Public statistics and private experience:
Varying feedback information in a take-or-pass game

Steffen Huck and Philippe Jehiel
University College London

September 27, 2004

Abstract

We study how subjects in an experiment use different forms of public information about their opponents' past behaviour. In the absence of public information, subjects appear to use rather detailed statistics summarizing their private experiences. If they have additional public information, they make use of this information even if it is less precise than their own private statistics—except for very high stakes. Making public information more precise has two consequences: It is also used when the stakes are very high and it reduces the number of subjects who ignore any information—public and private. That is, precise public information crowds in the use of own information. Finally, our results shed some light on unravelling in centipede games.

JEL codes:

Keywords: backward induction, analogy-based equilibrium, experiment

1   Introduction

Any form of belief-based learning requires that agents receive some information about their opponents' play. The information might be precise or noisy, gained through own experience or through other channels, regard all parts of opponents' strategies or only certain aspects thereof. In this paper we study how the provision of different forms of *public* information about others' behaviour affects subjects' play in a specific and highly stylized strategic interaction.

Specifically, we consider a variant of the so-called centipede game introduced by Rosenthal (1981). This game has the attractive feature that opponents have to take many similar decisions. At each decision node a

1

player must decide between Take or Pass. This allows us to vary the precision of information about others in a very natural way. Jehiel (2003) has proposed that boundedly rational agents who think about others who have to take similar decisions repeatedly might actually use coarse aggregate statics when forming their beliefs. In a centipede-like game a very coarse statistic would, for example, tell you that your opponent passed, *on average across all decision nodes*, $x\%$ of the time whereas a fine statistic would tell you the pass rate of your opponent at every decision node. The framework introduced in Jehiel (2003) permits the description of the interaction of players who base their beliefs on any such statistics, and the corresponding equilibrium called Analogy-Based Expectation Equilibrium is parameterized by the coarseness of the statistics used by the players.[1]

In our experiment, groups of subjects play several times a centipede-like game, and we provide subjects with public information about past behavior of their opponents. This information varies in its precision. It can either be based on averages across all nodes or it can be node-specific. Furthermore, the provided statistics are either averaged across the entire history of play or based on moving averages from the last five periods. This gives rise to a 2x2 design. In addition, we have a treatment without public information where agents have only their own experience.

When analysing the data from the treatment without public information we compute private statistics that are equivalent to the four different public statistics we provide in the other treatments. Analysing the decision data we find that subjects' behaviour is best explained by use of the most precise private statistic. That is, at any given node a subject's behavior is best explained as a reaction to what they learned from their own recent experience about what happens at the following node. The higher the pass rate that they experienced at the following node, the more likely they are to pass at the present node.

Surprisingly, we find that subjects' behavior is not much affected by their own past performance as measured, for example, by their payoff so far or the frequency with which they won games. Overall, this suggests that subjects who have just their own experience use their experience in a rather sophisticated manner. Their behaviour is much better described by a model with node-specific memory than by a model with a coarser memory structure where they just remember one average across all nodes or by a model in which they solely remember their past average performance.

---

[1]The statistic used by a player is referred to as an analogy class and defined as a partition of the decision nodes of the other player.

Using this result as a benchmark we can then analyse behaviour in the treatments with public information and we can study to what extent subjects use public information and private experience. We find that most of the time subjects who do make use of information make use of both types of information. However, the number of subjects who disregard any information (public as well as own) is considerably larger if the public information lacks precision. Thus, high-quality public information triggers the general use of information, private and public.

Our experiment also sheds new light on the issue of unravelling in take-or-pass games. From a theory viewpoint, more unravelling should be expected with the node-specific statistics than with the coarse statistics. With node-specific statistics, this is the classical insight: players can detect the exact node at which their opponent takes and as a consequence players should take earlier and earlier. With coarse statistics like average pass rates across all nodes players fail to identify when exactly their opponent takes. As a consequence a few pass decisions can stabilize play and prevent unravelling (see Jehiel 2003 for details on this).

As we have seen that subjects tend to rely on rather precise node-specific statistics when they have only their own experiences, one would, therefore, expect a lot of unravelling in the treatment without public information. However, as it turns out unravelling in that treatment is far from complete, which is due to the fact that without good public information there are a lot of "non-learners", i.e., subjects who always do the same regardless of what happens in the game. In fact, these non-learners pass so often that they help to stabilize play in the treatment without public information quite a bit above the Nash equilibrium. We also observe that there is much more pronounced unravelling in the presence of a precise public statistic than in the other treatments. This is because (1) a precise public statistic reduces the number of non-learners as we have stated before and non-learners appear to pass a lot; and (2) the pass rate at a given node falls dramatically when *both*, the private and public information about the pass rate at the next decision node are very low.[2] Since such a scenario can only occur in the treatment with node-specific statistics, it provides a further explanation for the more pronounced unravelling attached to this treatment.

