WZB

Wissenschaftszentrum Berlin
für Sozialforschung

David Danz
Steffen Huck
Philippe Jehiel

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

Discussion Paper

SP II 2016–201

February 2016

(WZB) Berlin Social Science Center

Research Area
Markets and Choice

Research Units
Market Behavior and Economics of Change

Wissenschaftszentrum Berlin für Sozialforschung gGmbH
Reichpietschufer 50
10785 Berlin
Germany
www.wzb.eu

David Danz, Steffen Huck, Philippe Jehiel
Public statistics and private experience: Varying feedback information in a take-or-pass game

Affiliation of the authors:

David Danz
WZB and TU Berlin

Steffen Huck
WZB and UCL

Philippe Jehiel
PSE and UCL

Wissenschaftszentrum Berlin für Sozialforschung gGmbH
Reichpietschufer 50
10785 Berlin
Germany
www.wzb.eu

Abstract

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

by David Danz, Steffen Huck and Philippe Jehiel*

We study how subjects in an experiment use different forms of public information about their opponents' past behavior. 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.

Keywords: Backward induction; analogy-based expectation equilibrium; learning; experiment

JEL classification: C72, C92, D83, D84

* E-mail: david.danz@wzb.eu, steffen.huck@wzb.eu, jehiel@pse.ens.fr.

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' behavior 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 player must decide between Take or Pass. This allows us to vary the precision of information about others in a natural way. Jehiel (2005) 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 (2005) 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 analyzing 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. Analyzing the decision data we find that subjects' behavior is best explained by use of the most precise private statistic. That is, at any given node a subject's behavior is best ex-

[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.

plained 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 behavior 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.

Using this result as a benchmark we can then analyze behavior 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 2005 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 the *public* (rather than the private) 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 observed in 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 (the player who does 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 the last decision node that was reached. 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. Figure 1 shows the winner's payoff for all possible last decision nodes. Slightly abusing standard terminology

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

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

we will refer to these in the following as *end nodes*.

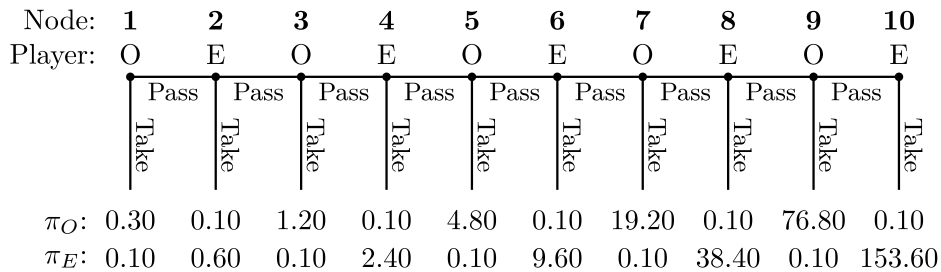

Figure 1: The take-or-pass game. The payoffs of player Odd (O) and Even (E) are denoted by π_O and π_E, respectively.

The game has a number of Nash equilibria in pure strategies and infinitely many in mixed strategies. But all these equilibria induce the same 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 (2005). Of particular relevance to this paper is the analogy-based expectation equilibrium approach introduced in Jehiel (2005): 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 at all nodes and the Odd player passes at all nodes except node 9

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

is an equilibrium under this assumed analogy grouping (see Jehiel 2005). 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 behavior 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 a few 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.

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 which were kept constant across treatments.

The experimental sessions lasted 50 rounds. In each round subjects were randomly paired to play the extensive-form take-or-pass game.[5] Roles were randomly assigned before the first round of the experiment and kept fixed during the entire course of a session. Once the game ended, they could infer their payoffs and a new round was started. 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.[6]

The five treatments we considered are listed in Table 1. In treatment NO, no public information was available, i.e., subjects did not receive any information about the past play of others other than their experience from own past play.[7]

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

[6] We paid two randomly selected rounds to decrease the variance in subjects' final payments. For a recent discussion of (improvements of) the random incentive system see Johnson et al. (2015).

