Managing growth to achieve efficient coordination in large groups

Roberto Weber*

Department of Social and Decision Sciences

Carnegie Mellon University

Pittsburgh, PA 15213

rweber@andrew.cmu.edu

August 8, 2002

*This paper benefited greatly from the author's discussions with Colin Camerer, David Grether, Chip Heath, George Loewenstein, Richard McKelvey, and Tom Palfrey, and from helpful feedback from Roberto Fernandez, Jacob Goeree, and participants at the Stanford Strategic Management conference, February 20-21, 1998; the March 1998 Public Choice/Economic Science Association meetings; and seminar participants at Caltech, University of Arizona, London Business School, University of Southern California, The Wharton School, Carnegie Mellon, Harvard Business School, Northwestern, University of Chicago, and Stanford University. Thanks to Chris Anderson, Angela Hung, Anthony Kwasnica, and Alex Markle for help in conducting some of the experiments. This research was funded by the Russell Sage Foundation Behavioral Economics Small Grant Program and by Carnegie Mellon.

Managing growth to achieve efficient coordination in large groups

Roberto Weber

Abstract

Several previous experiments using the minimum effort (weak link) coordination game reveal a striking regularity – large groups *never* coordinate successfully on the efficient equilibrium. Given the frequency with which large real-world groups, such as firms, face similarly difficult coordination problems, this poses an important question for economics and organization: Why do we observe large successfully coordinated groups of people in the real world when they are so difficult to create in the laboratory? This paper presents one reason. The experiments show that, even though efficient coordination does not occur in groups that start off large, efficiently coordinated large groups can be "grown." By starting with small groups that find it easier to coordinate, we can add people at a sufficiently slow rate to create efficiently coordinated large groups. This represents the first experimental demonstration of large groups regularly coordinated at high levels of efficiency.

Coordination is an important problem for economics and organization (Coase, 1937; March and Simon, 1958; Schelling, 1960; Arrow, 1974; Cooper, 1999). Coordination deals with situations where economic actors need to attempt to match the actions of others without knowing what these others will do – such as buyers or sellers searching for a market, workers with complementary production tasks, or consumers purchasing products with network externalities. In game-theoretic terms, coordination problems arise when there are multiple equilibria, and players must tacitly resolve which one to play.

This paper is about coordination problems in large groups, and the special difficulty that group size creates for efficient coordination. There is considerable experimental evidence – mostly from a game known as the minimum effort (or weak link) coordination game – that group size has a strong effect on the ability of groups to coordinate (e.g., Van Huyck, Battalio and Beil, 1990; Weber, et al., 2001). Large groups of people almost never coordinate successfully, and repetition alone does not solve the problem. However, this is inconsistent with the real world, where we observe groups much larger than those in experiments – such as firms and countries – where coordination plays a crucial role, but where these large groups have managed to coordinate successfully. To give a specific example, researchers have presented the airline industry as an industry where coordination problems are both difficult and ubiquitous and are very similar to those modeled by the minimum effort game (Knez and Simester, 2001; Gittell, 2001). However, we observe large airlines that operate smoothly and appear to have solved these coordination problems, even for very large groups. For another example, consider that problems like corruption in a given country can be modelled by an infinitely-repeated game, similar to the prisoner's dilemma, played between the country's inhabitants (Bicchieri and Rovelli, 1995). The infinite horizon gives rise to several equilibria, some more efficient than others, and inhabitants of the country are better off when they coordinate on a low-corruption equilibrium. If "large" laboratory groups cannot coordinate efficiently in this type of game, large countries should find it difficult to coordinate on an efficient equilibrium. But then why are some large countries, like Sweden, coordinated on efficient repeated-game equilibria with very little corruption? Equally important, however, is why other similarly large countries, like Nigeria, are stuck at less efficient repeated-game equilibria with much higher levels of corruption. This paper presents one possible reason.

This paper argues, and demonstrates, that the ability of large groups or organizations to coordinate successfully is critically affected by the group's growth process itself. Consistent with previous experimental research, coordination is much easier in smaller groups. Therefore, a small group, such as the founding members of a firm, do not face substantial difficulty in coordinating efficiently. Once they have done so, they can establish a set of self-reinforcing rules or norms – either tacit or formalized – governing what actions are appropriate. These norms allow the group or organization to continue to successfully coordinate activity. As the group grows, new entrants' exposure to these norms allows these entrants to be aware of the appropriate behavior, and creates an expectation for everyone in the group of what everyone else (including the new entrants) will do. Therefore, slow growth and exposure of new entrants to the group's previous norms can overcome large-group coordination failure.

In the rest of this paper, this is clearly demonstrated using experiments on the minimum effort coordination game. The experiments are motivated by a simple model that shows why growth should work (Weber, 2000). The intuition underlying this model is presented in the next section. In the experiments, efficiently coordinated large groups, which are impossible to obtain when a group starts out at a large size, are "grown" in the laboratory by starting with a small group and adding a few entrants at a time who are exposed to the group's history. In doing so, this paper produces the *only* laboratory demonstration of the regular occurrence of efficiently coordinated large groups.

1 The minimum effort coordination game

The minimum effort, or weak link, coordination game was first studied experimentally by Van Huyck, Battalio and Beil (1990) and has since been studied by several researchers (e.g., Cachon and Camerer, 1996; Weber, et al., 2001; Chaudhuri, Schotter, and Sopher, 2002). In the game, which is an n-person version of the stag hunt game, players choose numbers, which can be thought of as orderable strategies such as effort or contribution levels. Every player receives a payoff that is a function of her choice and the minimum choice of all other $n-1$ players (thus the term "weak link," since every player's payoff is partially determined by the lowest choice in the group). A seven-action version of the game is represented in Table 1, which gives the payoff to each player as a function of her choice and the minimum choice of all n players.

The diagonal cells correspond to outcomes in which the player is making the same choice as the group minimum. When everyone makes the same choice, and therefore receives the

Player's	Minimum choice of all players						
choice	7	6	5	4	3	2	1
7	.90	.70	.50	.30	.10	-.10	-.30
6		.80	.60	.40	.20	.00	-.20
5			.70	.50	.30	.10	-.10
4				.60	.40	.20	.00
3					.50	.30	.10
2						.40	.20
1							.30

Table 1: Payoffs (in dollars) for minimum-effort game

same payoff, the outcome is one of the game's seven pure-strategy Nash equilibria (each represented by a payoff along the diagonal). In any of these equilibria, no player has an incentive to unilaterally increase her choice (which would not affect the minimum but would lower her payoff) or decrease it (which would lower the minimum and also result in a lower payoff). Notice, however, that these equilibria differ because those corresponding to higher choices also yield higher payoffs. Therefore, more efficient coordination corresponds to all players making higher choices in equilibrium. The Pareto-optimal (or efficient) equilibrium results when all of the participants select the highest choice, 7, and receive \$0.90. It is in the players' mutual interest to reach this outcome and the players realize this.[1]

However, the efficient equilibrium may not be easy to achieve because players are faced with strategic uncertainty. Everyone may recognize the efficient equilibrium, but may be unsure of what others will do. Therefore, players may choose something other than 7,

[1]This game does not have the conflicting incentives present in the prisoner's dilemma. In the minimum effort game, as in the closely related stag hunt game, players want to reciprocate high actions of others with high choices and the best outcome for all players is the efficient equilibrium.

particularly when they think it is more likely that someone else will choose something other than 7. Simply being unsure about what others will do may lead players to take different actions.

Previous experiments with minimum effort coordination games have established clear regularities. Tacit coordination on the efficient equilibrium is impossible for large groups. Of the seven sessions initially conducted by Van Huyck, et al. (1990) (VHBB) with groups of size 14 to 16, after the third period the minimum in *all* games was the lowest possible choice. For small groups (n = 2), coordination on the efficient equilibrium was much easier – it was reached in 12 of 14 (86%) of the groups studied (a result replicated by Camerer and Knez (2000)). Table 2 summarizes the distribution of fifth-period minima in several different experiments, all using the Van Huyck, et al., game in which subjects choose integers from 1 to 7, and choosing 7 is efficient.

1	2	3	4	5	6	7	Group Size	N	Source
9	0	0	0	0	0	91	2	28	VHBB, 1990; Camerer & Knez 2000
37	15	15	11	0	4	18	3	60	Knez & Camerer 1994, Camerer & Knex 2000
80	10	10	0	0	0	0	6	114	Knez & Camerer 1994
100	0	0	0	0	0	0	9	18	Cachon & Camerer 1996
100	0	0	0	0	0	0	14-16	104	VHBB, 1990

Table 2: Fifth period group minima (by %) in various 7-action minimum-effort studies (1 = inefficient; 7 = efficient)

The effect of group size could hardly be stronger. Subjects in a group of size 2 are almost assured to coordinate on the efficient equilibrium. Subjects in larger groups (six or more) are almost assured to converge to the least efficient outcome.[2] Thus, there is a strong negative relationship between a group's size and the ability of its members to coordinate efficiently.[3]

The result that large groups fail to coordinate efficiently is even largely (though not quite always) immune to communication. The above experiments all used versions of the game in which there was no communication between people about to play the game. This leaves unaddressed one possible solution to large-group coordination failure: communication. However, a couple of studies have demonstrated that communication leads to only slight improvements. For instance, Weber, et al., (2001) had subjects play a four-action version of the game in groups of 8 to 15 where one player was randomly selected to address the other players and urge them to coordinate efficiently. This communication was public and common knowledge – everyone sat in the same room and together observed the "leader" addressing the entire group. However, only in 1 of 9 such groups was the group subsequently able to coordinate on a minimum above the lowest value.