Compared to other centipede games we find a much higher degree of unravelling, which is in part due to the payoff structure that we employ and that removes incentives for cooperation (we assume that the player who does

---

[2]Such an attitude toward extreme information is per se sufficient to generate unravelling.

end node	1	2	3	4	5
winner's payoff	£0.30	£0.60	£1.20	£2.40	£4.80
end node	6	7	8	9	10
winner's payoff	£9.60	£19.20	£38.40	£76.80	£153.60

Table 1: Winner's payoff in the TOL game.

not take gets invariably the same payoff irrespective of when the opponent takes). Thus, the non-unravelling observed in earlier centipede games (for example, McKelvey and Palfrey 1992 or Nagel and Tang 1998) is at least partly due to subjects' willingness to cooperate, perhaps, induced through "social preferences" (see, for example, Bolton and Ockenfels 2000).

The remainder of the paper is organized as follows. In Section 2 we introduce the game that we study in the experiment and offer a brief theoretical discussion. In Section 3 we introduce the experimental design and procedures. Section 4 contains the data and data analysis and Section 5 concludes.

2 The game

We study a version of Rosenthal's (1982) centipede game, a simple game of take-or-pass.[3] There are two players, called Even and Odd. The game has nine decision nodes. At each node one of the two players decides between Take or Pass. Odd decides at odd nodes 1, 3, 5, 7, and 9; Even at even nodes 2, 4, 6, and 8. If a player takes, the game is over. The game also ends if Odd passes at node 9. To make our terminology as simple as possible we will say that, if odd passes at node 9, node 10 is reached where Even automatically takes. Let the player who ends the game by taking be called the "winner" and the other player the "loser". This helps us to define the payoffs in a simple manner. The loser earns £0.10 regardless of what was the last decision node. The winner's payoff, on the other hand, depends on the last decision node. At node 1 it is £0.30. After that it doubles from node to node, reaching £153.60 at node 10. Table 1 shows the winner's payoff for all possible last decision nodes. Slightly abusing standard terminology we will refer to these in the following as *end nodes*.

The game has a number of Nash equilibria in pure strategies and infinitely many in mixed strategies. But all these equilibria induce the same

---

[3]Reny (1993) calls the same game Take-it-or-leave-it.

equilibrium *path* where Odd takes immediately at node 1. The unique subgame perfect equilibrium prescribes for both players to take at every node. Our game shares all these properties with Rosenthal's original game. However, there is one major difference: the payoff of the player who does not take is the same irrespective of when take occurs.[4] That is, one of the players is unambiguously the loser and the decision to pass cannot be interpreted as a cooperative move. Despite the pass move not being reducible to a cooperative move, there are still several reasons for why a player might wish to pass in our game. Intuitively, a player in a gambling mood (gambling here is about the behavior of the other player) may be thought of as being ready to pass at least in the early nodes of the game.

Several approaches have been proposed to explain why players may pass in the centipede game. A discussion of these approaches appears in Rubinstein (1998) and Jehiel (2003). Of particular relevance to this paper is the analogy-based expectation equilibrium approach introduced in Jehiel (2003): This approach assumes that players base their choice of strategy on the sole information about the average behavior of their opponent over bundles of nodes referred to as analogy classes. To illustrate the approach, assume that players use the coarsest analogy partition. That is, each player bundles all the nodes of the other player into a single analogy class and bases his choice of strategy on the sole information about the average pass rate of the other player throughout the game. A strategy profile in which the Even player passes in all nodes and the Odd player passes in all nodes except node 9 is an equilibrium under this assumed analogy grouping (see Jehiel 2003). Let us review the reasoning of the Even player. This player knows that the Odd player passes on average with probability 4/5 (this is the statistic that would emerge from the assumed strategy profile). Extrapolating that this is the Odd player's behaviors at each of his decision nodes 1-3-5-7-9, the Even player finds passing attractive even at decision node 8 ($4/5 \times 2 > 1$).

More generally, most coarse analogy grouping would allow players to pass for a few number of times. This is because such behaviors give rise to high pass rates, and, based on such an information, players would find it optimal to pass except toward reaching the end of the game. The logic of backward induction breaks down in this approach because players fail to identify exactly when their opponent stops passing.

---

[4]Such a specification is considered by Reny (1993).

3   Experimental design and procedures

In the experiment we vary the amount and type of public information players receive about other players whenever they have to make a decision. But before going into the details of the treatments, let us briefly describe those aspects of the design and procedures which were kept constant across treatments.

