

Do we follow others when we should? A simple test of rational expectations¹

Georg Weizsäcker

London School of Economics and Political Science

April 2008

Abstract: The paper presents a new meta data set covering 13 experiments on the social learning games by Bikhchandani, Hirshleifer, and Welch (1992). The large amount of data makes it possible to estimate the empirically optimal action for a large variety of decision situations and ask about the economic significance of suboptimal play. For example, one can ask how much of the possible payoffs the players earn in situations where it is empirically optimal that they follow others and contradict their own information. The answer is 53% on average across all experiments – only slightly more than what they would earn by choosing at random. The players' own information carries much more weight in the choices than the information conveyed by other players' choices: the average player contradicts her own signal only if the empirical odds ratio of the own signal being wrong, conditional on all available information, is larger than 2:1, rather than 1:1 as would be implied by rational expectations. A regression analysis formulates a straightforward test of

¹I am grateful for helpful comments by Mohit Bhargava, Stefano DellaVigna, Erik Eyster, Bernardo Guimaraes, Dorothea Kübler, Luke Miner, Charles Noussair, Matthew Rabin, Myung Seo and seminar audiences at Berkeley, Copenhagen and MPI Bonn. Mohit Bhargava and Luke Miner also provided excellent research assistance. The collection of the meta data set was made possible by Jonathan Alevy, Lisa Anderson, Antonio Guarino, Charles Holt, Angela Hung, John List, Markus Nöth, Clemens Oberhammer, Joerg Oechssler, Brian Rogers, Andreas Roider, Andreas Stiehler, Martin Weber and Anthony Ziegelmeyer who provided the individual data sets. Finally, I thank the ESRC (Award RES-000-22-1465) for funding.

rational expectations, which rejects, and confirms that the reluctance to follow others generates a large part of the observed variance in payoffs, adding to the variance that is due to situational differences.

Keywords: Social learning, information cascades, failure of rational expectations, meta analysis

JEL Classification: C72, C92, D82

Contact: g.weizsacker@lse.ac.uk

1 Introduction

The paper re-visits the experimental literature on social learning and addresses a significant gap in the existing data analyses: to what extent is it beneficial for a decisionmaker to learn from others and contradict his or her own private information – and if it is beneficial, do people do it? These questions lie at the heart of understanding social learning behavior but have been answered only very partially. The main difficulty in finding appropriate answers is the identification of conditions under which a decisionmaker should follow others. Most previous discussions in the literature measure a decisionmaker’s success under the assumption that other players obey a given model solution, e.g. Bayes Nash Equilibrium or Quantal Response Equilibrium. Any such theory implies an optimal response of the decisionmaker herself, which can be held against the data.² But despite the undisputable usefulness of these solutions, it is important to note that they are often inaccurate and thus their implications are imperfect benchmarks. It has never been established to what degree the decisionmakers choose the *empirically optimal* action, i.e. the action that is ex-post optimal most of the time under identical conditions. This paper’s premise is that the empirically optimal action is what we should measure the success of social learning by, because it reflects all relevant situational factors, including those that are influenced by the true behavior of other players. The paper combines the raw data from a large set of experiments (all of which follow Anderson and Holt, 1997) into a new meta data set and asks whether the average behavior can be viewed as approximately payoff maximizing, given the behavior of others. It finds that the success of learning from others is very modest. Conditional on being in a situation where it is empirically optimal for the participants to contradict their own private information, they choose this action (among two

²Anderson and Holt (1997), the influential first study in this experimental literature, contains a data analysis using both Bayes Nash Equilibrium and Quantal Response Equilibrium. Most subsequent studies have summarized behavior on the equilibrium path of Bayes Nash Equilibrium, see e.g. the survey in Kübler and Weizsäcker (2005). Variants of Quantal Response Equilibrium are discussed in Oberhammer and Stiehler (2003), Kübler and Weizsäcker (2004), Choi, Gale, and Kariv (2004), Drehmann, Oechssler and Roeder (2005), Krämer, Nöth and Weber (2006) and Goeree et al (2007).

alternatives) in less than half of the cases.³ On the other hand, in cases where their own signal happens to support the empirically optimal action, the participants more likely realize this and follow the signal nine out of ten times.

These and all other results in the paper are generated under a reduced-form approach, without imposing a behavioral model of the players' decisionmaking process or a model of how they take other players' actions into account. (For example, the analysis does not assume that players apply Bayesian updating.) A disadvantage of not imposing a behavioral model is that we learn little about the nature of the suboptimality. Yet the results are supportive of the model estimates by Nöth and Weber (2003) and Goeree et al (2007), which suggest that people assign too much weight to their own private information, relative to the publically observable choices of others.⁴

But can the high error rate perhaps be attributed to the lack of strong incentives? Especially in situations where a player receives a signal that is different from other players' signals, the expected gain from making correct inferences is reduced: with conflicting pieces of information, the total value of the available information is small and all choice options are unsuccessful with relatively high likelihood. This makes it less important to choose the best option. At the same time, players may also find it more difficult to identify the optimal action in situations with conflicting information. One may thus expect a positive correlation between simplicity and lucrativity (value of the available information) of choice situations. The mere choices frequencies may therefore be misleading indicators of success if they do not consider the differences in the importance of making the optimal choice – the players may do the right thing predominantly when it is important.⁵

³The frequency is about 0.44 and depends on the set of included cases – the identification of the empirically optimal action relies on small sample sizes for many decision problems so that it is better not to include these cases (see Section 3). But for any of a large selection of possible sets, the frequency lies below one half.

⁴These and other reports of biases relative to a rational-expectations benchmark (Oberhammer and Stiehler, 2003, Kübler and Weizsäcker, 2004, Kraemer, Nöth and Weber, 2006) are essentially confined to establishing statistical significance. The existing literature has largely sidestepped a detailed discussion of payoff outcomes – as opposed to behavioral outcomes – and most studies merely contain measures of overall earnings in the experiment. A notable exception is in Goeree et al's (2007) discussion of higher average earnings of players who choose later in the game.

⁵The complication is particularly severe in social learning games, where uncertainty about others' strategies prevails so that the value of the available information (here, the observation of others' actions) is unknown to the researcher.

But with enough data one can estimate the value of the available information and control for it in the analysis. Below, calculations will show that the success rate indeed increases with the importance of making the optimal choice, but still social learning is unsuccessful in payoff terms. In situations where participants should contradict their own signal, they only receive 53% of the higher of the two possible prizes on average (normalizing the low prize to zero). They could have earned 64% of the high prize on average, had they always realized that it is better to contradict their own information. In contrast, in the complementary set of situations, where it is empirically optimal for participants to follow their own information, they are much more successful and earn 73% of the high prize on average, out of 75% that they could earn from always behaving optimally. Indeed, the numbers show that the latter set of situations is more lucrative for the players, but also that the difference in success rates generates a large payoff gap.

Estimating the value of the players' information is particularly straightforward in games like the present ones, where there are only two payoff outcomes and two available actions for each player. One can simply count how often each of the two actions would have yielded the higher of the two possible prizes, across all instances where the available information is identical.⁶ The relative frequency of the two counts then yields the desired estimate. This exercise is applied to the new data set consisting of data from 13 different experimental studies on the game by Bikhchandani, Hirshleifer, and Welch (1992), which is widely viewed to capture the essence of social learning. All of these experiments follow the same format, with only minor modifications and differences in the experimental procedures, such as different precisions of signals. The data set contains a large variety of different decision situations – there are more than 10000 distinct decision problems if one counts decisions in different experiments separately, and almost 30000 decisions in total – but all of them are generated within the same basic game, with two possible states of the world and two actions per player, and they can therefore be summarized in meaningful ways. Section 3 describes the statistical connection between the participants' frequency of following their own information and the estimated value of this action. This description demonstrates a strongly inflated tendency of

⁶As will be made clear in Section 2, this estimation only views situations as identical if their informational conditions are the same and if they appear in the same treatment of the same experiment.

participants' to follow their private information, and indeed gives a negative answer to the question in the paper's title, as described above: when participants should optimally contradict their own signal, this action is observed in less than half of the cases. As the value of this action increases, its frequency increases, but at a slow rate. To observe a frequency of one half or more, the action's value needs to exceed $2/3$ of the higher prize. That is, to induce the average participant to act against her signal, the evidence conveyed by the other players' decisions needs to be so strong that even after accounting for the countervailing significance of the signal, the empirical likelihood ratio against the signal is larger than 2:1.

The reader may note that the above classification of situations – whether or not the own signal supports the empirical optimal action – makes for an unusual explanatory variable because it is unobservable to the decisionmaker. But of course, it is correlated with other, more transparent classifications. Section 4 considers a natural one, whether or not a player's signal is in contradiction to a strict majority of previous players' choices. This variable and the success rates under each of its categories are highly correlated with the above-described results,⁷ indicating again that people underestimate how much information is contained in other people's choices. There is also a positive effect of unanimity of previous choices: participants are much more successful in learning from others if the others all agree.