Chaudhuri, Schotter and Sopher (2002) examined coordination in the minimum effort game using several "advice" treatments in which each cohort of players gave advice to a subsequent generation on how to play the game (and subsequently received payoffs linked to the payoffs recieved by this "progeny"). Using groups of size 8, they ran several different treatments in which they varied who received the advice (whether each "parent's" message went to only one person or to everyone in the group) and the way in which the advice

[2]Once these groups reached the inefficient outcome, they were not able to subsequently increase the minimum.

[3]This result has also been replicated using slightly different versions of the game (e.g., Weber, et al., 2001).

was communicated (whether it was written or publicly announced). They found that in most of the treatments, advice had no positive effect on the ability of players to coordinate efficiently.[4] In fact, only one treatment produced efficient coordination (in a few sessions): when the advice given by all parents from the previous generation was read aloud publicly to a group of subjects all sitting in the same room. Therefore, while communication helps slightly, it in no way regularly solves the difficulty in coordination faced by large groups. Indeed, Caudhuri, et al., express surprise at how difficult it is to find a treatment in which inter-generational communication solves coordination failure.

The results of these experiments using the minimum effort game and communication indicate even more how difficult it is to obtain efficient coordination in large groups. The only way to obtain an efficiently coordinated large group is to start off with a large group that is efficiently coordinated and then allow all the members of this group to communicate their history and a recommendation to a subsequent group of players. Moreover, this history and communication must take place in a "commom knowledge" environment where everyone receiving the communication is together in the same room while hearing the announcement. This does not, of course, begin to answer the question of how we obtain efficiently coordinated large groups in the first place.

Given the link between coordination in minimum effort games and coordination problems faced by real-world groups (e.g, Knez and Simester, 2001; Camerer and Knez, 1997; Nanda, 1997), all of the above results points to the impossibility of obtaining efficient coordination

[4]In fact, when the advice given to the next generation was not public and common knowledge, coordination on the efficient equilibrium was less likely than if there was no advice.

in real large groups. However, we do observe efficiently coordinated large groups in the real world. To see one reason why this is the case, we need to begin by recognizing that few large groups (such as firms and communities) start off at a large size. Most groups in fact begin small, when solving coordination problems is easier according to the above experimental literature. Once successfully coordinated, then, it might be the case that these groups manage to remain coordinated as new entrants are added, growing into efficiently coordinated large groups.

Using a simple dynamic model based on Crawford's (1995) model of adaptive dynamics in order-statistic coordination games (including the minimum effort game), Weber (2000) shows how growth can produce large efficiently coordinated groups. Specifically, the model assumes that players are initially unsure of what actions other players will take, and that this strategic uncertainty is captured by a stochastic component of players' choices. The variance of this stochastic element decreases with experience – as players learn about the actions taken by others. The dynamic process determining players' choices is simply an adjustment from a player's choice in the previous period towards the minimum in the previous period. Using a model like this, Crawford (1995) shows why the minimum in large groups falls to 1, while small groups are more likely to coordinate efficiently.

Weber (2000) builds on Crawford's model by defining "growth paths" as sequences of plays of the minimum effort game in which the set of players weakly increases in size across periods and where the number of players at the end of the sequence is strictly greater than at the first period. Similarly to Crawford (1995), the model assumes that as players gain experience, their strategic uncertainty is reduced. The model also assumes that future entrants who are not yet playing the game – but will be playing in the future – have decreased strategic uncertainty, as long as they can observe the outcome of plays of the game.[5] The main result from this model is that growth stochastically produces large groups coordinated at higher levels of efficiency, after t periods, than groups that start off at a large size and play at that size for t periods. Moreover, "slower" growth paths result in stochastically higher efficiency than "faster" growth paths that always have at least the same number of players. The intuition behind these results is simple: smaller groups coordinate at higher levels of efficiency, and when the new entrants are added to achieve a group size of n, the overall strategic uncertainty (which produces lower minima in large groups) is much lower than it is for a group that starts off playing at size n.

The hypothesis in this paper is that initial successful coordination followed by growth that is not too rapid, and in which new entrants are aware of the previous history of successful coordination, can lead to efficient coordination in large groups. The experiments in the next section directly test this hypothesis.

2 Growing efficient coordination in the laboratory

The experiments are conducted similarly to previous experiments by Van Huyck, et al.. (1990) and Weber, et al. (2001). There are two almost identical experiments. In the

[5]The assumption that people can "learn" about how a game will be played by watching other players is supported by experiments (e.g., Duffy and Feltovich, 2002).

main experiment, the rate of growth was fixed and pre-determined by the experimenter. This experiment tests the main hypothesis that growth can produce efficiently coordinated large groups, relative to control groups that start off at a large size. The results are clear: while none of the control groups manages to coordinate on a minimum above one, the grown groups frequently manage to coordinate at higher levels of efficiency, including several groups coordinated on the highest possible minimum.

As a result of the successful, though fragile, growth in the first experiment, a second experiment looks at whether subjects themselves are aware of the need for careful growth. In this experiments, groups were again grown as before, but the rate of growth was endogenous and not determined until the experiment. In each period, a "manager" not playing the minimum effort game determined the size of the group. This experiment provides additional evidence of successful growth, even with endogenously determined growth paths. In addition, while limited by a very small number of observations due to the high cost of conducting a session (each data point involves using 13 paid subjects for 2 hours), the experiment reveals that subjects are not initially aware of the need for careful growth and tend to grow the groups too quickly, but that some of them learn to start over and are successful with slower growth.

2.1 Exogenous pre-determined growth

2.1.1 Experimental Design

Since large groups of ten or more subjects have never consistently coordinated efficiently in previous experiments – in fact, they almost always end up at the least efficient outcome – this first experiment was designed to explore whether a slow, controlled growth rate determined by the experimenter could create large groups that coordinated efficiently. In the experiment, groups of 12 students (at Stanford, UC Santa Cruz, or Carnegie Mellon) were assembled in one room. Subjects were presented the game in Table 1, though it was framed in the context of a report completion as in Weber, et al. (2001).[6] Subjects were informed about how growth would take place (see below). They then played 22 periods in which the group was grown from a size of 2 to 12 players, except in the control sessions, where they played 12 periods at a fixed group size of 12 and there was no mention of growth. Instructions were read aloud and, before playing the game, subjects answered several questions to check their comprehension of the instructions and the game. At the end of the experiment, subjects were paid their earnings in cash.

In each experimental session, subjects were anonymously assigned participant numbers. Each of the growth sessions consisted of 22 periods. In the first several periods, only participants 1 and 2 played the game while the other subjects sat quietly. Participants were

[6] This context presents each player's choice as the choice of a time at which to contribute a section of a report. Each player receives a bonus depending on when the report is completed (when the last section is contributed) and pays a cost (which is higher for contributing earlier). The game was presented both in this context and also by giving the subjects Table 1 with an explanation of the payoffs. Weber, et al. (2001) include experiments both with and without this context and find no difference in the results.

told that they would receive a fixed, positive, "fair" amount for periods in which they were not playing the game, but that the exact amount would not be revealed until the end of the experiment.[7] In each period, participating subjects recorded a number from 1 to 7 (indicating the contribution time for their section of the report) on a piece of paper and handed it to the experimenter.[8]

At various preannounced and commonly known periods, other participants (referred to only by number) joined the group of those actively playing the game. For each session, there was a schedule of such additions – referred to as a *growth path* – that was handed to all subjects at the beginning of the experiment. For example, in one of the growth paths a third participant (#3) was added in period 7, joining the first two participants (#1 and #2) who continued to participate. Subjects all knew the predetermined growth path, and they knew that earlier participants always continued to participate. At the end of each period, the minimum was announced to all the subjects, including those who were not actively participating yet. This was the only feedback provided. In all growth paths, all 12 subjects were participating by the last few periods.

---

[7]This was done because of the concern that announcing the per-period amount could create a focal point that might influence participants' choices. Subjects were told that the reason they were not informed of the amount was because "we do not want this to influence your choices." To ensure that they believed the amount was fixed – and not dependent on what happened in the experiment – the experimenter placed a large envelope at the front of the room and told subjects that the amount was written on a sheet inside the envelope. This sheet was shown to subjects at the end of the experiment.

