

WISE CROWDS OR WISE MINORITIES?

Christoph Brunner and Jacob K. Goeree*

January 1, 2012

Abstract

This paper reports results from social learning experiments where subjects choose between two options and each subject has a small chance of being perfectly informed about which option is correct. In treatment “sequence,” subjects observe the entire sequence of predecessors’ choices while in treatment “no-sequence” they only observe the number of times each option has been chosen. The theoretical predictions are that subjects follow their immediate predecessors in treatment sequence and follow the minority in treatment no-sequence (Callander and Hörner, 2009). The former prediction is borne out in the data, but subjects tend to follow the majority in treatment no-sequence. We observe substantial heterogeneity in levels of strategic thinking, as predicted by level- k and Cognitive Hierarchy. While these models reproduce some features of our data, their fit is poor because of the assumed best-response behavior. Allowing for some degree of “trembling” improves the fit significantly, especially if subjects are aware that others tremble, as in logit-QRE. The “noisy introspection” model proposed by Goeree and Holt (2004), which combines different levels of thinking with error-prone behavior, best describes the data.

JEL code: C92

Keywords: Social learning, experiments, quantal response equilibrium, level- k models, noisy introspection

*Brunner: Department of Economics, University of Heidelberg, Alfred-Weber-Institut, Bergheimerstrasse 58, 69115 Heidelberg, Germany. Goeree: Chair for Organizational Design, Department of Economics, University of Zürich, Blümlisalpstrasse 10, CH-8006, Zürich, Switzerland. We gratefully acknowledge financial support from the European Research Council (ERC Advanced Investigator Grant, ESEI-249433). We thank Kevin Watts for research assistance and seminar participants at National Taiwan University (Taipei), Ohio State University (Columbus), and the Economic Science Association Meetings (Tucson) for helpful comments.

1. Introduction

When people are imperfectly informed they may try to learn from others' choices. Prospective graduate students, for instance, often inquire which schools other students chose to apply to. Teenagers consult the charts before deciding what music to buy and tourists tend to prefer well occupied restaurants to half empty competitors. In many of these situations, only the choices made can be observed not the exact information others had when making their choices. And while in some instances the exact sequence of predecessors' choices is observed (as in the US primary elections) more often only aggregate statistics based on those choices are available (e.g. the number of diners in a restaurant).

When do others' decisions contain relevant information and what course of action do they suggest? Obviously, predecessors' choices matter only when their payoffs are correlated to some extent. For this reason, most of the social learning literature makes the simplifying assumption that agents have identical values for the available options, as is the case, for instance, when buying stocks.¹ In this common-value environment, folk wisdom suggests it would be best to follow the majority, an intuition that is formalized by theoretical models of social learning. Bikhchandani, Hirschleifer, and Welch (1992), for example, consider a model where agents are privately informed about which of two options is better and the quality of information is the same across agents. They demonstrate that after a few decisions, information cascades occur and all agents herd on the majority choice regardless of their private information.² Following the majority is also optimal in Banerjee's (1992) model where uninformed agents have the ability to signal that they have no information.

In contrast, Callander and Hörner (2009) consider a situation where it can be optimal to follow the *minority*. In their model, agents differ in terms of the quality of information they possess and observe only the number of decisions for each option. In this paper, we consider the following simplified version of their model: each agent has a small chance of being perfectly informed about which of two options is correct or gets no information at all (besides the prior information that puts equal weight on both options). The information agents receive and

¹Exceptions include Smith and Sørensen (2000) and Goeree et al. (2006, 2007) who study social learning with heterogeneous payoffs. Others' choices are also important when network externalities are present, e.g. using the same technology standard. See Hung and Plott (2000) and Drehmann et al. (2005b) for sequential decision making experiments with informational and payoff externalities. Cipriani and Guarino (2005), Drehmann et al. (2005a), and Bose et al. (2009) consider situations where predecessors' choices affect payoffs by changing the prices of the alternatives.

²Laboratory evidence provides partial support for these predictions in the sense that information cascades do occur but are often broken. In addition, subjects tend to overweigh their private information vis-à-vis that contained in publicly observable predecessors' choices (see, e.g., Anderson and Holt, 1997; Çelen and Kariv, 2004a,b; Goeree et al., 2006, 2007).

the order in which they move are exogenously determined.³ Finally, the probability of being informed is low enough such that, under common knowledge of rationality, it is always optimal to follow the minority. This result is explained in more detail below, but to glean some intuition consider an uninformed agent who learns that two predecessors have chosen restaurant *A* and one has chosen restaurant *B*. Such an outcome can occur only when the second or third agent were informed (or both, but this case is relatively unlikely when the probability of being informed is low). If the second agent was informed then the third agent faced a tie and chose randomly, in which case following the minority is no worse than following the majority. And if the third agent was informed, it is strictly better to follow the minority.

This simple logic extends to more general minority-majority divisions if common knowledge of rationality can be subsumed (see Callander and Hörner, 2009). But once we introduce the possibility of “trembles” or mistakes, it breaks down. Goeree et al. (2007), for example, find that in standard social learning experiments (based on Bikhchandani et al., 1992), cascades do form but almost never last as subjects frequently opt to follow their contrary information and break the cascade.⁴ While uninformed agents in our experiment do not possess any private information,⁵ trembles may still occur especially because the information conveyed by predecessors’ choices may be of low quality and value. Intuitively, the possibility of trembles greatly alters equilibrium predictions in the Callander–Hörner setup. In the example above, for instance, the 2-1 division between restaurants is more likely caused by an uninformed agent deviating when the probability of being informed is very low.

More generally, whether it is optimal to follow the majority or the “deviant minority” therefore depends on the likelihood of mistakes, the quality of others’ information, the correlation in tastes, etc. It would be hard to distinguish these confounding elements in data from the field, which is why we turn to the lab. We conducted two types of sequential decision-making experiments: in treatment “sequence,” agents can see the entire sequence of predecessors’ decisions and in treatment “no-sequence” they only see the number of predecessors’ choices for either option. Collecting data from both treatments allows us to connect our findings to prior literature, which mostly employs the sequence treatment, and enables us to evaluate the efficiency gains that may result from the additional information in treatment sequence.

³Kübler and Weizsäcker (2004) present a study where the information available to decision makers is endogenous, while Ottaviani and Sørensen (2001) allow the sequence of decision makers to be endogenous.

⁴In these experiments, cascade breaks are informative and prevent the learning process from getting stuck. As a result, full information aggregation becomes possible in the limit as the number of agents grows large. Goeree et al. (2007) demonstrate how a logit quantal response model can account for much of the dynamics in the experiments.

⁵Previous social learning experiments have documented the tendency of subjects to overweigh their private information. This is not possible in our design, since subjects get either no information or information that is perfectly informative.

We find that informed subjects always follow their signal, i.e. they always pick the correct option. On average, uninformed subjects tend to follow unanimous predecessors close to 90% of the time both in treatments sequence and no-sequence. Furthermore, the frequency with which unanimous predecessors are followed significantly rises (to levels between 90% and 100%) as the number of predecessors grows. This high percentage of rational choices is maybe not surprising given that the decision problem faced by an uninformed agent is relatively easy when all predecessors agree. When there is a deviator in treatment sequence, the rational choice is to follow such a deviator but subjects in the experiment tend to follow the deviator only 72% of the time. The frequency with which a deviator is followed is significantly higher when the deviator’s choice belongs to the majority (80%) than when it belongs to the minority of previous choices (58%). Finally, in treatment no-sequence, subjects tend to follow the minority only 28% of the time. This percentage significantly declines when the difference between the number of majority and minority choices grows.