Section 4 presents these results in a regression analysis that projects the observed behavior on explanatory variables that describe the nature of the situation. Importantly, one of the control variables is the value of the empirically optimal choice. If this value is held constant, and under the assumption of rational expectations, there is no reason for a systematic behavioral change in response to the nature of the situation. The significance test for the other explanatory variables therefore makes for a straightforward consistent test of the hypothesis that players respond to rational expectations. The test rejects strongly: the behavioral change depending on whether the own signal contradicts the majority of previous choices is highly significant and accounts for about

⁷The average earnings are 55% of the high prize in situations where the own signal contradicts the majority; were the participants to choose the empirically optimal action in each case, they would earn 65% in these situations. In the remaining situations, where the signal coincides with the majority of previous choices – so that the bias in favor of the own signal may help the participants – they earn 72% out of the maximum feasible payoff of 75%.

half of the total payoff differences between the two classes of situations.

Section 5 concludes with a discussion of how the methodology may be applied in other data contexts. It argues that the identification of empirically optimal actions can be used generally in tests of rational expectations models, if enough data is available.

2 A meta data set of cascade game experiments

Anderson and Holt's (1997) experiment was replicated and modified by numerous researchers, and the new meta data set is restricted to this class of experiments because the structure of the game is thereby held constant in all subsets of data.⁸ The following describes the game with Anderson and Holt's specific parameter values. At stage $t = 0$, Nature draws one of two states of the world, $\omega \in \Delta \equiv \{A, B\}$, with $\Pr(A) = 0.5$. The states are represented by two urns labelled A and B . Of the balls in urn A , a fraction $q_A = 2/3$ is labelled a and $1 - q_A$ is labelled b . Urn B , analogously, has fractions of $q_B = 2/3$ labelled b and $1 - q_B$ labelled a . Nature's draw is not revealed, so that it is unknown in the remainder of the game whether balls are chosen from A or B . A set of $T = 6$ players makes predictions about this event. At stage $t = 1$, the first player receives a private signal, in the form of a ball drawn from the true urn. The player then predicts the state of the world, i.e. chooses an action $d \in \Delta$. At $t = 2$, the next player receives a signal from the same urn, makes a prediction, and so on until the game ends after stage T . If a player's prediction coincides with the true urn, she gets a fixed amount U , here normalized to 1. Otherwise, the player gets 0.⁹ Social learning is possible because at $t \geq 2$ the players observe the predictions made by all previous players $1, \dots, t - 1$ (but not their signals). Hence, players can make their choices dependent on the choices

⁸Other social learning experiments are those on learning in networks (e.g. Choi, Gale, and Kariv, 2004), on cascade games that do not follow the Anderson/Holt format (Allsopp and Hey, 2000, Celen and Kariv, 2004, Guarino, Harngart and Huck, 2007) and games where agents are replaced by groups (Fahr and Irlenbusch, 2008). Several papers also address how experimental participants play social learning games against automated opponents with exogenously fixed behavior (Huck and Oechssler, 1999, Grebe, Schmidt and Stiehler, 2006, Kraemer, Nöth and Weber 2004).

⁹Note that since there are only two possible payoffs for each player, risk considerations would be absent under any expected-utility model.

of other players. Information cascades can arise in Bayes Nash Equilibrium: e.g., if the first two players make an identical prediction, say, A , then in sequential equilibrium the third player chooses A even if she has signal b . The fourth player, too, follows the prediction of the initial three players, and so on for later players. Quantal Response Equilibrium and its variants also prescribe that the decision makers learn from the predecessors' choices, and tend to follow them.

Anderson and Holt played the game with 18 participants, each of whom took part in 15 repetitions of the game, with randomly changing player positions between the repetitions. In a separate asymmetric treatment, 36 subjects played a modified version (again with 15 repetitions), where the proportion of a balls in urn A is $q_A = 6/7$, and the proportion of b balls in urn B is $q_B = 2/7$. Both treatments are used in the analysis, yielding a total number of $N = 810$ decisions. The following lists the other data sources:

- Willinger and Ziegelmeyer (1998, $N = 324$). Basic (symmetric) Anderson/Holt experiment, with $q \equiv q_A = q_B = 0.6$, 36 participants, and 9 repetitions.
- Anderson (2001, $N = 270$). Symmetric Anderson/Holt experiment, with $q \equiv 2/3$, 18 participants, and 15 repetitions.
- Hung and Plott (2001, $N = 890$). Replication of Anderson/Holt with $T = 10$ and $q = 2/3$, in three treatments with minor differences in the experimental implementation. 40 participants and 22.25 repetitions on average.
- Ziegelmeyer et al (2002, $N = 810$). 54 participants and 15 repetitions with $T = 9$, $\Pr(A) = 0.55$ and $q = 2/3$. Subjects in one of the two treatments also reported beliefs about the state of the world.
- Nöth and Weber (2003, $N = 9834$). Variant of the game with $T = 6$, where q is drawn for each player from $\{0.6, 0.8\}$, in independent draws with equal probabilities, and the realized signal precision is known to the player herself but unknown to other players. 126 participants and about 78 repetitions on average.
- Oberhammer and Stiehler (2003, $N = 876$). Variant where subjects also announce their

willingness to pay for playing the game, with $T = 6$, $q = 0.6$, 36 participants, and about 24 repetitions on average.

- Kübler and Weizsäcker (2004, $N = 482$). Variant where subjects also decide whether or not to receive a private signal, this decision not being revealed to other players. 36 participants and 15 repetitions, with $T = 6$ and $q = 2/3$. 68 observations are dropped from the original data, in cases where the corresponding participants requested no signal.¹⁰
- Dominitz and Hung (2004, $N = 2270$). 90 participants, games with $T = 10$ and $q = 2/3$. 30 participants played 20 regular repetitions, and the other 60 participants reported belief statements during the last 10 out of 20 repetitions.
- Cipriani and Guarino (2005, $N = 161$). Variant where players have an outside option, modelled after the possibility not to trade in financial markets. 48 participants and 10 repetitions, with $T = 12$ and $q = 0.7$. 319 observations are dropped from the original data because the participant or one of her predecessor chose the outside option.¹¹
- Drehmann, Oechssler and Roider (2005, $N = 2789$) 1840 participants played in 8 different Internet-based treatments with different signal precisions and different values for $\Pr(A)$. 267 participants were management consultants who played among each other, with $T = 7$. In addition, 1573 participants of different backgrounds (mostly students or graduates of universities) played in games with $T = 20$ (mostly). Participants played up to three repetitions.
- Alevy, Haigh and List (2006, $N = 1647$). Replication of both Anderson/Holt treatments, with $T = 5$ or $T = 6$, 15 repetitions, and with subjects of different backgrounds: 55 financial market professionals, and 54 undergraduate students. Treatments also differed with respect to gains/losses framing. Eight treatments in total.

¹⁰Receiving a signal is weakly dominant in this game.

¹¹For risk-neutral expected-utility maximizers, the outside option is strictly suboptimal for almost all subjective beliefs that a player may have about the state of the world. This game is the only game with three possible payoffs, so that risk considerations may become important. For the sake of completeness, I decided to include the (few) data despite these differences.

- Goeree et al (2007, $N = 8760$). 380 participants play in four long-game treatments, with $T = 20$ and $T = 40$, $q = 5/9$ and $q = 2/3$, and an average of 22.7 repetitions.

In total, the meta data set contains 29923 individual decisions $\{s_i\}_{i=1}^{29923}$, made by 2813 participants in 13 separate studies. All of them follow the observation of a private signal and a (possibly empty) string of previous choices made in analogous situations. In all of them there are two actions and two possible payoffs,¹² but the data set nevertheless contains decisions with a large variety in environments, instructions, and histories of other players' choices.

From here on, let a "decision" s_i be a row in the data matrix – a vector of variable realizations that describe a participant's choice problem in one particular repetition of a game. Each decision s_i in the data set contains the description of the game ($\Pr(A)_i, q_{A,i}, q_{B,i}, T_i$), the participant's player position t_i , her choice d_i and the true state of the world ω_i .¹³ In addition, the following variables are contained in s_i :

group_i: An identifying variable that denotes the group of participants with whom the participant played the current repetition of the game.¹⁴

treatment_i: A categorical variable that defines two decisions (s_i, s_j) to be in the same treatment if (i) the participants received the same instructions for the current game, and (ii) the participants (and their opponents) are drawn from the same pool.¹⁵

I_i (information): A set of variables describing the history of play in the current repetition of the game, including all previous choices as well as the participant's signal. E.g. if the participant

¹²But with the qualifier about Cipriani and Guarino (2005) – see the previous footnote.