[7] As in the other treatments, the computer screens in the NO treatment did not provide subjects with a history of their own previous choices or the past actions of their immediate opponents. Thus, in the NO treatment subjects had to memorize the actions of the other players throughout the experiment to form private statistics analogous to the public information provided in the AN and NS treatments.

Table 1: Overview of treatments and sessions.

Treatment	**NO** (No Informaton)	**AN** (All Nodes)	**NS** (Node specific)	**AN-MA** (All nodes, moving average)	**NS-MA** (Node specific, moving average)
Information about other player's average pass rate	./.	Average over all nodes, over all previous rounds	Subsequent node, over all previous rounds	Average over all nodes, over last five rounds	Subsequent node, over last five rounds
Sessions	3	2	2	3	3
Subjects	42 (14+12+16)	26 (12+14)	30 (12+18)	48 (18+18+12)	40 (12+14+14)

In the remaining four treatments, subjects received some information about the past play of others. More specifically, from round two on, odd subjects were informed about *pass rates* of even subjects in previous rounds, i.e., some relative frequency with which the group of even subjects had passed in the past (and vice versa). This public information was updated each round and presented to the subjects before they made their choice.

In treatment AN pass rates were aggregated over <u>a</u>ll <u>n</u>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 <u>n</u>ode-<u>s</u>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 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 <u>m</u>oving <u>a</u>verages from the last 5 rounds. Similarly, treatment NS-MA was identical to treatment NS with the exception that pass rates were calculated as moving averages from the last 5 rounds.

The experiment was conducted at the Experimental Laboratory of the ELSE Centre at the Department of Economics at UCL in 2002. Subjects were recruited from the ELSE experimental subject pool, which includes mainly UCL undergraduate students across all disciplines. Each session had an even number of subjects who were recruited via Email. (Actual numbers

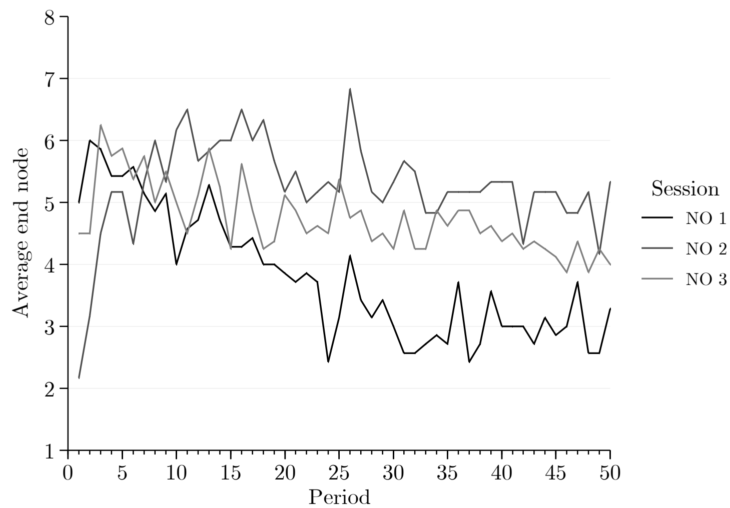

Figure 2: Evolution of mean end nodes in treatment NO.

varied from 12 to 18.) In all, 186 subjects participated. Subjects had no previous experience with this experiment and each subject participated in only one session. Table 1 shows how the total of subjects was allocated to the treatments. The experiment was computerized using EXPECON (Tomlinson 2005) but instructions were handed out on paper (see Appendix A). Subjects were paid immediately after each session. On top of their game payoff, subjects received a £8 as a show-up fee.

