University of Mannheim / Department of Economics

Working Paper Series

Herding differently: A level-k model of social learning

Stefan P. Penczynski

Working Paper 13-01

Januar 2013

Herding differently: A level-k model of social learning

Stefan P. Penczynski*
University of Mannheim

January 2013

This paper proposes a behavioral model of social learning that unifies various forms of inferential reasoning in one hierarchy of types. Iterated best responses that are based on uninformative level-0 play lead to the following of the private information (level-1), to the following of the majority (level-2), to a differentiated view on predecessors (level-3), etc. I present evidence from three sources that these are the prevalent types of reasoning in social learning: a review of social learning studies, existing data from Çelen and Kariv (2004) as well as new experimental data that includes written accounts of reasoning from incentivized intra-team communication.

Keywords: Social learning, levels of reasoning

JEL Classification: C91, D82, D83, D84

*Department of Economics, University of Mannheim, L7, 3-5, 68131 Mannheim, Germany, stefan.penczynski@uni-mannheim.de, Tel. +49 621 181 3656, Fax. +49 621 181 1893. I thank Ayala Arad, Konrad B. Burchardi, Dirk Engelmann, David Gill, Rosemarie Nagel, Anders Poulsen, David Rojo Arjona, Robert Sugden and seminar participants at ESA Chicago 2011, University of East Anglia, SABE Granada 2012, HeiMaX Heidelberg 2012, GfeW Karlsruhe 2012, M-BEES Maastricht 2012, and University of Mannheim for helpful comments. I also thank Shachar Kariv and Boğaçhan Çelen for the provision of the data from their 2004 study. Katharina Momsen, Allison Nuñez, Dennis Shiraev, and Julia Walecki provided excellent research assistance.

How does the observation of others' choices influence a person's vote for a political candidate, a CEO's placing of a takeover bid or a doctor's choice of a medical treatment. In other words, how do people use information about others' decisions when they have to decide on a similar matter? Can the iterated belief structure of the level-k model for strategic decisions usefully describe how others' decisions are fathomed in such a *social learning* context? This study aims to illuminate these questions with the use of experimental data.

A first theoretical perspective on social learning was given by models of Bikhchandani, Hirshleifer and Welch (1992) and Banerjee (1992). In a sequence of decision-makers that want to act according to the true but uncertain state of the world, players hold private information and observe their predecessors' decisions. A striking conformity and a wasteful neglect of private information are predicted to eventually emerge in form of an information cascade.

A large experimental literature, starting with Anderson and Holt (1997, AH henceforth), set out to test the assumption of Bayesian inference within error-rate or Quantal Response Equilibrium (QRE) models (for example, Çelen and Kariv, 2004; Drehmann, Oechssler and Roider, 2005; Goeree, Palfrey, Rogers and McKelvey, 2007). Kübler and Weizsäcker (2004) estimate the depth of reasoning in an error-rate model as they allow the error to depend on the level of reasoning. The studies generally find substantial deviations from Bayesian inference, mostly in that decision-makers overweight their own information.

Other studies try to explicitly model the deviations from Bayesian inference and propose various types of players. Hung and Plott (2001) consider "truthful, private-information revealers" and "naïve Bayesians", who think others are truthful. Eyster and Rabin (2005) analyze "cursed" players who underinfer from others actions and thus play truthful. Further, they propose "Best Response Trailing Naïve Inference", a best response to cursed players (Eyster and Rabin, 2010). Dominitz and Hung (2004) consider Bayesians that deem off-equilibrium play informative, an intuition already formulated by AH and predicted by the error-rate models.

In this paper, I propose a behavioral model of social learning that unites a variety of these types in one coherent framework. Inspired by the level of reasoning model of strategic thinking (Nagel, 1995; Stahl and Wilson, 1995; Costa-Gomes, Crawford and Broseta, 2001; Camerer, Ho and Chong, 2004), my model organizes types by the number of iterated best responses to uninformative level-0 play.In the resulting hierarchy, types are linked via the population belief so that lower level types describe the model that higher types form of other decision-makers.