¹³Information on ω_i is available for each decision – which is crucial for estimating the empirical value of each action – thanks to the "no-lying-to-subjects" policy in experimental economics: all urn and signal draws in the experiments were actually made according to the described process, and not made up by the experimenters.

¹⁴With the exception of the Internet-based experiment by Drehmann, Oechssler and Roeder (2005), the matching of participants into groups remained constant throughout all experiments, so that no participant was part of several groups.

¹⁵This does not require identical conditions for other parts of the experiment, e.g. for the belief elicitation in some studies or for the number of repetitions of the game. In all cases except one, this treatment definition coincides with that of the original studies' treatment definitions (the exception is Dominitz and Hung, 2004, who had two treatments with different belief elicitation procedures). But most studies included additional treatments.

in decision s_i acts as the first player in the game, she may have $I_i = a$, or if she acts as the third player she may have $I_i = ABa$. In the case of the Nöth/Weber dataset where the signal strength differs between the participants, the information about the own signal strength is also part of I_i .¹⁶

contradict_i: An indicator of whether or not the participant in decision s_i contradicts her own signal – e.g. chooses $d = A$ with own signal b .¹⁷

sitcount_i: A counting variable that contains the number of times that a decision with the same information I_i occurred within the same treatment. In treatments with A/B symmetry ($\Pr(A) = 0.5, q_A = q_B$), this includes situations with the symmetric information set. (E.g. $AABa$ is viewed as identical to $BBAb$.)

Let $\tilde{S}(s_i)$ be the set of all decisions that are identical to s_i , in the sense that the (treatment _{i} , I_i)-description is the same. By definition, $\text{sitcount}_i = |\tilde{S}(s_i)|$. Further, let $\tilde{S} \equiv \{\tilde{S}_k\}_k$ be the resulting set of all distinct sets of identical situations: each element of \tilde{S} is a set of identical decisions and each s_i appears in one and only one element of \tilde{S} . Notice that these definitions ignore information about history of play in previous repetitions of the game, i.e. they pool situations as identical across different histories. The subsequent analysis relies on this pooling and thereby makes the simplifying assumption that pre-current-game history does not enter into a participant's decisionmaking process.¹⁸ Under this assumption it holds that the decisionmaking process in decision s_i – whatever it may be – uses only information that is identical across $\tilde{S}(s_i)$ because all that the participants learn about the random draws or about the possible characteristics of other players is reflected in (treatment _{i} , I_i).¹⁹ The following variable estimates the value of this

¹⁶In effect, there are four different signals in this data set: $a - weak, a - strong, b - weak$ and $b - strong$.

¹⁷For this definition to apply also in the case of the Nöth/Weber experiment where there are four signals, the phrase "contradicts her own signal" is to be understood e.g. as a choice of B after observing a signal $a - weak$ or $a - strong$.

¹⁸The assumption is restrictive but still relatively mild: it is difficult to learn about other player's strategies in social learning games because their signals remain unknown and player positions change after each repetition of the game. Several studies (e.g. Kübler and Weizsäcker, 2004) looked for behavioral changes over the course of the experiments and found at most weak changes. The meta data set allows for further checks, which support the assumption. For example, the frequency of following the own signal changes by 0.6 percentage points between the first and the second half of the experiments (0.767 versus 0.761).

¹⁹This observation about the information being constant across $\tilde{S}(s_i)$ is correct only in the pure case of complete

information, using the ex-post knowledge about ω_i , the true underlying state of the world. It is the key ingredient of the analysis – the empirical value of contradicting one’s signal.

mean_pay|contradict_i: Averaging across situations $s_j \in \tilde{S}(s_i)$, the payoff that the participant earns if she chooses $\text{contradict}_i=1$. Due to the payoff normalization, this variable is identical to the frequency across $s_j \in \tilde{S}(s_i)$ of the participant’s signal being wrong, i.e. the state of the world ω_j not being indicated by the signal.

The variable $\text{mean_pay|contradict}_i$ is an average of a finite sample. As sitcount_i increases, this average approaches the mean of the corresponding random variable, i.e., it approaches the expected payoff from contradicting one’s signal, conditional on all available information. This conditional expected payoff is what that the player should care about: she should contradict her signal if the expected payoff from doing so exceeds 0.5.

For small sitcount_i , however, $\text{mean_pay|contradict}_i$ may be far away from its expected value. I will therefore restrict attention to cases where there are strictly more than 10 occurrences of identical situations in the data set, i.e. where $\text{sitcount}_i = |\tilde{S}(s_i)| > 10$. Almost half of the decisions are therefore not considered in the statistical analysis: out of the 29923 situations, 15734 remain in the restricted data set. To give the reader an impression of the resulting selection of the subsample as well as a general overview of the data, the following two tables present some features of the restricted data set (first entry in each cell) and the full data set (second entry, in italics). Table 1 organizes the data by decisions $\{s_i\}_i$, and Table 2 by distinct decisions $\{\tilde{S}_k\}_k$ (so that each distinct decision counts only once for the reported values in Table 2).

anonymity between the participants during the experiment. If participants can observe each other’s person-specific characteristics, there are informational differences even within $\tilde{S}(s_i)$ – but these should be minor. Another caveat is that unobservable differences between experimental sessions within the same treatment cannot be taken into account.

	# of decisions	ave. of sitcount _{<i>i</i>}	from largest data set	history of only A or only B
$t_i=1$	3198	456.8	51.6% (NW)	-
	<i>3306</i>	<i>442.1</i>	<i>49.6% (NW)</i>	-
$t_i=2$	3025	263.1	54.2% (NW)	-
	<i>3304</i>	<i>241.5</i>	<i>49.7% (NW)</i>	-
$t_i=3$	2727	141.8	60.1% (NW)	63.2%
	<i>3301</i>	<i>118.1</i>	<i>49.7% (NW)</i>	<i>60.6%</i>
$t_i=4$	2379	87.0	68.9% (NW)	49.4%
	<i>3298</i>	<i>64.2</i>	<i>49.7% (NW)</i>	<i>43.5%</i>
$t_i=5$	2024	63.6	80.1% (NW)	43.1%
	<i>3289</i>	<i>40.8</i>	<i>49.8% (NW)</i>	<i>34.8%</i>
$t_i=6$	1844	51.7	83.7% (NW)	40.0%
	<i>3120</i>	<i>32.1</i>	<i>52.5% (NW)</i>	<i>30.4%</i>
$t_i=7$	136	35.5	83.1% (DH)	100.0%
	<i>884</i>	<i>8.0</i>	<i>36.4% (GMPR)</i>	<i>31.0%</i>
$t_i=8$	134	33.8	82.1% (DH)	100.0%
	<i>880</i>	<i>7.2</i>	<i>36.6% (GMPR)</i>	<i>29.1%</i>
$t_i=9$	139	28.9	91.4% (DH)	91.4%
	<i>878</i>	<i>6.3</i>	<i>36.7% (GMPR)</i>	<i>26.3%</i>
$t_i=10$	128	32.8	82.0% (DH)	100.0%
	<i>788</i>	<i>6.8</i>	<i>40.9% (GMPR)</i>	<i>22.2%</i>
$t_i>10$	0	-	-	-
	<i>6875</i>	<i>1.2</i>	<i>80.6% (GMPR)</i>	<i>4.8%</i>
total/ave.	15734	196.5 (ave.)	61.9% (NW)	52.9% (ave.)
	<i>29923</i>	<i>104.5 (ave.)</i>	<i>32.9% (NW)</i>	<i>29.1% (ave.)</i>

Table 1: Decisions in $\{s_i\}_i$, split up according to player position t . Note: The first entry in each cell described the restricted data set where sitcount_{*i*} > 10, the second, italicized entry describes the full data set. Abbreviations: NW-Nöth/Weber 2003, DH-Dominitz/Hung 2004, GMPR-Goeree et al 2007.