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 2, 3, and 4 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 1 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 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_{it} 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 and for different spans of time. In particular, we analyzed the last 10, 20, 30, and more periods. Tables 2, 3, and 4 show the results of regressions, where we pool AN with AN-MA and NS with NS-MA.[8] 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. For the last 20 (or more) periods β is significantly negative for all treatments (one-sided tests). In treatment NO and the AN treatments the time coefficient β becomes small and becomes insignificant in the last 10 periods. In contrast, in the NS treatments β is always—including the last 10 periods—significantly negative. Furthermore, unravelling is strongest (in terms of the size of the β coefficient) in the treatments with node-specific information, a little weaker in treatments with information about aggregate nodes and still weaker in the treatment where subjects can only rely on feedback about their own play.[9] This pattern holds for all time intervals except for the comparison of the NO treatment and the AN treatments when the regression is based on the last 40 periods. We summarize our findings in

Result 1 Only when public information is precise, i.e., in treatments with node-specific information, 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, unraveling is less pronounced and

[8]We take all odd subjects as the repeatedly measured units. Results with even subjects are virtually identical.

[9]A joint regression over all treatments and periods of the end node on a constant, the period, and two interactions of the period with a dummy for the AN treatments and the NS treatments, respectively, reveals that the time trends in the NO treatment, the AN treatments, and the NS treatments are significantly different from each other ($p \leq 0.018$ for all comparisons).

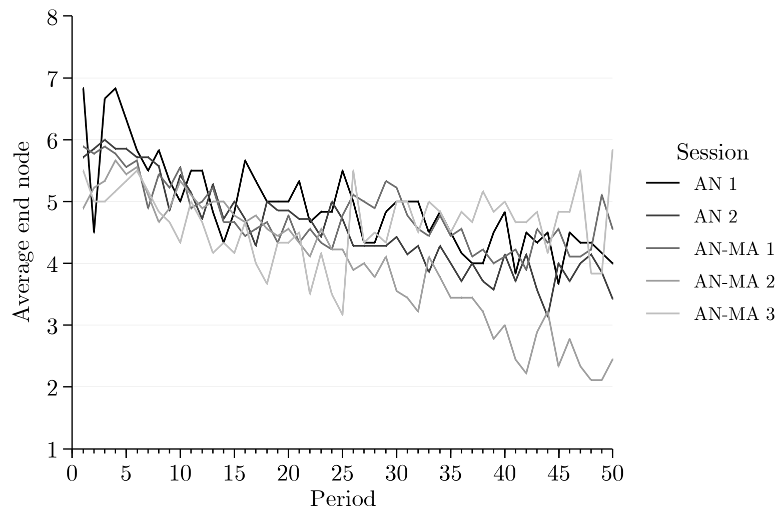

Figure 3: Evolution of end nodes in treatments AN and AN-MA.

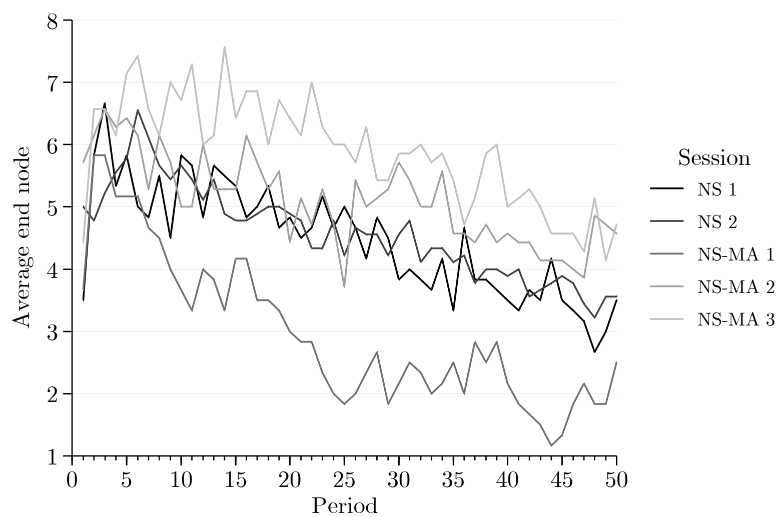

Figure 4: Evolution of end nodes in treatments NS and NS-MA.

