pareto-worst outcome and stick with it. Surprisingly, this advice is typically ignored in the first round of the subsequent generation, as subjects tend to take a chance and choose a number higher than they were advised. As

time goes on, however, they quickly learn to choose the worst action so that by the end of the session (after 10 rounds), they not only choose the worst action, 1, but suggest it to their successors. Hence, in the private advice setting, rather than experience teaching subjects the folly of their inefficient ways, it teaches them the wisdom of selfish choice which they pass on to their successors. Such advice is ignored, initially, and then its wisdom recognised in time to be passed on once again.

We proceed as follows. In Section 1 we present our experimental design. In Section 2 we discuss our results. In Section 3 we explore the beliefs behind the players' actions. Finally in Section 4 we offer some concluding remarks.

1. Experimental Design

In the experiments discussed in this article we present an inter-generational version of the original VBB experiments presented in Table 1 and described above. Groups of 8 subjects are recruited into the laboratory and play the minimum effort game for 10 rounds. After their participation is over each one is replaced by another agent, his laboratory descendent, who then plays the minimum game for 10 rounds with a new group of subjects. This group is then replaced by another 8 successors who take their place and play on. When generations change, after 10 rounds of repetition, outgoing agents are allowed to pass on advice in the form of free-form written messages to their successors. The successors, depending on the game, are able to view these messages and some subset of the history of play of previous generations before they make their choices. Payoffs are equal to the sum of what an agent earns during his lifetime plus what his successor earns in the next generation so there is complete inter-generational caring (more precisely, there is no one-period ahead discounting). Finally, before the first and last of the ten rounds of any agent's life, we ask them to state their beliefs (using a proper scoring rule). Specifically, we ask them to state the frequency with which they expect each of the 7 strategies will be played in the 8 person population. This allows us to investigate the relationship between observed actions and (normally) unobserved beliefs.

Prior to the first round of any generation, subjects are presented with a set of written instructions which are read out loud to them after they have finished reading them privately. After questions are answered, depending on the treatment, subjects are allowed to view the history of the previous generations and read the advice offered by their predecessor. In some treatments they are only allowed to read the advice offered by their own predecessor in the previous generation and cannot see the history of previous plays from earlier generations. After the last round, subjects write advice to their successors and leave. (When they write advice they know whether it is to be made public to all eight subjects in the next generation or simply to be read privately by their successor.) They are paid their earnings in cash upon completion of the session. At a later date each subject receives a second payment which is equal to the earning of his successor in the next generation. They are also paid separately for their predictions according to a quadratic scoring rule. (See the instructions to the experiment contained in the Appendix for a description of the scoring rule used.) The payment for the predictions is a small fraction of what they could earn in the game itself.

We conducted a total of 11 distinct games in two distinct blocks. We first describe Block I in some detail, and then go on to Block II. Block I was a set of five different games which varied according to the information available to subjects. In Game 1, the Replicator (No-Advice) Game, we aimed to replicate the VBB results (albeit with only 8 rather than 14 or 16 subjects), running the minimum effort game as they did with neither generations nor advice for 10 rounds.

In running our inter-generational experiments we started by running Game 2, a 'Garden of Eden' or Progenitor Game in which an initial group of eight subjects played the game for the first time with no advice. This generation was the progenitor of all generations in all games in Block I in the sense that the first generations of all other games that followed used the advice of this progenitor generation. For example, in Game 3 we had six generations of subjects play the minimum effort game in circumstances where each player, prior to starting the experiment, could view the history of all actions taken by his predecessors in all generations before him and could also receive advice from his immediate predecessor. Here advice was private so that while each agent knew that the others were receiving advice, they did not know the content of any advice other than their own. This was the Advice Plus History Game. Game 4 also had six generations of subjects playing the minimum effort game except here subjects could only receive advice from their immediate predecessors but could not see any history of past plays. The first generation here received the same progenitor advice as did the subjects in Game 3. This Game was called the Advice Only Game.

Finally, Game 5 was the Public Advice Game. This experiment was run for nine generations but it was run differently for the first five and last four generations. During the first five generations, before play started in the first round, subjects were given a sheet of paper upon which was written the advice offered by each of the subjects in the previous generation. (The first generation here received all the advice from the Progenitor generation of Game 2.) Each subject in this treatment knew that all other subjects were looking at a sheet containing exactly the same advice offered by all the subjects in the previous generation but an individual subject had no idea if the others in the session actually read the sheet or how carefully they read it. Hence, we call these first five generations the 'Public Advice Almost Common Knowledge' condition. Starting with Generation 6, however, we not only gave the subjects the advice sheets listing all the advice left by the previous generation but also read these pieces of advice out loud so the content of the advice on these sheets was common knowledge. This we call the 'Public Advice Common Knowledge' condition. In none of these generations did subjects receive any information about the history of plays from previous generations. Table 2(a) provides the details about the design of our Block I experiments.

While our initial results were interesting, they did not control for the quality of advice. In other words, two factors were closely at work determining when coordination was possible: the quality of the advice being offered and the manner in which it was distributed. To control for the effect of advice quality we ran another set of six games in Block II. In these games, subjects played the same minimum effort game for 10 rounds just as our subjects did in Block I. They were told that a previous generation had played the game before them and had left advice for them. (As we explain below in Section 2.2, this advice was different from that left by the subjects of Game 2 above.) They were also told that they would leave advice for a set of subjects who would replace them.

Table 2
 (a) *Experimental Design for the Inter-generational Games in Block I*

No.	Game	Number of Generations	Rounds Per Generation	Subjects Per Generation	Number of Subjects
1	Replicator: No-Advice or History	4	10	8	32
2	Progenitor: No History or Advice, but Advice left	1	10	8	8
3	Private Advice-Plus-History	6	10	8	48
4	Private Advice Only	6	10	8	48
5	Public Advice	9	10	8	71
	(Public Advice <i>Not</i> Read Aloud – Almost Common Knowledge)	(5)	(10)	(8)	(40)
	(Public Advice Read Aloud – Common Knowledge)	(4)	(10)	(8)	(31)
	Total				207

(b) *Experimental Design for Block II Games*

No.	Game	Number of Groups	Rounds Per Group	Subjects Per Group	Advice Quality	Number of Subjects
6	Common Knowledge of Advice	4	10	8	Good	32
7	Almost Common Knowledge of Advice	4	10	8	Good	32
8	Advice projected on Overheads for all subjects to see	3	10	8	Good	24
9	Common Knowledge of Advice	4	10	8	Very Good	32
10	Almost Common Knowledge of Advice	4	10	8	Very Good	32
11	Private Knowledge of Advice	3	10	8	Very Good	24
	Total					176

Note. In the very last session of the Public Advice game, we only had 7 subjects instead of the 8 that we used in every other session.

Hence as far as they were concerned they were playing as one generation in an inter-generational game, which they were. These games differ from the Block I games in that we are focusing on the effect of a particular set of advice on the play of several groups of subjects, rather than following the evolution of play over many generations. Subjects in the Block II games received either ‘good’ or ‘very good’ quality advice. Depending on the game, this advice was distributed to subjects either privately, as almost common knowledge, or as common knowledge. Block II then defined a design where the factors were the advice quality (good and very good) and the way it was disseminated (privately, as almost common knowledge, or as common knowledge). We filled out all the cells in this design except that we did not combine the ‘good’ advice with private distribution since it was clear that this would not spur efficient outcomes, since neither the ‘almost common knowledge’ nor the ‘common knowledge’ conditions resulted in efficiency with ‘good’ advice. We also considered, for ‘good’ advice only, a game intermediate between almost common knowledge and common knowledge. In this game subjects received a sheet of paper containing all the advice of their predecessors and this same advice sheet was projected on the laboratory wall using an overhead projector for all to see. We only ran this game for ‘good’ advice for reasons that will become clear shortly.

The advantage of our Block II games is that we are able to control the quality of advice that subjects get as well as the way it is distributed. The disadvantage is that these

games are no longer strictly inter-generational so we are unable to see how generational advice influences behaviour over time. This is not a serious drawback since one clear and consistent result we will demonstrate is that once advice in any generation falls below a certain quality threshold, no matter how that advice is distributed, cooperation is very unlikely to be achieved. Our focus of attention in these experiments then shifts to the impact of advice quality and distribution method on the first generation receiving it. What we are interested in is how the quality of advice and the manner in which it is distributed affect the beliefs and actions of the subjects receiving it as they enter the first round of their 10 round interactions.

In our analysis of the Block II data we concentrate to an extent on first round behaviour because if coordination is not achieved in round 1 of a group's interaction there is no precedent for it ever being achieved in a later round. Coordination in the first round is therefore a necessary condition for coordination in later rounds and in this sense behaviour in these games is extremely path dependent where the initial impact of advice on beliefs and actions is the key determinant of the future outcome.

The exact experimental design of the Block II experiments is summarised in Table 2(b).

The experimental sessions were conducted at Washington State University, at the Center for Experimental Social Science (C.E.S.S) at New York University, the University of Auckland and Rutgers University. (No subject pool effects were noted so all observations are pooled). Inexperienced subjects were recruited from undergraduate courses and participated for about one and one half hours. Average payoffs were approximately \$19.00.

2. Results

We present our results by first describing the behaviour of the minimum choices in the inter-generational Block I games and then moving on to the Block II games. By comparing these results we highlight the main result of our article which is that it is extremely hard to get agents to coordinate their actions in this game and that the necessary conditions for coordination are that the advice offered be not only extremely positive but also offered publicly. If advice is less exhortative (only of the 'good' variety), however, then the manner in which it is distributed does matter at least for the behaviour of subjects in the first round. However, because behaviour in succeeding rounds is so path dependent, first round behaviour becomes crucial to eventual coordination.