[8]To prevent players from knowing which others were participating, all players handed in slips of paper; non-participants simply checked a box saying they were not participating.

2.1.2 Results: Control sessions

Five control sessions (n = 60) were conducted using undergraduates at Stanford (sessions C1 & C2), Caltech (sessions C3 & C4), and Carnegie Mellon (session C5). The results are reported in Figure 1, which presents the minimum choice across all 12 periods for each session.[9] In addition, the two thicker lines present the average of the session minima and the average choice across all sessions.

(Figure 1 about here)

Overall, the results replicate previous experimental results on large groups playing the minimum effort coordination game. The minimum converges to 1 in all five control groups. Both the average choice and the average of the minima consistently decrease and end up at or near one by the final periods. Note also that the average choice is initially high, indicating that many subjects select high actions in early periods, but that the minimum is nonetheless low since it is sensitive to outliers.

2.1.3 Results: Growth sessions

Nine sessions with pre-determined growth paths (n = 108) were conducted at Stanford (sessions 1 - 4), UC Santa Cruz (sessions 5 - 7), and Carnegie Mellon (sessions 8 & 9).

The purpose of this experiment is to explore the possibility that growth may lead to more efficiently coordinated large groups. Therefore, each of the growth paths used was

[9] The entire dataset of choices can be obtained from the author's website: http://www.andrew.cmu.edu/user/rweber.

selected with the hope that it would be slow enough to create a large group that coordinated more efficiently than large groups had in earlier experiments and in the control sessions. The principles driving the choice of these growth paths were based on the intuition that slow, regular growth, and successful experience with initial growth probably give the best chance of obtaining efficient coordination (as shonw in Weber (2000)). This implies that good candidates for successful growth paths should only add a few players at a time and should allow time between growth periods or "spurts," particularly initially. Therefore, the growth paths were designed to first allow the establishment of successful coordination in the group of size two (by allowing them to play several periods before adding more participants) and then to add players in a slow and regular manner. Thus, with one exception (the last two players added in growth path 3), the growth paths add only one player at a time. Figure 2 shows the three growth paths used in these experiments.

(Figure 2 about here)

Figures 3a through 3c present the minimum choices for sessions 1 through 9. Each figure also presents the corresponding growth path.[10] The marker-less solid line in each figure shows the growth path (measured on the left vertical axis), while the remaining lines all present the minima in a session (measured on the right vertical axis). The data are also

[10] The choices in the final two periods (21 and 22) are not reported here and not included in the analysis because there is often a strong end of experiment effect in these games, in which subjects change their choices in the final period or two (perhaps to punish or do better that others, or perhaps because they believe that others will do so – in which case doing so is a best response). The minima are particularly sensitive to only one subject changing her behavior in this manner. While this phenomenon is interesting and frequently shows up in previous experiments, this paper is not concerned with what occurs in the final periods (after growth is completed), but rather with coordination during and immediately after growth.

presented in Table 3.

(Figure 3 about here)

An examination of behavior in individual sessions helps shed light on several behavioral regularities. First, the small groups were able to coordinate efficiently. In all but one of the sessions, the group of size two was able to coordinate on a minimum of 7 for at least two consecutive periods. In the remaining one (session 7) the group was able to coordinate on a high level of efficiency (minimum = 6).

	Period																			
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
Growth path 1	*2*	*2*	*2*	*2*	*2*	*2*	*3*	*3*	*4*	*4*	*5*	*6*	*7*	*8*	*9*	*10*	*11*	*12*	*12*	*12*
Session 1	7	7	7	7	7	7	6	6	5	5	4	4	3	3	2	1	1	1	1	1
Session 2	6	5	7	6	7	7	5	6	5	6	5	5	5	5	5	5	5	5	5	5
Growth path 2	*2*	*2*	*2*	*2*	*2*	*3*	*3*	*3*	*3*	*4*	*4*	*5*	*6*	*7*	*8*	*9*	*10*	*11*	*12*	*12*
Session 3	7	7	7	7	7	5	5	5	5	5	5	5	5	5	5	5	1	1	1	1
Session 4	7	7	7	7	7	7	7	7	7	7	7	7	7	7	7	7	7	7	7	7
Growth path 3	*2*	*2*	*2*	*2*	*2*	*3*	*3*	*3*	*3*	*4*	*4*	*5*	*6*	*7*	*8*	*9*	*10*	*12*	*12*	*12*
Session 5	6	6	7	7	7	7	7	7	7	7	7	7	7	4	3	3	3	3	3	3
Session 6	7	7	7	7	7	7	7	7	7	7	7	7	7	5	4	3	2	1	1	1
Session 7	6	6	6	6	6	6	6	6	6	6	6	6	1	4	4	4	4	4	3	1
Session 8	7	7	7	7	7	7	7	7	7	7	7	7	7	7	7	7	7	7	7	7
Session 9	7	7	7	7	7	7	7	7	7	7	7	7	7	7	7	7	7	7	7	7

Table 3: Minima by period for “growth” sessions

A second observation is that growth does not always work. In four of the sessions (1, 3, 6, and 7) the minimum by period 20 is 1. Therefore growth does not always produce efficiently coordinated groups, even with growth paths designed to do so.

An equally strong regularity, however, is the fact that growth just as often produces large efficiently coordinated groups. In three sessions (4, 8, and 9), the minimum remains at 7 throughout 20 periods. In another session (2), the minimum never falls below 5, while in another (5) it does not fall below 3. Therefore, in five out of nine sessions, the large groups coordinate on an equilibrium with a minimum greater than 1, and in three of those they coordinate on the efficient equilibrium. Recall that in previous experiments and in the control sessions, large groups (of 6 or more) *never* remained coordinated on a minimum greater than 1.

In addition, in all the sessions that end up at a minimum of 1, the minimum is higher at least through a group size of 9. This higher level of efficiency for groups of size 9 is surprising in light of the fact that the minimum was always 1 for the large groups (nine or larger) in Table 3.[11] Thus, there is clear support for the hypothesis that starting with a two-person group, which reliably reach efficiency, and then adding players at a slow rate enables much better coordination than starting with a large group.[12]

To directly test whether or not growth results in successful coordination, we need to compare choices in the control sessions with those in the growth sessions. Table 4 compares the distribution of subject choices in the five control sessions and the nine growth sessions

[11]While Table 2 reports the fifth period minima, the first period minima in previous experiments were not as high as in the sessions reported here and there was never a minimum of 7. Note also that in the control sessions the minimum was never above 4 in any period.

[12]The results of the experiments by Knez and Camerer (1994), which showed that "merging" two three person groups leads to coordination failure (the minimum fell to 1 80 percent of the time), indicate that growth can be too rapid. This provides additional evidence that *controlled* growth can play an important role in obtaining successful coordination in large groups. Note that the minimum for a group of size 6 (which corresponds to the size of Knez and Camerer's rapidly grown groups) was 1 in only one of the seven sessions in experiment 1.

in the fourth period in which participants played at a group size of 12.[13]

		Control		Growth	
Choice	7	3	(5.0 %)	32	(29.6 %)
	6	0	(0.0 %)	3	(2.8 %)
	5	6	(10.0 %)	13	(12.0 %)
	4	22	(36.7 %)	8	(7.4 %)
	3	7	(11.7 %)	10	(9.3 %)
	2	9	(15.0 %)	9	(8.3 %)
	1	13	(21.7 %)	33	(30.6 %)
Total		60		108	
Minima		1,1,1,1,4		1,1,1,1,3,4,5,7,7	

Table 4: Distribution of subject choice in fourth period as 12-person groups

The frequency of subjects choosing 1 is high in both conditions (13 of 60 in the control sessions and 33 of 108 in the growth sessions). While the proportion is higher for the growth sessions, this is not surprising since in three of the nine grown groups the minimum was 1 even before the group reached a size of 12. In these groups, therefore, there were more previous periods for subjects' choices to converge towards the inefficient equilibrium than in the control treatment, where initial choices are usually high and dispersed but then converge downwards.[14]

Just as interesting, however, is the difference in the distribution of subjects choosing

[13]Selecting a comparison period is somewhat tricky. The control groups played as large groups for 12 periods. The grown groups all started off at small sizes and did not reach a group size of 12 for several periods. The earliest period in which a grown group reached the maximum size was period 18. The key question is when the comparison should be made. A reasonable comparison is to compare the control groups in period t with the t-th period in which the grown groups played as groups of size 12. In this case, subjects in both treatments have $t-1$ periods of play in 12-person groups and therefore share a similar history. The analysis here sets $t=4$ because this is the greatest value of t for which there is data in all of the growth sessions (growth path 2 has a group size of 12 only in periods 19 through 22), and it allows the groups several periods in which to coordinate.