In particular, a level-1 player believes that all predecessors' actions stem from uninformative level-0 play; her best response is therefore to follow the private signal. Believing all others to play according to their signal, the level-2 player views all previous actions as equally informative as her own private signal. Level-3 players realize that some level-2 actions do not reflect private signals and assess them as uninformative. So does level-4 for level-3 play, etc.

The model as a whole is related to the incomplete information level-k model which Crawford and Iriberri (2007) introduce to analyze data on common value auctions. The most closely related model is by Brocas, Carillo, Wang and Camerer (2010) who provide a level-k model in the context of zero-sum betting. Both applications feature inference of private information as well as strategic interaction. In contrast, looking at social learning puts the focus entirely on inference and abstracts from strategic interaction.

Most of the experimental literature in social learning relates to the seminal paper by Anderson and Holt (1997). In their experiment, subjects observe a binary private draw that reveals the state of the world correctly with a probability of $\frac{2}{3}$. Before subjects choose a binary action to match the state, they can observe the actions of all predecessors. Due to the coarse binary structure, this setting provides only limited information about the reasoning underlying the observed decisions. For example, level-2 or higher and Bayesian reasoning are predicted to always play the same actions. For Eyster and Rabin (2010, p. 236), "the existing experimental literature is generally not well-designed to differentiate among likely hypotheses about the nature of observational learning." In my analysis, I turn to a different framework of social learning and a new experimental method in order to be able to differentiate types.

First, I reconsider various experimental studies in the AH tradition and relate the findings to features of the proposed model. Second, I analyze the rich data of a social learning experiment with continuous actions from Çelen and Kariv (2004, ÇK henceforth). Finally, I use the original AH and ÇK setups in two experiments with intra-team communication as introduced by Burchardi and Penczynski (2012).

In all instances, I find evidence in favor of the modeled heterogeneity in inferential reasoning. In all analyses on the individual level, subjects of all types are detected and the modal behavior turns out to be level-2. The arguments observed in the team communication mirror the decision rules in the proposed hierarchy of types.

I conclude that the level-k framework is a valid behavioral model of social learning which explains occurring deviations from Bayesian play and rational expectations (see Weizsäcker, 2010). The extension of the level-k concept to social learning has the potential to contribute a similarly good understanding of inferential reasoning as it achieves for strategic reasoning. Furthermore, this extension highlights the empirical importance of the cognitive processes that are modelled in the level-k literature and provides another view on how people form models of others in their minds.

The paper is structured as follows. Section 1 introduces the general model. The empirical part is started with a new look at past studies in section 2. The analysis of ÇK's data (section 3) and of the experiments with intra-team communication (section 4) follows. Section 5 discusses the results before section 6 concludes.

1. The model

After a general introduction to the model, I apply it in sections 1.2 and 1.3 to the social learning frameworks used in the studies by AH and ÇK.[1]

1.1. The general model

Consider a general social learning framework, in which players take actions sequentially and aim to match an uncertain state of the world $\omega \in \Omega$. A player t receives information about the state of the world via a private signal $s_t \in \mathcal{S}$. Furthermore, she observes the history of actions up to round t, denoted by $H_t = (a_1, a_2, \ldots, a_{t-1})$. The strategy of a level-k type at position t in the sequence is $\sigma^k(s_t, H_t)$, mapping the observed history H_t and the signal s_t to an action $a \in \mathcal{A}$ in a way that the action maximizes the expected payoff under the information s_t and H_t and the belief that others follow the strategy σ^{k-1}. I introduce the model using a degenerate population belief on level $k-1$, $p^{k,l} = 1$ for $l = k-1$, and discuss unconstrained population beliefs afterwards. In order to talk about the inference from the action to the observed signal under a given strategy, I denote the inverse image of a strategy as $\sigma^{k\leftarrow}(a_t, H_t) = \{s_t \in \mathcal{S} : \sigma^k(s_t, H_t) = a_t\}$, which maps histories and actions into sets of signals.