Table 2: Estimated time trends of end nodes in treatment without public information (NO).

Dependent var.:	End node				
Time interval:	All Periods	$t > 10$	$t > 20$	$t > 30$	$t > 40$
Period β	-0.030^{***}	-0.033^{***}	-0.021^{***}	-0.016^{**}	-0.019
	(0.003)	(0.003)	(0.005)	(0.009)	(0.025)
Constant α	5.341^{***}	5.451^{***}	4.986^{***}	4.735^{***}	4.858^{***}
	(0.165)	(0.215)	(0.266)	(0.418)	(1.143)
$\chi^2_{(1)}$	116.740	106.407	20.263	3.099	0.580
N	1050	840	630	420	210

Note: Linear regressions with individual random effects. Values in parentheses represent standard errors. Stars represent p-values of one-sided tests for the time trends and two-sided tests for the constant: $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 3: Estimated time trends of end nodes in treatment with public information aggregated over nodes (AN and AN-MA).

Dependent var.:	End node				
Time interval:	All Periods	$t > 10$	$t > 20$	$t > 30$	$t > 40$
Period β	-0.038^{***}	-0.031^{***}	-0.031^{***}	-0.031^{***}	0.007
	(0.002)	(0.002)	(0.004)	(0.007)	(0.020)
Constant α	5.517^{***}	5.242^{***}	5.247^{***}	5.256^{***}	3.529^{***}
	(0.101)	(0.123)	(0.181)	(0.311)	(0.932)
$\chi^2_{(1)}$	427.454	154.423	67.569	21.731	0.106
N	1850	1480	1110	740	370

Note: Linear regressions with individual random effects. Values in parentheses represent standard errors. Stars represent p-values of one-sided tests for the time trends and two-sided tests for the constant: $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 4: Estimated time trends of end nodes in treatment with node-specific public information (NS and NS-MA).

Dependent var.:	End node				
Time interval:	All Periods	$t > 10$	$t > 20$	$t > 30$	$t > 40$
Period β	-0.049^{***}	-0.050^{***}	-0.044^{***}	-0.056^{***}	-0.036^{**}
	(0.002)	(0.003)	(0.004)	(0.007)	(0.021)
Constant α	5.898^{***}	5.909^{***}	5.696^{***}	6.567^{***}	5.656^{***}
	(0.183)	(0.203)	(0.235)	(0.331)	(0.951)
$\chi^2_{(1)}$	533.198	347.756	129.727	55.873	3.018
N	1750	1400	1050	580	290

Note: Linear regressions with individual random effects. Values in parentheses represent standard errors. Stars represent p-values of one-sided tests for the time trends and two-sided tests for the constant: $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

behavior eventually settles down such that 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 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-strategy types* (**AP**) These are subjects whose behavior is consistent with a fixed pure strategy in 80% of all rounds, i.e., we allow for ten deviations over 50 periods.

3. *Non-rationalizable types* (**N**) These are subjects (in the role of Odd) 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 5 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 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.[10] 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 more on this below). With

[10]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 (two-tailed Pearson: $p = 0.047$).

Table 5: Types of subjects in the different treatments.

Treatment	Types P	AP	N	A
NO	0	9	3	30
	(0%)	(21.4%)	(7.1%)	(71.4%)
AN	0	7	0	19
	(0%)	(26.9%)	(0%)	(73.1%)
AN-MA	1	9	2	36
	(2.1%)	(18.8%)	(4.2%)	(75%)
NS	1	3	0	26
	(3.3%)	(10%)	(0%)	(86.7%)
NS-MA	0	3	3	34
	(0%)	(7.5%)	(7.5%)	(85%)
All	2	31	8	145
	(1.1%)	(16.7%)	(4.3%)	(78%)

Note: Absolute number in first row, percentage in second row. Types are denoted by P (Pure), AP (Almost Pure), N (Non-rationalizable), and A (Adaptive).

fewer subjects reacting to information, behavior is more likely to settle down when public information is 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

Table 6: Estimated decision rules at nodes 3, 4, 5, and 6 in treatment NO.