2.1. Block I Games – Minimum Choices

Table 3(a) shows the behaviour of the minimum in the Block I games of our experiment, while Table 3(b) does the same for the VBB experiment.

In Figure 1 we present the round-by-round minimum choices of subjects for each of the five games in Block I along with those of the subjects in the seven groups run by VBB. Here we show the pattern of choices made by each generation in the different games over all 10 rounds. Each block in Figure 1 represents one generation and shows the minimum choices that particular generation made over all 10 rounds.

Table 3
 (a) *Observed Minimum Choices in Block I games*

	Rounds									
	1	2	3	4	5	6	7	8	9	10
Game 1: Replicator										
Group 1	2	1	1	1	1	1	1	1	1	1
Group 2	2	1	1	1	1	1	1	1	1	1
Group 3	1	1	1	1	1	1	1	1	1	1
Group 4	1	1	1	1	1	1	1	1	1	1
Game 2: Progenitor										
Progenitor Group	1	1	1	1	1	1	1	1	1	1
Game 3: Advice-Plus-History										
Generation 1	1	1	1	1	1	1	1	1	1	1
Generation 2	1	1	1	1	1	1	1	1	1	1
Generation 3	4	4	4	4	4	4	4	4	4	4
Generation 4	1	2	2	2	2	2	2	2	2	1
Generation 5	1	1	1	1	1	1	1	1	1	1
Generation 6	1	1	1	1	1	1	1	1	1	1
Game 4: Advice Only										
Generation 1	1	1	1	1	1	1	1	1	1	1
Generation 2	1	1	1	1	1	1	1	1	1	1
Generation 3	1	1	1	1	1	1	1	1	1	1
Generation 4	1	1	1	1	1	1	1	1	1	1
Generation 5	1	1	1	1	1	1	1	1	1	1
Generation 6	1	1	1	1	1	1	1	1	1	1
Game 5: Public Advice										
Generation 1	1	1	1	1	1	1	1	1	1	1
Generation 2	1	1	1	1	1	1	1	1	1	1
Generation 3	6	5	5	4	1	1	1	1	1	1
Generation 4	4	4	3	2	1	1	1	1	1	1
Generation 5	4	3	1	1	1	1	1	1	1	1
Generation 6	7	7	7	7	7	7	7	7	7	7
Generation 7	6	6	6	6	6	6	6	6	6	5
Generation 8	7	7	7	7	7	7	7	7	7	7
Generation 9	7	7	7	7	7	7	7	7	7	7

(b) *Group Minima – Van Huyck et al. (1990) Experiments*

Group	Rounds									
	1	2	3	4	5	6	7	8	9	10
1	2	2	2	1	1	1	1	1	1	1
2	2	1	1	1	1	1	1	1	1	1
3	4	2	2	1	1	1	1	1	1	1
4	4	2	3	1	1	1	1	1	1	1
5	3	2	1	1	1	1	1	1	1	1
6	1	1	1	1	1	1	1	1	1	1
7	1	1	1	1	1	1	1	1	1	1

As can be seen from Tables 3(a) and 3(b) and Figure 1, the behaviour of the group minima in the VBB and our Replicator (No-Advice) games is striking. First, note that in the VBB experiment (Table 3(b)), none of the 7 groups managed to achieve a minimum greater than 4 and in no group did the minimum remain above 1 for more than




Fig. 1. *Behaviour of the Minimum Across Block I Games*

three rounds. In two of the seven groups (Groups 6 and 7) we observe 1 chosen in each round. This behaviour is even more dramatic in our four replicator experiments (Game 1 in Table 3(a)). Here in no round was the minimum greater than 2 and that only occurred in two sessions and only during the first round. Beyond that all minima were equal to 1.

It was exactly this behaviour that we expected would disappear when we introduced advice into our design. To our surprise we found just the opposite to be true. For example, in the Advice Only Game (Game 4 in Table 3(a)) the behaviour exhibited was less coordinated than that in either the VBB or our replicator experiments. In no round of any generation did we ever observe a minimum above 1. Hence, it appears that a treatment where subjects cannot observe previous history but can only pass on advice is hostile to efficiency in a risky environment.

When advice and history are both available, as in Game 3, we do observe somewhat more improved coordination. Here, as in the Advice Only Game, in four of the six generations (generations 1, 2, 5, and 6) we observed a minimum of 1 in each round. However, in generation 3, we observe the first instance of an interior equilibrium existing for all 10 rounds. Here the minimum starts out at 4 in round 1 and stays at that level for the remainder of the game. In the next generation we also see an elevated minimum of 2 existing from round 2 to round 9. These two generations exhibit behaviour different from any seen in the VBB or Replicator games since they exhibit the first instances of a minimum above 1 lasting past round 3. Still, these results cannot be considered evidence of any strong impact of advice on behaviour in our inter-generational set up. Finally, in spite of the fact we see a minimum higher than 1 for 2 generations in this game, subjects find it impossible to sustain that level of coordination and by Generations 5 and 6, the minimum is again 1 for all 10 rounds.

In order to find truly different behaviour one must look at the results of Game 5, the Public Advice Game from Generation 6 on. In Generations 1–5, while all subjects read all of the advice from their eight predecessors, this advice was not read out loud; so in our terminology it was distributed as ‘almost common knowledge’. As one can see,

despite some attempts at improved coordination – for instance in Generation 3 the minimum is 6 in round 1 and then 5 for the next two rounds and 4 in round 4 before collapsing to 1 for the rest of the session – subjects in the first five generations were not successful in sustaining coordination. In Generation 6 we introduced a treatment change and we read out loud the advice from all the Generation 5 subjects. As one can see this had a large impact, raising the minimum to 7 in all ten rounds. From this generation on in no round was the minimum choice below 5 (which occurred only once) and was 7 for 30 of the total 40 rounds. (Of the other 10 rounds it was 6 in nine of them and 5 in one.) In short, this treatment, public advice with common knowledge, was successful in breaking the stranglehold that the all-1 equilibrium had on behaviour up until this point.

In order to quantify the differences in choice behaviour among the treatments more precisely we estimate the probability of making various choices. We estimate an ordered probit model with a random effects specification of the error term to take account of individual-specific components in choice behaviour. An ordered choice model is appropriate in this context as there is a natural ordering of the choices in the game. The intuition is that the factors that determine choice ought to work in a similar way on all the choices, so there is an efficiency gain in estimation from estimating a single coefficient for each independent variable for each of the choices (unlike, say, an (unordered) multinomial model, where there are separate estimated effects for each choice). In other games there need not be an ordering in available stage game choices but, as we have seen, there are Pareto-ranked equilibria in the minimum effort game and there is an evident monotonic structure in the choices. There are seven possible ‘responses’ an individual can make in each period in the minimum effort game. The choice of $n = \{1, \dots, 7\}$, are coded as n . Separate ‘cut-points’, which are essentially separate constant terms for each choice, are estimated. The cut-point is an (estimated) threshold parameter between choosing one of two adjacent choices.³ We use period (round of a game), generation of a treatment, sex (‘0’ if male, 1 otherwise) and current cash balance of a subject, as well as coefficients on dummy variables for the different game treatments, as independent variables.⁴

Table 4 contains estimates from the ordered probit procedure and Tables 5 and 6 contain the fitted (predicted) values for the probabilities of each choice (1 to 7). The estimated coefficients in Table 4 indicate the direction of the effect of a variable on the estimated choice probabilities but are difficult to quantify, in terms of the magnitude of the effect on the estimated probabilities of choice. The actual probabilities are calculated via a progressive decumulation of the standard normal distribution, so the estimated probability is a nonlinear transformation of the linear summation of the independent variables weighted by their estimated coefficients. Many variables are significant, including period, balance and the public advice treatment dummies.

The period variable has a negative coefficient, consistent with the general tendency for subjects to choose lower numbers in later periods. The generation variable is negative but small and insignificant, consistent with the fact that, in a given treatment, there is generally no trend in choice behaviour over generations. Instead, there is a

³ See Wooldridge (2002), for example, for details of the ordered probit model.

⁴ We used the REOPROB procedure, an ‘Ado’ procedure in *Stata* authored by Guillaume Frechette.

Table 4
Ordered Probit Estimates of Choice Behaviour in Block I Games

Dependent Variable: Number chosen {1...7} by a subject in a particular period			
Variable	Coefficient	z-statistic	P > z
Period	-0.53	-13.44	0.00
Generation	-0.04	-0.63	0.53
Sex ('0' for male and '1' for female)	-0.10	-0.61	0.54
Balance	0.36	6.67	0.00
No Advice Treatment	-0.27	0.90	0.37
Advice Only Treatment	-0.56	-1.72	0.09
Advice-Plus-History Treatment	0.07	0.19	0.85
Public Advice Not Read Aloud Treatment	0.75	2.39	0.02
Public Advice Read Aloud Treatment	2.96	5.89	0.00
Cut1	-2.08	-6.92	0.00
Cut2	-1.46	-4.90	0.00
Cut3	-1.13	-3.81	0.00
Cut4	-0.54	-1.84	0.07
Cut5	-0.19	-0.65	0.51
Cut6	0.19	0.64	0.53
<i>Rho</i>	0.55	15.73	0.00
Likelihood Ratio $\chi^2(9) = 848.19$	Prob > $\chi^2 = 0.00$	N = 1860	Log Likelihood = -2086.81

Table 5
Estimated Probabilities of Choosing a Particular Number in Block I Games
 (at treatment averages of other variables)

Choice	Progenitor Generation	No Advice	Advice Only	Advice-Plus-History	Public Advice Almost Common Knowledge	Public Advice Common Knowledge
1	0.56	0.65	0.76	0.49	0.22	0.00
2	0.22	0.20	0.15	0.24	0.22	0.00
3	0.08	0.07	0.04	0.10	0.13	0.00
4	0.09	0.07	0.04	0.11	0.21	0.01
5	0.03	0.02	0.01	0.03	0.09	0.01
6	0.01	0.01	0.00	0.02	0.07	0.02
7	0.01	0.00	0.00	0.01	0.07	0.96

similar pattern within each generation (typically more high choices in early periods, lower choices in later periods).⁵

The fitted values in Tables 5 and 6 provide a more intuitive picture of the experiment. Table 5 contains the estimated choice probabilities with all variables except the treatment dummies held at the treatment average. Thus, these will most closely replicate the raw data averages. Table 6 contains estimates with variables held at the full sample averages. The estimated probabilities are somewhat less dramatically different

⁵ We tried a variable for the generation interacted with the treatment, in order to capture treatment-specific generational effects, but found that we were clearly over-fitting the model. The likelihood improved only modestly (using the Akaike Information Criterion, we reject the model with cross effects) and the parameter estimates became unstable and difficult to interpret, though the estimated choice probabilities were largely unaffected.