As is common in the level-k literature, the model postulates a hierarchy of

[1] The application to the framework of Eyster and Rabin (2010) is shown in section A of the supplementary material.

heterogeneous types. In my model, the hierarchy's starting point is the level-0 player that plays uninformatively and randomizes uniformly over the action space,

$$\sigma^0(s_t, H_t) \sim U(\mathcal{A}).$$

Smith and Sørensen (2000) model a related, uninformative type which I view as nested in my specification. Their "crazy type" always chooses one and the same action independently of signal and history. Intuitively, a level-0 player lacks either an understanding of the game or the motivation to use any information for his action.

The level-1 player assumes others to be level-0 players and best responds accordingly. Given the uninformativeness, the recovery of the private signals of preceding level-0 players is not possible, $\sigma^{0\leftarrow}(a_t, H_t) = \mathcal{S}$ is not informative. Solely the own private signal is informative to level-1 players and enters their strategy,

$$\sigma^1(s_t, H_t) = \sigma^1(s_t).$$

This player understands how her private signal is informative but either expects others to play randomly or does not understand how others' decisions can be useful. This type has been discussed among others as "Truthful, private-information revealer" (Hung and Plott, 2001), "Follow your own signal"-type (Kübler and Weizsäcker, 2004, $\lambda_1 \to \infty$, $\lambda_2 = 0$)[2], and "Fully cursed player" (Eyster and Rabin, 2005).

While level-1 players disregard the history of random actions, level-2 strategies depend on it heavily since all earlier actions are believed to be informative of the private signals, with only a strict subset of $\mathcal{S}$ leading to each action, $\sigma^{1\leftarrow}(a_t, H_t) \subsetneqq \mathcal{S}$. Call the set of predecessors whose actions are informative about the signal under a level-k strategy $\mathcal{I}^k = \{i : \sigma^{k\leftarrow}(a_i, H_i) \subsetneqq \mathcal{S} \wedge i < t\}$. It follows that $\mathcal{I}^1 = \{i : i < t\}$, the set of all predecessors. The level-2 strategy simply aggregates the own signal and the detailed information about the other signals,[3]

$$\sigma^2(s_t, H_t) \quad = \quad \sigma^2(s_t, \{\sigma^{1\leftarrow}(a_i, H_i)\}_{i\in\mathcal{I}^1}).$$

This player expects others to play according to their signal and does not see how they might be like him: influenced by their history. This type has been discussed

[2] λ_i is the precision in a logistic choice function for the ith iterated best response.

[3] With slight abuse of notation, the argument H_t is replaced by the inferred set of signals to highlight the informational input of the strategy.

as “Naive Bayesian” (Hung and Plott, 2001), “counting heuristic”-type (Kübler and Weizsäcker, 2004, $\lambda_1 = \lambda_2 \to \infty$, $\lambda_3 = 0$) and BRTNI (Eyster and Rabin, 2010, Best Response Trailing Naïve Inference).

A level-3 player can infer the private signals from observed actions if the action of level-2 players is dependent of their signal, i. e. the signal is not made irrelevant by others' actions and $\sigma^{2\leftarrow}(a_i, H_i) \subsetneqq \mathcal{S}$. The level-3 strategy is based on an inference from those predecessors' actions that are informative,

$$\sigma^3(s_t, H_t) \quad = \quad \sigma^3(s_t, \{\sigma^{2\leftarrow}(a_i, H_i)\}_{i \in \mathcal{I}^2}).$$

Level-3 players understand that others are influenced by their predecessors and take this into account when inferring private signals. They are the lowest type in the hierarchy that necessarily differentiate individual predecessors in their inference, like Bayesian players do. Unlike Bayesians, they attribute previous play to level 2, and do not expect others to be like themselves. The level-3 type is somewhat related to the Bayesians that infer from off-equilibrium play as in Dominitz and Hung (2004) and AH's (p. 850) intuition.