Dependent var.:	Pass			
	Node 3	Node 4	Node 5	Node 6
Private Statistic β	0.450***	0.612***	0.546***	0.560***
(node-specific, last 5 periods)	(0.041)	(0.067)	(0.094)	(0.156)
Constant α	0.527***	0.407***	0.312***	0.221***
	(0.051)	(0.036)	(0.043)	(0.046)
R^2	0.283	0.182	0.128	0.100
$\chi^2_{(1)}$	118.310	84.248	33.555	12.860
N	594	478	352	114

Note: Linear probability model with individual random effects. Values in parentheses represent standard errors. Stars represent (two-sided) p-values: $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

using random-effects linear probability models with just one independent variable capturing *subjects' own experience*. Specifically, we estimate the model

$$p_{n,i,t} = \alpha_n + \beta_n S_{n+1,i,t} + v_{n,i} + \varepsilon_{n,i,t} \tag{1}$$

where $S_{n+1,it}$ is subject i's private statistic for the average pass rate at node $n+1$ previous to the decision period t, $v_{n,i}$ is a subject-specific random effect and $\varepsilon_{n,i,t}$ the remaining error. 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.[11] We estimate this model separately for nodes 3, 4, 5, and 6.[12]

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 one that corresponds to the NS-MA treatment (see Table 6 for the estima-

[11]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.

[12]In the table and for any comparisons of the descriptive accuracy of specifications with different private statistics, we exclude the first six periods since here the NS and NS-MA statistics (as well as the AN and AN-MA statistics) are the same.

tion results for this specification).[13] Thus, it appears that in the absence of information provided by the experimenter subjects memorize their own experiences in a rather subtle way. We further observe that *unconditional passing* becomes less pronounced at later nodes, reflected by the falling constants, 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).

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 unambiguous. None of the extra variables has a consistent significant effect while the estimates for the coefficient of the private statistic remain significant and similar in size throughout all nodes when one includes the other variables. Furthermore, for each node, a model selection based on the Bayesian information criterion (BIC) selects the specification including the private statistic (1) over a competing model containing the effects of aspiration levels, learning direction theory, and time (but not the private statistic).[14] We summarize this in

Result 4 Subjects' behavior is not consistently affected by their past performance nor by the outcome of their last game nor by the passing of

[13]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 16.3%. For the models with NS statistics averaged over all periods the average R^2 is 7.2%. Finally, we the two models that examine private statistics averaged over all nodes the average R^2 are 1.4% (AN-MA) and 2.0% (AN).

[14]The precise specification of the alternative model is $p_{n,i,t} = \alpha_n + \gamma_n \pi_{i,t} + \lambda_n L_{n,i,t} + \kappa t + v_{n,i} + \varepsilon_{n,i,t}$, where all variables are defined as in Appendix B. The BIC for model (1) [for the alternative model] is $221.59, 583.48, 460.17, 143.50$ $[248.77, 614.83, 461.56, 157.57]$ for node 3, 4, 5, and 6, respectively. The results are qualitatively the same, i.e., for each node the model with the private statistic is selected over the alternative model, when (i) a time trend is included in (1) or excluded in the alternative specification, (ii) the alternative specification includes only the effects of aspiration levels, or (iii) the alternative specification includes only learning direction theory.

time itself.

Let us now turn our 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

$$p_{n,i,t} = \alpha_n + \beta_n S_{n+1,i,t} + \gamma_n P_{n+1,i,t} + v_{n,i} + \varepsilon_{n,i,t}$$

where $P_{n+1,i,t}$ 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 7 and 8.[15]

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. However, at node 6 when stakes reach £9.60 subjects start to ignore public imprecise public information but still rely on it when it is node-specific. 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 approach £10 and the public statistic is imprecise the public statistic is ignored.

4.3 Unravelling

From a theory viewpoint, if players rely on the node-specific statistic 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 2005). Thus, less unravelling should be expected in this case.