Table 6
Estimated Probabilities of Choosing a Particular Number in Block I Games
 (at full sample averages of other variables)

Choice	Progenitor Generation	No Advice	Advice Only	Advice-Plus- History	Public Advice Almost Common Knowledge	Public Advice Common Knowledge
1	0.45	0.56	0.67	0.42	0.19	0.00
2	0.24	0.22	0.19	0.24	0.21	0.01
3	0.10	0.08	0.06	0.11	0.13	0.01
4	0.13	0.09	0.06	0.13	0.22	0.04
5	0.04	0.03	0.01	0.04	0.10	0.06
6	0.02	0.01	0.01	0.03	0.07	0.09
7	0.02	0.01	0.00	0.02	0.08	0.79

here (e.g., the probability of choosing 7 is 0.79 in the Public Advice Common Knowledge treatment in Table 6 vs. 0.96 in the same treatment in Table 5) across the games, probably due to the lower average balance subjects held in the games where advice was not read aloud. Overall, however, the estimation exercise contains no great surprises, confirming what one sees in the raw data. There is a large and significant effect when advice is read aloud, dramatically increasing the probability of choosing 7.

One possible concern here is that we should also take account of potential dependencies arising from the fact that individuals in a given generation play together for 10 consecutive periods and thus their actions may become correlated over time.⁶ To account more formally for generational dependencies in the error structure of the model, in addition to the individual-specific dependencies that the random effects structure addresses would evidently be very difficult and unlikely to yield a large improvement in the present context, simply because the generational effects do not loom large. We take some assurance from the finding of Robinson (1982) that dependencies in discrete choice model estimation are not such a critical problem, in that the maximum likelihood estimator is still consistent, as it is with independent observations, under less stringent assumptions. Robinson notes the extreme computational difficulties associated with formally modelling dependencies in the discrete choice context.

There is, in fact, very little interaction, as such, among players in a given generation. At the outset they have, depending on the treatment, either a single piece of advice from one player in the previous generation, different for each member of the group, or else the same collection of advice as all the other players in the group. The only feedback they receive after making a choice in any period is the minimum of all choices made by players in that group (generation) for that period. They do not learn, for example, the full set of choices made by the members of a group. So we believe the individual-specific component is the most important thing to account for in the error structure.

Presumably the concern is that what we are interpreting as a treatment effect (i.e., the much higher propensity of subjects in the public advice-read aloud treatment to choose 7) is possibly an artefact of a situation in which some of the groups of players who composed the generations of that treatment, by happenstance, end up 'getting along' and choosing 7 a lot. To address this concern we estimated an alternative version to the

⁶ We would like to thank an anonymous referee for bringing this concern to our attention.

model in Table 4 in which we include dummy variables for each group (i.e., each generation of a treatment), rather than dummy variables for the different treatments as a whole. If our finding of a larger treatment effect for the public advice-read aloud treatment is robust, then the dummy variables for those generations should, in general, be larger than the dummy variables associated with other treatments. We do not simply want to average the generation dummies by treatment and compare them, as this is essentially what the treatment dummies in Table 4 give us. Instead, we treat the collection of generation dummies for a given treatment as a random sample from a distribution and then test whether the distributions of generation dummies differ between treatments.

A Kruskal-Wallis test of the null hypothesis that the samples of dummy variable coefficients across treatments are from the same distribution is rejected (χ^2 (3 d.f.) = 9.643, $p = 0.02$). Pairwise rank-sum tests between the sample of coefficients from the public advice-read aloud (common knowledge) treatment and all other treatments all show that the 'read aloud' coefficients come from a distribution that lies to the right of the distributions from which the coefficients for other treatments are drawn. For instance using pairwise rank-sum tests we can reject the null hypothesis that the dummies for the public advice read aloud (common knowledge) treatment are drawn from the same distribution as

- (1) the dummies in the public advice not read aloud (almost common knowledge) treatment ($z = 2.24$, $p = 0.03$) or
- (2) the dummies in the advice-plus-history treatment ($z = 2.23$, $p = 0.02$).

We interpret this as confirmation of the findings from the model reported in Table 4.⁷

Thus while our experimental agents were finally capable of 'talking themselves to efficiency', the process was much harder than we expected and occurred only when advice was public and common knowledge.^{8,9}

⁷ Because the alternative model with the treatment dummies essentially reinforces the results presented in Table 4 we have not reported those results here. We have also not provided the detailed results of the pairwise rank-sum tests. These results are available from the authors upon request.

⁸ Note that including advice in our experiments is different from either including cheap talk or allowing free communication among decision makers, both of which have been known to increase efficiency. See Cooper *et al.* (1989, 1992), Dawes *et al.* (1977), Isaac and Walker (1988), Charness (2000), Burton and Sefton (2004) and Blume and Ortmann (2007). Cheap talk statements are public, non-binding and payoff irrelevant statements made by the players who are actually going to play the game and not their predecessors. In contrast, except for the public advice treatment, our advice statements are private and made by predecessors. Even when we made advice public and common knowledge, these statements are still not made by the people who are about to play the game. Our advice treatments are different from the communication treatments found in public goods experiments since we only permit one-sided statements to be made and not bilateral or multilateral non-binding discussions. Blume and Ortmann (2007) find that such costless messages do increase efficiency in the minimum effort game but subjects still find it difficult to coordinate to the Pareto-dominant outcome. These authors conduct 8 sessions of the minimum effort game using the same payoff matrix as in the VBB study with 9 subjects in each session. Out of 8 sessions there is only one session where all subjects choose 7 for all 8 periods. (This is because a choice of 7 by all members of the group will yield a payoff of \$1.30 to each subject in any period. Over 8 periods this would amount to a total earning of \$10.40. This was true of only one session – session M8Min.) The other sessions achieved various degrees of cooperation but none of them achieved coordination at the Pareto-dominant outcome for all 8 periods.

⁹ To demonstrate that the result that common knowledge of advice facilitates coordination to the payoff dominant outcome is not caused by the fact that in our study a professor reads the advice out loud Antonopoulos *et al.* (2002), in a follow-up paper to this one, replicates this same discontinuity in behaviour when advice is common knowledge and when it is not, using only undergraduate students as experimental administrators.

2.2. Block II Games – Minimum Choices

Because each generation is free to give advice to the next, our experimental design in Block I does not control for advice quality and thus we cannot say precisely how the way that advice is distributed might affect efficiency. In Block II we isolate the advice quality variable in order to study the interaction between choices and how advice is distributed. To investigate this we ran an additional set of 6 games (see the experimental design in Table 2(b)). Note that, here, our design is different from Block I since here we repeatedly bring in new groups and present them with the same advice distributed in a fixed manner. This allows us to compare behaviour directly in groups that received the same advice distributed in different ways, and to see how advice quality matters for a given way of distributing advice. Whereas the Block I games essentially provide us with a single instance of each treatment (albeit with several related generations), the Block II games allow us to collect repeated observations for each ‘advice quality’-‘information condition’ combination.

More precisely, each new group of subjects in the Block II games was given one of the two sets of advice contained in Table 7. The ‘good’ advice was written by Generation 3 of the public advice game (where advice was almost common knowledge) and failed to spur cooperation in Generation 4, while the ‘very good’ advice was written by Generation 7 after we changed to common knowledge and led to consistent coordination in Generation 8.

Note that while both sets of advice strongly encourage actions that lead to efficient play, the second set (on the right-hand side of Table 7) is unequivocal both on what rule to use in playing the game and on what to do in the first round. All subjects in the ‘very good’ advice category urge their successors to choose seven. However in the ‘good’ advice category (on the left-hand side of Table 7) two people equivocate. The first subject says to start with 6 in round 1 and then go down to 5 in round 2 while subject number 6 does not give clear advice as to what to do in the first round. Still, all other subjects unambiguously urge choosing 7.