The level-3 strategy will lead to a set of players $\mathcal{I}^3$ with informative actions which level-4 players will exclusively use to infer private signals. Depending on the informativeness of actions and signals, higher types' strategies are more or less different from lower types and strategies can be identical to Bayesian play. The strategy of a general level-k type can be expressed as

$$\sigma^k(s_t, H_t) \quad = \quad \sigma^k(s_t, \{\sigma^{k-1\leftarrow}(a_i, H_i)\}_{i \in \mathcal{I}^{k-1}}).$$

The population beliefs in the level-k literature are specified as either degenerate on the next lower level $k-1$ (e.g. Nagel, 1995; Costa-Gomes and Crawford, 2006) or as a non-degenerate distribution on levels $0, \ldots, k-1$ (e.g. Stahl and Wilson, 1995; Camerer et al., 2004). They usually are not specified to differentiate between individual opponents because in the common applications of the level-k model players mostly face single individuals or homogeneous groups, like in the beauty contest game, normal form games, or multi-player auctions.

In social learning environments, differentiation enters naturally because players differ in their location in the sequence and thus in the information they hold. Since players know the entire public information available to their predecessors they will differentiate between them when they expect predecessors' strategies to depend on this information. It follows that level-1 and level-2 players do not differentiate predecessors because they believe them to play independently of their idiosyncratic history. However, from level-3 onwards, the population belief, $p_i^{k,l}$, depends on the position of the predecessors i because, starting with level-2, the players' strategies depend on their history H_i.

For ease of exposition, the model is so far formulated with a degenerate and homogeneous population belief. In the non-degenerate model, the strategy σ^k would simply depend on inferences from lower level players $\sigma^{l\leftarrow}$, $l = 1, \ldots, k-1$ weighted by the population belief $p_i^{k,l}$. The complete specification that I want to put forward posits an initial homogeneous population belief (without subscript i), but allows that the observed history can be used to update the beliefs in a way that differentiates predecessors. Furthermore, I posit nearly degenerate population beliefs with most of the probability mass on level $k-1$ and small probabilities on the remaining levels. In particular, I assume that $p_i^{k,k-1} = 1 - \sum_{d=1}^{k-1} \varepsilon^d$ and $p_i^{k,l} = \varepsilon^{(k-1)-l}$ for $l = 0, \ldots, k-2$, for a small and positive ε.

This specification is chosen for two reasons. First, if an observed action and history rules out that the player is of level-$k-1$, the update with Bayes' rule for this individual leads again to a nearly degenerate population belief on level $k-2$.[4] Second, believing with small probability that a predecessor plays independently of his signal and history, $p_i^{k,0} = \varepsilon^{k-1}$, works as a tie-breaker and rationalizes that players follow their private information in case they would be indifferent between two actions under $p_i^{k,0} = 0$.

1.2. Anderson and Holt (1997)

The framework used by AH has a simple binary information structure with states of the world $\omega \in \{A, B\}$, private signals $s_t \in \{A, B\}$, and actions $a_t \in \{A, B\}$. Both states of the world are equally likely. The private signal is informative with $Pr(s = \omega) = \frac{2}{3}$. The payoff of player t is positive when $a_t = \omega$ and 0 otherwise. The level-0 player randomizes over the action space $\mathcal{A} = \{A, B\}$,

$$\sigma^0(s_t, H_t) = \begin{cases} A, & \text{with Pr} = 0.5, \\ B, & \text{with Pr} = 0.5. \end{cases}$$

As in the general case, the level-1 player best-responds by simply following his signal,

$$\sigma^1(s_t, H_t) = \sigma^1(s_t) = s_t \subset \mathcal{S}.$$

[4] For example, a level-3 player observes an action that cannot come from a level-2 player. Then the updated population belief is again nearly degenerate: $p_i^{3,0} = \frac{\varepsilon}{\varepsilon+1}$ and $p_i^{3,1} = \frac{1}{\varepsilon+1}$. This example reflects the intuition that breaks in a cascade are most likely informative, i. e. they result from level-1 play.