While observed choices deviate from theoretical predictions they are approximate best responses to the empirical distribution of play in the following sense. Given the choices of others, and given the signals used in the experiment, the cost of not following unanimous predecessors in the experiment is \$1.53 on average. Likewise, the cost of not following a deviator in treatment sequence is \$1.23 on average, and the cost of not following the minority is $-\$0.36$ on average. In other words, subjects who are (imperfect) profit maximizers would nearly always follow unanimous predecessors, would more likely than not follow a deviator in treatment sequence (although not as frequently as they would follow unanimous predecessors), and would follow the *majority* in treatment no-sequence. The canonical model that captures this type of imperfect maximization behavior is the (logit) QRE model. We show that, on an aggregate level, logit-QRE is able to reproduce the main features of our data quite well.

In the logit-QRE model, however, agents are assumed to be *ex ante* symmetric, which is clearly not true in our data. While some subjects make rational choice in all ten periods of the experiment, others do so in less than half the periods. This type of heterogeneity is captured by models that allow for different levels of strategic thinking, such as the level- k model (e.g. Crawford and Iriberry, 2007a,b) and Cognitive Hierarchy (Camerer, Ho, and Chong, 2004). When we apply level- k and Cognitive Hierarchy to the data, we find they produce a worse fit both on an individual and aggregate level. The reason for their poor performance is that these models subsume best response behavior (given beliefs), except for level-0 who randomizes when uninformed. The best-response assumption often conflicts with intuitive comparative statics observed in the data, e.g. subjects tend to follow unanimous predecessors more frequently when the group of predecessors is large. In addition, the best-response assumption implies that a

subject who mostly but not always makes a rational choice, is classified as level-0 even though most of her choices suggest a higher level of thinking.

Goeree and Holt (2004) propose a “noisy introspection” model that blends the idea of different levels of strategic thinking with noisy responses (trembling). In particular, the noisy introspection model replaces the strict best responses of the level- k model with logit “better responses.” Importantly, agents in the noisy introspection model are assumed to be aware that others tremble. For example, when computing the probability that the minority choice is correct in treatment no-sequence, agents take into account the possibility that the minority arose because of trembles. As a result, the model can predict why subsequent choices favor the majority (not the minority) even for agents with high levels of strategic thinking.

We find that the noisy introspection model fits the data best – it provides a significant improvement in fit relative to logit-QRE and a dramatic improvement relative to level- k and Cognitive Hierarchy. To illustrate the importance of the “common-knowledge-of-trembling” assumption that underlies noisy introspection, we also estimate a noisy version of the level- k model in which agents tremble but assume others don’t. We find that the noisy level- k model provides a better fit than level- k and Cognitive Hierarchy, but not nearly as good as noisy introspection.

The paper is organized as follows. In the next section, we briefly discuss the theoretical predictions for the two treatments. The experimental design is presented in Section 3. The results of the experiment can be found in Section 4. In section 5 we apply alternative models of bounded rationality to explain individual and aggregate outcomes. Section 6 concludes.

2. Theoretical Predictions

Treatment sequence is a simple variant of the social learning model proposed by Bikhchandani, Hirschleifer, and Welch (1992). There are two options, A and B , that are equally likely to be correct and a finite set of n agents labeled $t = 1, 2, \dots, n$. Each agent chooses either A or B after having observed a private signal s_t and the decisions of predecessors. Signals in the experiment are either fully informative or not informative at all: if ω denotes the correct option then $s_t = \omega$ with probability $q > 0$ and $s_t = \emptyset$ with probability $1 - q$.

Given that some agents may be fully informed, the perfect Bayesian equilibria of treatment sequence are easy to derive. First, an agent with $s_t = \omega$ chooses ω . Second, if all predecessors agree, then an uninformed agent follows the majority since either all predecessors were uninformed and the agent is indifferent or some predecessors were informed and the agent strictly

prefers to follow the majority. Third, if predecessors were not unanimous, i.e. choices switched from one option to the other, then the first predecessor who “deviated” by not following her predecessors must have been informed. In this case, the agent should follow the deviator. Note that all three cases can be succinctly summarized as follows: under common knowledge of rationality, agents follow their immediate predecessor. Finally, if predecessors switched from one option to the other and then switched back, play is off the equilibrium path. In this case, the agent cannot infer anything from prior play and simply randomizes when $s_t = \emptyset$ and chooses ω when $s_t = \omega$.

In treatment no-sequence, informed agents with $s_t = \omega$ again simply choose ω . Also, uninformed agents simply follow the majority if all predecessors agree. The surprising result for treatment no-sequence is that uninformed agents may be better off following the *minority*, a result due to Callander and Hörner (2009). To glean some intuition, consider their illustrative example in which uninformed agent $t = 4$ has to choose after two predecessors have chosen option A and 1 predecessor has chosen option B . In this case, agent $t = 1$ must have been uninformed and must have chosen the wrong option since otherwise predecessors would have been unanimous. Furthermore, either (i) agent 2 was uninformed in which case she followed agent 1 and agent 3 was informed and choose differently from her predecessors, or (ii) agent 2 was informed and choose differently from agent 1 and agent 3 was uninformed and randomly followed agent 1, or (iii) both agents 2 and 3 were informed and chose differently from agent 1. Under scenario (i) the minority choice is correct, under (ii) both options are equally likely to be correct, and under (iii) the majority choice is correct. For $q < \frac{1}{2}$, the situation in which 2 agents are informed is less likely, and, hence, the minority is more likely to be correct. A simple calculation shows that

$$P_{1,2}(\text{minority is correct}) = \frac{3 - 3q}{4 - 2q} \quad (1)$$

Moreover, for small q , the probability that the minority is correct *grows* as the majority increases in size. For example, consider the case where three predecessors have chosen option A and 1 predecessor has chosen option B . The fourth agent must have faced (i) unanimous predecessors in which case the minority choice B is correct (if option A were correct the fourth agent should pick A whether or not informed), or (ii) a 2-1 majority for A in which case the minority is incorrect (now if option B were correct the fourth agent should pick B whether or not informed, resulting in a tie). Note that situation (i) only requires the fourth agent to be informed while situation (ii) requires the fourth agent and at least one other agent to be




Figure 1: Probability Minority is Correct in Treatment No-sequence with $q = 0.2$ (Left) and $q = 0.1$ (Right) as a Function of the Number of Predecessors (x -axis) and the Minority Size (Label Next to a Line).

informed. The chance that the minority is correct is now

$$P_{1,3}(\text{minority is correct}) = \frac{2(1-q)^2}{2(1-q)^2 + q(1+q)} \quad (2)$$

and it is straightforward to verify that $P_{1,3} > P_{1,2}$ for small q . The minority wisdom becomes even stronger as the majority grows larger than 3: the outcome in which $m \geq 4$ predecessors chose A and only 1 predecessor chose B requires at least $m - 1$ agents to be informed when the majority choice A is correct, but requires only the final agent to be informed when the minority choice B is correct. Furthermore, this logic extends to minorities of sizes different than 1.