Our finding in the treatment without public information – 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 more pronounced if there is a precise public statistic (i.e., in the NS treatments). To understand this differential degree of unravelling, we introduce two additional slices of the data.

[15] As in Table 6, we exclude the first six observations in the regressions.

Table 7: Estimated decision rules at nodes 3, 4, 5, and 6 in AN treatments.

Dependent var.:	Pass			
	Node 3	Node 4	Node 5	Node 6
Private Statistic β	0.540***	0.507***	0.644***	0.375**
(node-specific, last 5 periods)	(0.032)	(0.055)	(0.062)	(0.158)
Public Statistic γ	1.202***	3.130***	1.105**	−1.780
	(0.218)	(0.633)	(0.435)	(1.192)
Constant α	−0.455***	−2.125***	−0.683**	1.615*
	(0.163)	(0.502)	(0.342)	(0.969)
R^2	0.286	0.181	0.187	0.059
$\chi^2_{(2)}$	467.446	166.909	126.175	9.616
N	1288	797	613	155

Note: Linear probability model with individual random effects. Values in parentheses represent standard errors. Stars represent (two-sided) p-values: $^*p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 8: Estimated decision rules at nodes 3, 4, 5, and 6 in NS treatments.

Dependent var.:	Pass			
	Node 3	Node 4	Node 5	Node 6
Private Statistic β	0.372***	0.450***	0.522***	0.408***
(node specific, last 5 periods)	(0.041)	(0.046)	(0.057)	(0.077)
Public Statistic γ	0.196***	0.652***	0.564***	0.374**
	(0.068)	(0.089)	(0.126)	(0.146)
Constant α	0.444***	0.048	0.056	0.077
	(0.043)	(0.049)	(0.056)	(0.063)
R^2	0.263	0.301	0.229	0.131
$\chi^2_{(2)}$	256.679	378.519	169.918	44.274
N	1130	953	630	296

Note: Linear probability model with individual random effects. Values in parentheses represent standard errors. Stars represent (two-sided) p-values: $^*p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 9: Pass rates of non-adaptive subjects.

Treatments	Node 3	Node 4	Node 5	Node 6
NO	0.913	0.708	0.400	0.254
AN/AN-MA	0.852	0.695	0.321	0.439
NS/NS-MA	0.960	0.816	0.592	0.452

Note: Pass rates computed excluding the first six periods to make them comparable with estimates for adaptive subjects.

First, Table 9 shows the pass rates of the non-adaptive subjects for nodes 3 to 6. The table reveals that non-adaptive subjects tend to pass a lot which does put a bound on the possibilities for unravelling.[16] Of course, in the NS treatments there are only roughly half as many non-adaptive subjects than in the two others (14.3% as compared to 28.6% in NO and 25.7% 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 10 shows how adaptive 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 the private or the public statistic about the next node were below 10%.

Table 10: Pass rates of adaptive subjects in treatment NS-MA.

	Node 3	Node 4	Node 5	Node 6
Private statistic < 0.1	0.500	0.305	0.356	0.263
Public statistic < 0.1	0.222	0.143	0.143	0.056

Note: Pass rates computed excluding the first six 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 their own experience has been bad they are much more optimistic and incidentally these numbers are basically the same in all other treatments. Hence, we should expect if adaptive subjects were on

[16]Over all treatments, the average pass rate of non-adaptive players is significantly higher than the average pass rate of adaptive players (t-test yields $p < 0.001$; node-specific t-tests yield $p = 0.004$ and $p = 0.023$ for odd players at node 3 and 5, respectively, and $p < 0.001$ and $p = 0.001$ for even players at node 4 and 6, respectively).

their own they would converge very closely to Nash equilibrium play with precise public information. With less precise public information this is less likely as extreme information is, by design, less frequent.