Each session had an even number of subjects who were recruited via Email. (Actual numbers varied from 12 to 18.) Roles were randomly assigned before the first round of the experiment and kept fixed during the entire course of a session. Sessions lasted 50 rounds. In each round subjects were randomly paired to play the extensive-form take-or-pass game.[5] Once the game ended they could infer their payoffs and a new round was started. Interaction took place via computer terminals using Tomlinson's (2003) package but instructions were handed out on paper (see Appendix A).

For actual monetary payoffs, we selected two rounds, one from the first 25, one from the second 25, which was, of course, known by subjects. On top of their game payoff, subjects received a £8 as a show-up fee. They were paid immediately after each session and no subject participated more than once.

The various treatments we considered are listed below. In the first four treatments, after the first round, subjects received some information about the past play of others whenever they had to make a decision. More specifically, odd subjects were informed about past *pass rates* of even subjects, i.e., some relative frequency with which the group of even subjects had passed in the past (and vice versa). In the last treatment labelled NO, no public information was available to subjects.

- In **treatment AN** pass rates were aggregated over **a**ll **n**odes and all previous rounds. Whenever Odd had to make a decision, a message was displayed saying that Even subjects had previously passed in x % of all instances. As average pass rates were only updated between rounds, the number x would not change in the course of a single game.

- In **treatment NS** pass rates were **n**ode-**s**pecific but still aggregated over all previous periods. Whenever Odd had to make a decision at node k, a message was displayed saying that Even subjects had

---

[5]Randomizations were done on the spot, i.e., the matching did not follow any predetermined pattern.

AN	NS	AN-MA	NS-MA	NO
$12 + 14$	$12 + 18$	$18 + 18 + 12$	$12 + 14 + 14$	$14 + 12 + 16$
$= 26$	$= 30$	$= 48$	$= 40$	$= 42$

Table 2: Subjects participating in treatments and sessions.

previously passed at node $k + 1$ in x % of all instances where node $k + 1$ was reached. If $k + 1$ had not been reached before, subjects were told so. Again, pass rates were updated after each round.

- **Treatment AN-MA** was identical to treatment AN with the exception that pass rates were calculated as moving averages from the last 5 rounds.

- **Treatment NS-MA** was identical to treatment NS with the exception that pass rates were calculated as moving averages from the last 5 rounds.

- Finally, in **treatment NO** subjects did not receive any information about average pass rates other than their experience from own past play.

In all, 186 subjects participated. Table 2 shows how the total of subjects was allocated to treatments and sessions.

4 Experimental results

4.1 Evolution of end nodes

A first approach to how subjects played the take-or-pass games is obtained by looking at the end nodes they reached. Figures 1, 2, and 3 show average end nodes for all sessions. Casual inspection of the figures reveals that without feedback information there is some rather light unravelling in the beginning but rather stable play over the last 30 or so rounds. In the AN treatments there seems to be some unravelling during the first half, followed by rather stable play towards the end. Only in session AN-MA 2 unravelling appears to continue until the very end. In general, there is not much difference between AN and AN-MA. Finally, in the NS treatments we observe more consistent downward trends in all sessions. In NS-MA 3 the unravelling is almost complete after the first half of the experiment after which average end nodes fluctuate between 1 and a little over 2. Again, there is no
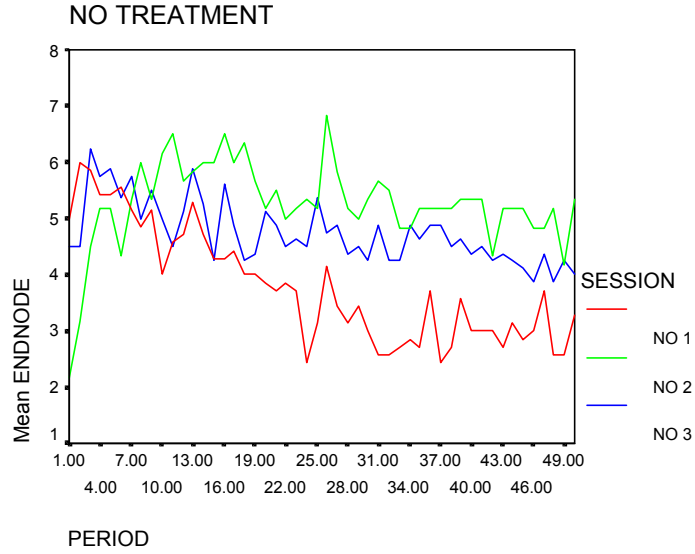
Figure 1: Evolution of mean end nodes in NO.

perceptible difference between information that is aggregated over all rounds and moving averages.