The level-2 player thus deems previous decisions similarly informative as his own signal and aggregates the information by counting the evidence C^1 for a given state of the world,

$$C^1(\omega|s_t, H_t) = \mathbb{1}(s_t = \omega) + \sum_{i<t} \mathbb{1}(\sigma^{1\leftarrow}(a_i, H_i) = \omega).$$

The level-2 player plays according to the state with the most evidence,

$$\sigma^2(s_t, H_t) = \begin{cases} A, & \text{if } C^1(A|s_t, H_t) > C^1(B|s_t, H_t), \\ B, & \text{if } C^1(A|s_t, H_t) < C^1(B|s_t, H_t), \\ s_t, & \text{otherwise.} \end{cases}$$

Consequently, the level-3 player can infer the private signal from actions only if a different signal had caused the level-2 player to choose a different action. A level-3 player sums the evidence C^2 in favor of a given state of the world as follows

$$C^2(\omega|s_t, H_t) = \mathbb{1}(s_t = \omega) + \sum_{i\in\mathcal{I}^2} \left(\mathbb{1}(\sigma^{2\leftarrow}(a_i, H_i) = \omega)\right).$$

The level-3 strategy is then

$$\sigma^3(s_t, H_t) = \begin{cases} A, & \text{if } C^2(A|s_t, H_t) > C^2(B|s_t, H_t), \\ B, & \text{if } C^2(A|s_t, H_t) < C^2(B|s_t, H_t), \\ s_t, & \text{otherwise.} \end{cases}$$

In this environment, level-3 actions and beliefs about the state of the world are identical to Bayesians' for a given history. Strategies σ^2 and σ^3 yield the same actions, hence higher level players only differ in the population belief and not in their actions or beliefs about the state of the world.[5]

1.3. Çelen and Kariv (2004)

The setting implemented in ÇK has discrete actions $a_t \in \{A, B\}$, but continuous, uniformly distributed, bounded signals $s_t \sim U[-10, 10]$, with the payoff depending on the match of the action with the sign of the sum of all players' signals $\sum_{t=1}^{T} s_t$. For a positive sum, action A yields a positive payoff and action B nothing, and vice versa for a negative sum.

Usually, a strategy would map from the space of the observed history and the signal $\{A, B\}^{t-1} \times [-10, 10]$ into the action space $\{A, B\}$. Since the expected payoffs are monotonic in the private signal, the optimal strategies are summarized

[5] This shows why action data in this framework is not well-suited to distinguish certain forms of behavior, as noted by Eyster and Rabin (2010).

in a threshold $\theta \in [-10, 10]$. In the experiment, subjects choose a threshold before seeing their own signal. After the realization of the signal, the computer derives actions from the thresholds as follows,

$$\sigma(s_t, \theta(H_t)) \quad = \quad \begin{cases} A & \text{if } s_t \geq \theta(H_t), \\ B & \text{if } s_t < \theta(H_t). \end{cases}$$

Since the action only indicates that the signal is above or below the threshold, the inference about the signal from the action requires above all a belief about the threshold θ. Then, in conjunction with the observed action, a range of the possible signal realization can be inferred and leads to an expected value of the signal.

A rational Bayesian player uses previous actions as well as beliefs about previous cutoffs to craft her optimal cutoff strategy,[6]

$$\theta^b(H_t) \quad = \quad -E\left[\sum_{i=1}^{t-1} s_i \middle| H_t, \theta^b\right].$$

This kind of inference assumes rational expectations with respect to the thresholds θ^b used by predecessors.