	# of distinct decisions	ave. of sitcount _{<i>i</i>}	from largest data set	history of only A or only B
$t_i=1$	38	84.2	31.6% (AHL)	-
	<i>62</i>	<i>53.3</i>	<i>35.5% (DOR)</i>	-
$t_i=2$	52	58.2	19.2% (AHL)	-
	<i>116</i>	<i>28.5</i>	<i>31.0% (DOR)</i>	-
$t_i=3$	59	46.2	25.4% (GMPR)	64.4%
	<i>206</i>	<i>16.0</i>	<i>29.1% (AHL)</i>	<i>50.5%</i>
$t_i=4$	59	40.3	27.1% (NW)	54.2%
	<i>338</i>	<i>9.8</i>	<i>27.8% (AHL)</i>	<i>29.0%</i>
$t_i=5$	53	38.2	60.4% (NW)	43.3%
	<i>534</i>	<i>6.2</i>	<i>25.7% (AHL)</i>	<i>17.2%</i>
$t_i=6$	64	28.8	76.6% (NW)	28.1%
	<i>629</i>	<i>5.0</i>	<i>17.0% (GMPR)</i>	<i>10.5%</i>
$t_i=7$	5	27.2	60.0% (DH)	100.0%
	<i>434</i>	<i>2.0</i>	<i>37.3% (GMPR)</i>	<i>8.3%</i>
$t_i=8$	5	26.8	60.0% (DH)	100.0%
	<i>506</i>	<i>1.7</i>	<i>39.5% (GMPR)</i>	<i>7.1%</i>
$t_i=9$	6	23.2	83.3% (DH)	83.3%
	<i>543</i>	<i>1.6</i>	<i>41.8% (GMPR)</i>	<i>5.7%</i>
$t_i=10$	5	25.6	60.0% (DH)	100.0%
	<i>528</i>	<i>1.5</i>	<i>46.2% (GMPR)</i>	<i>3.9%</i>
$t_i>10$	0	-	-	-
	<i>6400</i>	<i>1.1</i>	<i>79.4% (GMPR)</i>	<i>1.7%</i>
total/ave.	346	45.5 (ave.)	32.1% (NW)	51.2% (ave.)
	<i>10296</i>	<i>2.9 (ave.)</i>	<i>59.7% (GMPR)</i>	<i>5.8% (ave.)</i>

Table 2: Distinct decisions in \tilde{S} , split up according to player position t . Note: The first entry in each cell describes the restricted data set where sitcount_{*i*} > 10, the second, italicized entry describes the full data set. Abbreviations: AHL-Alevy/Haigh/List 2006, DOR-Drehmann/Oechssler/Roieder 2005, NW-Nöth/Weber 2003, DH-Dömnitz/Hung 2004, GMPR-Goeree et al 2007.

The tables show the large variety of decision situations in the meta data set. For example, the first entries in the cells of the second column of Table 2 show that for each of the player position up to $t = 6$, there are between 38 and 64 distinct decisions among those with $\text{sitcount}_i > 10$. The second column of Table 1 shows that these sets of distinct situations comprise several thousand individual decisions at each player position. For player positions 7, ..., 10, which appear in only six experiments, there are much fewer suitable decisions available, and not a single decision with $t > 10$ satisfies the requirement that $\text{sitcount}_i > 10$.

The third columns of the two tables show the strong decrease of sitcount_i for increasing player positions: due to the fact that later positions can have many more different histories, their average number of identical decisions is much lower. Nevertheless, there are large numbers of identical decisions even in later positions. E.g. for $t = 6$, each of the 1844 decisions that are included in the restricted data set appears 51.7 times on average. Averaging across all decisions in the restricted data set, decisions appear identically in 196.5 cases each. Table 2 shows that averaging across distinct decisions, each decision appears 45.5 times – this is necessarily lower because the decisions with many identical occurrences do not get larger weight in this average.

The fourth columns show the respective largest proportions of observations from an individual data set. Table 1 reveals that in the restricted data set, the experiment by Nöth and Weber (2003), which had many identical games, is rather dominant particularly for player positions $t = 5$ and $t = 6$. This is the main reason for not increasing the cutoff value for sitcount_i to a higher value than 10: requiring an even higher number of identical decisions would have increased the precision of $\text{mean_pay}|\text{contradict}_i$, but it would also have restricted the meta data set even more to its largest homogeneous subset of decisions. While it is not clear that this or any other kind of selection would introduce problematic biases – because all included experiments are selected situations anyway, and all of them are carefully designed and are equally suitable to detect social learning patterns – it would certainly run counter to the idea of a meta analysis and would reduce the variability of explanatory variables that are included in the regressions of Section 4.²⁰

²⁰In parallel to the contents of the next sections, robustness checks were run where the cutoff was increased to $\text{sitcount}_i > 30$ as well as a separate analysis where the data by Nöth and Weber (2003) were excluded. All linear regressions were also repeated with a probit model. Unless noted otherwise, the results were similar.

The fifth columns of the tables show the proportions of decisions that follow a history where all predecessors agree in their choices. These are potential decisions where social learning should be most prevalent, and they are also relatively simple decisions. The tables show that the selection effect on this variable due to the restriction $\text{sitcount}_i > 10$ is not very strong for player positions up to $t = 6$, which make up most of the restricted data set. For player positions $t = 1$ and $t = 2$, the variable is not defined because there can be no disagreement yet, and for positions $3, \dots, 6$, about half of the decisions follow a history where all previous players agreed.

3 Reluctance to contradict one’s own information

This section shows the main result concerning the question in the title of the paper. Figure 1 plots $\text{mean_pay|contradict}_i$ against the proportion of contradicting one’s signal among all identical situations. That is, for each decisions s_i with $\text{sitfreq}_i > 10$, the figure contains a marker with x -value $\text{mean_pay|contradict}_i$ and y -value given by the following variable.

prop_contradict_i: Frequency of $\text{contradict}_j = 1$ across all $s_j \in \tilde{S}(s_i)$.

Every marker in the scatterplot reflects one of the 364 distinct decisions. The regression uses all 15734 individual decisions and not just each distinct situation once, thereby weighting according to sitfreq_i . It includes an intercept plus linear, squared and cubed terms of the regressor variable.

The figure shows a large discrepancy between decisions where $\text{mean_pay|contradict}_i$ is small versus large: for the decisions in the left half of the figure, where the own signal is correct in more than half of the cases, the participants mostly make the correct choice and follow their signal. In contrast, in the decisions depicted in the right half of the figure the frequency of the signal being correct is smaller than one half but the participants often fail to contradict their signal. Had they made the payoff-maximizing choice in each case, or even approximately so, the regression line would be an S-shaped line through $(0.5, 0.5)$, which is a possible shape of the fitted line with squared and cubed regressors, and it would lie close to 1 in the right half of the figure. But the proportions in the right half are mostly far away from this optimum. Averaging over all decisions with $\text{mean_pay|contradict}_i \geq 0.5$, the frequency of the optimal choice is 0.44. In contrast, for the decisions in left half of the figure, they make the optimal choice with a frequency of 0.91.




Figure 1: Proportion of contradicting own signal across all decisions in the restricted data set with $\text{sitcount}_i > 10$. Note: 346 distinct decisions, 15734 individual decisions in total. Regression includes linear, squared and cubed x-variable.

The correspondence at least has the right slope as the proportion of contradicting the signal increases in $\text{mean_pay|contradict}_i$. A natural question to ask is how large does $\text{mean_pay|contradict}_i$ have to be for the average participant to optimally contradict their own signal with more than probability 0.5. The answer is most easily seen by observing that the regression line reaches the level of 0.5 at $\text{mean_pay|contradict}_i = 0.68$. (Considering a series of subsamples of the data, conditioning on prop_contradict_i , would yield the same result.) The empirical likelihood of the own signal being correct therefore has to be very low before the typical participant discards the signal. On average, only in decisions where the signal is wrong more than twice as often as it is correct (i.e., at an empirical odds ration of 2:1 against the signal) do the participants contradict this signal with a frequency of more than one half.

Without making any assumptions on the process of decisionmaking, one can conclude from these observations that choices cannot be best responses to rational expectations about the underlying distributions. Whatever the true structure of the data generating process, it cannot be that the participants optimally respond to it because otherwise the proportions in the right half of the figure would lie closer to 1. One can even reject that participants only imperfectly respond to rational expectations but in a way that is unbiased between following and contradicting the own signal. Under such a model, rational expectations would still imply that the unsystematic disturbances would only attenuate the S-shape in the estimated correspondence: for large samples the regression consistently estimates the frequency of contradicting the own signal as a function of this action's value (consistent for $\text{sitcount}_i \rightarrow \infty$). Under rational expectations the turnover point therefore lies at $\text{mean_pay|contradict}_i = 0.5$. But the vertical distance between the regression line and $(0.5, 0.5)$ is highly significant.²¹ ²² Section 4 will provide further tests and a more precise formulation of the

²¹The t -value for the vertical distance between the regression line and $(0.5, 0.5)$ is $t = -24.15$ in a regression where the errors are clustered by group $_i$.

²²A straightforward argument rules out the possibility that the failure of the regression line to go through $(0.5, 0.5)$ is driven entirely by random deviations of $\text{mean_pay|contradict}_i$ from its true expected value. One can estimate the standard deviation of $\text{mean_pay|contradict}_i$ for each decision s_i , using the fact that $\text{mean_pay|contradict}_i$ is a counting variable with a known sample size. The mean of these estimated standard deviations is 0.05, far smaller than the horizontal distance of the regression line from $(0.5, 0.5)$. Another possible check of the potential impact of limited sample sizes is to consider only cases where sitcount_i exceeds more stringent thresholds, so that more

rational expectation hypothesis.