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 nodes 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.[17]

5 Concluding remarks

Our experimental results can be summarized 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 stakes are 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

[17] There is also much more unravelling in high-stakes *three*-player centipede games as studied by Rapoport, Stein, Parco, and Nicholas (2003). 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.

economic applications and deserves further study. In that sense we advocate the methodology 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. (2005). Analogy-based expectation equilibrium. *Journal of Economic Theory* 123, 81-104.

[4] Johnson, C. A., A. Baillon, H. Bleichrodt, Z. Li, D. van Dolder, and P. P. Wakker. (2015) Prince: An Improved Method for Measuring Incentivized Preferences. Available at SSRN 2504745.

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

[6] 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.

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

[8] Rapoport, A., W.E. Stein, J.E. Parco, and T.E. Nicholas (2003). Equilibrium play and adaptive learning in a three-person centipede game. *Games and Economic Behavior* 43, 239-265

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

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

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

[12] 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.

[13] Tomlinson, C. (2005). EXPECON: A flexible architecture for building networked experiments. *Centre for Economic Learning and Social Evolution (ELSE)*. Available online at: http://else.econ.ucl.ac.uk/newweb/software/expecon.pdf.

A Experimental instructions

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 consists 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 be doubled in each and every stage, finally, reaching £153.60 in stage 10. The following table shows you all the numbers.

Payoffs of participant who chooses TAKE									
1	2	3	4	5	6	7	8	9	10
£0.30	£0.60	£1.20	£2.40	£4.80	£9.60	£19.20	£38.40	£76.80	£153.60

[*NS Treatments:*
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 [*NS-MA:* 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 [*NS-MA:* 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 [*NS-MA:* five rounds], you will be told there is no data available. [*NS-MA:* (During the first five rounds, averages will be based on all previous decisions.)]]

[*AN Treatments:*
When making a decision you will have some information about the past. More specifically, Odd participants will be told how often, on average, Even participants have chosen to PASS in the past [*AN-MA:* five rounds]. Similarly, Even participants will be told how often, on average, Odd participants have chosen to PASS in the past [*AN-MA:* five rounds]. These averages will be updated after each round and will be available to you whenever you make a choice. [*AN-MA:* (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, $\pi_{i,t}$, subject i's average payoff up to period $t-1$; a variable capturing the possible effect of learning direction theory, $L_{n,i,t} = 1$ if subject i lost in period $t-1$ on node $n+1$ (and 0 otherwise); and a variable capturing the passing of time, t itself. Table 11 shows the estimates for the random-effects model

$$p_{n,i,t} = \alpha_n + \beta_n S_{n+1,i,t} + \gamma_n \pi_{i,t} + \lambda_n L_{n,i,t} + \kappa t + v_{n,i} + \varepsilon_{n,i,t}, \qquad (2)$$

where $n = 3, 4, 5, 6$ 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.

Table 11: Estimated decision rules at nodes 3, 4, 5, and 6 in treatment NO.

Dependent var.:	Pass			
	Node 3	Node 4	Node 5	Node 6
Private Statistic β	0.423***	0.550***	0.616***	0.578***
(node specific, last 5 periods)	(0.049)	(0.077)	(0.101)	(0.167)
Aspirations γ	0.015*	0.008	0.014	−0.001
	(0.009)	(0.007)	(0.010)	(0.009)
Learning dir λ	0.064**	−0.046	−0.088	0.017
	(0.030)	(0.043)	(0.057)	(0.125)
Period κ	−0.003**	−0.003*	−0.002	0.008**
	(0.001)	(0.002)	(0.002)	(0.003)
Constant α	0.539***	0.492***	0.329***	0.054
	(0.063)	(0.071)	(0.080)	(0.107)
R^2	0.285	0.189	0.139	0.151
$\chi^2_{(2)}$	167.043	102.590	55.791	19.449
N	594	478	352	114

Note: Stars represent (two-sided) p-values: $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Discussion Papers of the Research Area Markets and Choice 2016

All discussion papers are downloadable:
http://www.wzb.eu/en/publications/discussion-papers/markets-and-choice