Figure 1 shows the relevant probabilities for the experimental setup discussed below: the left panel corresponds to the case where the probability of being informed is $q = 0.2$ and the total number of agents is $n = 7$, and the right panel corresponds to the case where $q = 0.1$ and $n = 13$. In the left panel, the number of predecessors (on the x -axis) is at most 6, and the minority size can be either 0, 1, or 2, as indicated by the labels next to the three lines that represent the probabilities the minority is correct for these cases. A minority of size 0 means that predecessors were unanimous in which case the minority choice is more likely to be incorrect, as indicated by the decreasing line that is everywhere below 0.5. The first two points of the line labeled “1” can be computed from (1) and (2) above for $q = 0.2$: note that $P_{1,m}$ for $m = 2, 3, 4, 5$ is everywhere above 0.5, i.e. the minority is more likely to be correct, and is increasing in m , i.e. the minority wisdom becomes stronger as the majority grows. The same is true for $P_{2,m}$ for $m = 3, 4$ as shown by the line labeled “2.” The right panel establishes the same properties for our second treatment in which $q = 0.1$ and $n = 13$.

3. Experimental Design

The experiments were conducted in the Social Science Experimental Laboratory (SSEL) at Caltech using undergraduate and graduate students as subjects. Each subject was allowed to participate only once. For each session, we randomly determined the correct option in each of the ten periods.⁶ And for each period, we randomly determined for each subject whether she was informed or not. Since the specific sequence of informed and uninformed agents can affect efficiency and behavior, we used the same draws for both treatment sequence and treatment no-sequence. We ran ten sessions for both treatments.

The experiments were run by hand. After reading the instructions out loud, subjects had an opportunity to ask clarifying questions. We used an urn to select the subject who had to guess first. The experimenter then went to that subject’s seat and indicated on the subject’s record sheet whether that subject was informed, and, if so, what the correct option was. Then another draw without replacement determined which subject had to guess next, etc. Each time, the experimenter revealed the correct option to only informed subjects. In addition, all subjects, informed or uninformed, were told the sequence of predecessors’ choices in treatment sequence and the number of choices for both options in treatment no-sequence.

The main design parameters are the probability that an agent is informed (q) and the number of subjects participating in a session (n). We chose q and n such that it was optimal to follow the minority in treatment no-sequence in all possible situations (assuming common knowledge of rationality). For $q = 0.1$ the maximum number is $n = 13$ and for $q = 0.2$ the maximum number is $n = 7$, see Figure 1.⁷ At the end of each period, the correct option was revealed to everyone and subjects earned \$4 if they had picked the correct option (and 0 otherwise). At the end of the experiment, subjects were paid, in cash, their cumulative earnings plus a \$5 show-up fee. Average earnings for the different treatments are shown in the rightmost column of Table 1. Sessions with 13 subjects usually took about one hour while sessions with only seven subjects typically lasted about 35 minutes.

⁶To avoid a bias for one of the options, we denoted option A with a circle (\circ) and option B with a cross ($+$) since these symbols have no obvious ordering. Also, there was no practice period that might have allowed subjects to coordinate their guesses at no cost. To test whether these measures were sufficient to avoid a bias, we considered the 143 cases where both options had the same number of choices in treatment no-sequence. In 82 of these cases, subjects chose A . Using a two-tailed test and a 5% confidence level, we cannot reject the null hypothesis that the number of choices of A follows a binomial (143, 0.5) distribution.

⁷More generally, the larger is q the lower is the largest n for which following the minority is optimal (Callander and Hörner, 2009).

Treatment	Prob Informed (q)	# Subjects (n)	# Sessions	Earnings
Sequence	0.1	12-13	5	\$31.3
Sequence	0.2	7	5	\$29.8
No-sequence	0.1	11-13	5	\$27.1
No-sequence	0.2	7	5	\$26.5

Table 1: Experimental Design Parameters and Subjects' Earnings.

4. Results

We first discuss the extent to which subject behavior conforms to theoretical predictions. We pool the data from the $q = 0.1$ and $q = 0.2$ sessions since there were no significant differences between them.

4.1. Subject Behavior

Not surprisingly, informed subjects in the experiment always follow their signals, and, hence, always select the correct option. In contrast, uninformed subjects do not always choose according to their optimal strategy under common knowledge of rationality. We differentiate among the following situations:

- (S1) all predecessors chose the same option.
- (S2) one predecessor deviated from his unanimous predecessors in treatment sequence.⁸
- (S3) one option is chosen by a minority of predecessors in treatment no-sequence.

As explained above, rational agents would follow unanimous predecessors in both treatments, a deviator in treatment sequence, and the minority in treatment no-sequence. Figure 2 summarizes the behavior of uninformed subjects in treatment sequence. For each of the ten sessions we compute the fraction of situations in which subjects follow their optimal strategy under common knowledge of rationality. The box plot summarizes these ten independent data points.⁹

The leftmost box in Figure 2 summarizes behavior in treatment sequence when all predecessors chose the same option. The median frequency with which uninformed subjects follow their

⁸We are not considering situations in which two or more agents deviated from their immediate predecessor. In these situations, the assumption of common knowledge of rationality is clearly violated.

⁹The boundaries of the box correspond to the first and third quartile, the line within the box marks the median. Observations that are more than 1.5 times the interquartile range higher (lower) than the third (first) quartile are considered to be outliers and shown separately. The whiskers show the position of the lowest (highest) observation that is not an outlier.

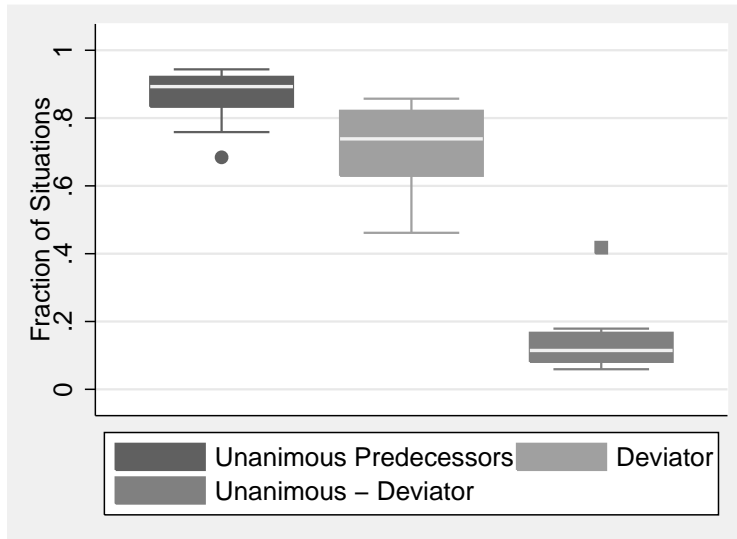


Figure 2: Subject Behavior in Treatment Sequence

unanimous predecessors (across ten sessions) is 0.89. In other words, 89% of the time, subject behavior coincides with the optimal strategy (under common knowledge of rationality) for this specific situation. The middle box concerns situations in which one predecessor deviated from his unanimous predecessors. The median frequency with which uninformed subjects follow such a deviator is 74%. The rightmost box summarizes the difference between the frequency with which subjects follow unanimous predecessors and the frequency with which they follow a deviator. In each one of ten sessions, this difference is positive. Treating each session as an independent observation, this difference is therefore significant using any conventional test. Subjects in our experiments are more likely to follow unanimous predecessors than to follow a deviator in treatment sequence.