But recall the discussion in the introduction, where it was argued that the frequency of making the optimal choice is only a partial indicator of success because the participants earn more from making the optimal choice in some situations than in others. To calculate the payoff of a participant in decision s_i , let $p(\text{contradict}_i)$ be her probability of choosing $\text{contradict}_i = 1$. Her expected earnings are

$$E[\pi_i] = p(\text{contradict}_i) \cdot E[\pi_i | \text{contradict}_i = 1] + (1 - p(\text{contradict}_i)) \cdot E[\pi_i | \text{contradict}_i = 0].$$

Here, $E[\pi_i | \text{contradict}_i = 1]$ is the expected payoff from contradicting the signal – which varies between decisions s_i, s_j if they are not identical, and cannot be influenced by the participant’s action. Due to the fact that one action yields 1 iff the other action yields 0, it holds that $E[\pi_i | \text{contradict}_i = 0] = 1 - E[\pi_i | \text{contradict}_i = 1]$, so the expression can be rewritten as

$$E[\pi_i] = p(\text{contradict}_i) \cdot (2E[\pi_i | \text{contradict}_i = 1] - 1) + 1 - E[\pi_i | \text{contradict}_i = 1]. \quad (1)$$

Substituting the sample means into this expression – prop_contradict_i for $p(\text{contradict}_i)$ and $\text{mean_pay} | \text{contradict}_i$ for $E[\pi_i | \text{contradict}_i = 1]$ – yields the actual average earnings in decisions that are identical to s_i . Applying this calculation to all decisions where $\text{mean_pay} | \text{contradict}_i \geq 0.5$ shows that in these situations the participants earn 0.53 of the normalized pie size. In contrast, had they chosen $p(\text{contradict}_i) = 1$ in all of these decisions, they would have earned 0.64, and they would have earned 0.36 from doing the opposite, $p(\text{contradict}_i) = 0$. Had they simply randomized between A and B , they would have predicted the correct urn in half of the cases and therefore would have earned 0.5. Hence, their actual earnings are much closer to the payoff from simple randomization than they are to the payoff-maximizing strategy. In contrast, the same calculation for the decisions with $\text{mean_pay} | \text{contradict}_i < 0.5$ shows that in these situations the success of the participants is much higher: when following the own signal is optimal, the participants earn 0.73 out of the 0.75 that they could have earned if they had always made the optimal choice.

reliable values of $\text{mean_pay} | \text{contradict}_i$ are used. The pattern is unchanged when such additional restrictions are made ($\text{sitcount}_i > 30$ and $\text{sitcount}_i > 50$, respectively).

But from this analysis one cannot tell whether there is a bias in one class of situations, or the other class, or in both – one can only establish a discrepancy between the two. Evidently, the data pattern suggests that relative to rational expectations, the participants have a general tendency to follow their signal. To the extent that such a bias exists also in the left half of Figure 1, this bias would help the participants to make the optimal choice. In the right half of the figure, such a bias would be harmful and participants could earn more if they were to overcome it. But little can be said about where the anomaly occurs.

In the absence of a structural model of behavior, the analysis therefore does not reveal much about the nature of the deviations from rational expectations. The next section will yield some insights on this issue, by looking for additional systematic patterns across different subsets of the data. It will demonstrate that the participants' ability to learn from others depends on whether the other players' actions are in contradiction to the own signal and whether the other participants show an unambiguous string of actions.

I also wish to point out clearly that other contributions to the literature have formulated models and other arguments that speak to the evidence presented above. Willinger and Ziegelmeyer (1998) were the first to raise the tendency to follow the own signal, which was later discussed most explicitly in Nöth and Weber (2003) and Goeree et al (2007). In the model estimates of Goeree et al (2007), the weight on the own signals is estimated to be significantly higher than the weight of a player's belief before observing the signal. In the error-rate analyses of Nöth and Weber (2003), Oberhammer and Stiehler (2003), Kübler and Weizsäcker (2004) and Kraemer, Nöth and Weber (2006), related arguments are given in the observations that the participants appear to attribute too low precision to other players. But all of these earlier results rely on much more restrictive assumptions about the decisionmaking process, and none has assessed the economic significance of the deviations from optimal play in much detail. This is further pursued in Section 4.

4 Accounting for losses: Behavioral versus situational determinants of payoffs

This section investigates co-variables of behavior, with the initial goal of identifying conditions under which social learning is more or less successful, measured by the frequency with which the participants choose the empirically optimal action. The section will then demonstrate that a substantial proportion of the payoff variation between these different conditions can be attributed to behavioral deviations from a rational-expectations benchmark. For this demonstration, the value of the available information will be held constant across different situations, so that the remaining payoff differences can be attributed to changes in the participants' behavior.

The first explanatory variable of interest is a proxy for the amount of available information. It describes whether the participants observe more or fewer other participants in the current repetition of the game:

late_{*i*}: An indicator variable that takes on the value of 1 if the participant has player position $t \geq 5$.

One may expect that social learning is more successful in cases where $\text{late}_i = 1$ because people may be more easily persuaded to pay attention to others if the amount of observable actions is large. On the other hand, they may also find it harder to interpret a large set of actions relative to a small set.

The second variable of interest indicates whether or not the actions of other players tend to contradict the player's own signal.

counter_majority_{*i*}: An indicator of whether or not the player's own signal is in contradiction to a strict majority of the predecessors' actions. (For example, if predecessors's actions are *AAB* and the participant's signal is *b*.)

Given that the previous section has already demonstrated a systematic tendency towards following the own signal, it is natural to suspect that the participants are more successful in situations where other players' actions indicate that they received the same signal ($\text{counter_majority}_i = 0$), because in these situations the signal tends to be correct.²³

²³Not surprisingly, $\text{counter_majority}_i$ is highly correlated with the separation in Section 3, according to

The third co-variate is whether the previous players in the current game unanimously chose the same action.

full_agreement_i: An indicator that is 1 if all predecessors chose A or all predecessors chose B , in the current repetition of the game.

Perfect agreement among predecessors can potentially make learning more successful because the string of observations is easier to interpret. On the other hand, if a player makes the mistake to follow the own signal too often, then this mistake may be more costly in cases with $\text{full_agreement}_i = 1$ because in these cases it is more likely that the other players are correct. Another possibility – which will be confirmed below – is that full agreement among other players has an especially strong effect on the success of learning when it is interacted with the length of the observed string of actions (late_i). Such an interaction effect may arise e.g. due to the difficulty of interpreting long strings of observations that are not unanimous.

The following is the dependent variable of all regressions in this section. It measures whether the participant chooses the action that is empirically optimal.

optimal_i: An indicator of whether or not the participant chooses the action with higher empirical value. If $\text{mean_pay|contradict}_i > 0.5$, then optimal_i is 1 iff the participant contradicts her signal. If $\text{mean_pay|contradict}_i \leq 0.5$, optimal_i is 1 iff the participant follows her signal.

The following variable is the empirical value of the optimal action (paralleling the calculation of $\text{mean_pay|contradict}_i$):

mean_pay|optimal_i: The frequency across $s_j \in \tilde{S}_i$ of receiving the high payoff from choosing $\text{optimal}_i = 1$.

The variable $\text{mean_pay|optimal}_i$ is included as a control in some of the subsequent regressions. Controlling for the value of the actions is relevant because the participants may exhibit a payoff-sensitive precision of play. For decisions where the empirical value of making the correct choice is low (close to 0.5) the participants may have a relatively low success rate, by either of three

 $\text{mean_pay|contradict}_i <> 0.5$. In the restricted dataset with $\text{sitcount}_i > 10$, there are 5146 decisions with $\text{counter_majority}_i = 1$, of which 4012 also have $\text{mean_pay|contradict}_i > 0.5$. The conceptual differences are that the value of $\text{counter_majority}_i$ is observable to the participants and that it does not rely on estimates of underlying values, which makes the interpretation of results simpler.

mechanisms: (i) a larger difficulty of finding the optimal action, (ii) a consciously higher likelihood to deviate in situations where less is at stake, (iii) the fact that $\text{mean_pay|optimal}_i$ measures the "true" underlying expected value with measurement error, so that the optimal action may be misclassified for cases where $\text{mean_pay|optimal}_i$ is close to 0.5. In either case, given that the dummy variables are correlated with $\text{mean_pay|optimal}_i$, the latter should be included.

With such a control, the dummy coefficients in the regressions have a distinct interpretation: they describe the size of behavioral changes that are associated with the nature of the situation (as captured by the dummies), but not with the importance of making the respective choices. Under the assumption that the participants have rational expectations about the underlying distributions and best respond to these expectations with equal precision across the dummy categories, all of the dummy coefficients are predicted to approach zero as $\text{mean_pay|contradict}_i$ approaches the true expected payoff. The regressions with controls therefore represent a simple test of the hypothesis that people respond to rational expectations.