In order to quantify the quality of advice offered by our subjects we took the written advice and coded them according to what we think they implied about the suggested course of action for the first round of any generation’s life. These codes are listed at the bottom of each column in Table 7. For example, if a subject said, ‘Choose 7 in round 1 and then choose the round $t - 1$ minimum in round t ’, we coded this with a 7 since it indicated that in the first round 7 should be chosen and then the subject should see what happens and then follow the minimum thereafter. Such advice was fairly typical in that most advice messages suggested dynamic rules for subjects to follow. The problem is that in certain cases they did not specify what action to take in round 1. For example, a subject may write, ‘Choose the minimum of last round’. This rule is well specified for all rounds but the first. In other cases a subject may write, ‘Choose pretty high in round 1 and then follow the minimum of last round’. To consistently code this data we took all statements that offered no advice for the first round and initially coded them with a zero. We then imputed a value of 5.5 to all advice suggesting a ‘high’ action in round 1, i.e. advice that did not specify a number but implied a range of high numbers. Similarly, we imputed a value of 2.5 to all advice suggesting a ‘low’ action in round 1 without being specific. (5.5 is the midpoint of 4, 5, 6 and 7, while 2.5 is the midpoint of 1, 2, 3

Table 7
The Exact Advice Given to the Subjects in the Block II Games

Subject	Good Advice	Very Good Advice
1	Be ahead of everybody, start with a 6, then go down to 5, etc. You will be able to make the most money that way.	Pick 7 every time, EVERY TIME. If everyone picks 7 every time, everyone will make the max per round \$1.30×10 \$13.00), plus you can make the full \$1.28 for each of the predictions rounds. Don't be stupid. Pick 7. Honestly, you're here for the money anyway, right?
2	Choose #7. Don't be tempted to deviate but everyone must choose #7.	If you don't start the first round with '7' then the pattern thereafter will be '7' or lower. Bottom line – you must begin the first period with a '7'...Or else!!!!
3	True, if everyone selects #7, you have max profit. But when you see the smallest # move down, you should follow.	Pick 7 for crying out loud! But if there is a weirdo who picks lower, pick that number too. Pick 7!!! Trust each other it will help you too!
4	If <i>everyone</i> continues to pick 7 you will maximise your profit. Anything else and profit maximisation is not possible.	For the first round, you must trust the other participants & choose 7. Choosing 7 gives the maximum payoff. The (<i>sic</i>) adjust your choice by following the trend after the first round. Be consistent!
5	Start with 7! Everyone agree at least once. Once someone starts using <i>one</i> join them.	It would be best for everyone to choose 7 each time. However, if one person consistently chooses a lower number, you will make more profitably conforming to them.
6	Follow the trend. Ideally you want to maximise at 7 but inevitable someone doesn't get it.	Picking 7 will yield the maximum payoff pick 6 if everyone picks 7. So start out picking 7, however, some people are very untrusting and will or 5 – if this happens, 6, follow the trends, if everyone starts picking start picking that also.
7	Stay with 7 unless someone won't use that number. If they insist on a lower number go with it.	Chose 7 & and hope everyone else does. But it is important to follow any trends you notice.
8	Pick #7 until after it is apparent that the number declines and then follow it down.	The thrill of not choosing 7 leads only to a smaller payoff than both you and everyone else could earn.
Advice coded as	6, 7, 7, 7, 7, 0, 7, 7	7, 7, 7, 7, 7, 7, 7, 7

and 4, so this coding imputes values in a symmetric fashion.) Pieces of advice that were nonsensical or that otherwise made no reference to specific choices or range of choices were left coded with a zero.

Using this coding convention we coded our 'good' advice as (6, 7, 7, 7, 7, 0, 7, 7) and our 'very good' advice as (7, 7, 7, 7, 7, 7, 7, 7). Again, we stress that to most people the good advice listed above would be considered a strong statement in favour of coordination. However, since this is a coordination game subjects need to know more than that in order to cooperate. They need to know that everyone else knows that everyone else considers this set of advice to be sufficiently strong to act cooperatively and this must be considered common knowledge.

What our Block-II experiments do then is to hold advice constant and ask whether the way this advice is presented to subjects, the information condition, makes any

difference in their ability to coordinate on efficient actions. While even the good advice is quite strongly supportive of such coordination, there are differential degrees of success at coordination over different information conditions for the two sets of advice.

Table 8 and Figures 2(a) and 2(b) present the results of the Block II games. Table 8 presents the minima of each group from round 1 to 10. Figures 2(a) and 2(b) present the minimum chosen in each round for the 10 round lifetime of a group. Figures 2(a) and 2(b) are the Block II analogs of Figure 1 which presents similar information for the Block I games except that here we do not have generations since each group brought into the laboratory receives the same advice. So we have labelled them as ‘Groups’ rather than ‘Generations’. Moreover any temporal sequence of these results is arbitrary since there is no inter-generational structure connecting the groups.

Looking at Figure 2 one can see that behaviour differs dramatically as we vary the quality of advice offered to subjects. More precisely, when subjects are offered good

Table 8
Observed Minimum Choices in Block II Games

	Rounds									
	1	2	3	4	5	6	7	8	9	10
Good Advice										
Game 6: Common Knowledge of Advice and ‘Good’ Advice										
Group 1	7	7	7	7	7	7	7	7	7	7
Group 2	5	1	1	1	1	1	1	1	1	1
Group 3	4	1	1	1	1	1	1	1	1	1
Group 4	7	1	1	1	1	1	1	1	1	1
Game 7: Almost Common Knowledge of Advice and ‘Good’ Advice										
Group 1	1	1	1	1	1	1	1	1	1	1
Group 2	1	1	1	1	1	1	1	1	1	1
Group 3	7	7	7	7	7	7	7	7	7	7
Group 4	1	1	1	1	1	1	1	1	1	1
Game 8: Advice on Overhead Projector for all subjects to see and ‘Good’ Advice										
Group 1	1	2	3	4	1	1	1	1	1	1
Group 2	4	1	1	1	1	1	1	1	1	1
Group 3	7	7	7	7	7	2	1	1	1	1
Very Good Advice										
Game 9: Common knowledge of Advice and ‘Very Good’ Advice										
Group 1	7	7	7	7	7	7	7	7	7	7
Group 2	7	7	7	7	7	7	7	7	7	7
Group 3	7	7	7	7	7	7	7	7	7	7
Group 4	5	4	1	1	1	1	1	1	1	1
Game 10: Almost Common Knowledge of Advice and ‘Very Good’ Advice										
Group 1	7	7	7	7	7	7	7	7	7	4
Group 2	7	7	7	7	7	7	7	7	7	1
Group 3	7	7	7	7	7	7	7	7	7	7
Group 4	7	7	7	7	7	7	7	7	7	7
Game 11: Private Knowledge of Advice and ‘Very Good’ Advice										
Group 1	4	4	4	4	4	4	4	4	4	4
Group 2	7	5	3	1	1	1	1	1	1	1
Group 3	3	2	1	1	1	1	1	1	1	1




Fig. 2. (a) Behaviour of the Minimum Across Block II Games with 'Good' Advice and (b) Behaviour of the Minimum Across Block II Games with 'Very Good' Advice

advice it appears that no matter how it is presented it is not sufficient to get them to cooperate over much of the horizon of the experiment. Looking at Figure 2(a) we find that when we have 'good' advice the preponderant choice for the minimum is 1. Only one group in the almost common knowledge game and one group in the common knowledge game manage to sustain a minimum choice of 7 for all 10 rounds. When the advice is 'very good', on the other hand, all four groups in the almost common knowledge game manage to achieve efficiency for all 10 rounds and 3 out of the 4 common knowledge groups manage to do so as well.

The Block I results imply that in order to achieve coordination at the payoff dominant outcome, one must not only present subjects with strong advice but one must also give it to them in a way that makes the advice 'common knowledge'. Our Block II results

suggest that in order to achieve such cooperation, the advice being passed on must be unanimous and unequivocal in exhorting successors to choose 7. Once the advice meets this unanimity threshold it does not matter whether it is distributed in a way that makes this advice ‘common’ or ‘almost common’ knowledge as long as the advice is made public. However when the advice quality falls below that threshold – i.e. advice is ‘good’ but not good enough – with a small number of subjects offering equivocal advice regarding the optimal choice – agents fail to coordinate no matter how that advice is given to them – either as ‘common’ or ‘almost common’ knowledge. Thus it would appear at first glance that our Block II results minimise the importance of the ‘common knowledge’ aspect of advice.

However upon closer inspection we find that there is no contradiction here. Our results from Block I and Block II indicate that behaviour in this game is extremely path dependent with the choices made (and the minimum chosen) in subsequent rounds being crucially dependent on the first round minimum. For example, there is never a case where the tenth round minimum winds up being 7 when 7 was not the minimum of the first round. In addition, if the first round minimum is 1 then the tenth-round minimum is certain to be 1. Put differently, choosing 7 in round 1 is a necessary condition for a round 10 minimum of 7. When the first round choice is neither 1 nor 7, the evidence is less clear. While it is highly likely that in this case the tenth round minimum will be 1, there are several cases where a group winds up at an intermediate level of cooperation in round 10 having started with an intermediate minimum in round 1.

Since round 1 is so crucial to achieving coordination we want to focus on it more intensively here and perform a simulation. In each of the Block II games we had either three or four groups of 8 subjects (i.e. either 24 or 32 subjects per game). Note, however, that while group effects (and the choices made in round 1) may affect the choices in later rounds, when subjects enter the laboratory for the first time and before they make their round 1 choice, the group that they are randomly assigned to is of no consequence. Hence, all first round choices are independent of each other. This being the case, if a game has 32 subjects in it, then we are free to form all possible groups of 8 from those 32 subjects (and not just the four groups these subjects were randomly assigned to) and then look at what the minimum would have been if those groups had formed given their first-round choices in Block II. The minimum chosen in the first round of all the possible hypothetical groups provide us with clues about the path of choices in subsequent rounds. If these hypothetical first round minima differ, we can expect that the tenth-round behaviour evolving from them would be different as well.