The differences in the degrees of unravelling that appear to be obvious from looking at the three figures can be validated statistically. We estimated random-effect panel regressions of the type

$$n_{it} = \alpha + \beta t + v_i + \varepsilon_{it}$$

where n is the end node reached by player i, t is time, v_i is the subject-specific error term and ε_{it} the residual. We estimated this model separately for all treatments (either for odd or for even players to avoid double counting) and for different spans of time. In particular, we analyzed the last 10, 20, 30, and more periods. For the last 30 (or more) periods β is significantly negative for all treatments. In treatment NO the time coefficient β becomes small and loses its high significance in the last 20 periods. And it becomes insignificant on all levels in the last 10 periods. In both, AN and AN-MA, β is significantly negative for the last 20 (or more) periods but loses its significance in the last 10 periods. In contrast, in both, NS and NS-MA, β is always—including the last 10 periods—significantly negative. Tables 3, 4, and 5 show the results of regressions for the first 30, the last 20, and the last

8

Figure 2: Evolution of end nodes in AN & AN-MA.

Figure 3: Evolution of end nodes in NS & NS-MA.

9

	NO		
	first 30	last 20	last 10
Const α	5.367***	4.735***	4.858***
	(.161)	(.418)	(1.142)
Time β	−.031***	−.016*	−.019
	(.006)	(.009)	(.025)
R^2	.03	.004	.002
N	630	420	210

Table 3: Estimated time trends in NO. Three stars indicate 1 percent significance level, two stars 5 percent, and one star 10 percent.

10, where we pool AN with AN-MA and NS with NS-MA.[6] In case of the latter two we do, however, exclude session NS-MA 1 when we estimate the last 20 and the last 10 periods because there no further unravelling can be expected. The tables show a clear pattern. Unravelling is strongest (in terms of the size of the β coefficient) and longest lasting in the treatments with node-specific information, a little weaker and finally fading in treatments with information about aggregate nodes and still weaker in the treatment where subjects can only rely on feedback about their own play.

We summarize our findings in

Result 1 Only when public information is very precise, i.e., both node-specific and based only on the immediate history, we find continuous unravelling until the very end of the experiment. In all other treatments, i.e., when public information is less precise or absent, behaviour eventually settles down and unravelling stops before the last period is reached.

4.2   Individual behavior

To understand what is driving the above result and to gain insight into how subjects use public and private information we shall now turn to the analysis of individual strategies. First of all, we shall classify subjects into four different categories:

1. *Pure-strategy types* (**P**) These are subjects whose behavior in all rounds is consistent with a fixed pure strategy, i.e., they always take at the

---

[6] Again we take all odd subjects as the repeatedly measured units. Results with even subjects are virtually identical.

	AN & AN-MA		
	first 30	last 20	last 10
Const α	5.647***	5.256***	3.529***
	(.098)	(.311)	(.932)
Time β	−.055***	−.031***	−.007
	(.004)	(.007)	(.020)
R^2	.09	.02	.0002
N	1690	740	370

Table 4: Estimated time trends in AN and AN-MA. Three stars indicate 1 percent significance level, two stars 5 percent, and one star 10 percent.

	NS & NS-MA		
	first 30	last 20	last 10
Const α	5.802***	6.567***	5.656***
	(.134)	(.331)	(.951)
Time β	−.066***	−.056***	−.036*
	(.004)	(.007)	(.025)
R^2	.06	.06	.01
N	1670	580	290

Table 5: Estimated time trends in NS and NS-MA. Three stars indicate 1 percent significance level, two stars 5 percent, and one star 10 percent.

	Types			
Treatments	P	AP	N	A
NO	–	6	2	34
	–	14.3%	4.8%	81%
AN	–	6	–	20
	–	23.1%	–	76.9%
AN-MA	2	6	2	38
	4.2%	12.5%	4.2%	79.2%
NS	1	2	–	27
	3.3%	6.7%	–	90.0%
NS-MA	2	1	1	36
	5.0%	2.5%	2.5%	90.0%
all	4	21	6	155
	2.2%	11.3%	3.2%	83.3%

Table 6: Types of subjects in the different treatments. Absolute number in first row, percentage in second row.

same node (or lose the game because their opponent takes at an earlier node). Among those would also be subjects who follow the backward induction solution and take at the earliest possible node. However, there are no subjects who exhibit such behavior.

2. *Almost pure-srategy types* (**AP**) These are subjects whose behavior is consistent with a fixed pure strategy in 90% of all rounds, i.e., we allow for five deviations over 50 periods.