Figure 3 summarizes the behavior of uninformed subjects in treatment no-sequence. As in Figure 2, the leftmost box summarizes behavior when subjects face unanimous predecessors. Not surprisingly, the results are very similar to those reported for treatment sequence (since there is no reason why subjects should behave differently in the two treatments when all predecessors chose the same option). The median frequency with which subjects follow unanimous predecessors is now 86%. The middle box concerns situations in which there is a minority choice among predecessors. Obviously, behavior deviates very strongly from the optimal strategy under common knowledge of rationality: subjects only follow the minority about 29% of the time (median over all ten sessions). In fact, in all but one session, subjects follow the majority *more* frequently than the minority. Finally, for each of the ten sessions of treatment no-sequence, we compute the difference between the frequency with which subjects

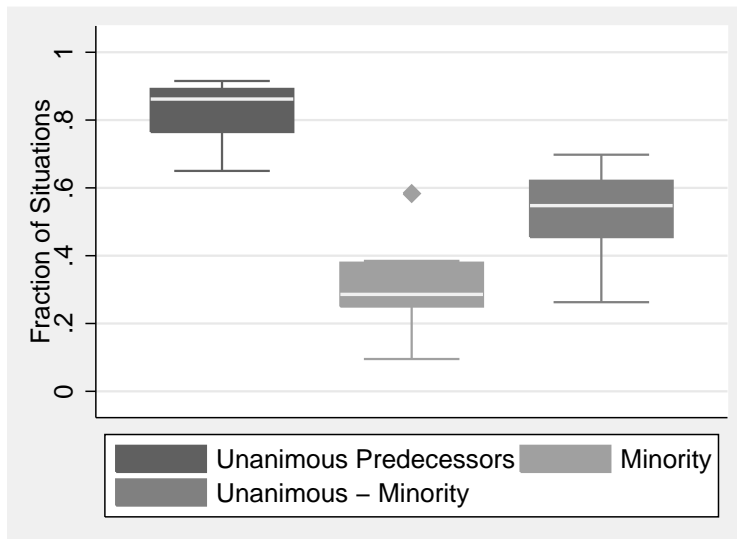


Figure 3: Subject Behavior in Treatment No-Sequence

follow unanimous predecessors and the frequency with which they follow the minority. The results are captured by the rightmost box in Figure 3. Subjects in our experiments are clearly more likely to follow unanimous predecessors than to follow the minority.

Summing up our findings for the three situations:

- (S1) While subjects do not always follow their unanimous predecessors, they do so more than 50% of the time in all 20 sessions.¹⁰ Furthermore, the tendency of subjects to follow their unanimous predecessors becomes stronger as the number of unanimous predecessors grows. Figure 4 shows the frequency with which unanimous predecessors are followed when the number of predecessors varies from 1 to 12. A simple Probit regression in which the dependent variable is the choice to follow unanimous predecessors, P_{follow} , and the independent variable is the number of predecessors, n , results in significant estimates: $P_{\text{follow}} = \Phi(\beta_0 + \beta_1 \cdot n)$ yields $\beta_0 = 0.7(0.1)$ and $\beta_1 = 0.12(0.04)$, where the numbers in parentheses denote the robust standard errors (clustered by subject).
- (S2) In nine out of ten sessions, subjects follow a deviator in treatment sequence more than 50% of the time.¹¹ Furthermore, the tendency to follow a deviator is stronger if the choice of the deviator coincides with the majority choice. Figure 5 shows the frequency with which a deviator is followed when the deviator's choice is in the minority (left circle), when there

¹⁰We can easily reject the null hypothesis that subjects follow their unanimous predecessors in 50% of all instances using a Wilcoxon signed-rank test ($p < 0.001$).

¹¹The null that subjects follow a deviator half of the time in treatment sequence can be rejected using a Wilcoxon signed-rank test ($p = 0.004$).




Figure 4: Frequency with which unanimous predecessors are followed as a function of the number of unanimous predecessors. The size and color of each circle reflects the number of occurrences.

is tie (middle circle), and when the deviator’s choice is in the majority (right circle). To verify whether the observed differences are statistically significant, we estimate a simple Probit model in which the chance of following a deviator, P_{follow} , is explained by whether there is minority, tie, or majority: $P_{\text{follow}} = \Phi(\beta_{\text{min}} + \beta_{\text{tie}} + \beta_{\text{maj}})$ yields $\beta_{\text{min}} = 0.20(0.23)$, $\beta_{\text{tie}} = 0.34(0.31)$, and $\beta_{\text{maj}} = 0.65(0.29)$, where the numbers in parentheses denote the robust standard errors (clustered by subject). The difference between β_{maj} and β_{min} is significant at the 5% level.

(S3) In all but one session of treatment no-sequence, subjects follow the majority more frequently than the minority.¹² Furthermore, subjects are increasingly *less* inclined to follow the minority, the smaller the minority is relative to the majority. (Recall that under common knowledge of rationality, the wisdom of the minority becomes *stronger* the smaller the minority is relative to the majority, see Figure 1.) Figure 6 shows the frequency with which subjects follow the minority as a function of the difference between the size of the majority and the size of the minority. A Probit analysis in which the chance of following the minority, P_{follow} , is regressed on the difference between majority and minority sizes, Δn , results in significant estimates: $P_{\text{follow}} = \Phi(\beta_0 + \beta_1 \cdot \Delta n)$ yields $\beta_0 = -0.26(0.14)$ and $\beta_1 = -0.098(0.039)$, where the numbers in parentheses denote the robust standard errors (clustered by subject).

¹²The null that subjects are equally likely to follow the minority or the majority in treatment no-sequence is rejected using a Wilcoxon signed-rank test ($p = 0.004$).

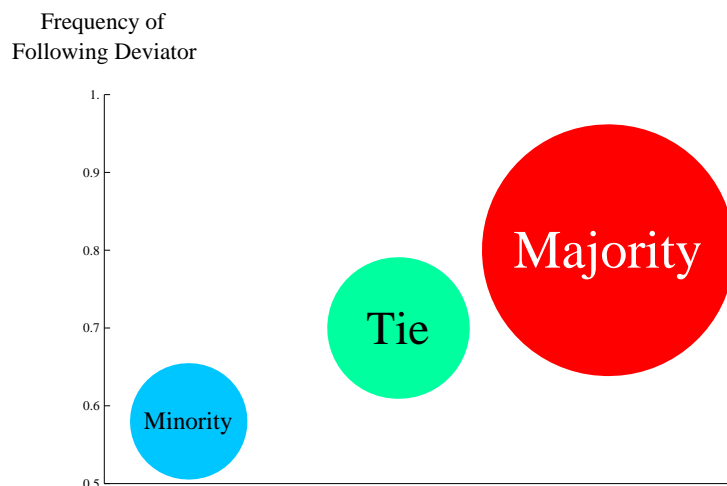


Figure 5: Frequency with which a deviator’s choice is followed when it is in the minority, creates a tie, or is in the majority. The size and color of each circle reflects the number of occurrences.

The observed choice frequencies across the three situations can be ranked as follows:¹³

$$P(\text{follow unanimous predecessors}) > P(\text{follow deviator}) > \frac{1}{2} > P(\text{follow minority}) \quad (3)$$

In other words, subjects are more likely to follow unanimous predecessors than a deviator, are more likely to follow unanimous predecessors than the minority, and are more likely to follow deviators than the minority. It is interesting to compare the ranking of choice frequencies with that of the associated average payoff differences, computed using the actual signals and choices of the experiment:

$$\Delta\pi(\text{unanimous predecessors}) > \Delta\pi(\text{deviator}) > 0 > \Delta\pi(\text{minority})$$

where $\Delta\pi$ measures the payoff difference between following and not following for each of three situations. To summarize, given the observed play of others, not following a deviator is a costly mistake (\$1.23), but it is less of a mistake than not following unanimous predecessors (\$1.53). Furthermore, not following the minority is *not* a mistake, i.e. in the no-sequence sessions, following the majority yields higher earnings (\$0.36).