We form a ‘hypothetical’ data set with either $32!/24!8! = 10,518,300$ or $24!/16!8! = 735,471$ hypothetical groupings for each game and see what behaviour would be forthcoming from each of these groups had they formed. The point of this exercise is to be sure that the small number of groups in each game does not mask larger tendencies in the population. This could work in two ways. For example, if there were a large number of players initially choosing small numbers in one of three or four groups of a game but not in the others, this will show up as a lower expected minimum for round 1 in the hypothetical data set than is evident from the actual groupings observed. Alternatively, if there are only one or two players choosing small

numbers in another game, then there will be many more hypothetical groups with higher minima than in the actual groupings observed. There is too much interdependence to be able to do formal statistical tests with this hypothetical data but the exercise is useful as a straightforward calculation of what the full set of possible outcomes, in fact, is.

Further, if we can establish that holding advice constant at its good-quality level (i.e. below the threshold) increases cooperation only when advice is distributed as common knowledge, then we would have reinforced our Block I results. We need not be concerned about the very good advice since when the advice is very strong and above the unanimity threshold it leads to cooperation no matter how it is distributed, provided it is distributed publicly. We do this and present the results in Figure 3 and Table 9.

As we see in Table 9 and Figure 3, holding advice constant at the good quality level, the manner of information transmission matters. First notice (from rows 4 and 5 of Table 9) that in the almost common knowledge and overhead projector games with ‘good’ advice, 59% and 56.5% of the hypothetical first round minima would have been 1.

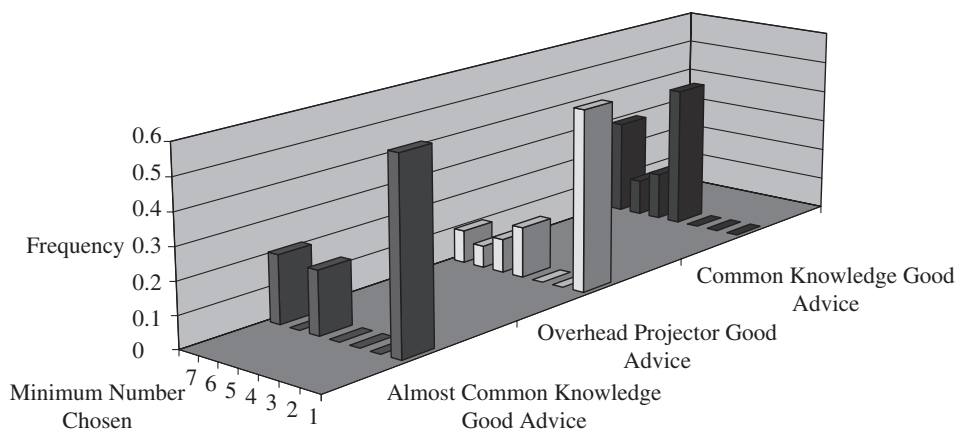


Fig. 3. Probability of Observing a Particular Minimum Value in Block II Games with ‘Good’ Advice using ‘Hypothetical’ Dataset

Table 9

Probability of Observing a First Round Minimum of 1, 2, 3, 4, 5, 6 or 7 Across the Various Games in Block II

Game	Minimum						
	1	2	3	4	5	6	7
Private Knowledge and ‘Very Good’ Advice	0	0	0.33	0.39	0.105	0.14	0.03
Almost Common Knowledge and ‘Very Good’ Advice	0	0	0	0	0	0	1
Common Knowledge and ‘Very Good’ Advice	0	0	0	0	0.44	0	0.56
Almost Common Knowledge and ‘Good’ Advice	0.59	0	0	0	0.20	0	0.21
Advice on Overheads and ‘Good’ Advice	0.565	0	0	0.16	0.105	0.07	0.10
Common Knowledge and ‘Good’ Advice	0	0	0	0.44	0.15	0.11	0.295

However *none* of the first round minima would have been 1 (last row of Table 9) in the common knowledge good advice game. Since round 1 minima of 1 are certain indicators of round 10 minima of 1, this means that almost 60% of the groups in either the overhead projector or almost common knowledge game would have spent their entire 10-round life choosing 1. While approximately 30% of first round minima would have been 7 in the common knowledge good-advice game only 21% would have been so in the almost common knowledge good-advice game, and 10% in the Overhead Projector game. Finally while 100% of the first round minima in the common knowledge good advice game would be 4 or more (last row of Table 9), only 41% of the first round minima would have been 4 or more in the almost common knowledge good-advice game (Row 5 of Table 9). Figure 3 highlights the fact that with common knowledge and ‘good’ advice the probability of observing a first round minimum of 1 is zero.

The effect of the information treatment is less dramatic for very good advice. Even for the private advice game with very good advice the probability of a first round minimum less than 3 is zero. At the same time, the probability of a first round minimum of 7 with private advice is essentially zero as well (Row 2 of Table 9), so this game seems doomed to suboptimal outcomes. Both the almost common knowledge and the common knowledge games with very good advice have zero probability of a first round minimum less than 5. We do not illustrate these probabilities for the ‘very good’ advice games as the differences among games are less dramatic.

Projecting the ‘good’ advice on a screen, in the overhead projection game, does not result in an average minimum different from the almost common knowledge game. Projecting the advice on a screen for all to see evidently does not induce the same degree of common knowledge as reading the advice out loud for all to hear. In short, information affects first round choices holding advice constant.

These results lead to the following conclusions:

Conclusion 1:

On the basis of our Block II results, we can say that if advice is strong enough (with all subjects in a group strongly exhorting their successors to choose 7) then efficient coordination is achieved regardless of the manner in which the advice is distributed, as long as it is public.

Conclusion 2:

Also on the basis of our Block II results we find that when advice is insufficiently strong, then efficient coordination is likely to be established only if that advice is distributed as common knowledge.

Of course there is no guarantee that *any* kind of advice will lead to efficient coordination as long as it is delivered in a way that ensures common knowledge. We have studied the effect of two distinct sets of advice, both generally strong in encouraging efficient actions, but one distinctly stronger than the other, and have uncovered a differential effect of the information condition. Presumably truly bad advice, uniformly encouraging subjects to make choices of 1 in the game, would not be likely to lead to efficient coordination. But we have found that, in practice, the quality of advice is endogenous to the information condition as in the Block I games, so our conclusions

are relevant for the generally strong advice that is left in those games with high degrees of common knowledge.^{10,11}

3. Beliefs and Common p-beliefs

The problem faced by our subjects is simple. When they hear a piece of advice offered in either one of our public advice treatments, (the issues are different for private advice) they must ask themselves whether they believe that all the others heard what they heard and also if they interpreted it to mean what they did. For example, if it was announced that ‘you all should choose 7’ whether you will actually do so depends on whether you think the other seven people heard what was said, i.e., no one was day dreaming, and whether they all took this statement to be strong enough to lead them to choose seven. You will do so if you believe with sufficiently high probability, p , that all others believe with sufficiently high probability that that statement means choose 7. For example, I may not think the statement was strong enough since I would only choose 7 if I heard ‘you should *really* all choose 7’. For me the insertion of the word ‘really’ is crucial. So what is not common knowledge across subjects is whether they are all speaking the ‘same language’ i.e., interpreting statements to have the same meaning. However, if we all share a strong enough common belief (a high enough p-belief) that these statements will be interpreted identically, then that common p-belief will be enough to get us to coordinate. This is what we think public and strongly exhortative advice achieves.

Monderer and Samet (1989, p. 170) motivate their paper by posing a similar problem. They state the following:

¹⁰ Chaudhuri *et al.* (2006) study a voluntary contributions mechanism using a very similar inter-generational paradigm with advice and also find that contributions are significantly higher and closer to the social optimum when the advice from a previous generation is made public and common knowledge.

¹¹ Van Huyck *et al.* (1992) study the role of credible assignments in a two-person coordination game with three Pareto-ranked equilibria. There is an external arbiter who instructs the participants to choose a particular strategy in the game. They find that when the arbiter instructs participants to choose a strategy commensurate with the payoff dominant outcome, 98% of the pairs playing the game do coordinate to that outcome. It is possible to think of advice playing a role similar to such external assignments in our study. There are a number of ways in which our results generalise those reported by Van Huyck *et al.* First, the external assignments in Van Huyck *et al.* are ‘common knowledge’ announcements since the arbiter’s message is projected on the wall of the laboratory and also read out loud. We have shown in our Block I games that even minor departures from such a common and shared comprehension of the message (especially if coupled with small equivocations in the advice) can lead to seriously sub-optimal outcomes. Second, Van Huyck *et al.* find that when the arbiter instructs subjects to choose a strategy that conflicts with payoff dominance, such assignments are *not* credible in that 25% or fewer pairs actually coordinate to the assigned outcome and the majority of subjects continue to choose the strategy that leads to the payoff dominant outcome. We have shown that when announcements in the form of advice are less than common knowledge, assignments to Pareto-dominated outcomes are quite credible. Thus in the presence of strategic uncertainty the prospects of sub-optimal outcomes are actually more likely than the Van Huyck *et al.* results suggest. Finally, the arbiter in the Van Huyck *et al.* study is very much in the nature of a *deus ex machina*, whereas in our study the strongly exhortative advice that leads to efficiency arises endogenously as a function of the information conditions implemented. In a related study Brandts and Cooper (2006) analyse the ‘corporate turn-around game’ which is a modified version of the Van Huyck *et al.* minimum effort game. These authors look at the role and impact of financial incentives in the form of a bonus in resolving coordination failures and report that such incentives are indeed successful in fostering successful coordination. What is surprising is that the magnitude of the bonus does not seem to matter in that larger bonuses were no more successful in enhancing coordination than smaller bonuses.