A natural choice for level-0 play is a uniform randomization of the cutoffs,

$$\theta^0(H_t) \sim U[-10, 10].$$

Since the cutoff and the signal are automatically translated into an action, this would be an atypical level-0 belief: the resulting action is as informative as a cutoff of 0. As opposed to the level-0 *action*, I will therefore model the level-0 *belief* as randomizing uniformly[7] over $\{-10, 10\}$,

$$\widehat{\theta^0}(H_t) = \begin{cases} -10, & \text{with } \Pr = 0.5, \\ 10, & \text{with } \Pr = 0.5. \end{cases}$$

This level-0 belief reflects a uniform random choice of actions $a_t \in \{A, B\}$[8] and the best response for a level-1 player is to follow the private signal and therefore play the cutoff strategy

$$\theta^1(H_t) = 0.$$

[6]See equations 1 and 2 in ÇK.

[7]Any strategy is uninformative that exclusively involves the thresholds -10 and 10.

[8]Since the information content of the level-0 action is not obvious, it is potentially not uncovered by a higher level player. Then, the level-0 action has the same effect in the minds of higher level players as the modeled level-0 belief.

For a level-2 reasoner, the observation of one action resulting from such a level-1 strategy leads to an expected signal of 5 after action A and -5 after action B. It follows that $E\left[\sum_i^{t-1} s_i \,|H_t, \theta^1\right] = 5 \cdot \#A - 5 \cdot \#B$, where $\#a$ gives the number of previously played A or B actions. It can be seen that level-2 players reach the limits of the action space very quickly and play -10 or 10 once one action was played in two more occasions than the other one. Here, the level-k model predicts the information cascades in which own signals are entirely disregarded ($\theta \in \{-10, 10\}$) as analyzed in ÇK. Only five different actions are predicted to be observed from a level-2 player,

$$\theta^2(H_t) = -E\left[\sum_i^{t-1} s_i \middle| H_t, \theta^1\right] \in \{-10, -5, 0, 5, 10\}.$$

The best response to level-2 play fully discounts actions that resulted from an uninformative threshold of -10 or 10 and uses only the remaining actions to infer signals. Therefore, the level-3 uses her own signal in addition to the information from previous informative play,

$$\theta^3(H_t) = -E\left[\sum_i^{t-1} s_i \middle| H_t, \theta^2\right].$$

Due to the continuous strategy and signal space in this framework, higher level players can be differentiated. They form different beliefs about the informativeness of observed actions and therefore choose different thresholds. Denoting the strategy of a level-k player by θ^k, the strategies can be recursively expressed as

$$\theta^k(H_t) = -E\left[\sum_i^{t-1} s_i \middle| H_t, \theta^{k-1}\right].$$

2. A new look at the literature

The large amount of data available from various implementations of the AH game enables a first crude analysis of heterogeneity in terms of level-1 vs. level-2 and higher. For that purpose, using the meta-dataset of Weizsäcker (2010), I analyze decision situations in which the private signal s_t goes against the majority in (H_t, s_t). Thus, the choice is between "following the own signal" (level-1) and "following the majority" (level-2 or higher, 2+). For experiments with $Pr(\omega = A) = \frac{1}{2}$ and $Pr(s_t = \omega) = \frac{2}{3}$, the last column of Table 1 reports the fraction of players in the respective study that follow the own signal, from lowest to highest. This fraction varies between 36.9% and 63.4%.

Such heterogeneity resembles the results of "beauty contest" games where the average level-k estimated in a cognitive hierarchy model depends on the subject pool (Table II in Camerer et al., 2004). There, students from highly-ranked or technical universities or professors and graduate students achieve on average higher levels of reasoning than others. The same pattern is found here. The last column implies a ratio of level-1 to level-2+ decisions between 0.59 and 1.74. For a Poisson-distributed level of reasoning, this implies average level estimates in the range $\tau \in [0.85, 1.7]$.[9] In Camerer et al. (2004), estimated average levels have a mean of 1.3 and a median of 1.61, the obtained range is therefore a plausible one.