[14]This aspect of the comparison works against the hypothesis that growth works, since choices in unsuccessfully grown groups are usually almost all 1's even prior to the first period as a 12-person group, while in the control groups many subjects' choices do not converge to 1 until later periods.

seven between the two treatments. In the control sessions, only 3 of 60 subjects (5%) did so, while this is true of 32 of 108 subjects (30%) in the growth sessions. Therefore, the number of subjects playing the efficient strategy is much higher in the grown groups than in the control sessions. The distributions of choices in Table 4 are significantly different when compared using a Chi-Square test ($\chi^2_{(6)} = 34.5$, $p < 0.001$) and a Kolmogorov-Smirnov one-tailed test ($\chi^2_{(2)} = 13.4$, $p < 0.01$).

Both of the above tests, however, rely on the assumption that all the observations in each treatment are independent. Since this assumption is not satisfied because the choices of subjects in a session in a particular period are affected by a shared history, the level of significance reported by the statistics is exaggerated. One way to address this problem is to treat each session as the independent unit of analysis and, for instance, examine only the minima. These minima are reported in the final row of Table 4. Note that the fourth-period minimum in all but one of the control sessions is 1 and that, while the minimum in four of the growth sessions is also 1, the minimum is greater than 1 in the remaining five sessions. Moreover, the growth sessions produce minima of at least 5 three times, which never occur in the control sessions. A Mann-Whitney U test of the minima yields the test statistic $z = 1.33$. While the corresponding one-tailed p-value of 0.10 is not highly significant, it must be noted that this test and p-value are *extremely* conservative since they treat each group of twelve subjects as just one observation.

Another observation is that it appears that early experience with growth is important in determining subsequent success. In sessions 2 through 9, the minimum did not drop when the fourth person was added compared to what it was when the third participant entered. In these groups, the minimum was 5 or greater through at least period 12, indicating that subjects may have learned that controlled growth was possible (at least for a while). In session 1, however, the minimum dropped both of the initial times the group grew and continued to drop with growth, indicating that the initial experience with growth led subjects to believe that the group could not grow successfully.

In fact, the results of session 1 shed light on an interesting phenomenon (which also surfaces – much more strikingly – in the experiment reported in the next section). In session 1, when a third participant was added in period 7, the minimum dipped below 7. The minimum was 6 in periods 7-8. When a fourth participant was added, in period 9, the minimum fell further to 5 and stayed there in period 10. When a fifth participant was added, in period 11, the minimum fell to 4. This pattern suggests an interesting conjecture. Earlier research showed that precedents often matter dramatically in this kind of game, in the sense that a group expects the minimum in one period to be the minimum in an upcoming period. That is, previous choice establishes a strong precedent that is reinforced by subsequent actions. In session 1, however, players seem to be inferring a precedent from the relation between *changes in structure* (group size) and changes in behavior. The fact that the minimum fell by one when group size increased from 2 to 3, and from 3 to 4, seems to create a precedent that "when we grow the minimum falls by 1," which is largely self-fulfilling in later periods. This suggests that players may form "higher-order" precedents based on not just levels of previous play (e.g., expect the previous minimum to be the minimum again), but also on the relation between levels of previous play and group sizes or transitions. The fact that the minimum falls by one unit when a third person is added, and falls again by that same amount when a fourth person is added, seems to create a belief that adding a person leads the minimum to fall by one (which is self-fulfilled when the fifth, seventh, ninth and tenth people are added, though not when the sixth and eighth are added). This kind of behavior – of which an even more striking example appears in the next section – had not been observed in previous work because nobody had changed structural variables repeatedly from period to period in a way that allows formation of higher-order relational precedents.

The above analysis convincingly shows that controlled growth can solve the problem of large group coordination failure. The growth sessions produce several groups coordinated at higher minima than 1 (including minima of 7), a result that has never been found in previous experiments on minimum effort coordination games. The growth sessions in this experiment all use a slow, regular growth path intended to give growth the best chance of producing large, efficiently coordinated groups. It is worth asking how the success of growth might be affected by growth paths that differ. An equally important question is whether subjects themselves are aware of the need for slow, controlled growth.[15] The following experiment provides exploratory evidence addressing both of these issues.

[15] It has been noted that while coordination problems play an important role for firms in several industries, they are often neglected by both managers and organizational researchers (March and Simon, 1958; Lawrence and Lorsch, 1967; Heath and Staudenmayer, 2000).

2.2 Endogenous dynamic growth

2.2.1 Experimental Design

In a second experiment one participant was randomly selected to act as a "manager" and determine the growth path during the experiment. Otherwise, the experiment was conducted in much the same way as the first experiment. The manager was placed in a separate room from the remaining 12 subjects and an experimenter carried information between the two rooms. The minimum effort coordination game that the other 12 participants were to play was described to the manager (again framed in the context of a project completion). In all periods after the first – in which group size was fixed at 2 – the manager selected the size of the group that would play the game. The manager's earnings in each period were determined by the number of active participants and by the group minimum. Table 5 describes the possible earnings for the manager.[16] For any group size, the manager is better off when the group coordinates efficiently. Also, the manager's payoff is higher when efficiently coordinated groups are larger, but the opposite is true for inefficient groups. Therefore, the manager has an incentive to create a large group, but only if it is efficiently coordinated.

Following the manager's determination of the group size in each period, a group of up to 12 subjects played the game in the same way as in the first experiment. The instructions

[16]The earnings are determined according to the following formula (rounded to the nearest cent):

$$\pi = \frac{n(min - 3.5)}{100} + 0.05$$

except for the payoff when the group size is twelve and the minimum is 7. Since the goal was for managers to attempt to reach this outcome, a large bonus was awarded for achieving it.

Number of participants	Group Minimum 1	2	3	4	5	6	7
2	0.00	0.02	0.04	0.06	0.08	0.10	0.12
3	-0.03	0.01	0.04	0.07	0.10	0.13	0.16
4	-0.05	-0.01	0.03	0.07	0.11	0.15	0.19
5	-0.08	-0.03	0.03	0.08	0.13	0.18	0.23
6	-0.10	-0.04	0.02	0.08	0.14	0.20	0.26
7	-0.13	-0.06	0.02	0.09	0.16	0.23	0.30
8	-0.15	-0.07	0.01	0.09	0.17	0.25	0.33
9	-0.18	-0.09	0.01	0.10	0.19	0.28	0.37
10	-0.20	-0.10	0.00	0.10	0.20	0.30	0.40
11	-0.23	-0.12	-0.01	0.11	0.22	0.33	0.44
12	-0.25	-0.13	-0.01	0.11	0.23	0.35	1.00

Table 5: Manager's payoffs (in dollars)

for these subjects were the same as in the previous experiment, except they were informed that the number of active participants was not pre-determined and would be determined at the beginning of each period by the manager. To determine the group size in a period, the manager simply wrote down a number between 2 and 12, and then the subjects with participant numbers from 1 to that number played the game in that period.[17]

The experiment lasted 35 periods to give the managers plenty of time to experiment with growth. Four sessions (n = 52) were conducted at Caltech. The experiment lasted about 2 hours.

[17]The one other difference with experiment 1 was that the manager was given the option, at the beginning of each period, to randomly reassign participant numbers. This was intended to allow the manager to "restart" the group in case the first few participants became stuck at a bad equilibrium, since previous results indicate that this does occasionally happen (though rarely) in small groups. This option was rarely used by managers.

2.2.2 Results

The growth paths employed by the four managers (sessions E1 through E4) are presented in Figure 4. The most important thing to note is that all four managers initially grew the groups quickly. In the first period, managers were constrained to a group size of two, but in the next period, all of them added at least three new players. It is also worthwhile to note that two of the four managers (sessions E1 and E4) subsequently implemented much slower growth paths.

In order to more closely analyze the behavior of individual managers, we examine the individual session data. Figures 5a through 5d report the results (growth path and minima) for each of the endogenous growth sessions. The left vertical axis presents the group size, while the right vertical axis presents the group minimum.[18]

(Figure 4 & 5 about here)

While there are only four sessions, an examination of the data reveals some interesting regularities. First, initially rapid growth by managers resulted in coordination failure. While group size in the first period was constrained to be 2, all four managers added at least 3 more players in the next period. In three of the four sessions (E1, E2, and E4), the minimum was initially high (6 or 7), but this minimum fell when the managers reached a group size of 9 or larger by the third period. In the other session (E3), the minimum was initially low (3) and did not increase as the manager grew the group. The fact that none of the managers allowed

[18]As with the first experiment, choices in the final two periods are not reported because of end-of-session effects.

small groups time to coordinate efficiently before growing points to a lack of cognition of both the difficulty of coordinating large groups – which is consistent with previous research (see Weber, et al., 2001). Moreover, the fact that the minimum fell when these managers grew quickly indicates the need for slow, controlled growth to solve coordination failure.

Following this initial failure, however, some of the managers learned to grow using a slower and more regular approach. Two of the managers (E2 and E3) continued to grow too quickly – obtaining higher minima at small group sizes, but adding too many players right away, causing the minimum to fall again. The other two managers (E1 and E4), however, started over with small groups and then grew slowly (never adding more than 2 participants at a time) to create large groups coordinated on minima of 6 and 7. It is interesting to note that the growth path used by the manager in session E1 is very similar to the growth paths used in the sessions with pre-determined growth paths (see Figure 2).