¹³Using a Wilcoxon matched-pairs signed-rank test, we can reject the null that subjects follow unanimous predecessors equally frequently as a deviator in treatment sequence ($p = 0.002$). Likewise, we can reject the null that subjects are equally likely to follow unanimous predecessors as the minority in treatment no-sequence ($p = 0.002$) and the null that subjects follow deviators equally frequently as the minority ($p < 0.001$).




Figure 6: Frequency with which the minority is followed as a function of the difference between majority and minority size. The size and color of each circle reflects the number of occurrences.

4.2. Efficiency

For the specific draws used in the experiment, the theoretically expected fraction of correct choices is 78% in treatment sequence and 72% in treatment no-sequence.¹⁴ Actual efficiency was 69% in treatment sequence and 62% in treatment no-sequence. Finally, the predicted fraction of correct choices when uninformed agents choose randomly and informed agents choose correctly is 58%.

Treating each session as an independent observation ($n = 10$), we test the null hypothesis that the distribution of the fraction of correct choices is identical under both treatments. A Wilcoxon matched-pairs signed-rank test allows us to reject the null ($p = 0.039$). Therefore, efficiency in treatment sequence is significantly higher than in treatment no-sequence. We also test whether observed efficiency is significantly higher than when uninformed subjects choose randomly and informed subjects choose correctly. Using the same test, we can reject the null for treatment sequence ($p = 0.006$) but not for treatment no-sequence ($p = 0.432$). This result suggests that information about predecessors’ choices improves efficiency only when the entire sequence of decisions is known – subjects are not able to improve their decisions when observing only the number of prior choices for each option.

¹⁴We computed the fraction of correct choices separately for each session and then took the average across sessions.

5. Explaining Subject Behavior

In this section, we analyze the observed deviations through the lens of alternative models of bounded rationality, including logit-QRE as well as two models in which agents exhibit different levels of strategic thinking (level- k and Cognitive Hierarchy). We find that a “noisy introspection” model (Goeree and Holt, 2004) that combines heterogeneity in strategic thinking with error-prone behavior fits our data best.

5.1. Logit-QRE

Consider again the example discussed in Section 2, where two predecessors chose option A and one predecessor chose option B . Whether the minority or majority is more likely to be correct depends on whether the sequence {uninformed, uninformed, informed}, for which the minority is correct, is more likely to occur than the sequence {uninformed, informed, informed}, for which the majority is correct. For $q < 0.5$, subjects should follow the minority under common knowledge of rationality. If subjects tremble, however, then the minority is not necessarily correct even for sequence {uninformed, uninformed, informed} since the second uninformed subject may not have followed the first one (such a tremble is not unlikely since the associated payoff loss is small); following the majority may be better as a result.

The above intuition can be formalized using a logit-QRE model in which subjects’ choice probabilities are positively but not perfectly related to expected payoffs, i.e. subjects are “better responders” not necessarily “best responders.” For subject t , let H_t denote the profile of predecessors’ choices and s_t the subject’s signal. Then the probability that subject t chooses c_t is

$$P(c_t|H_t, s_t) = \frac{1}{1 + \exp(\lambda\pi(1 - 2P(\omega = c_t|H_t, s_t)))} \quad (4)$$

where λ is a “rationality” parameter that determines how sensitive choice probabilities are with respect to expected payoffs and $\pi = \$4$ is the payoff of picking the correct option. Note that, since $\lambda\pi > 0$, agents more often choose the option that is more likely to be correct.

The curves in Figure 7 show the predicted probability of following unanimous predecessors, a deviator, or the minority respectively for different values of λ . (To generate this figure we averaged predicted logit choice probabilities over all occurrences of the particular situation in the experiment.) For low levels of λ , logit-QRE predicts that uninformed subjects are more likely to follow the majority than to follow the minority. The predicted probability of following a deviator is higher than 0.5 for all λ , and the same is true for the predicted probability of following unanimous predecessors.

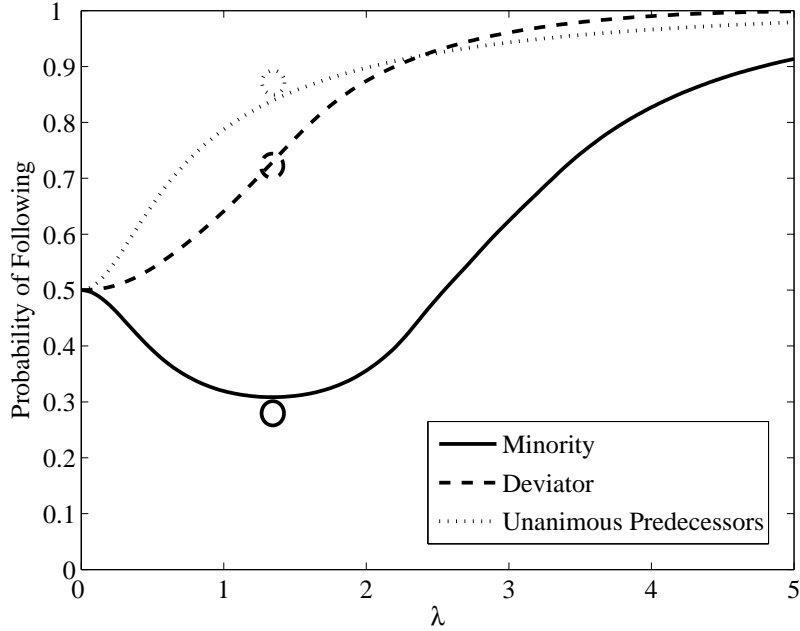


Figure 7: Uninformed Subjects’ Behavior as a Function of λ in a Logit QRE Model for Situations that Occurred in the Experiment.

Using maximum-likelihood techniques (treating each decision as an independent observation) we estimate $\lambda = 1.4$ (0.1) for treatment no-sequence, $\lambda = 1.3$ (0.1) for treatment sequence, and $\lambda = 1.3$ (0.1) for the pooled data, where the number in parentheses denotes the robust standard error (clustered by subject). The estimates are listed in the topmost panel of Table 2 below (second column labeled “ λ ”), together with the associated likelihood (third column labeled “ LogL_{obs} ”), the likelihood that results if all uninformed subjects choose randomly (fourth column labeled “ $\text{LogL}_{\text{random}}$ ”), and the best possible likelihood (fifth column labeled “ $\text{LogL}_{\text{best}}$ ”) that results by using the observed fraction for either option as the predicted fraction (separate for each different situation that occurred in the experiment).¹⁵ The rightmost column in Table 2 labeled “% Explained” provides a “goodness-of-fit” defined as follows:

$$\% \text{ Explained} = \frac{\text{LogL}_{\text{obs}} - \text{LogL}_{\text{random}}}{\text{LogL}_{\text{best}} - \text{LogL}_{\text{random}}} \times 100\%$$

The circles in Figure 7 show the average frequency with which subjects in the experiment follow unanimous predecessors, a deviator, or the minority respectively. The data averages are shown

¹⁵We treat two situations as different iff at least one of the models we estimate makes a different prediction for the two situations. More specifically, (1) we distinguish between informed and uninformed agents, but all situations informed agents encounter are considered equivalent, (2) the situation where n (m) predecessors have chosen A (B) is considered equivalent to the situation where m (n) predecessors have chosen A (B), and (3) situations with an equal number of choices for both options are equivalent.