Source	Subject pool	N	Level-1 (%)
Goeree et al. (2007)	Caltech/UCLA	1575	36.9
Dominitz and Hung (2004)	Carnegie Mellon	622	38.1
Hung and Plott (2001)	Caltech	266	41.0
Ziegelmeyer et al. (2008)	Strasbourg	202	43.1
Nöth and Weber (2003)	Mannheim	1479	43.1
Kübler and Weizsäcker (2004)	Harvard	139	51.1
Alevy et al. (2007)	Maryland/ Professionals	168	51.8
Anderson (2001)	Virginia	81	51.9
Cipriani and Guarino (2005)	NYU, UCL	34	55.9
Oberhammer and Stiehler (2003)	Humboldt Berlin	270	56.3
Drehmann et al. (2005)	Well-educated public	95	62.1
Anderson and Holt (1997)	U of Virginia	80	62.5
Willinger and Ziegelmeyer (1998)	Strasbourg	93	63.4

Notes: Observations feature $Pr(\omega = 0) = 0.5$ and $Pr(s_t = \omega) = \frac{2}{3}$. Exceptions are Ziegelmeyer et al. with $Pr(\omega = 0) = 0.55$, Nöth and Weber with $Pr(s_t = \omega) = 0.6$, Cipriani and Guarino with $Pr(s_t = \omega) = 0.7$ and Willinger and Ziegelmeyer with $Pr(s_t = \omega) = 0.6$.

Table 1: Fraction of decisions following the own signal, not the majority

Apart from the subject pool, the specific experimental implementation seems to have an impact on the average level of reasoning as well. In particular, the two studies from Strasbourg occupy very different ranks. One major difference is that Ziegelmeyer, Koessler, Bracht and Winter (2008) elicit beliefs about true states of the world and predecessors' signals, while Willinger and Ziegelmeyer (1998) do not. It is possible that this procedure fosters a good understanding of

[9]This range is in accordance with the limited depth of reasoning found in Kübler and Weizsäcker (2004, $\lambda_3 = 0$).

the situation and allows more subjects to evaluate and make use of the public information.[10]

Overall, this table provides first suggestive evidence of a heterogeneity in inferential reasoning analogue to the heterogeneity which the level-k literature has identified in strategic settings.

Some progress in terms of differentiating between various ways of reasoning behind one of the binary actions A and B has been made by collecting further information on beliefs. Oberhammer and Stiehler (2003), Dominitz and Hung (2004) and Ziegelmeyer et al. (2008) elicit beliefs and find that the beliefs become more extreme during a cascade, a behavior that – in my model – would be predicted by level-2 play but at odds with level-1, level-3 and Bayesian play. This points to the important role of heterogeneity as established in my model and in particular to the potential relevance of level-2 play.

3. Data from Çelen and Kariv (2004)

ÇK implemented their social learning game as explained in section 1.3 in an experiment at New York University with 40 undergraduate subjects that played 15 independent rounds in sequences of 8 players. After each completed sequence, participants were informed whether their action coincided with the state of the world and were paid \$2 if it did.

I will use the 15 choices of each subject to investigate their individual type. Following Costa-Gomes and Crawford (2006), I compare the chosen cutoffs $\{\theta_r\}_{r=1}^{15}$ to the cutoffs $\{\theta_r^k\}_{r=1}^{15}$ that an ideal level-k or rational type is predicted to choose in the exact same 15 situations. I distinguish between the 4 types: level-1, level-2, level-3, and Bayesian. Figure 1 illustrates "fingerprints" of ideal types and selected subjects.

As the distance metric between fingerprints of subjects and ideal types I use the sum of squared differences, $SSD = \sum_{r=1}^{15}(\theta_r - \theta_r^k)^2$. I use this measure because calculating the expected payoff for a cutoff in a given situation is extremely complex due to the payoff function and the uniformly distributed signals. It can be observed, however, that the probability of playing the correct action is quadratic in the cutoff because the cutoff linearly influences i) the probability of playing a certain action and ii) the probability of playing the correct action given

[10] The experiments in Drehmann, Oechssler and Roider (2005) and Cipriani and Guarino (2005) are framed in a more complex financial market context. This might be a reason for the relatively low ranks despite their highly educated subject pool.