There is also further support in these experiments for the main hypothesis of this paper that slow, regular growth can lead to successful coordination in large groups. In the two sessions in which the managers started over at a small size and grew slowly (sessions E1 and E4), the result was large groups that coordinated on high minima (6 and 7). In addition, the failure to succeed of the other two managers indicates that the rate of growth is important in obtaining efficient coordination in large groups.

Finally, the results of session E4 provide strong additional support for the previously mentioned conjecture, based on the results of session 1, that experience with growth plays an important role in subsequent reactions to growth. To see why, consider what happened in session E4. The minimum was 6 for the group of size 2 in the first period. The manager then tried to grow the group too quickly (increasing the size to 10) and the minimum fell to 1 in the next period. The minimum remained at 1 until the manager set group size at 3 and the minimum moved up to 3. The manager then increased the group size by one and allowed the group of size 4 enough time to raise the minimum to 6. When group size again increased by one, the minimum again fell to 3, but then rose to 7 as group size remained the same for several periods. For the remainder of the experiment, the manager always added only one or two participants at a time and the minimum fell to either 3, 4, or 5 every time the group grew. However, after the initial drop following growth in each of these instances, the minimum increased by *exactly 1 in every period in which the group size remained the same*. This continued until period 33, when the entire group of 12 participants coordinated successfully on 7. Thus, after initially growing too quickly, the manager in session E4 discovered a growth path slow enough to lead to efficient coordination.

Similarly to session 1, session E4 reveals a self-enforcing norm concerning how the group reacts to growth. In session 1, the norm was "the minimum falls by one when we grow." In session E4, a precedent was established indicating that every time the group grew, the minimum fell, but that in every period that the group did not grow, the minimum went up by exactly one. The strength of this precedent and the extent to which it was a self-reinforcing belief held by *all* players is evident in the last few episodes of growth. When the group grew to a group size of 10, the minimum fell to 5. In the next two periods, *all* 10 players played first 6 and then 7. When the group grew again to a size of 12, the minimum fell to 4.[19] In the next three periods, *all* 12 players played first 5, then 6, and then 7. This points to a strong coordinating effect of previous experience with growth, and a very strong norm prescribing how subjects should react not only to the previous minimum, but also to what happened in previous experience with growth.

3 Conclusion

Previous studies reveal the difficulty involved with obtaining efficient coordination in large groups playing the minimum effort game. In fact, no previous study has regularly produced independent, efficiently coordinated large groups. The results of both experiments in this paper clearly show that efficient coordination in large groups is possible when groups start off small and then grow slowly. Over one-half of the 12-person groups in the growth sessions coordinate on a minimum greater than 1, as do the two groups in the second experiment that grow slowly. However, growth does not always work, and even the groups that were grown slowly and carefully did not always remain coordinated efficiently.

The results of the second experiment also indicate that subjects may not be aware of the importance of slow, regular growth. All four of the managers grew quickly initially, and as a result none were initially successful in maintaining efficient coordination while growing the

[19]Since this was the first time that the group had grown successfully by more than one, there was noticeable agitation (e.g., fidgeting, longer response time) by several participants in the experiment. This points to the importance of the precedent since players were nervous because they had never experienced this particular kind of growth before and were therefore uncertain about what the outcome would be. As a result, the minimum fell below (to 4) what it had fallen the last few times the group grew (to 5).

group. However, following this early negative experience, two of the subjects in the role of manager were able to subsequently grow large groups coordinated on high levels of efficiency, using the same principles employed in the growth paths of experiment 1.

Another interesting result of these experiments was unintended. In two sessions, there is clear evidence that subjects are forming norms about how to react, not only to the previous minimum, but also to what happened in previous experience with growth. This is most striking in session E4, where all of the subjects changed their choices in exactly the same complicated way in several of the latter periods.

Taken together, the results indicate that growth (coupled with the exposure of new entrants to a group's history) is one reason why we observe large, successfully coordinated groups in the real world, but also that the growth process itself – and previous experience with growth – may be incredibly important in determining the success or failure of a growing group. Firms, which are often efficiently coordinated large groups, may only get to be so through careful growth.

There are also specific ways in which the results can be extended to provide real-world implications. For instance, the success of slow growth is linked to the exposure of new entrants to the group's history, which can be thought of as a form of the training or acculturation that new members of an organization frequently undergo. Moreover, the early failure of groups in the endogenous growth sessions appears to produce an instance of the common view in the business world that firms can "grow too fast." While – as with all experimental results – the value of these kinds of connections lies in the eye of the beholder, the main contribution of this paper is clear: by managing growth, it is possible to create efficient coordination in large groups.

References

[1] Arrow, K. J. *The Limits of Organization.* New York, New York: W. W. Norton & Company.

[2] Bicchieri, C. and C. Rovelli. 1995. Evolution and revolution: The dynamics of corruption. *Rationality and Society*, 7(2): 201-224.

[3] Cachon, G. P. and C. F. Camerer. 1996. Loss avoidance and forward induction in experimental coordination games. *Quarterly Journal of Economics* 111: 165-194.

[4] Camerer, C. F. and M. Knez. 1997. Coordination in organizations: A game-theoretic perspective. in *Organizational Decision Making* ed. Zur Shapira. Cambridge: Cambridge Series on Judgment and Decision Making.

[5] Camerer, C. F. and M. Knez. 2000. Increasing cooperation in social dilemmas by establishing a precedent of efficiency in coordination games. *Organizational Behavior and Human Decision Processes*, 82: 194-216.

[6] Chaudhuri, A., A. Schotter and B. Sopher. 2002. Talking ourselves to efficiency: Coordination in inter-generational minimum games with private, almost common and common knowledge of advice. Unpublished manuscript.

[7] Coase, R. H. 1937. The nature of the firm. *Economica* 4: 386-405.

[8] Cooper, R. W. 1999. *Coordination Games: Complementarities and Macroeconomics*. Cambridge: Cambridge University Press.

[9] Crawford, V. P. 1995. Adaptive dynamics in coordination games. *Econometrica* 63(1): 103-143.

[10] Duffy, J. and N. Feltovich. 2002. Do actions speak louder than words? An experimental comparison of observation and cheap talk. *Games and Economic Behavior*, 39(1): 1-27.

[11] Gittell, J. H. 2001. Supervisory span, relational coordination and flight departure performance: A reassessment of post-bureaucracy theory. *Organization Science*, 12(4): 467-482.

[12] Heath, C. and N. Staudenmeyer 2000. Coordination neglect: How lay theories of organizing complicate coordination in organizations. *Research in Organizational Behavior*, 22: 153-191.

[13] Knez, M. and C. F. Camerer. 1994. Creating expectational assets in the laboratory: Coordination in weakest-link games. *Strategic Management Journal* 15: 101- 119.

[14] Knez, M. and D. Simester. 2001. Firm-wide incentives and mutual monitoring at Continental airlines. *Journal of Labor Economics*, 19(4): 743-772.

[15] Lawrence, P. R. and J. W. Lorsch. 1967. *Organization and Environment: Managing Differentiation and Integration*. Boston: Harvard Business School Press.

[16] March, J. G. and H. A. Simon. 1958. *Organizations*. New York: John Wiley & Sons.

[17] Nanda, A. 1997. Implementing Organizational Change. *Comportamento Organizacional E Gastao*, 3(1). .

[18] Schelling, T. 1960. *The Strategy of Conflict*. Boston: Harvard University Press.

[19] Van Huyck, J. B., R. C. Battalio and R. O. Beil. 1990. Tacit coordination games, strategic uncertainty and coordination failure. *The American Economic Review* 80(1): 234-248.

[20] Weber, R. A. 2000. *Interdependence in Laboratory Groups and Organizations*. Doctoral Thesis. Division of Humanities and Social Sciences, California Institute of Technology.

[21] Weber, R. A., C. F. Camerer, Y. Rottenstreich and M. Knez. 2001. The illusion of leadership: Misattribution of cause in coordination games. *Organization Science*, 12(5): 582-598.