The following gives a precise definition of the rational-expectations hypothesis that is tested here. Let $\mathbf{x}_i \equiv (\text{treatment}_i, I_i)$ be the vector of variables that is observable to the decisionmaker in situation s_i , including the explanatory dummies that are defined above. The hypothesis is that the player uses the relevant payoff information that is contained in \mathbf{x}_i , but disregards \mathbf{x}_i otherwise: a player exhibits (probabilistic) best responses to rational expectations if action A is chosen iff

$$E[\pi_i|A, \mathbf{x}_i] - E[\pi_i|B, \mathbf{x}_i] + \epsilon_i \geq 0$$

where ϵ_i is i.i.d. across i , with a distribution F that is symmetric around zero, and independent of \mathbf{x}_i . The choice probability is therefore

$$\begin{aligned} p_i(A) &= 1 - F(E[\pi_i|A, \mathbf{x}_i] - E[\pi_i|B, \mathbf{x}_i]) \\ &= 1 - F(2E[\pi_i|A, \mathbf{x}_i] - 1). \end{aligned}$$

Since the labelling choice of actions A and B is arbitrary in this formulation, one can write the probability of making the empirically optimal choice as $1 - F(2E[\pi_i|\text{optimal}_i, \mathbf{x}_i] - 1)$, which depends on $E[\pi_i|\text{optimal}_i, \mathbf{x}_i]$ alone. Therefore, if a regression controls for $E[\pi_i|\text{optimal}_i, \mathbf{x}_i]$ but includes elements of \mathbf{x}_i as separate regressors (the dummies) then the latter's coefficients are predicted to be

zero. Since $\text{mean_pay|optimal}_i$ approaches $E[\pi_i|\text{optimal}_i, \mathbf{x}_i]$ for large sample sizes ($\text{sitcount}_i \rightarrow \infty$), the t-tests for the dummies in a regression with controls for $\text{mean_pay|optimal}_i$ are consistent tests of rational expectations.

In addition, the following tables will also contain regressions where $\text{mean_pay|optimal}_i$ is not included. These regressions only contain dummy variables, thus allowing to read off the average success rates for each category from the estimated coefficients.

Table 3 considers in columns (1) to (8) the main effects of the above explanatory variables. In all of the subsequent analysis, attention is restricted to player positions $t \geq 2$, i.e. to those players who observe at least one predecessor in the current repetition of the game.

	optimal _i			
	(1)	(2)	(3)	(4)
late _i	0.116 (.009)	-0.002 (0.008)	-	
counter_majority _i	-	-	-0.343 (.013)	-0.199 (.015)
full_agreement _i	-	-	-	-
mean_pay optimal _i	-	24.055 (4.583)	-	23.077 (4.277)
mean_pay optimal _i ²	-	-26.864 (6.043)	-	-26.145 (5.623)
mean_pay optimal _i ³	-	10.176 (2.618)	-	9.918 (2.427)
constant	.705 (.009)	-6.370 (1.140)	0.887 (0.009)	-5.861 (1.068)
# of obs.	12536	12536	12536	12536
R ²	0.016	0.217	0.150	0.255

Table 3: Frequencies of making the empirically optimal choice. Note: Data includes cases with $\text{sitcount}_i > 10$ and $t \geq 2$. Robust standard errors in parentheses, clustered by group_i .

	optimal _i			
	(5)	(6)	(7)	(8)
late _i	-	-	0.126 (.010)	0.039 (0.009)
counter_majority _i	-	-	-0.348 (0.013)	-0.211 (0.016)
full_agreement _i	-0.010 (0.015)	0.001 (0.010)	0.062 (0.013)	0.034 (0.012)
mean_pay optimal _i	-	24.101 (4.582)	-	24.128 (4.163)
mean_pay optimal _i ²	-	-26.929 (6.054)	-	-27.644 (5.469)
mean_pay optimal _i ³	-	10.205 (2.626)	-	10.580 (2.363)
constant	0.752 (.011)	-6.381 (1.135)	0.805 (0.013)	-6.117 (1.037)
# of obs.	12536	12536	12536	12536
R ²	0.000	0.217	0.170	0.257

Table 3 (ctd.): Frequencies of making the empirically optimal choice.

	optimal _i			
	late _i =0		late _i =1	
	(9)	(10)	(11)	(12)
full_agreement _i	-0.048 (0.021)	-0.024 (0.013)	0.119 (0.014)	0.080 (0.013)
mean_pay optimal _i	-	26.689 (6.444)		1.345 (5.398)
mean_pay optimal _i ²	-	-28.364 (8.485)		-1.372 (7.054)
mean_pay optimal _i ³	-	10.039 (3.671)		0.839 (3.024)
constant	0.740 (0.017)	-7.382 (1.602)	0.764 (0.010)	0.157 (1.352)
# of obs.	8131	8131	4405	4405
R ²	0.002	0.285	0.024	0.084

Table 3 (ctd.): Frequencies of making the empirically optimal choice.

	optimal _{<i>i</i>}			
	late _{<i>i</i>} =0		late _{<i>i</i>} =1	
	(13)	(14)	(15)	(16)
counter_majority _{<i>i</i>}	-0.246 (0.032)	-0.052 (0.032)	-0.176 (0.021)	-0.158 (.023)
full_agreement _{<i>i</i>}	0.109 (0.020)	0.067 (0.016)	0.146 (0.011)	0.102 (0.011)
(counter_majority _{<i>i</i>} × full_agreement _{<i>i</i>})	-0.221 (0.034)	-0.178 (0.030)	-0.075 (0.032)	-0.035 (0.030)
mean_pay optimal _{<i>i</i>}	-	27.603 (5.700)		5.423 (5.269)
mean_pay optimal _{<i>i</i>} ²	-	-31.046 (7.480)		-5.943 (6.878)
mean_pay optimal _{<i>i</i>} ³	-	11.614 (3.232)		2.325 (2.937)
constant	0.804 (0.018)	-7.244 (1.426)	0.835 (0.012)	-0.867 (1.315)
# of obs.	8131	8131	4405	4405
R^2	0.208	0.315	0.010	0.126

Table 3 (ctd.): Frequencies of making the empirically optimal choice.

The overriding feature of the regressions is the strongly negative effect of counter_majority_{*i*}. In situations where most previous players choose a different action than is indicated by a participant's signal, her frequency of making the empirically optimal choice is much lower. The difference in frequencies is 0.343, relative to an average success rate of 0.887 in the situations with counter_majority_{*i*} = 0, see column (3). The result corresponds to the tendency described in Section 3: people far too often follow their own signals, and this is harmful in situations where the majority of previous choices indicates that the signal is likely wrong. In part, the result can be explained by the different incentives that the participants face depending on the value counter_majority_{*i*}. As column (4) shows, controlling for the incentives by including mean_pay|optimal_{*i*} reduces the size of the coefficient on counter_majority_{*i*}. This reflects the fact that mean_pay|optimal_{*i*} is lower in situations where counter_majority_{*i*} = 1, and that participants make the empirically optimal choice more often when it is important. But even when mean_pay|optimal_{*i*} is controlled for, the partial effect of counter_majority_{*i*} is -0.199, a large and highly significant deviation from the hypothesis

of best responding to rational expectations.

The variable late_i has a positive coefficient in column (1), indicating that participants in positions $t \geq 5$ have a higher success rate on average, by 11.6 percentage points. However, controlling for mean_pay|opt_i in column (2) shows that the higher success is entirely associated with the higher incentive to make the optimal choice. Holding mean_pay|opt_i constant, the success rates are virtually identical between earlier and later positions in the game. Higher earnings for later players are therefore associated with the higher lucrativity of situations. The rate of optimal play improves for later players in the game, but only in correspondence with the higher incentives.

The variable full_agreement_i is not correlated with success rates, unless the other dummy variables are introduced as well (columns (5) to (8)).²⁴ This is surprising because one would have expected learning to be easier when $\text{full_agreement}_i = 1$. But the difference between the results in columns (5) and (6) versus (7) and (8) points at a potential interaction between the dummy variables, which is examined next. Columns (9) to (16) interact the three variables of interest, by separating the data according to late_i and, in columns (13) to (16), by including an interaction term for $(\text{counter_majority}_i \times \text{full_agreement}_i)$. The regressions show that a full agreement among the previous players can indeed strongly increase the success rate, but especially for late positions in the game. This is most directly seen in columns (9) to (12), where full_agreement_i has a significant and positive effect only in the subsample with $\text{late}_i = 1$.