3. *Non-rationalizable types* (**N**) These are subjects (in the role of Even) who sometimes pass at the very last node. We do not want to speculate about what drives them.

4. *Adaptive types* (**A**) All other subjects, i.e., subjects who do different things at different times, presumably because of different information or experiences.

Table 6 shows the absolute and relative frequencies of all four types in the different treatments.

A couple of observations are in order. While non-rationalizable types are extremely rare, the vast majority of subjects are adaptive players. Moreover, the share of adaptive players is highest in the treatments with node-

specific information which is statistically significant.[7] More precise information about different nodes makes experimentation and learning more attractive. This increase in adaptive behavior might be important for explaining the differences in unravelling observed above (see also further insight on this below). With fewer subjects reacting to information, behavior is more likely to settle down in treatment when public information is rather imprecise or absent. We summarize this in:

Result 2 The number of adaptive players, i.e. players who react to public information and/or private experience is increasing in the quality of the public information.

In the following we shall try to understand the behavior of adaptive players. In particular, we want to analyze how adaptive types react to their own experiences and the information provided. For that we can eliminate subjects who always do the same—in particular as their presence would cause a selection bias when we compare behavior at different nodes. (Those who always take at node 4 are never present at node 6.) Moreover, we focus on nodes 3, 4, 5, and 6. This is because there is very little variation at the first two nodes (with pass rates above 90% in all treatments) and not enough data for the last three nodes since they are reached too rarely.

To understand how own experiences enter subjects' decision rules, we first analyze treatment NO, where subjects have no additional information they can base their decision on. In a first step we estimate decision rules using random-effects linear probability models with just one independent variable capturing *subjects' own experience*. For reasons of parsimony we only examine the explanatory power of private statistics that are constructed in the same way as the statistics we provide in the other treatments, i.e., we examine models with private statistics for own experienced pass rates at specific nodes for the last five periods; at specific nodes for the entire history; averaged over all nodes and the last five periods; and, finally, averaged over all nodes and the entire history.[8]

The estimation results draw an extremely clear picture. For each node the R^2 is, by far, highest for the model with the most precise statistic, the

---

[7]Pooling the treatments with node-specific information and pooling all others, we find that the share of adaptive players compared to other types is significantly higher in the NS treatments at a level of 5.8% (two-tailed Pearson).

[8]In some cases the most detailed statistic that corresponds to the NS-MA treatment might be missing, simply because the next node has not been reached in the last five periods. Since we have found that pass rates given the statistic is missing are indistinguishable from pass rates when the statistic is zero we replaced missing values by zeros.

| | Treatment NO | | | |
	node 3	node 4	node 5	node 6
Const α	.486***	.400***	.282***	.193***
	(.040)	(.027)	(.027)	(.034)
Private Stat β	.605***	.950***	1.377***	2.048***
	(.040)	(.068)	(.144)	(.330)
R^2	.37	.30	.20	.21
N	676	558	373	150

Table 7: Estimated decision rules at nodes 3, 4, 5, and 6 in treatment NO. Three stars indicate 1 percent significance level.

one that corresponds to the NS-MA treatment.[9] Thus, it appears that in the absence of information provided by the experimenter subjects memorize their own experiences in a rather subtle way. Table 7 shows estimation results for the random-effects model

$$p_{it} = \alpha + \beta S_{it} + v_i + \varepsilon_{it} \tag{1}$$

where S_{it}^j is subject i's private statistic for the average pass rate at node $j+1$ in the last five periods previous to the decision period t, v_i^j is a subject-specific random effect and ε_{it}^j the remaining error. We estimate this model separately for nodes 3, 4, 5, and 6.

A few observations are immediate. Reflected by the increasing values of β, pass rates become more information sensitive at later nodes; reflected by the falling constants, *unconditional passing* becomes less pronounced at later nodes. However, there is significant unconditional passing at all nodes we estimated—even towards the end of the game.

Result 3 In the absence of public information players appear to store their own experiences in a rather sophisticated way. They appear to memorize node-specific information for the more recent past as opposed to more aggregate information about the past behavior of their opponent(s). Furthermore, we find that subjects' play becomes increasingly information sensitive as one moves further down the game tree.

---

[9] The average R^2 for models with the private NS-MA statistic for the first six nodes (where we have more than 100 observations for each) is 25.6%. For the models with NS statistics averaged over all periods the average R^2 is 12.7%. Finally, we the two models that examine private statistics averaged over all nodes the average R^2 are 0.9% (AN-MA) and 1.6% (AN).