Figure 1. Choices in control sessions

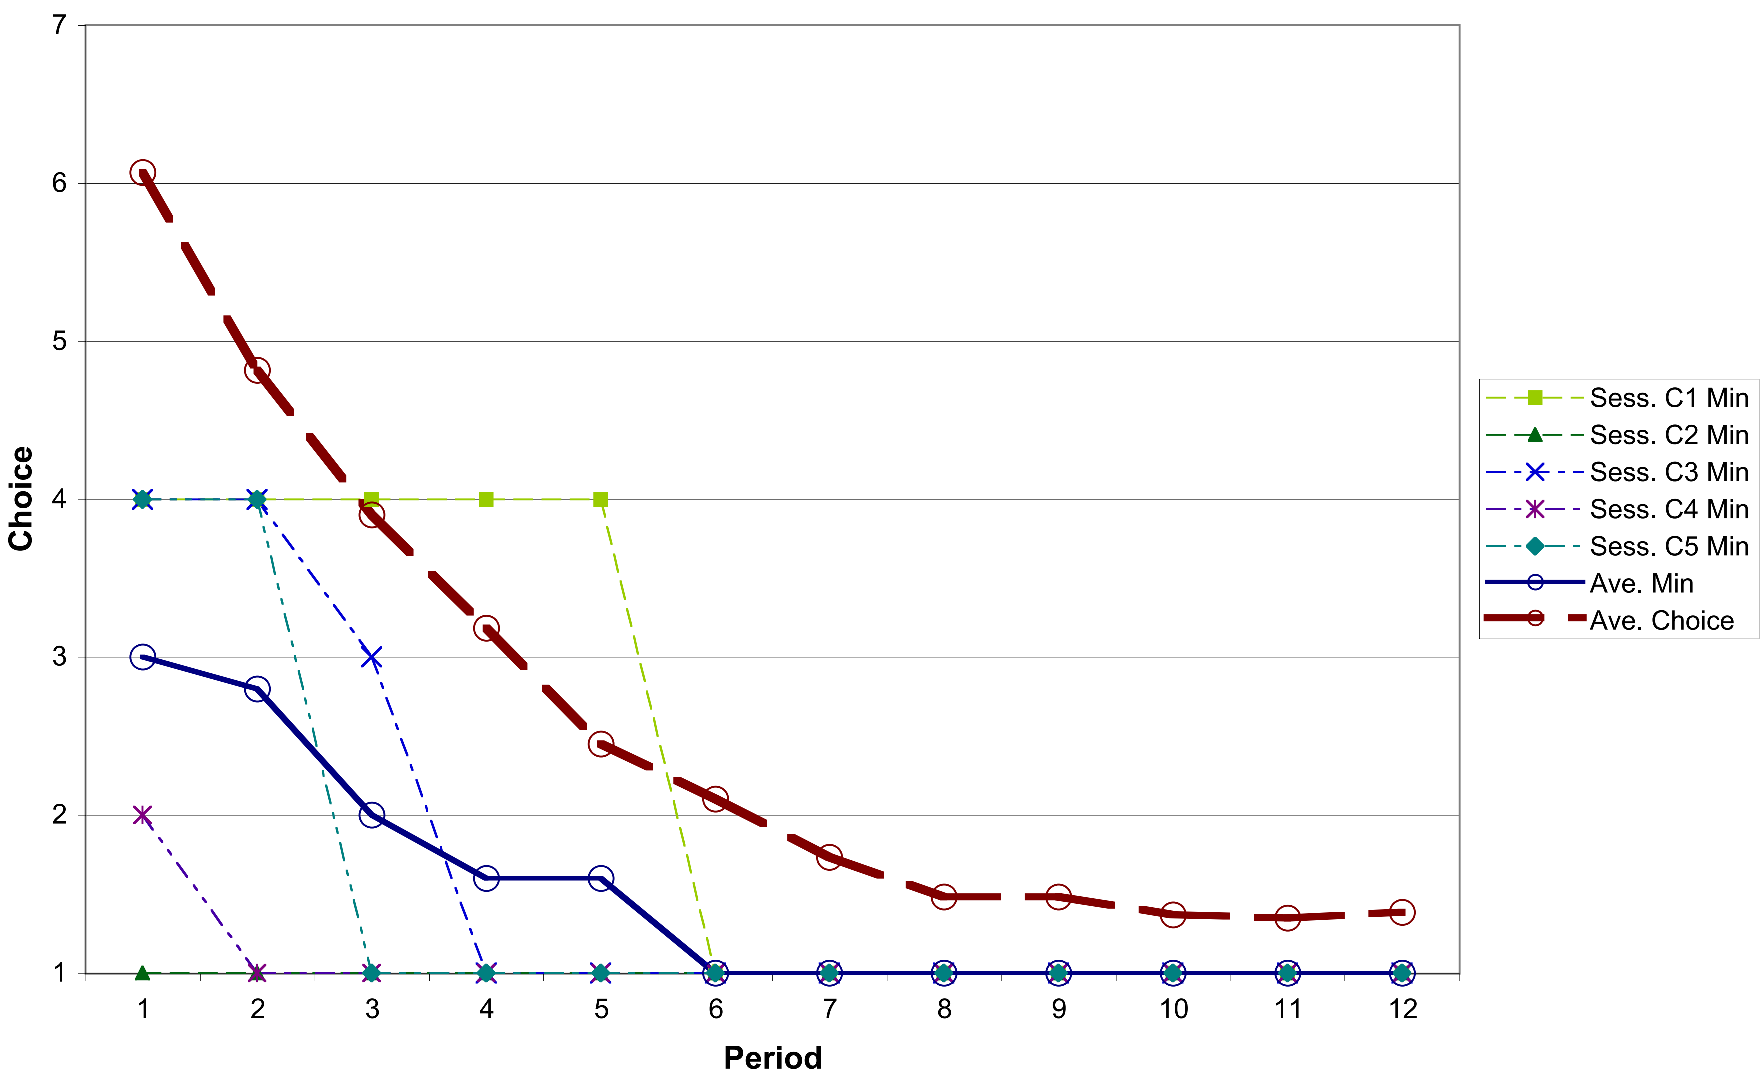

Figure 2. Pre-determined growth paths

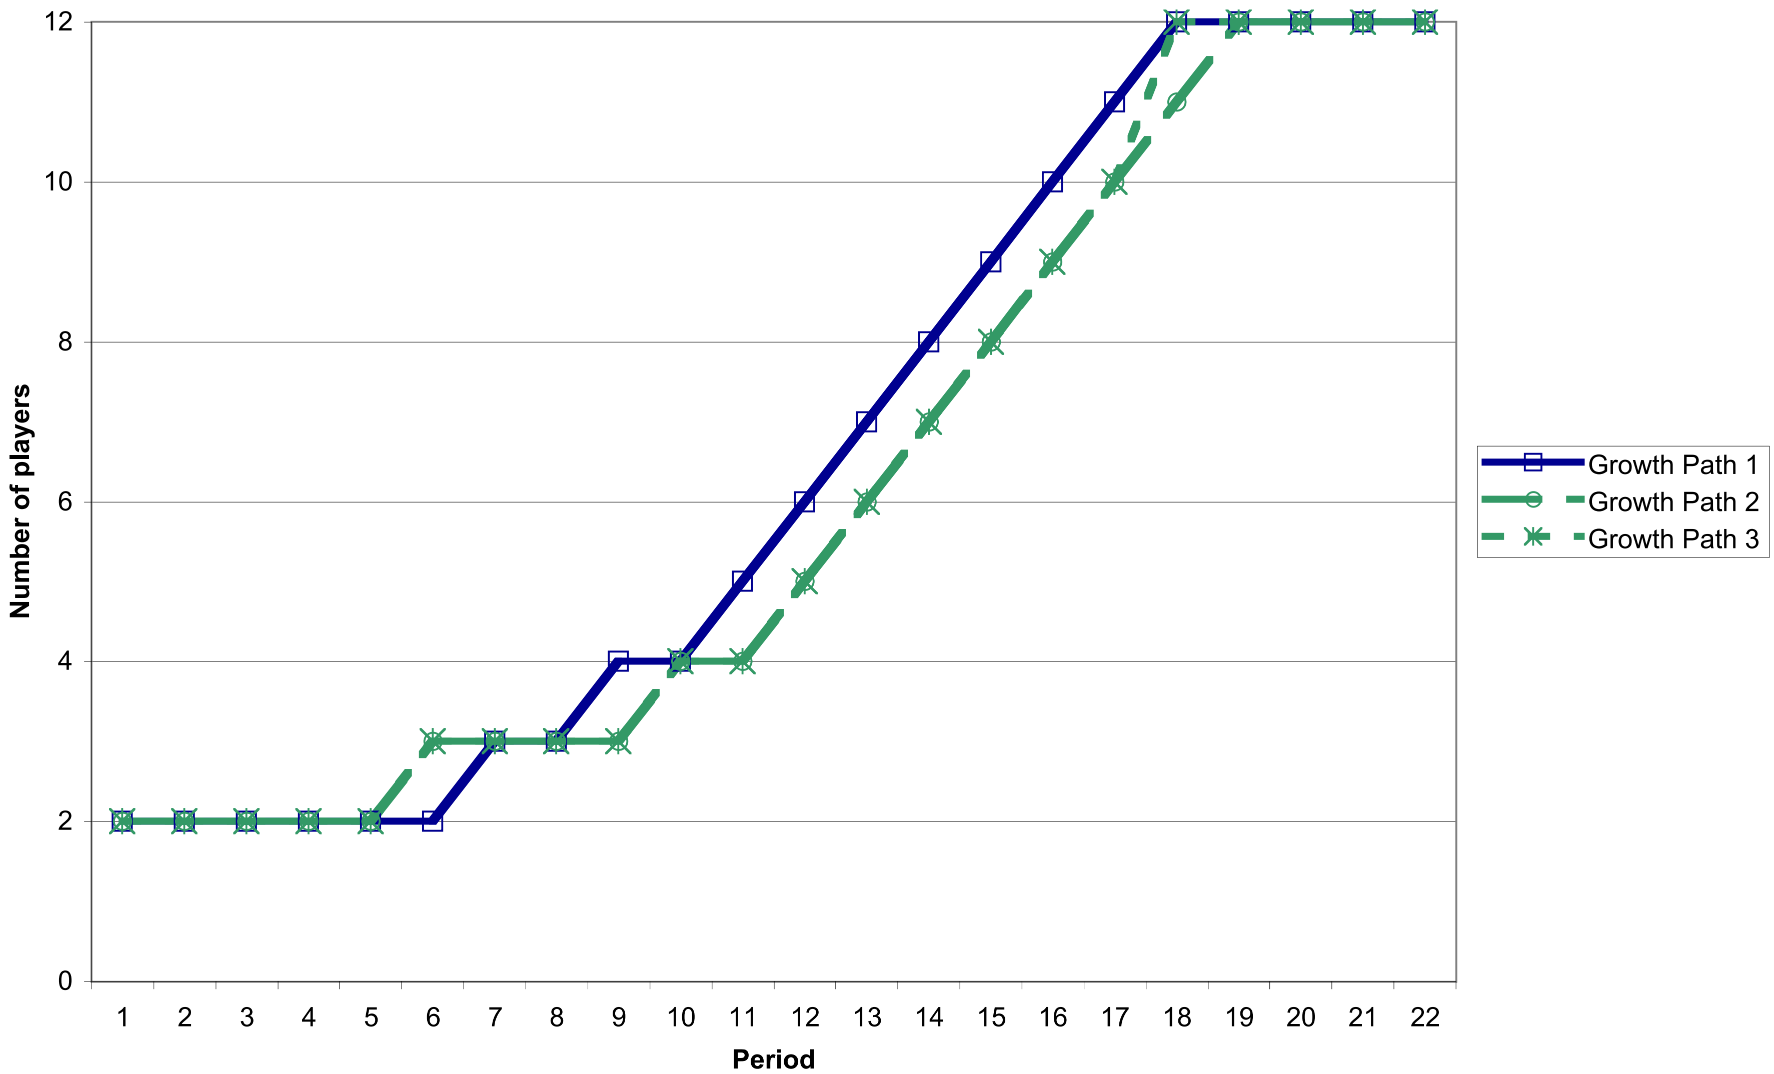

Fig 3a. Growth path 1 and period minima for sessions 1 & 2

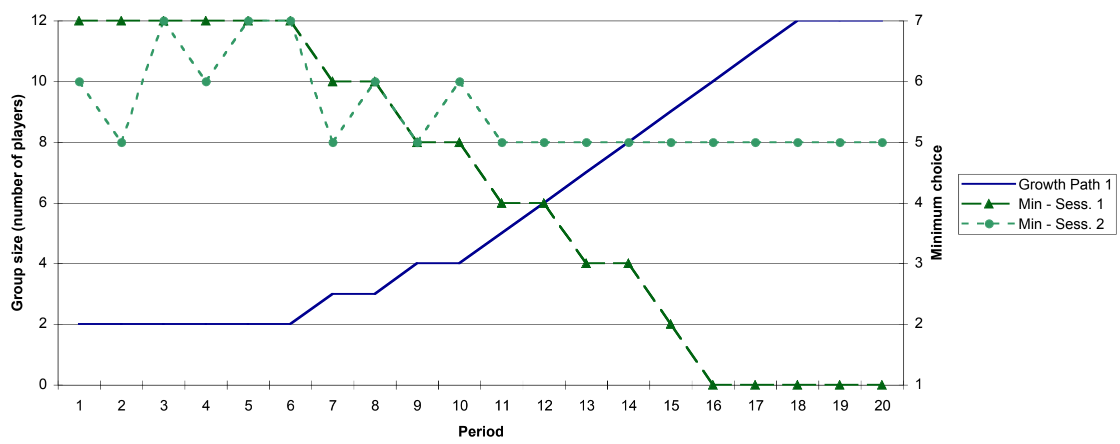

Fig 3b. Growth path 2 and period minima for sessions 3 & 4