The most frequent cases of common knowledge are public announcements. Consider, for example, an auction. Once the auctioneer has publicly announced a price, it is assumed by most suppliers of auction models to be common knowledge to the participants of the auction. But is it really? One should always allow for some small probability that a participant was absent-minded or deaf at the time of the announcement. No matter how small that probability is, the price is not common knowledge. Nobody knows for sure that the others know the price. We cannot even build the first story of this formidable tower of hierarchies of knowledge, ‘I believe that you believe that I believe. . .,’ that is required for common knowledge.

Later in their paper (pp. 176–7) they expand on the same example as follows:

To illustrate the notions of evident p -belief and common p -belief let us reconsider auctions. Let C be the event ‘the price of the picture is \$1000.’ Let E be the event ‘the auctioneer announces that the price of the picture is \$1000.’ If everyone must hear the announcements of the auctioneer, then E is evident knowledge. That is, whenever E occurs everyone knows that it has occurred. However, if there is some positive (possibly small) probability ε that not all the audience are hearing, then E is not evident knowledge. Moreover, E may not be evident p -belief for high p . This is the case, for example, if E occurs and one of the agents who assigns a low probability to C does not hear the announcement. However, if F is the event ‘the audience are all hearing, and the auctioneer announces that the price of the picture is \$1000,’ then F is evident $(1 - \varepsilon)$ -belief. Therefore, C is common $(1 - \varepsilon)$ -belief at each state of the world in which F occurs.

While we do not model these ideas formally in this article, in this Section we try to give some evidence that subjects’ beliefs are altered by the quality of advice and the manner in which it is disseminated. While the beliefs we discuss below are not *common beliefs* or common p -beliefs (such beliefs are unobservable) the fact that people’s actions are sensitive to the beliefs discussed below implies strongly that when advice is public and strong people have enough confidence (share a strong enough common belief) to take the risky action of choosing 7.

With this in mind we now turn to the analysis of an aspect of strategic interaction that one generally is not able to observe: the beliefs of players in the game about what is going to happen in the game. As we collected information on the beliefs of players (before the first and tenth rounds of play) after they received advice, we are able to say something about this. One of the main avenues for advice to influence behaviour is through beliefs. In this experiment we have only measured first-order beliefs (i.e., point estimates) but they give us insight into the more significant higher order beliefs that are essential to get cooperative behaviour. We only analyse the beliefs elicited prior to the first round of play, as it is at this juncture that we may be able to detect the influence of advice on belief. The tenth round beliefs are highly consistent with the actual play of the game.

Prior to round 1 of the actual game, we asked each subject to report the distribution of actions that they believed would result in the subsequent round of play. Figures 4(a)

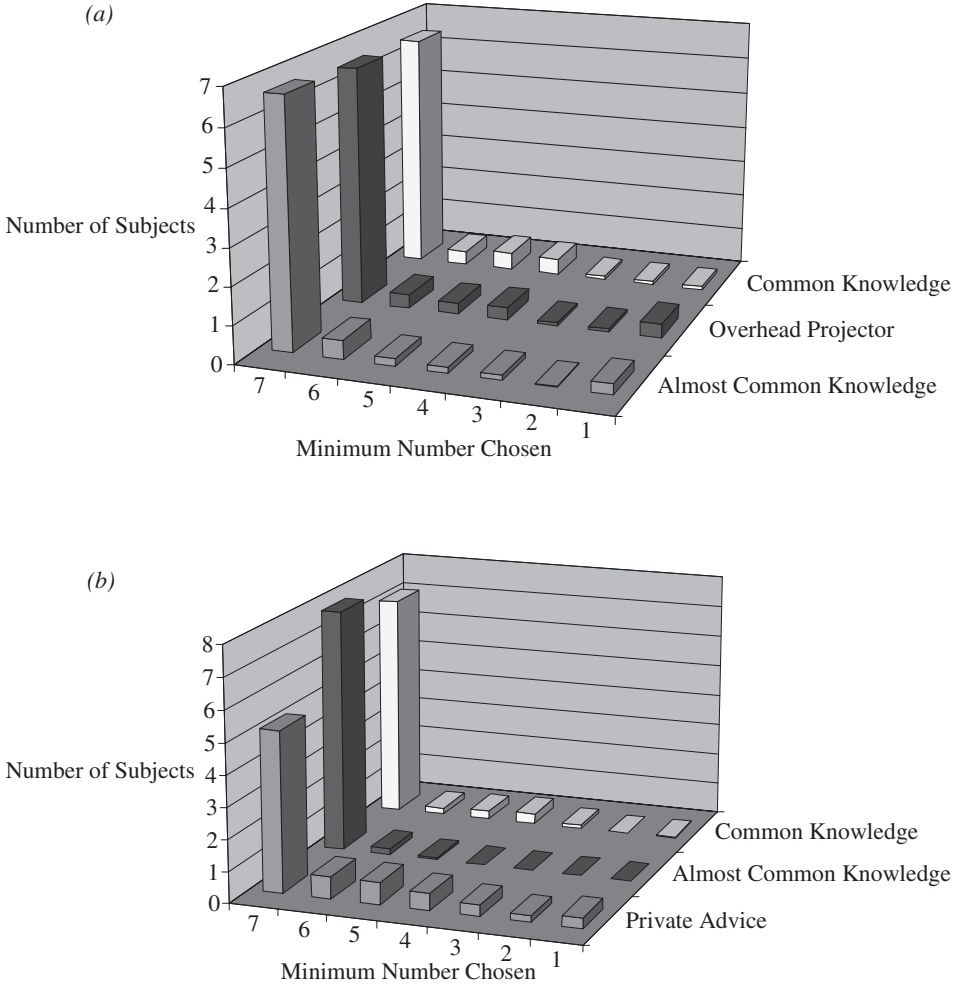


Fig. 4. (a) Comparison of Beliefs Across Block II Games with 'Good' Advice and (b) Comparison of Beliefs Across Block II Games with 'Very Good' Advice

and 4(b) represent these beliefs by presenting the mean number of subjects expected to choose each number between 1 and 7 in each game in Block II with 'good' and 'very good' advice respectively. These diagrams place the various actions 1, 2, ..., 7 along the horizontal axis and along the vertical they display the mean number of subjects (aggregated over all groups) predicted to choose that action.

Note that there is a small but noticeable difference in the beliefs of subjects across advice quality treatments. In the good-advice games (Figure 4(a)) the mean number of subjects expected to choose 7 was 6.63, 6.46, and 6.4 for the Almost Common Knowledge, Overhead Projector and Common Knowledge treatments respectively. In the games with 'very good' advice (Figure 4(b)) looking only at the Common Knowledge and Almost Common Knowledge conditions we find that no one expected anyone to choose 1, and the mean number of subjects expected to

choose 7 was 7.09 and 7.75 for the Common Knowledge and Almost Common Knowledge games respectively. While these differences may seem small, remember that the best response functions of our subjects are very sensitive to even small deviations from unanimous cooperation. Hence, even a small weight placed on low choices by others may spark a sub-optimal choice by a subject in round 1, which is a sure prediction that the group will end its 10-round experience with each other sub-optimally.

To demonstrate how beliefs affect behaviour look at the *private advice game* with ‘very good’ advice in Figure 4(b). Here the aggregate advice is very good yet because of the way it is distributed the beliefs of subjects are relatively pessimistic. On average subjects expect only 5 out of 8 group members to choose 7 in round 1 of the private advice game. Referring to the actions of subjects in this game (Table 8 and Figure 2(b)) we see that these diminished beliefs have a dramatic impact on the actions taken both in round 1 and in succeeding rounds as well.

The lesson to be learned from these Block II belief distributions is simple. Common Knowledge plus Public Advice seems to be a necessary condition for beliefs to be sufficiently positive so as to lead to efficient outcomes. All other settings hold the prospect of someone choosing 1 which leads to a spiral of actions towards 1. Even in the common knowledge version of the games with good advice there is a tiny amount of pessimism – a small amount of belief that someone will choose 1 – and there is greater such pessimism for the almost common knowledge and overhead projector games. These small amounts of pessimism in the less-than common knowledge games seems to have made large differences in behaviour.

We now return to the Block I games to help us answer a different but related set of questions. First we are interested in comparing the beliefs of subjects in our replicator (No-Advice) game to those of subjects in our private and public advice games. We are interested in answering two questions. First, does the existence of advice change the distribution of subjects’ beliefs from what it would be if no advice existed? Second, does advice increase the minimum action upon which there is positive probability placed? The second question is important since if subjects best respond then the best response rule is trivial: choose that action which you think is the minimum to be chosen by your cohort. Hence, if any subject believes that there will be even one other who will choose 1, then 1 is their best response. If advice can raise this expected minimum, it can succeed in raising subjects’ choices.

The answer to both of these questions can be seen in Figure 5 which presents the results game by game. This diagram places the various actions 1, 2, ..., 7 along the horizontal axis and along the vertical axis they display the mean number of subjects (aggregated over all generations) predicted to choose that action. For example, we see that on average in the No-Advice game subjects expected 3.6 people in their group of eight to choose 7, 1 to choose 6 etc.

Looking across these games there are some interesting results. First when comparing the beliefs of subjects in the No Advice game to those in the Private Advice games (Advice Only and Advice Plus History) we see that while there appears to be no significant difference (at the 5% level) between these distributions, private advice does seem to lower expectations in the sense that subjects expect more people to choose 7 in




Fig. 5. Comparison of Beliefs Across Block I Games

the No Advice game (mean = 3.6) than in either the Advice Only game (mean = 2.5) or the Advice Plus History game (mean = 2.6).¹²

In addition, subjects in the private advice games predict at least one person will choose 1 (mean = 1.2 for the Advice Only game and 1.4 for the Advice Plus History game) while in the No Advice game subjects expect an average of only 0.80 subjects will choose 1. Obviously, the advice offered damaged beliefs in the private advice games and led subjects to think that fewer people are likely to choose 7 and more are likely to choose 1.