	AN Treatments			
	node 3	node 4	node 5	node 6
Const α	−0.078	−1.001**	−.555*	1.400
	(.149)	(.436)	(.332)	(.908)
Private Stat β	.703***	1.045***	1.559***	2.065***
	(.031)	(.070)	(.126)	(.481)
Public Stat γ	.674***	1.689***	.952**	−1.542
	(.200)	(.549)	(.423)	(1.118)
R^2	.38	.28	.23	.12
N	1332	858	637	169

Table 8: Estimated decision rules at nodes 3, 4, 5, and 6 in AN treatments. One star indicates 10 percent significance level, two stars 5 percent, and three stars 1 percent.

How robust is this result when one controls for other aspects of subjects' learning? We have added variables for past performance to capture possible effects of aspiration levels (see, for example, Selten 1998, or Oechssler 2002) or learning direction theory (Selten and Buchta 1998) as well as a variable counting the periods to check for unravelling induced by the vanishing time horizon. The estimation results shown in Appendix B are unambigous. None of the extra variables has a consistent significant effect and the R^2's remain virtually unchanged. Moreover, the estimates for both, constant and coefficient of the private statistic, remain almost identical when one includes the other variables. Thus, any model selection criteria would select the most parsimonious model in (1). We summarize this in

Result 4 Subjects' behavior is not directly affected by their past performance nor by the outcome of their last game nor by the passing of time itself.

Let us now turn attention to estimating decision rules for the same nodes in the treatments where a public statistic is available. We shall estimate the following random-effects linear probability model

$$p_{it} = \alpha_0 + \beta_0 S_{it} + \gamma_0 P_{it} + v_i + \varepsilon_{it}$$

where P_{it}^j is the public statistic and all other variables are defined as before. We estimate this model separately for the AN and NS treatments and again for nodes 3, 4, 5, and 6. The results are shown in Tables 8 and 9.

15

	NS Treatments			
	node 3	node 4	node 5	node 6
Const α	.455***	.188***	.115***	.093
	(.039)	(.043)	(.048)	(.060)
Private Stat β	.448***	.710***	1.138***	1.098***
	(.039)	(.051)	(.077)	(.152)
Public Stat γ	.162***	.381***	.369***	.336***
	(.061)	(.080)	(.110)	(.127)
R^2	.33	.37	.35	.23
N	1180	1035	669	330

Table 9: Estimated decision rules at nodes 3, 4, 5, and 6 in NS treatments. One star indicates 10 percent significance level, two stars 5 percent, and three stars 1 percent.

Some observations can be made immediately. In general, subjects make use of all available information, private experience and publicly provided statistics. However, comparing the tables, it is apparent that the coefficients on the information variables are much bigger in the AN treatments than in the NS treatments. This may partly explain the comparatively limited unravelling in the AN treatments. (The second reason for less unravelling is, of course, as we have discussed above, the greater number of subjects who do not change their behaviour at all and always pass at the same nodes.) On the other hand, there appears to be some significant unconditional passing in the NS treatments which also puts a bound on unravelling.

Furthermore, we see that the public statistic has a more systematic effect when it is more precise. While it loses its significance in later rounds in the AN treatments it remains highly significant throughout all nodes in the NS treatments—even when the stakes are getting very high. On the other hand, less precise public information is ignored when the stakes are high.

We summarize our findings on the treatments with public information in

Result 5 In general, subjects use all types of information that is available. However, when the stakes are high and the public statistic is imprecise the public statistic is ignored.

16

Treatments	node 3	node 4	node 5	node 6
NO	.965	.738	.419	.269
AN	.862	.706	.306	.473
NS	.978	.788	.610	.449

Table 10: Pass rates of non-adative subjects. (Computed for last 45 periods to make them comparable with estimates for adaptive subjects.)

4.3 Unravelling

In the previous subsection we have offered two possible reasons for the different degrees of unravelling observed in our treatments. Here we want to reexamine the issue of unravelling in a little more detail. From a theory viewpoint, if players rely on the node-specific statistic, they should take as soon as they can. That is, from a learning perspective, a lot of unravelling should be expected. If players rely on a coarse statistic (average pass rate across all nodes) then a few Passes may take place before the system stabilizes (this is because players would fail to identify exactly when their opponent takes, see Jehiel 2003). Thus, less unravelling should be expected in this case.

Our finding in the no public information treatment that subjects mostly rely on their own most precise node-specific information would suggest that a lot of unravelling should appear in all treatments. But, as shown in Result 1 unravelling is much more pronounced if there is a precise public statistic (i.e. in treatment NS-MA). To understand this differential degree of unravelling, we introduce two additional slices of the data.