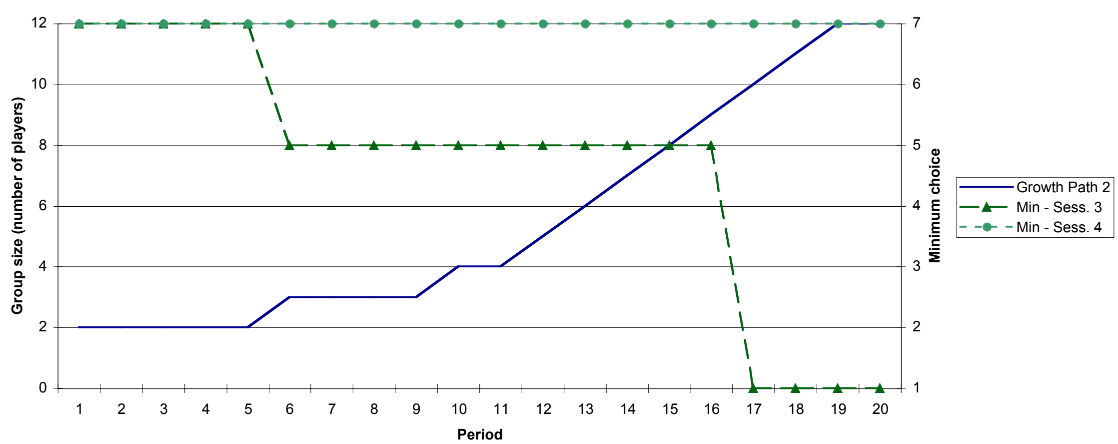

Fig 3c. Growth path and period minima for sessions 5 - 9

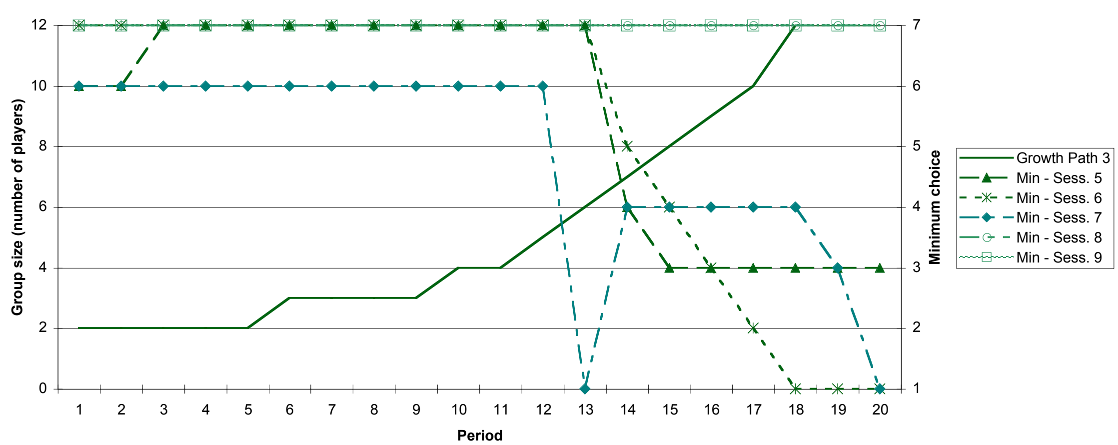

Figure 4. Endogenously determined growth paths

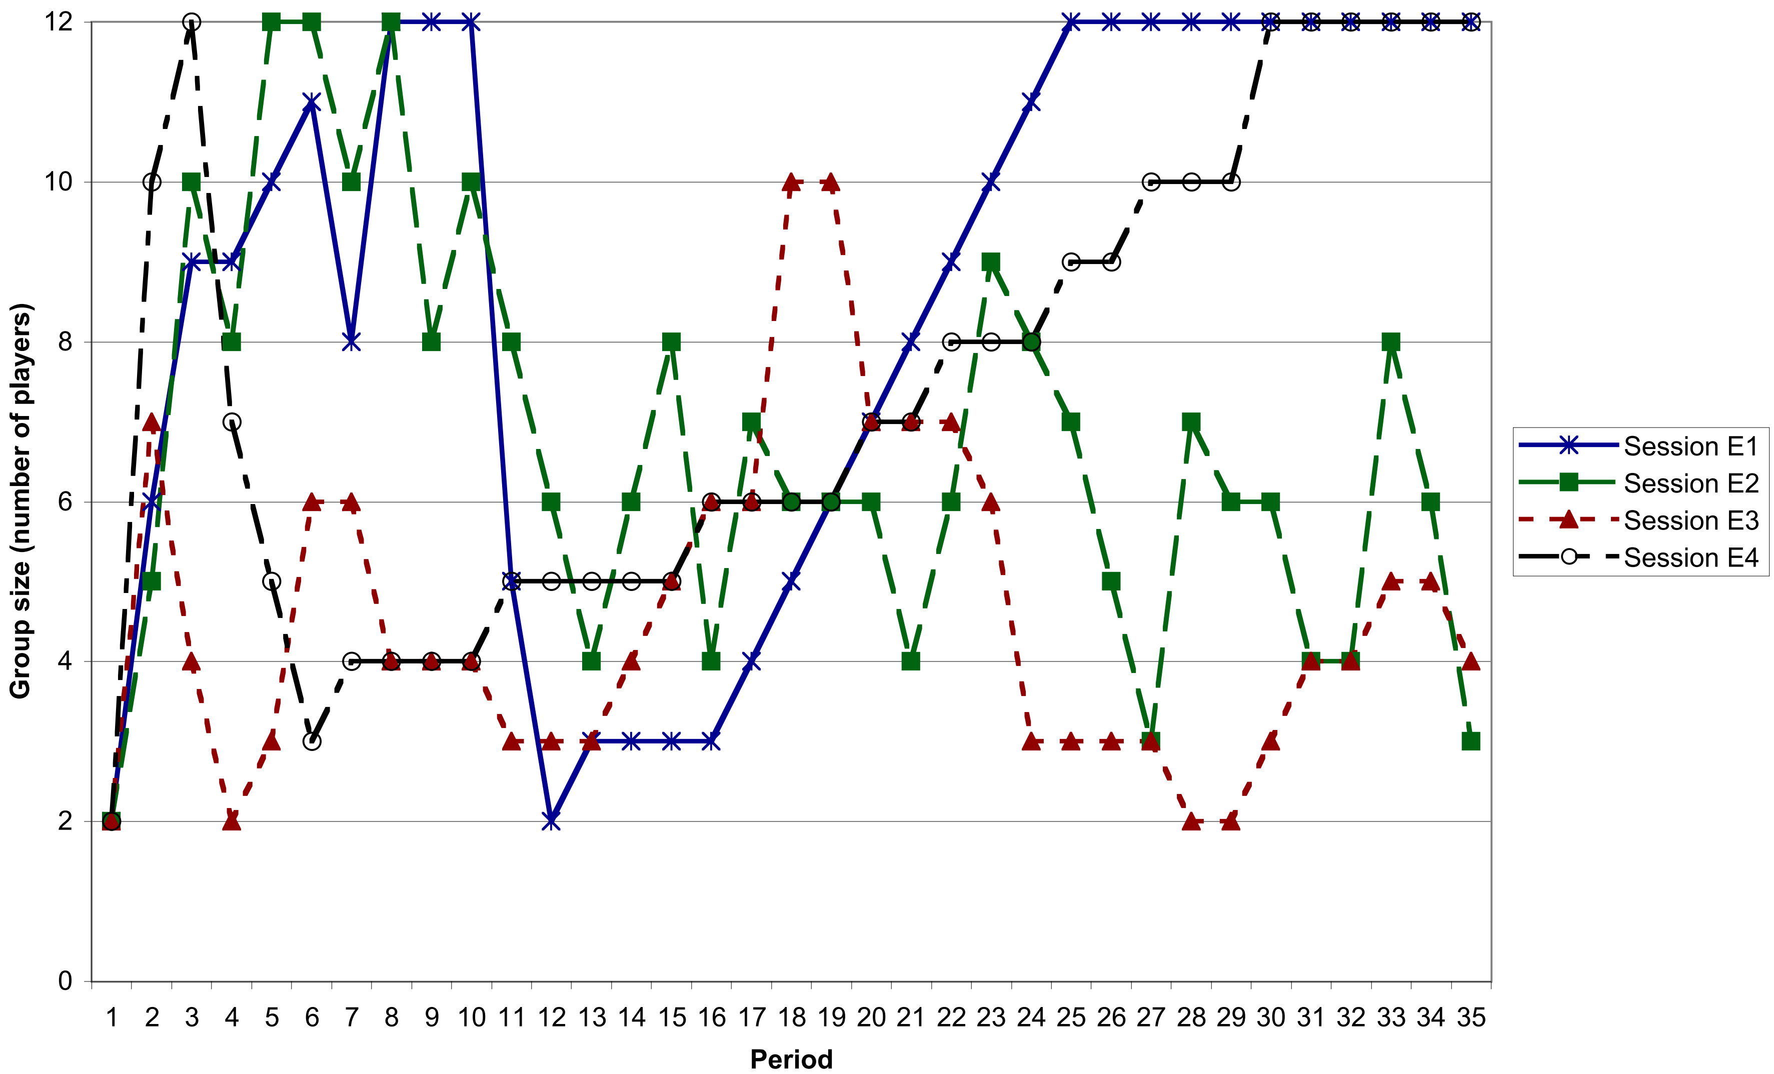

Figure 5a. Endogenous growth path and period minima in Session E1