Figure 8: Fraction of Rational Choices in Sequence (left) and No-sequence (right)

at the pooled estimate $\lambda = 1.3$ to facilitate comparison with logit predictions. Obviously, Logit-QRE can explain the qualitative patterns of our data as summarized in (3) extremely well.

5.2. Subject Heterogeneity

The homogeneous logit-QRE model assumes that agents are *ex ante* symmetric, i.e. they are equally likely to make mistakes. To test whether this is a reasonable assumption, we compute the fraction of rational choices for each subject by treatment.¹⁶ If all subjects were equally rational, this fraction would not vary much across subjects.

Figure 8 illustrates the distribution of rational choices separately for treatments sequence (left) and no-sequence (right). Clearly, there is substantial subject heterogeneity. While the standard deviation of the fraction of rational choices is very similar in both treatments (0.20 in sequence and 0.21 in no-sequence), the shape of the distribution is quite different. In treatment sequence, a fairly high proportion of subjects make choices that are always consistent with their optimal strategy (under common knowledge of rationality), but this is not the case in treatment no-sequence.

5.3. Level- k and Cognitive Hierarchy

To account for the observed heterogeneity, we next consider a “level- k model” (e.g. Crawford and Iriberri, 2007a,b) that allows for different levels of strategic thinking. We first describe

¹⁶We only consider situations in which uninformed subjects face a rational and an irrational choice. Hence, we drop choices made by the first subject in the sequence. We also drop choices made by subjects who face the same number of predecessors’ choices for both options. In treatment sequence, we drop (“off-the-equilibrium-path”) choices that occur after two or more predecessors deviated from their immediate predecessor’s choice.

how different types are defined in this model and then discuss the model’s predictions and the estimation of the type distribution.

In the level- k model, type k best responds believing others are of type $k - 1$. We assume that type 0 chooses randomly when uninformed and follows her signal when informed.¹⁷ In treatment sequence, an uninformed type 1 then follows the majority of predecessors – if all others are of type 0, the majority is more likely to be correct since type 0 picks the correct option when informed. Uninformed type 2 follows unanimous predecessors, but since type 1 deviates from the majority only when informed, uninformed type 2 follows such deviations. Note that it is possible that type 2 observes a sequence that is inconsistent with her beliefs¹⁸ in which case type 2 is assumed to randomly pick an option when uninformed and to follow her signal when informed. Uninformed type 3 behaves the same as type 2 except that there are more situations in which type 3’s beliefs are contradicted, and, hence, type 3 randomizes more often.¹⁹ All higher types (4 and above) are identical to type 3.

In treatment no-sequence, uninformed type 1 follows the majority of predecessors. Uninformed type 2 follows unanimous predecessors, and since type 1 only deviates from the majority when informed, uninformed type 2 follows the minority when predecessors are not unanimous. Callander and Hörner (2009) show that it is optimal to follow the minority when all others do, so uninformed types 3 and higher are identical to type 2. In treatment no-sequence, there are no situations in which beliefs about lower-types’ behavior are contradicted.

To facilitate comparison with the one-parameter logit-QRE model, we assume that types follow a Poisson distribution with parameter τ , which we truncate at the highest type that can be distinguished (2 in treatment no-sequence and 3 in treatment sequence).²⁰ Figure 9 summarizes the predictions of the level- k model for the three situations of interest for different levels of τ . For low τ , the level- k model is able to reproduce the main features of our data as summarized in (3) although the fit is not as good as that of logit-QRE (cf. Figure 7).

¹⁷Under the assumption that type 0 chooses randomly even when informed, type 1 would choose randomly when uninformed but would follow his signal when informed. Therefore, type 1 would correspond to type 0 under our assumptions. The only change to the model would be an introduction of an additional type that can only rarely be distinguished from type 1 (since informed agents are rare).

¹⁸For instance, if the majority choice switches from A to B after a deviation, then a next deviation (back to A) would violate the assumption that all other agents are of type 1.

¹⁹Observing two or more deviations from the immediate predecessors’ choice violates the assumption that all agents are of type 2, but if these deviations occur in favor of the majority’s choice then they are compatible with the assumption that all other agents are of type 1. Therefore, the set of situations that violate the assumptions made by type 2 is a strict subset of the set of situations that violate the assumptions made by type 3.

²⁰We also estimated the type distribution non-parametrically, see footnote 22.

We estimate τ separately for treatment sequence and treatment no-sequence by maximizing

$$L(\tau) = \prod_{t=1}^n \sum_{k=0}^m \frac{e^{-\tau} \tau^k}{k!} \text{Prob}(c_t | \text{type} = k) \quad (5)$$

where n is the number of subjects, m the number of different types (2 in treatment no-sequence and 3 in treatment sequence), and c_t the ten choices made by subject t , i.e. we assume that a subject's type is the same for all ten periods.

The estimation results for the level- k model are shown in the second panel of Table 2, and the third panel provides results for the closely related Cognitive Hierarchy model (Camerer, Ho, and Chong, 2004).²¹ Note that level- k and Cognitive Hierarchy fit about the same and substantially worse than logit-QRE.²²

The worse fit could have been expected from Figure 9, which shows that observed choice frequencies for the three situations of interest (represented by the circles) cannot be matched. More importantly, because of the best-response assumption underlying level- k , it cannot reproduce the intuitive data patterns of Figures 4–6. For example, in level- k , types 1 and 2 follow unanimous predecessors and they do so irrespective of whether the number of predecessors is 1 or 12. Clearly, this prediction is refuted by the data, see Figure 4 (and a similar argument applies to Figures 5 and 6).

The data patterns of Figures 4–6 are consistent with models in which choice probabilities vary continuously with expected payoff differences, i.e. when best responses are replaced by (logit) better responses. Introducing trembles has the additional benefit that the type distribution can be determined more robustly. Compared to previous experiments, we find a high fraction of type 0: 70% in treatment no-sequence and 63% in treatment sequence. The reason is that a subject is classified as type 0 even if her behavior is incompatible with type 1 or 2 in only a single period. The model presented in the next section avoids this problem.

²¹Cognitive Hierarchy differs from level- k only in that type k believes that others have types lower than k , not necessarily that they are all of type $k - 1$. Choice probabilities in Cognitive Hierarchy do not vary continuously with τ since a change in τ not only alters the type distribution but it can also shift the best response (since expected payoff calculations depend on the type-distribution). The likelihood function for the Cognitive Hierarchy model is not differentiable, and we introduce very small logit trembles (using $\lambda = 20$) to facilitate estimation. As for the level- k model, we assume type 0 picks randomly when uninformed and chooses the correct option when informed.

²² This remains true when we estimate the type distribution non-parametrically. In treatment sequence, the estimated fractions of types (0, 1, 2, 3) are (0.68, 0.17, 0.09, 0.05) and the associated loglikelihood is -512 . In treatment no-sequence, the estimated fractions of types (0, 1, 2) are (0.70, 0.26, 0.04) and the associated loglikelihood is -504 . In other words, allowing for an arbitrary type distribution hardly improves the fit compared to the assumed Poisson distribution of types.