The interaction effect of late_i and full_agreement is further clarified if it is also interacted with

²⁴The results on full_agreement are not all robust to excluding subsamples. When restricting the sample to $t \geq 3$ instead of $t \geq 2$, there is a small but significant effect of full_agreement_i : the dummy coefficient in regression (5) changes to 0.050 (std. err. 0.014). But in column (7), the analogous coefficient is 0.001 (0.010). When excluding the largest individual data set, Nöth and Weber (2003), full_agreement_i is significantly and strongly positively correlated with optimal_i , with dummy coefficients of 0.160 (0.038) and 0.20 (0.038) in the regressions corresponding to columns (5) and (7), and 0.120 (0.039) and 0.176 (0.085) in columns (9) and (11). However, when $\text{mean_pay|optimal}_i$ is controlled for, parallel to the regressions in columns (6), (8), (10) and (12), the coefficients are reduced to 0.053 (0.032), 0.111 (0.035), 0.050 (0.033) and -0.006 (0.076), respectively. All other results of Tables 3 are qualitatively unchanged when the data are restricted to subsets that exclude $t = 2$ situations or the data from Nöth and Weber (2003), or to observations with $\text{sitcount}_i > 30$. All conclusions from Table 3 are also supported by repeating the regressions as a probit instead of linear model.

counter_majority_{*i*}, in columns (13) to (16). First consider the case that counter_majority_{*i*} = 0. For these cases, the coefficients of full_agreement_{*i*} show that observing an unanimous set of actions increases the success rate by 0.109 for late_{*i*} = 0 and by 0.146 for late_{*i*} = 1. This confirms the results of columns (9) to (12) to a somewhat weaker degree, for those situations where the majority of previous actions and the participants's own signal are in agreement. But for counter_majority_{*i*} = 1 and late_{*i*} = 0, the participants in situations with full_agreement_{*i*} = 1 have a *lower* success rate than with full_agreement_{*i*} = 0, by a difference of 0.109 - 0.221 = -0.112.²⁵ For late_{*i*} = 1, on the other hand, the participants in situations with full_agreement_{*i*} = 1 have a success rate that is higher by 0.146 - 0.075 = 0.071, as one would expect.

In sum, the discussion shows that for later players, the positive effect of having previous players who all agree is larger than for early players. The effect counteracts to a substantial degree the negative effect of counter_majority_{*i*}: in situations where all three dummy variables are 1 (e.g. histories with information *AAAAb* or *AAAAAb*), the average success rate is 0.730, which is higher than in any other constellation with counter_majority_{*i*} = 1. In other words, the participants are fairly successful in learning from others only when the evidence conveyed by the others' choices is strong (late_{*i*} = 1) and unambiguous (full_agreement_{*i*} = 1).

Controlling for the empirical value of the optimal action (in regression in even-numbered columns) does not qualitatively change these conclusions but only reduces the size of the effects somewhat. This shows again that the hypothesis of best responding to rational expectations is violated – the participants exhibit systematically lower success rates in those situations where the actions of others are not unambiguous.

As an illustration, Figure 2 contains the same variables as Figure 1, plotting the empirical value of contradicting one's signal against the corresponding frequency, but highlights with solid markers

²⁵In part, this is driven by the observation that players at positions $t = 2$ and $t = 3$ have a very low success rate if counter_majority = 1 (0.404 and 0.453, respectively) and that by definition of the variables, in these cases it is always true that full_agreement_{*i*} = 1 holds because there cannot be a strict but ambiguous majority with only one or two predecessors. But it also reflects that there is no positive effect of full_agreement_{*i*} even for players at $t = 4$: conditional on counter_majority_{*i*} = 1 and $t_i = 4$, the frequency of optimal_{*i*} is 0.558 for full_agreement_{*i*} = 0 and 0.564 for full_agreement_{*i*} = 1.




Figure 2: Proportion of contradicting own signal. Note: See Figure 1. Highlighted situations are those with $late_i = 1$ and $fullagreehis_i = 1$.

the set of situations with $late_i = 1$ and $full_agreement_i = 1$. The solid line is the regression line of this subsample, whereas the dashed line is for the remaining subsample. As the figure shows, the rate of contradicting the own signal when it is optimal are much higher if the history of previous choices is long and unanimous. These situations are the typical cascade situations, where Bayes Nash Equilibrium predicts that players disregard their signals. The fact that the success rates are much lower in the remainder of the situations was never emphasized in the previous literature, to my knowledge. Instead, the high success rates along paths where cascades would theoretically arise was often pointed out (see e.g. Anderson and Holt, 1997, Kübler and Weizsäcker, 2005).

The remainder of the section describes the payoffs of participants in the different classes of situations that are captured by the dummy variables, and discusses the relative effect of behavioral

versus situational influences on the payoff variance. I will restrict attention to the two identified significant effects of $\text{counter_majority}_i$ (columns (3), (4) of Table 3) and of full_agreement_i for later player positions (columns (11), (12)). The relevant distinction between the two former classes of situations will be addressed in Table 4, and the latter in Table 5.

The first step is to calculate the obtained payoffs for the different classes of situations. Parallel to expression (1) one can calculate the average realized payoff in all situations that are identical to s_i as

$$\text{mean}(\pi_i) = \text{mean}(\text{optimal}_i) \cdot (2 \cdot \text{mean_pay}|\text{optimal}_i - 1) + 1 - \text{mean_pay}|\text{optimal}_i. \quad (2)$$

The row labelled "realized earnings" in Table 4 shows the corresponding values (averages of $\text{mean}(\pi_i)$ across all situations) for each of the categories of $\text{counter_majority}_i$. Across all situations with $\text{counter_majority}_i = 0$, the participants receive an average payoff of 0.748, and they receive an average payoff of 0.545 if $\text{counter_majority}_i = 1$.

Had they always chosen the empirically optimal action, how much would they have earned? This can be calculated by replacing optimal_i by 1 in expression (2). The row labelled "maximum earnings" shows the corresponding calculation results. Comparing realized with maximum earnings, it shows that for $\text{counter_majority}_i = 1$, the realized earnings are much further from their maximum than for $\text{counter_majority}_i = 0$. But the entries also show that situations with $\text{counter_majority}_i = 1$ are much more lucrative: choosing the optimal action yields 0.785 on average, compared to 0.648 in situations with $\text{counter_majority}_i = 1$. Overall, the calculations illustrate that the difference in the realized earnings is driven by both the difference in behavior and the difference in incentives. Both components have a substantial impact.

To explore this further, one can ask the following counterfactual question: by how much would the difference in the earnings decrease if the difference in behavior was absent? For example, how much would earnings under $\text{counter_majority}_i = 1$ increase if the participants were equally successful in making optimal choices as under $\text{counter_majority}_i = 0$? The answer can be generated by replacing $\text{mean}(\text{optimal}_i)$ in expression (2) by $(\text{mean}(\text{optimal}_i) - \text{dummy coefficient})$, where the dummy coefficient is taken from the regressions of column (3) or (4) of Table 3. This coefficient measures the difference in the success rates, so that subtracting it generates the desired

counterfactual behavioral change. The difference to the realized payoffs is then given by

$$\Delta \widehat{E}[\pi_i] | 1 \rightarrow 0 = -\text{dummy coefficient} \cdot (2 \cdot \text{mean_pay|optimal}_i - 1),$$

which is reported in the last row of Table 4. Similarly, the term $\Delta \widehat{E}[\pi_i] | 0 \rightarrow 1$ in the preceding row measures the counterfactual change in payoffs under $\text{counter_majority}_i = 0$ that would arise if behavior was equally (un)successful as in situations with $\text{counter_majority}_i = 1$. For each of the two rows, the first entry measures the counterfactual payoff increase when using only the simple dummy regression of column (3) of Table 3 – hence uses the simple difference in the data average as the dummy coefficient – whereas the second entry uses the regression that controls for $\text{mean_pay|optimal}_i$ when describing the behavioral difference (column (4)). Recall that under the hypothesis that participants best respond to rational expectations, the inclusion of such controls should capture any impact of the nature of the situation, including the dummy values. Hence, the counterfactual payoff changes in Table 4's columns (2) and (4) are estimates of the payoff changes that are associated with the observed deviations from rational expectations. For example, the entries in column (4) of Table (4) show that if behavior was equally successful under $\text{counter_majority}_i = 1$ as it is under $\text{counter_majority}_i = 0$, payoffs under $\text{counter_majority}_i = 1$ would increase by almost 6 percentage points, to about 0.61. The observed payoff differences between the two classes of situations would be substantially reduced if behavior was identical.