The results in the Public Advice games are quite different as can be seen from Figure 5. For example, in the Public Advice Common Knowledge (CK) game (fifth set of histograms from the left in Figure 5), we see that on average subjects expected 5.2 people to choose 7 in round 1 and *no one* to choose either 1, 2, or 3. It is only in the Public Advice Common Knowledge (CK) game that we get beliefs which place zero value on someone choosing 1. In all the other games subjects place a positive (albeit small) probability on someone in the group choosing 1. The firm belief that no one will choose a low number in the Public Advice Common Knowledge game seems to give subjects the confidence which allows them to choose 7 and maintain efficiency. In addition, the distribution of beliefs for this game is significantly different from the distribution of beliefs in all other games at the 1% level. When we compare the distribution of beliefs in the Public Advice Common Knowledge (CK) and Almost Common Knowledge (ACK) games, we see that common knowledge has an impact on beliefs. In Figure 5 the fourth set of histograms from the left presents the results of all five generations of the Almost Common Knowledge (ACK) game, the fifth set of histograms presents the beliefs of subjects in the four generations of the Common Knowledge (CK) game and the last (sixth) set presents the beliefs of subjects in the last

¹² The difference between the distributions of beliefs in the No Advice and Advice Plus history treatments is significant at the 10% level using a χ^2 test (χ^2 (7 d.o.f.) = 14.32). A χ^2 test fails to detect any difference in the distributions of beliefs in the No Advice and Advice Only treatments even at the 10% level (χ^2 (7 d.o.f.) = 8.88). See Table 10.

Table 10
 χ^2 Test on Belief Distribution for Block I Games

Treatment	Advice Only	Advice-Plus-History	Public Advice Almost Common Knowledge (ACK) All Generations	Public Advice Almost Common Knowledge (ACK) Generations 3 – 5 only	Public Advice Common Knowledge (CK)
No Advice	8.88	14.32*	11.84	26.91***	34.24***
Advice Only	—	3.84	31.2***	62.45***	64.16***
Advice-Plus-History	—	—	31.84***	63.12***	103.92***
Public Advice Almost Common Knowledge All Generations	—	—	—	—	37.92***
Public Advice Almost Common Knowledge Generations 3 – 5 only	—	—	—	—	31.19***

Note. *** indicates significant at 1%, ** indicates significance at 5% and * indicates significance at 10%.

three generations (Generations 3 to 5) of the Public Advice Almost Common Knowledge game. The first two generations are eliminated here because the advice offered here came initially from the progenitor game and was particularly pessimistic leading the minimum to quickly converge to 1. While the Public Advice Almost Common Knowledge (ACK) game (last three generations only) fostered beliefs that were more like the Common Knowledge game than either of the private advice games, it did not succeed as completely in eliminating the possibility that some subject might choose a low number. As can be seen by the actions taken (Figure 1 and Table 3(a)), even this small bit of suspicion that someone might choose 1 was enough to ruin cooperation.

The lesson learned here is that any doubt mentioned in advice statements may be enough to lead subjects to fear that others will not act cooperatively and hence cause them to deviate. The mapping from advice to actions seems to be a discontinuous one where the discontinuity arises even when the advice set seems to be rather strong in urging cooperation. The mapping from advice to beliefs, however, is more continuous. Small amounts of doubt lead to small probability weight placed on bad outcomes and this leads to big changes in behaviour.

4. Conclusions

This article is motivated by the conjecture that if we allow subjects playing a coordination game with Pareto-ranked equilibria to leave advice for their successors then they would manage, over time, to achieve an efficient outcome. In this sense we think people can ‘talk themselves to efficiency’ through advice. We find something rather different. If the advice offered by one generational agent to his or her descendent is private, in the sense that no other agent can hear it although it is common knowledge that all agents are receiving advice from their predecessors, then just the opposite occurs. Private advice between a predecessor and his successor, no matter how positive, fails to lead to efficient results.

When advice is public we find that subjects act as if there is a quality threshold with advice. If the advice offered to subjects is sufficiently strong in urging them to

cooperate (i.e., above this advice-quality threshold) then as long as that advice is offered in a public manner (either as common knowledge or as what we call ‘almost common knowledge’) we can expect cooperation to follow. However, if the advice quality is below the threshold then advice is far more likely to result in cooperation if it is not only public but also distributed in a manner that makes it common knowledge.

It is useful to distinguish between *establishing and maintaining* efficiency on the one hand, and *improving* efficiency on the other. We find that it is essential that play in the minimum effort game get started at an efficient level if there is to be any hope of efficiency being maintained. To establish efficiency, it seems that subjects have both to get the right advice and to take the right actions. When advice is public, subjects seem to leave better advice for their successors and subjects are more likely to follow this advice. Improving efficiency in the minimum effort game, either within a given generation or group, or over the course of several generations, seems to be almost impossible. Once play is at a suboptimal level (the minimum less than 7), play can be counted on to deteriorate to the worst outcome with a minimum of 1 within a few rounds. Similarly, once a generation of players has failed to establish and maintain efficiency, play in subsequent generations can be counted on to fail to coordinate as well.

We find the difficulty of achieving efficiency under these circumstances surprising but instructive since it indicates that if we expect policy makers (like central bankers) to be able to coordinate a move from a sub-optimal (underemployment) equilibrium to a Pareto improving one, they will need to assure all agents in the economy that all others in the economy are hearing the same message and that fact must be common knowledge. As Chwe (2001) points out, however, we may need to think of how we are going to structure our institutions so as to achieve the common knowledge necessary for the attainment of coordinated action since reaching common knowledge is not an easy task.

Appendix

PLAYER ID # _____

Instructions

(These instructions are for use in the Almost Common Knowledge, Overhead Projector and Common Knowledge Games)

This is an experiment in the economics of market decision making. Various funding agencies including the National Science Foundation, New York University and Rutgers University have provided funds to conduct this research. The instructions are simple. If you follow them closely and make appropriate decisions, you may make an appreciable amount of money. These earnings will be paid to you in cash at the end of the experiment.

You will be in a market with 7 other people. In this experiment there will be a number of periods. In each period every participant will pick a value of X. The values of X you may choose are 1, 2, 3, 4, 5, 6 or 7. The value you pick for X and the smallest value picked for X by any participant, including your choice of X, will determine the payoff you receive.

You are provided with a table which tells you the potential payoffs you may receive. Please look at the table (on Page 2) now. The earnings in each period may be found by looking across from the value you choose on the left hand side of the table and down from the smallest value chosen by any participant from the top of the table. For example if you choose a 4 and the smallest value chosen is a 3 you earn 80 cents that period.

At the beginning of every period each participant will write down the value of X they have chosen on the Record Sheet. (Page 5) The smallest value of X chosen will be announced and each participant will then calculate his/her earnings for that period.

If you will now look at your record sheet you will see the following entries. MARKET PERIOD, BALANCE, YOUR CHOICE OF X, SMALLEST VALUE OF X CHOSEN, and YOUR EARNINGS. In the first period your BALANCE is zero. In the second period your BALANCE is the value of your earnings in the first period. In the third period your BALANCE is the value of your BALANCE in the second period plus the value of your earnings in the second period. And so on. Please keep accurate records throughout the experiment. All payoffs in this experiment are designated in dollars and cents.

Unless you are in the first group to participate in this experiment, when you start the experiment you will receive advice on how to make your decisions from a group of players who participated in the experiment *prior to you*. Each of you will get to see the advice left by all the players in this group prior to you. (See Page 7) So each of you is looking at the exact same set of advice as everybody else. (*In common knowledge treatment add: 'Besides providing you with a sheet with the advice, this advice will also be read aloud by the experimenter.'* *In overhead transparency treatment replace previous sentence with: 'Besides providing you with a sheet with the advice, this advice will also be displayed on the screen in front of the classroom.'*) At the end of this session you will be asked to leave advice to the next group of players in the experiment. Please write your advice on the sheet provided (Page 6). Please write or print legibly.

Each of you is paired with another player, who you do not know and who will participate in the experiment *immediately after you*. You will receive a second payment, equal to the amount that this player, who will participate in the experiment immediately after you, makes in his or her session. You will be told how to collect this second payment after the instructions have been read.

To be sure that everyone understands the instructions please fill out the sheet labelled questions on Page 3 now. If there are any mistakes on any question sheet the experimenter will go over the instructions again. **IF YOU HAVE ANY QUESTIONS PLEASE ASK THEM AT THIS TIME!!!**

PAYOFF TABLE

		Smallest value of X chosen						
		7	6	5	4	3	2	1
Your choice of X	7	1.30	1.10	.90	.70	.50	.30	.10
	6	—	1.20	1.00	.80	.60	.40	.20
	5	—	—	1.10	.90	.70	.50	.30
	4	—	—	—	1.00	.80	.60	.40
	3	—	—	—	—	.90	.70	.50
	2	—	—	—	—	—	.80	.60
	1	—	—	—	—	—	—	.70

QUESTIONS

Please look at your payoff table and fill in the following blanks.