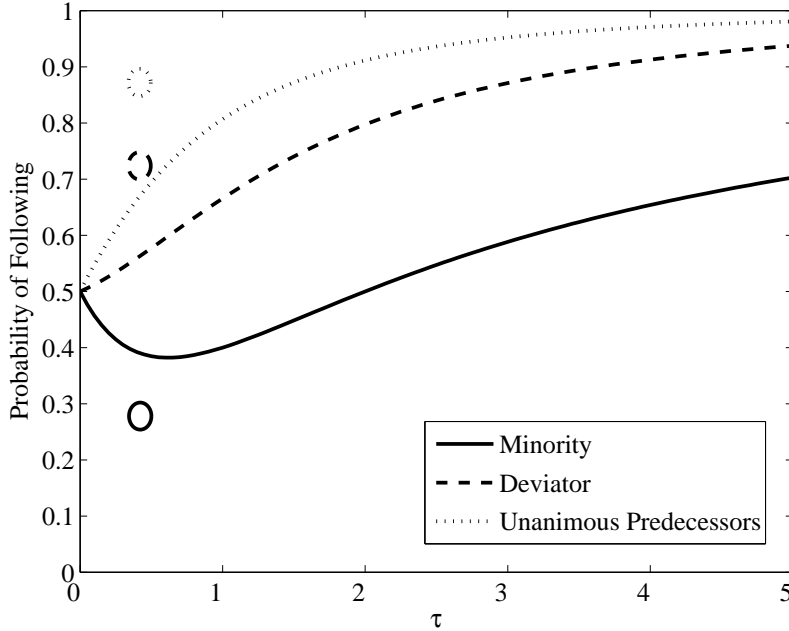


Figure 9: Uninformed Subjects’ Behavior as a Function of τ in a level- k Model for Situations that Occurred in the Experiment.

5.4. Noisy Introspection

The noisy introspection model proposed by Goeree and Holt (2004) combines subject heterogeneity and error-prone behavior by assuming that (i) subjects differ in levels of strategic thinking and (ii) subjects realize others are “better responders” and not necessarily “best responders.” Let $\phi_\lambda(\cdot)$ denote the logit response function, see (4), then type k ’s choice probabilities can be recursively defined as $P_k = \phi_\lambda \circ P_{k-1}$, i.e. type k better responds believing that observed choices are generated by better responders of type $k - 1$. Equivalently:²³

$$P_k = \underbrace{\phi_\lambda \circ \cdots \circ \phi_\lambda}_{k \text{ times}} \circ \phi_0 \quad (6)$$

So type k better responds to type $k - 1$, who better responds to type $k - 2$, \dots , and type 1 better responds to type 0 who chooses randomly (since ϕ_0 results in uniform choice probabilities). The noisy introspection model in (6) reduces to the level- k model of the previous section when $\lambda = \infty$, i.e. when logit better responses are replaced by standard best responses.

For finite levels of λ , predicted behavior is different for all types (unlike in the model

²³Goeree and Holt (2004) define choice probabilities for the noisy introspection model by considering an infinite sequence of logit responses: $P = \lim_{n \rightarrow \infty} \phi_{\lambda_n} \circ \cdots \circ \phi_1 \circ \phi_0$, where $\lambda_1 \leq \cdots \leq \lambda_n$. Note that this is equivalent to (6) when $\lambda_n = \cdots = \lambda_{n-k+1} = \lambda$ and $\lambda_{n-k} = \cdots = \lambda_1 = 0$.

	τ	λ	LogL _{Obs}	LogL _{Random}	LogL _{Best}	% Explained
Logit-QRE						
Sequence		1.3 (0.1)	-417	-686	-364	83.5%
No-sequence		1.4 (0.1)	-438	-679	-407	88.6%
Pooled		1.3 (0.1)	-856	-1366	-771	85.7%
Level-k						
Sequence	0.5 (0.1)		-516	-686	-364	52.8%
No-sequence	0.4 (0.1)		-504	-679	-407	64.3%
Pooled	0.4 (0.1)		-1021	-1366	-771	58.0%
Cognitive Hierarchy						
Sequence	1.2 (0.1)		-489	-686	-364	61.2%
No-sequence	0.6 (0.1)		-504	-679	-407	64.3%
Pooled	1.0 (0.1)		-1003	-1366	-771	61.0%
Noisy Level-k						
Sequence	1.9 (0.3)	2.0 (0.2)	-402	-686	-364	88.2%
No-sequence	1.0 (0.1)	1.6 (0.2)	-457	-679	-407	81.6%
Pooled	1.3 (0.1)	1.7 (0.1)	-869	-1366	-771	83.5%
Noisy Introspection						
Sequence	2.3 (0.3)	1.8 (0.2)	-385	-686	-364	93.5%
No-sequence	2.1 (0.4)	1.7 (0.1)	-429	-679	-407	91.9%
Pooled	2.2 (0.3)	1.8 (0.1)	-815	-1366	-771	92.6%

Table 2: Overview of Different Models, Robust Standard Errors are in Parentheses.

without trembles). To keep the model parsimonious and comparable to level- k , we again assume that the types follow a Poisson distribution. The estimated parameter values for τ and λ are shown in the bottom panel of Table 2. They are highly significant, and are higher than for the corresponding estimates in models that contain either λ or τ . Hence, compared to the level- k model, the introduction of trembles leads to a right-shift of the type distribution because subjects whose choices are almost always compatible with a type higher than 0 are now classified as such. Likewise, introducing types into logit-QRE increases the estimated rationality parameter because some of the randomness in observed choices is accounted for by the presence of type 0. Note that noisy introspection provides a significant improvement in fit relative to logit-QRE and a dramatic improvement relative to level- k and Cognitive Hierarchy.

Previous papers that have allowed for trembles within the level- k framework typically have assumed that subjects are unaware that others tremble (i.e. only the econometrician assumes trembles). In other words, subjects compute their expected payoffs under the assumption that others are best responders (even though they are better responders themselves). To gauge the importance of the “common-knowledge-of-trembles” assumption, we also estimate a noisy level- k model that follows by replacing others’ logit responses by best responses:

$$\tilde{P}_k = \phi_\lambda \circ \underbrace{\phi_\infty \circ \dots \circ \phi_\infty}_{k-1 \text{ times}} \circ \phi_0 \quad (7)$$

Note that in this model, there can be situations in which beliefs are inconsistent with observed

play. Also, there are now 4 (instead of 3) distinct types in treatment sequence²⁴ (the number of distinct types does not change for treatment no-sequence).

The estimation results in Table 2 demonstrate that this model fits slightly worse than logit-QRE, much better than level- k and Cognitive Hierarchy, but significantly worse than noisy introspection.

6. Conclusion

The “wisdom of the crowds” typically refers to the observation that a group may produce better decisions than any of its members could have. The common explanation is that groups aggregate diverse opinions and preferences, yielding more accurate information or more widely acceptable policies and rules (e.g. Surowiecki, 2004). The wisdom of the crowds is a central outcome of most social learning models in which imperfectly informed agents infer valuable information from predecessors’ choices. In the canonical social learning model developed by Bikhchandani et al. (1992), for instance, herding occurs frequently and almost immediately after a few decisions have been observed.