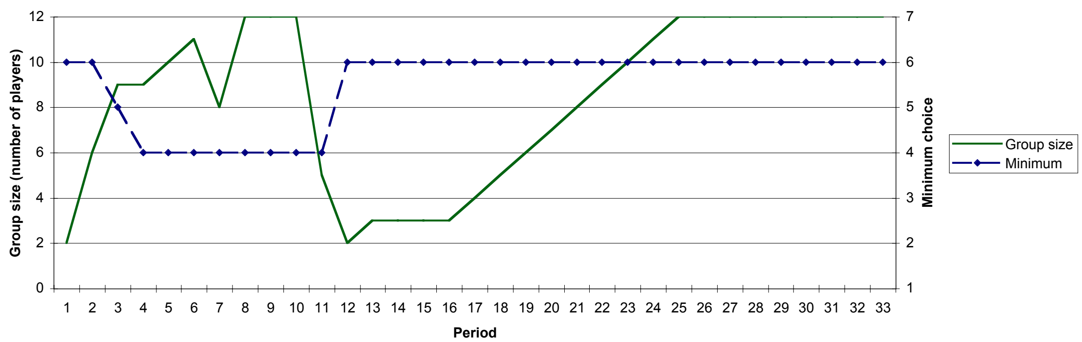

Figure 5b. Endogenous growth path and period minima in Session E2

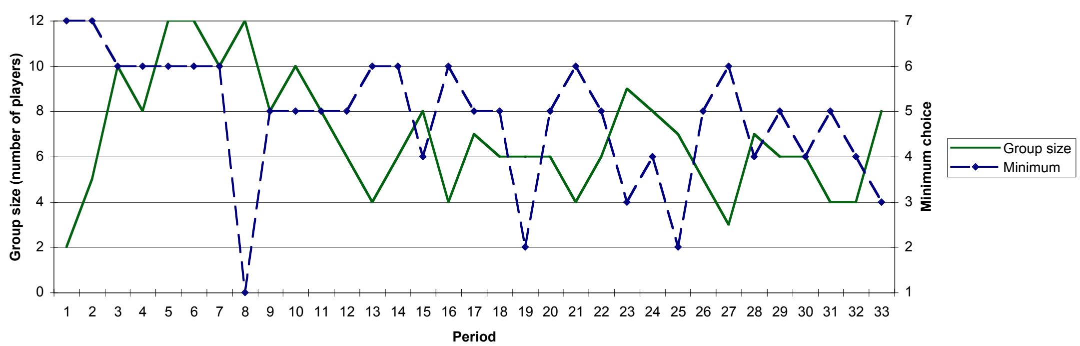

Figure 5c. Endogenous growth path and period minima in Session E3

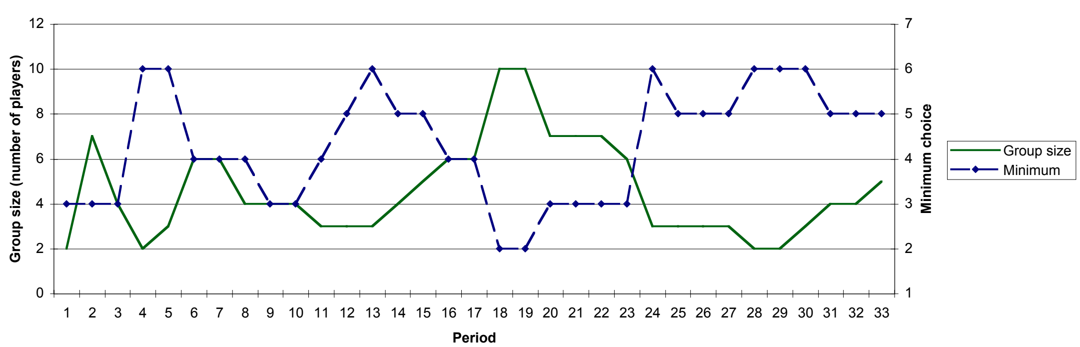

Figure 5d. Endogenous growth path and period minima in Session E4