Table 5 repeats the above exercises for the distinction between situations where $\text{full_agreement}_i = 1$ versus $\text{full_agreement}_i = 0$, both conditional on $\text{late}_i = 1$. Here, too, it shows that the sizable differences in realized payoffs between the two classes of situations (0.759 versus 0.663) would be substantially reduced – to at most half of the actual difference – if behavior was equally successful in both categories. This confirms that also for this second "irrational" reaction to the nature of situations, the payoff was strongly affected.

	counter_majority _i = 0		counter_majority _i = 1	
	(1)	(2)	(3)	(4)
controls for				
mean_pay optimal	no	yes	no	yes
realized earnings		0.748		0.545
maximum earnings		0.785		0.648
$\Delta \widehat{E}[\pi_i] 0 \rightarrow 1$	-0.195	-0.114	-	-
$\Delta \widehat{E}[\pi_i] 1 \rightarrow 0$	-	-	0.101	0.059

Table 4: Payoff results in response to behavioral variations
depending on counter_majority_i

	full_agreement _i = 0		full_agreement _i = 1	
	late _i = 1		late _i = 1	
	(1)	(2)	(3)	(4)
controls for				
mean_pay optimal	no	yes	no	yes
realized earnings		0.663		0.759
maximum earnings		0.758		0.813
$\Delta \widehat{E}[\pi_i] 0 \rightarrow 1$	0.061	0.041	-	-
$\Delta \widehat{E}[\pi_i] 1 \rightarrow 0$	-	-	-0.074	-0.050

Table 5: Payoff results in response to behavioral variations
depending on full_agreement_i

5 Conclusions

The paper presents a large meta dataset from thirteen cascade game experiments, and applies a simple counting technique to infer that the success rate is low in decisions where it would be optimal to follow the choices of others. This demonstrates a violation of rational expectations about the conditional likelihoods of the states of the world. This conclusion can be drawn without imposing any structural model of what players might believe about the behavior of the players that acted before her.

The paper also discusses the payoff impact that the deviations from rational expectations have. In situations where the own signal runs counter to a strict majority of previous choices, the participants earn much less than in the remaining situations, and it is estimated that the "behavioral" deviations from rational expectations account for a large share of the payoff difference. If the participants were as successful in identifying the optimal choice in these decisions as they are in the remaining decisions, the payoff difference would be reduced by a proportion between one quarter and one half. A similar result occurs for cases where the previous players were not in full agreement, which causes players to mis-read the available information and forgo substantial payoffs.

The methodology of the test for rational expectations in Section 4 is simple but novel in experimental economics, so the question arises whether it can be applied to other experimental (or other) data sets. It surely can, and in fact aspects of it are already applied in numerous studies: it is often calculated – though not often enough – how much the participants in an experiment give up due to their failure to choose a best response against the empirical distribution of payoff-relevant random variables. The innovative component of the methodology here is to not only estimate the empirical value of the available actions, but to also use it as an explanatory variable in regressions that describe behavior. This yields a test of rational expectations that does not rely on a structural behavioral model.²⁶ The key requirement is to have sufficient data so that the empirical value of an action is a useful estimate of the underlying "true" expectation. With a meta data set, which typically has lots of data, one can apply this approach to a large variety of decisions. As in other

²⁶See Weizsäcker (2003) and Goeree and Holt (2004) for Quantal-Response-Equilibrium-based tests of rational expectations in normal form games.

meta analyses, the wealth of data thus allows multiple different perspectives on the data – including some that are unanticipated by the original studies’ authors – without giving up statistical reliability. This general advantage of meta analyses is probably underrated, too, in the literature.²⁷

References

Alevy, Jonathan E., Michael S. Haigh and John A. List (2007). "Information cascades: Evidence from a field experiment with financial market professionals." *Journal of Finance* 62, 151-180.

Anderson, Lisa R. (2001). "Payoff Effects in Information Cascade Experiments." *Economic Inquiry* 39, 609-614.

Anderson, Lisa R., and Charles A. Holt (1997). "Information Cascades in the Laboratory." *American Economic Review* 87, 847-862.

Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch (1992). "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades." *Journal of Political Economy* 100, 992-1026.

Celen, Bogachan, and Shachar Kariv, Distinguishing Informational Cascades From Herd Behavior in the Laboratory." *American Economic Review* 94, 484-497.

Choi, Syngjoo, Douglas Gale, and Shachar Kariv (2004). "Learning in Networks: An Experimental Study." Working Paper, New York University.

Cipriani, Marco, and Antonio Guarino (2005). "Herd Behavior in a Laboratory Financial Market." *American Economic Review* 95, 1427-1443.

Dominitz, Jeff, and Angela A. Hung (2004). "Homogeneous Actions and Heterogeneous

²⁷Harless and Camerer (1994), Zelman (2003), Oosterbeek, Sloof and van de Kuilen (2004), Huck, Normann and Oechssler (2004) and Engel (2007) are earlier examples of meta analyses in experimental economics that tackle novel questions.

Beliefs: Experimental Evidence on the Formation of Information Cascades." Working Paper, Carnegie Mellon University.

Drehmann, Mathias, Jörg Oechssler, and Andreas Roider (2005). "Herding and Contrarian Behavior in Financial Markets - An Internet Experiment." *American Economic Review* 95, 1403-1426.

Engel, Cristoph (2007). "How Much Collusion? A Meta Analysis of Oligopoly Experiments." *Journal of Competition Law and Economics* 3, 491-549.

Fahr, Rene, and Bernd Irlenbusch (2008). "Who follows the crowd – groups or individuals?" Working Paper, University of Cologne.

Goeree, Jacob, and Charles Holt (2004). "A Model of Noisy Introspection." *Games and Economic Behavior* 46, 365-382.

Goeree, Jacob K., Thomas R. Palfrey, Brian W. Rogers, and Richard D. McKelvey (2007). "Self-correcting Information Cascades." *Review of Economic Studies* 74, 733-762.

Grebe, Tim, Julia Schmidt, and Andreas Stiehler (2006) "Do individuals recognize cascade behavior of others? An experimental study." Working Paper, Humboldt University Berlin.

Guarino, Antonio, Heike Harmgart, and Steffen Huck (2007). "When half the truth is better than the truth: Theory and experiments on aggregate information cascades." ELSE Working Paper, University College London.

Harless, David, and Colin F. Camerer (1994). "The Predictive Utility of Generalized Expected Utility Theories." *Econometrica* 62, 1251-1290.

Huck, Steffen, Hans-Theo Normann and Jörg Oechsler (2000). "Two are Few and Four are Many: Number Effects in Oligopoly Experiments." *Journal of Economic Behavior and Organization* 53, 435-446.

Huck, Steffen, and Jörg Oechsler (2000). "Information Cascades in the Laboratory: Do They Occur for the Right Reasons?" *Journal of Economic Psychology* 21, 661-667.

- Hung, Angela, and Charles Plott (2001). "Information Cascades: Replication and an Extension to Majority Rule and Conformity Rewarding Institutions." *American Economic Review* 91, 1508-1520.
- Kraemer, Carlo, Markus Nöth and Martin Weber (2006). "Information Aggregation with Costly Information and Random Ordering: Experimental Evidence." *Journal of Economic Behavior and Organization*, 59, 423-432.
- Kübler, Dorothea and Georg Weizsäcker (2004). "Limited depth of reasoning and failure of cascade formation in the laboratory." *Review of Economic Studies* 71, 425-441.
- Kübler, Dorothea, and Georg Weizsäcker (2005). "Are longer cascades more stable?" *Journal of the European Economic Association* 3, 330-339.
- Nöth, Markus, and Martin Weber (2003). "Information aggregation with random ordering: Cascades and overconfidence." *Economic Journal* 113, 166-189.
- Oberhammer, Clemens, and Andreas Stiehler (2003). "Does cascade behavior in information cascades reflect Bayesian Updating?" *Papers on Strategic Interaction*, Max Planck Institute for Research into Economic Systems.
- Oosterbeek, Hessel, Randolph Sloof and Gijs van De Kuilen (2004). "Cultural Differences in Ultimatum Game Experiments: Evidence from a Meta-Analysis." *Experimental Economics* 7, 171-188.
- Weizsäcker, Georg (2003). "Ignoring the Rationality of Others: Evidence from Experimental Normal-Form Games." *Games and Economic Behavior* 44, 145-171.
- Willinger, Marc, and Anthony Ziegelmeyer (1998). "Are more Informed Agents able to Shatter Information Cascades in the Lab?" In: Cohendet, Patrick, Patrick Llerena, Hubert Stahn, Gisele Umbhauer (eds.), *The Economics of Networks: Interaction and Behaviors*, Springer Verlag.
- Zelmer, Jennifer (2003). "Linear Public Goods Experiments: A Meta-Analysis." *Experimental Economics* 6, 299-310.

Zieglmeyer, Anthony, Jürgen Bracht, Frédéric Koessler, and Eyal Winter (2002). "Fragility of Information Cascades: An Experimental Study Using Elicited Beliefs." Working Paper, Max Planck Institute for Research into Economic Systems.