First, Table 10 shows the pass rates of the non-adaptive subjects for nodes 3 to 7. The table reveals that non-adaptive subjects tend to pass a lot which does put a bound on the possibilities for unravelling. Of course, in the NS treatments there are only roughly half as many non-adaptive subjects than in the two others (10% as compared to 19% in NO and 21.6% in the AN treatments). This basically halves the impact of non-adaptive play in the NS treatments and increases the room for unravelling. This is our first explanation for why unravelling is more pronounced in the NS treatments.

Second, Table 11 shows how subjects in treatment NS-MA react to extreme information, i.e., to very small observed *and* experienced pass rates. In principle, the reaction to extreme information alone can explain (different degrees of) unravelling. The table shows the pass rates of adaptive subjects for those cases where either private or public statistic (or both) about the next node were below 10%.

information	node 3	node 4	node 5	node 6
own < 0.1	.473	.324	.300	.239
public < 0.1	.208	.111	.207	.204
both < 0.1	.198	.000	.127	.089

Table 11: Pass rates of non-adative subjects. (Computed for last 45 periods to make them comparable with estimates for adaptive subjects.)

The table illustrates that subjects are much more careful when the public statistic is very bleak. If there own experience has been bad they are much more optimistic and incidentally these numbers are basically the same in all other treatments. However, if both statistics point to low pass rates of others at the next node subjects do become very careful indeed and at node 4 we even observe that nobody passes (out of 31 cases). Hence, we should expect that with very precise public information the adaptive subjects on their own would converge very closely to Nash equilibrium play. With less precise public information this is far less likely.

It is interesting to compare our data on unravelling with previous experiments on centipede games in which also the loser's payoff increases over time. Both payoffs increasing over time generates, of course, strong incentives for cooperation. McKelvey and Palfrey (1992) study such centipede games with four and six decision nodes. The games are repeated over ten rounds and while there is some unravelling it is very limited. In the games with six decision node the average end node falls from 4.29 in the first half of the experiment to 3.98 in the second half. Nagel and Tang (1998) study a reduced normal-form version of a centipede game with 14 decision nodes. Subjects play this game repeatedly for 100 periods and, quite amazingly, there is no unravelling at all. In fact, in some sessions the average end node even increases over time (see their Figure 4, p.362). Thus, a comparison with our data suggests that the different payoff structure that we employ makes a big difference even if it induces the same best reply correspondence. Taking away the possibility for mutually beneficial cooperation, there is much more unravelling.[10]

---

[10]There is also much more unravelling in high-stakes *three*-player centipede games as studied by Rapoport, Stein, Parco, and Nicholas (2000). Remarkably, the same pattern has also been observed in quite different games. Huck, Normann, and Oechssler (2004), for example, report that subjects in Cournot markets only manage to collude if there are no more than two competitors.

5   Concluding remarks

Our experimental results can be summarised as follows. When faced with their sole experience subjects seem to use their memory in a rather sophisticated way: They do not rely much on their past performance, and rather rely on some quite sophisticated estimate of their opponent's behavior that varies from one decision node to the other.

When public statistics are introduced subjects make use of both, their own experience and the public statistic—even if the latter is rather imprecise. Only when the stakes are very high coarse public statistics are ignored. Another effect of providing precise public information is that more subjects start using information—public and private. That is, precise public information crowds in the use of private information and reduces the number of non-learners. This may have important consequences for a variety of games and economic applications and deserves further study. In that sense we advocate the methodoly of varying public information in experiments.

References

[1] Bolton, G.E., and A. Ockenfels (2000). ERC: A Theory of Equity, Reciprocity, and Competition. *American Economic Review* 90, 166-193.

[2] Huck, S., H.T. Normann, and J. Oechssler (2004). Two are few and four are many: Number effects in experimental oligopoly. *Journal of Economic Behavior and Organization* 53, 435-446.

[3] Jehiel, P. (2003). Analogy-based expectation equilibrium. *Journal of Economic Theory* forthcoming.

[4] McKelvey, R.D., and T.R. Palfrey (1992). An experimental study of the centipede game. *Econometrica* 60, 803-836.

[5] Nagel, R., and F.F. Tang (1998). Experimental results on the centipede game in normal form: An investigation on learning. *Journal of Mathematical Psychology* 42, 356-384.

[6] Oechssler, J. (2002). Cooperation as a result of learning with aspiration levels. *Journal of Economic Behavior and Organization* 49, 405-409.

[7] Rapoport, A., W.E. Stein, J.E. Parco, and T.E. Nicholas (2000). Equilibrium play and adaptive learning in three-person centipede game. University of Arizona Working Paper.

[8] Reny, P. (1993). Common belief and the theory of games with perfect information. *Journal of Economic Theory* 59, 257-274.