In a clever and novel contribution, Callander and Hörner (2009) discuss settings where the *minority* is predicted to be correct. Necessary conditions for this prediction to hold are that agents are differentially informed and only observe the number of times each option is chosen (and not the exact sequence of prior choices as in Bikhchandani et al., 1992). In this paper, we test the Callander–Hörner model in a controlled laboratory setting. We employ a simplified version of their model where each agent is either perfectly informed (with small probability) or not informed at all (with complementary high probability). We report data from two treatments: in treatment “sequence,” subjects could observe the entire sequence of predecessors’ choices, while in treatment “no-sequence” they could see only how many times either option had been chosen. For our setup, the predictions are that subjects follow their immediate predecessors in treatment sequence and follow the minority in treatment no-sequence.

In a nutshell, the data may be characterized as follows: subjects tend to overwhelmingly follow unanimous predecessors in both treatments (87%), they more frequently than not follow a deviator in treatment sequence (72%), and they *don’t* follow the minority in treatment sequence

²⁴Like type 2, type 3 follows the minority. While type 2 assumes that everybody else follows the majority, type 3 assumes all other agents follow the minority. Therefore, type 2 and type 3 arrive at different expected probabilities that option A is correct and as a consequence also exhibit different trembles. Like type 3, type 4 assumes that everybody follows the minority. Since agents in the simple noisy level- k model are unaware of the fact that other agents tremble, type 4 is identical to type 3.

(28%). The observed deviations from theoretical predictions are approximate best responses to the empirical distribution of play. Given the choices of others, and given the signals used in the experiment, not following unanimous predecessors is very costly (\$1.53 on average), not following a deviator in treatment sequence is less costly (\$1.23 on average), and not following the minority *pays* (\$0.36 on average).

We analyze the deviations observed in our data using alternative models of bounded rationality. In the logit Quantal Response Equilibrium (QRE), for instance, agents are predicted to tremble, which can overturn the logic for why the minority is correct. Intuitively, if the probability of being informed is very low then a “deviant minority” is more likely the result of a tremble that should rationally be ignored. When applying the logit-QRE to our data we find that it is able to reproduce the main (aggregate) features. Zooming in on individual level data, however, reveals a substantial amount of heterogeneity that cannot be explained by the symmetric logit-QRE.

Heterogeneity in levels of strategic thinking is more naturally explained by the Cognitive Hierarchy model (Camerer, Ho, and Chong, 2004) or the closely related level- k model (e.g. Crawford and Iriberri, 2007). The level- k model, for example, assumes that agents of type k best respond to their beliefs that others are of type $k - 1$. We apply both models to our data and find they are also able to reproduce the main aggregate features, although not as well as logit-QRE. Furthermore, their fit of the individual data is substantially worse. The main reason for these shortcomings is the underlying best-response assumption, which conflicts with several intuitive comparative statics observed in our data (e.g. subjects tend to follow unanimous predecessors more frequently when the group of predecessors is large). In addition, the best-response assumption skews the estimated type distribution towards lower types.

The noisy introspection model proposed by Goeree and Holt (2004) combines heterogeneity in strategic thinking with noisy behavior by replacing level- k 's best response with a logit “better response.” Importantly, agents in the noisy introspection model know that others tremble. For example, when computing the probability that the minority choice is correct in treatment no-sequence, agents take into account the possibility that the minority arose because of trembles. As a result, the model can predict why subsequent choices favor the majority even for agents with high levels of strategic thinking. We illustrate the importance of the “common-knowledge-of-trembling” assumption by also estimating a noisy version of the level- k model in which agents tremble but assume others don't (as is typically done in the literature) and show that it fits significantly worse.

The noisy introspection model fits our data significantly better than any of the other models. It naturally captures heterogeneity in strategic sophistication as well as the endogenous effects

of noisy behavior. Without these endogenous effects, others' trembles are easily mistaken for informed behavior and sophisticated agents would have to conclude that the minority is wise. In contrast, realizing that others can be error prone, sophisticated agents in the noisy introspection model benefit from following the wisdom of the crowds.

References

- Anderson, L. R. and Holt, C. A. (1997) "Information Cascades in the Laboratory," *American Economic Review*, 87, 847-62.
- Banerjee, A. V. (1992) "A Simple Model of Herd Behavior," *Quarterly Journal of Economics*, 107 (3), 797-817.
- Bikhchandani, S., Hirschleifer, D. and Welch, I. (1992) "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades," *Journal of Political Economy*, 100(5), 992-1026.
- Bose, S., Orosel, G., Ottaviani, M. and Vesterlund, L. (2009) "Monopoly Pricing in the Binary Herding Model," *Economic Theory*, forthcoming.
- Callander, S. and Hörner, J. (2009) "The Wisdom of the Minority," *Journal of Economic Theory*, forthcoming.
- Camerer, C. F., T.-H. Ho, and J.-K. Hong (2004) "A Cognitive Hierarchy Model of Games," *Quarterly Journal of Economics*, 119(3), 861-898.
- Çelen, B. and Kariv, S. (2004a) "Distinguishing Informational Cascades from Herd Behavior in the Laboratory," *American Economic Review*, 94(3).
- Çelen, B. and Kariv, S. (2004b) "Observational Learning Under Imperfect Information," *Games and Economic Behavior*, 47(1), 72-86.
- Cipriani, M. and Guarino, A. (2005) "Herd Behavior in a Laboratory Financial Market," *American Economic Review*, 95(5), 1427-1443.
- Crawford, V. P. and N. Iriberrri (2007a) "Fatal Attraction: Salience, Naivete, and Sophistication in Experimental Hide-and-Seek Games," *American Economic Review*, 97, 1731-1750.
- Crawford, V. P. and N. Iriberrri (2007b) "Level- k Auctions: Can a Non-Equilibrium Model of Strategic Thinking Explain the Winner's Curse and Overbidding in Private-Value Auctions?" *Econometrica*, 75, 1721-1770.
- Drehmann, M. Oechssler, J. and Roeder, A. (2005a) "Herding and Contrarian Behavior in Financial Markets - An Internet Experiment," *American Economic Review*, 95(5), 1403-1426.
- Drehmann, M. Oechssler, J. and Roeder, A. (2005b) "Herding with and without Payoff Externalities - An Internet Experiment," University of Heidelberg Discussion Paper No. 420.

- Goeree, J. K. and C. A. Holt (2004) "A Model of Noisy Introspection," *Games and Economic Behavior*, 46(2), 2004, 365-382.
- Goeree, J. K., T. R. Palfrey, T., and B. Rogers (2006) "Social Learning with Private and Common Values," *Economic Theory*, 28(2), 245-264.
- Goeree, J. K., Palfrey, T., Rogers, B. and McKelvey, R. (2007) "Self-Correcting Information Cascades," *Review of Economic Studies*, 74(3), 733-762.
- Hung, A. A. and Plott, C. R. (2000) "Information Cascades: Replication and an Extension to Majority Rule and Conformity Rewarding Institutions," *American Economic Review*, 2001, 91(5), 1508-1520.
- Kübler, D. and Weizsäcker, G. (2004) "Limited Depth of Reasoning and Failure of Cascade Formation in the Laboratory," *Review of Economic Studies*, 71, 425-441.
- Ottaviani, M. and Sørensen, P. (2001) "Information Aggregation in Debate: Who Should Speak First?" *Journal of Public Economics*, 81, 393-421.
- Smith, L. and Sørensen, P. N. (2000) "Pathological Outcomes of Observational Learning," *Econometrica*, 68(2), 371-398.
- Surowiecki, J. (2004) *The Wisdom of Crowds: Why the Many Are Smarter Than the Few and How Collective Wisdom Shapes Business, Economies, Societies and Nations*, Little, Brown.