Your choice of X	The smallest value of X chosen	Your earnings
4	2	_____
2	2	_____
5	5	_____
6	4	_____

EXTRA INSTRUCTIONS

Occasionally you will be asked to predict what every participant will choose for X. When you are asked to do so, please write down your prediction of how many people will pick 7, 6, 5, 4, 3, 2, and 1. When you add your predictions of the number of people that will pick each value i.e. 7, 6, 5, 4, 3, 2, 1, they should add to 8.

You will be paid for each of your correct predictions as follows. Your earnings will equal 128 cents less the sum of squared differences between your predictions and the actual choices.

EXAMPLES: Suppose that 8 people each had a red ball and a blue ball, and that they were all asked to put one and only one of the balls into an urn. At the same time they each were asked to predict the number of red balls and the number of blue balls that would end up in the urn. With a payment rule like that above they would find their earnings as follows:

Diff	Predict	Actual	Sq. Diff	Predict	Actual	Sq. Diff
Blue	8	0	64	8	8	0
Red	0	8	64	0	0	0
			Total 128			Total 0
		128-128 = 0			128-0 = 128	

Diff	Predict	Actual	Sq. Diff	Predict	Actual	Sq. Diff
Blue	4	4	0	6	2	16
Red	4	4	0	2	6	16
			Total 0			Total 32
		128-0 = 128			128-32 = 96	

You will be told the actual choices made for the periods you were asked to make predictions at the end of the experiment.

IF YOU HAVE ANY QUESTIONS PLEASE ASK THEM NOW!!!

RECORD SHEET

SEX M F

PARTICIPANT #

circle one

MARKET PERIOD	1	2	3	4	5	6	7	8	9	10	11
---------------	---	---	---	---	---	---	---	---	---	----	----

BALANCE 0

YOUR CHOICE OF X

SMALLEST VALUE OF X CHOSEN

YOUR EARNINGS

Period	Predict	Actual	Sq.Diff	Period	Predict	Actual	Sq.Diff	Period	Predict	Actual	Sq.Diff
7	—	—	—	7	—	—	—	7	—	—	—
6	—	—	—	6	—	—	—	6	—	—	—
5	—	—	—	5	—	—	—	5	—	—	—
4	—	—	—	4	—	—	—	4	—	—	—
3	—	—	—	3	—	—	—	3	—	—	—
2	—	—	—	2	—	—	—	2	—	—	—
1	—	—	—	1	—	—	—	1	—	—	—
	Total	—	—	Total	—	—	—	Total	—	—	—
	128	- =	—	128	- =	—	—	128	- =	—	—

University of Auckland
 New York University
 Rutgers University

Submitted: 12 April 2006

Accepted: 22 February 2008

References

- Anderlini L. and Lagunoff R. (2005). 'Communication in dynastic repeated games: "whitewashes" and "coverups"', *Economic Theory*, 26(2) (August), pp. 265–99.
- Anderlini, L., Gerardi, D. and Lagunoff, R. (2007). 'A "super" folk theorem for dynastic repeated games', Working Paper, Georgetown University, January.
- Antonopoulos, C., Chau, L., Chaudhuri, A., Hahn, E. and Min, E. (2002). 'Efficiency breakthrough: conversion from almost common to common knowledge in the minimum effort game', Working Paper, Department of Economics, Wellesley College, May.
- Blume, A. and Ortmann, A. (2007). 'The effects of costless pre-play communication: experimental evidence from games with Pareto-ranked equilibria', *Journal of Economic Theory*, vol. 132(1) (January), pp. 274–90.
- Brandts, J. and D. Cooper (2006). 'A change would do you good... an experimental study on how to overcome coordination failure in organizations', *American Economic Review*, vol. 96(3) (June), pp. 667–93.
- Bryant, J. (1983). 'A simple rational expectations Keynes-type model', *Quarterly Journal of Economics*, vol. 98(3) (August), pp. 525–8.
- Burton, A. and Sefton, M. (2004). 'Risk, pre-play communication and equilibrium', *Games and Economic Behavior*, vol. 46(1) (January), pp. 23–40.
- Carlsson, H. and van Damme, E. (1993). 'Global games and equilibrium selection', *Econometrica*, vol. 61(5) (September), pp. 989–1018.
- Charness, G. (2000). 'Self-serving cheap talk: a test of Aumann's conjecture', *Games and Economic Behavior*, vol. 33(2) (November), pp. 177–94.
- Chaudhuri, A., Graziano S. and Maitra P. (2006). 'Social learning and norms in an experimental public goods game with inter-generational advice', *Review of Economic Studies*, vol. 73(2) (April), pp. 357–80.
- Chwe, M. (2001). *Rational Ritual: Culture, Coordination and Common Knowledge*, Princeton, NJ: Princeton University Press.
- Cooper, R. (1999). *Coordination Games: Complementarities and Macroeconomics*, Cambridge: Cambridge University Press.
- Cooper, R., DeJong, D., Forsythe, R., and Ross, T. (1989). 'Communication in the battle of the sexes game: some experimental results', *Rand Journal of Economics*, vol. 20(4) (Winter), pp. 568–87.
- Cooper, R., DeJong, D., Forsythe, R., and Ross, T. (1992). 'Communication in coordination games', *Quarterly Journal of Economics*, vol. 107(2) (May), pp. 739–71.
- Crawford, V. (1991). 'An "evolutionary" interpretation of Van Huyck, Battalio, and Beil's experimental results on coordination', *Games and Economic Behavior*, vol. 3(1) (February), pp. 25–60.
- Crawford, V. (1995). 'Adaptive dynamics in coordination games', *Econometrica*, vol. 63(1) (January), pp. 103–43.
- Cremer, J. (1986). 'Cooperation in ongoing organizations', *Quarterly Journal of Economics*, vol. 101(1) (February), pp. 33–49.
- Dawes, R., McTavish, J. and Shaklee, H. (1977). 'Behavior, communication, and assumptions about other people's behavior in a commons dilemma situation', *Journal of Personality and Social Psychology*, vol. 35(1) (January), pp. 1–11.
- Diamond, D., and Dybvig, P., (1983). 'Bank runs, deposit insurance, and liquidity', *Journal of Political Economy*, vol. 91(3) (June), pp. 401–19.
- Ichniowski, C., Shaw, K. and Prennushi, G. (1997). 'The effects of human resource management practices on productivity: a study of steel finishing lines', *American Economic Review*, vol. 87(3) (June), pp. 291–313.
- Isaac, M. and Walker, J. (1988). 'Communication and free-riding behaviour: the voluntary contribution mechanism', *Economic Inquiry*, vol. 26 (4) (October), pp. 585–608.
- Kandori, M. (1992). 'Repeated games played by overlapping generations of players', *Review of Economic Studies*, vol. 59(1) (January), pp. 81–92.
- Knez, M. and Simester, D. (2001). 'Firm-wide incentives and mutual monitoring at Continental Airlines', *Journal of Labor Economics*, vol. 19(4) (October), pp. 743–72.

- Kremer, M. (1993). 'The O-ring theory of economic development', *Quarterly Journal of Economics*, vol. 108(3) (August), pp. 551–75.
- Kuran, T. (1987). 'Preference falsification, policy continuity and collective conservatism', *ECONOMIC JOURNAL*, vol. 97(387) (September), pp. 642–5.
- Kuran, T. (1995). *Private Truths, Public Lies: The Social Consequences of Preference Falsification*, Cambridge, MA: Harvard University Press.
- Monderer D. and Samet, D. (1989). 'Approximating common knowledge with common beliefs', *Games and Economic Behavior*, vol. 1(2) (June), pp. 170–90.
- Morris, S., Rob, R. and Shin, H.S. (1995). 'p-dominance and belief potential', *Econometrica*, vol. 63(1) (January), pp. 145–57.
- Morris, S., and Shin, H.S. (1998). 'Unique equilibrium in a model of self-fulfilling currency attacks', *American Economic Review*, vol. 88(3) (June), pp. 587–97.
- Murphy, K., Shleifer, A. and Vishny, R. (1989). 'Industrialization and the big push', *Journal of Political Economy*, vol. 97(5) (October), pp. 1003–26.
- Robinson, P. (1982). 'On the asymptotic properties of estimators of models containing limited dependent variables', *Econometrica*, vol. 50(1) (January), pp. 27–41.
- Rubinstein, A. (1989). 'The electronic mail game: strategic behavior under "almost common knowledge"', *American Economic Review*, vol. 79(3) (June), pp. 385–91.
- Salant, D.J. (1991). 'A repeated game with finitely lived overlapping generations of players', *Games & Economic Behavior*, vol. 3(2) (May), pp. 244–59.
- Schotter, A. and Sopher, B. (2003). 'Social learning and coordination conventions in intergenerational games: an experimental study', *Journal of Political Economy*, vol. 111(3) (June), pp. 498–529.
- Schotter, A. and Sopher, B. (2006). 'Advice, trust and trustworthiness in an experimental intergenerational game', *Experimental Economics*, vol. 9(2) (June), pp. 123–45.
- Schotter, A. and Sopher, B. (2007). 'Advice and behavior in intergenerational ultimatum games: an experimental approach', *Games and Economic Behavior*, vol. 58(2) (February), pp. 365–93.
- Van Huyck, J., Battalio, R. and Beil, R. (1990). 'Tacit coordination games, strategic uncertainty, and coordination failure', *American Economic Review*, vol. 80(1) (March), pp. 234–48.
- Van Huyck, J., Gillette, A. and Battalio, R. (1992). 'Credible assignments in coordination games', *Games and Economic Behavior*, vol. 4(4) (October), pp. 606–26.
- Wooldridge, J. (2002). *Econometric Analysis of Cross Section and Panel Data*, Cambridge, MA: MIT Press.

Copyright of Economic Journal is the property of Blackwell Publishing Limited and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.