[9] Rosenthal, R. (1981). Games of perfect information, predatory pricing, and the chain store paradox. *Journal of Economic Theory* 25, 92-100.

[10] Selten, R. (1998). Aspiration adaptation theory. *Journal of Mathematical Psychology* 42, 191-214.

[11] Selten, R., and J. Buchta (1998). Experimental sealed bid first price auctions with directly observed bid functions. *Games and Human Behavior: Essays in Honor of Amnon Rapoport* (eds. D. Budescu, I. Erev, and R. Zwick). Hillsdale, NJ: Erlbaum.

[12] Tomlinson, C. (2003).

A   Instructions for treatment NS-MA

Welcome to our experiment!

Please read these instructions carefully! Do not speak to your neighbours and keep quiet during the entire experiment! If you have a question, raise your hand and the experimenter will see you.

In this experiment you will repeatedly make decisions. Doing this you can earn real money. How much you earn depends on your decisions and on the decisions of other participants. All participants receive the same instructions. You will stay anonymous to us and to the other participants.

The experiment will have 50 rounds. In each round you will be randomly matched with one other participant with whom you are going to interact. Your payoff in one particular round depends solely on the decisions taken by yourself and by the other participant that you were matched with. After the experiment we will randomly select two rounds that will be paid off for real. One payoff round will be selected randomly from rounds 1-25 and one payoff round will be selected randomly from rounds 26-50.

There are two groups of participants in this experiment, Odd participants and Even participants. In each round each Odd participant will be randomly matched with an Even participant and vice versa.

Each round consist of up to 10 stages. Odd participants have to make decisions in odd stages, Even participants have to make decisions in even stages. In each stage, the decision is between TAKE and PASS. If a participant chooses PASS the round continues into its next stage. If a participant

chooses TAKE the round is over. Finally, if the 10th stage is reached, PASS is no longer an option and the participant has to TAKE.

The payoffs are as follows. The payoff of the participant who has chosen TAKE depends on the stage in which he has done so. The payoff of the other participant is 10p regardless of the stage in which TAKE has been chosen.

The payoff for the participant who has chosen TAKE follows a simple rule. In stage 1 it is 30p. After that it will doubled in each and every stage, finally, reaching £153.60 in stage 10. The following table shows you all the numbers. [Here they are shown a table similar Table 1.]

When making a decision you will have some information about the past. More specifically, you will be told how often, on average, the other participants have chosen to PASS at the next stage in the past five rounds. For example, if you are an Even participant deciding at stage 4 you will be told how often Odd participants passed in stage 5 in the past five rounds. Averages will be updated after each round and will be available to you whenever you make a choice. If the next stage has not been reached in the past five rounds, you will be told there is no data available. (During the first five rounds, averages will be based on all previous decisions.)

These are the rules. Everything will happen exactly as specified by them. Enjoy.

B Estimation results with further variables

Here we show further estimation results for treatment NO. next to own experienced pass rates, we include a variable capturing the possible effect of aspiration levels, π_{it}, subject i's total payoff up to period $t-1$; a variable capturing the possible effect of learbning direction theory, $L_{it}^j = 1$ if subject i lost in period t on node $j+1$ (and 0 otherwise); and a variable capturing the passing of time, t itself. Table 12 shows the estimates for the random-effects model

$$p_{it}^j = \alpha^j + \beta^j S_{it}^j + \gamma^j \pi_{it} + \lambda^j L_{it}^j + \kappa t + v_i^j + \varepsilon_{it}^j. \qquad (2)$$

where $j = 3, 4, 5, 6$ is the decision node and all other variables are defined as before. Conventional wisdom would let us expect a positive sign for γ (if aspiration levels matter), a negative sign for λ (if subjects immediately react to bad experiences in the way of moving towards better responses), and a negative sign for κ (if there is unravelling because of a vanishing shadow of the future). Strikingly, none of these expectations turns out to be true.

	Treatment NO			
	node 3	node 4	node 5	node 6
Const α	.418***	.428***	.288***	.057
	(.050)	(.058)	(.070)	(.091)
Private Stat β	.701***	.942***	1.283***	2.177***
	(.043)	(.073)	(.156)	(.343)
Aspirations γ	−.006	.003	.013	.000
	(.007)	(.006)	(.009)	(.008)
Learn dir λ	.093***	−.008	−.036	.085
	(.027)	(.036)	(.055)	(.102)
Period κ	.001	−.001	−.002	.005
	(.001)	(.001)	(.002)	(.003)
R^2	.39	.30	.20	.23
N	676	558	373	150

Table 12: Estimated decision rules at nodes 3, 4, 5, and 6 in treatment NO. Stars inbdicate significance levels as before.