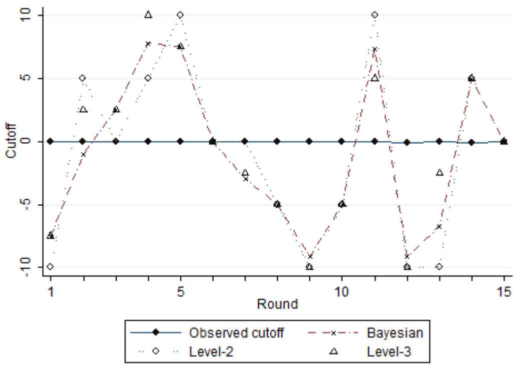

(a) Level-1 play (subject 17)

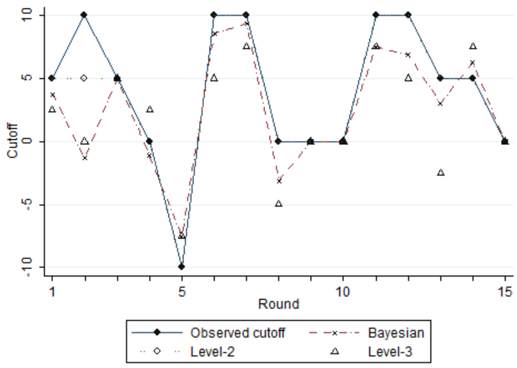

(b) Level-2 play (subject 8)

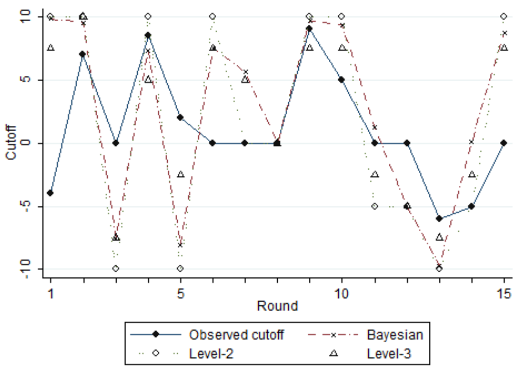

(c) Level-3 play (subject 25)

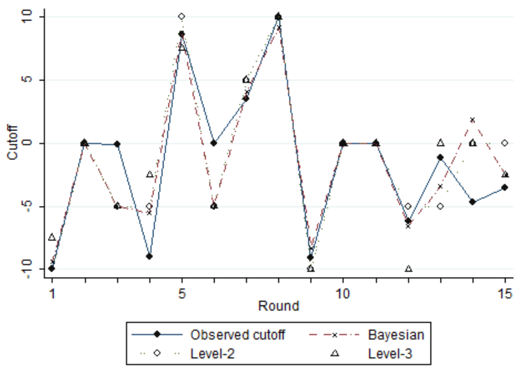

(d) Bayesian play (subject 21)

Figure 1: Subjects' and types' "fingerprints" in the data of ÇK

the action played.

Because of this complexity, it is difficult to craft an objective probability distribution of choosing the correct threshold for different ideal types. I can, however, calculate the distribution of SSDs resulting from random play and evaluate how likely it is that random play leads to smaller SSDs than the subject's SSD. With random play under the null hypothesis, this tests whether the observed play is significantly closer to the ideal type than random play. The distribution of SSDs is type-specific since, for example, random play achieves lower SSDs with respect to the level-1 strategy of always playing the midpoint of the action space ($\theta^1 = 0$) than under the level-2 strategy with frequent choices of $\theta^2 = |10|$.[11]

Using this data, I classify subjects into the type that gives the lowest significant SSD (p-value ≤ 0.05).[12] The six players that do not differ significantly from random play for any type are reported as unclassified (NA). The resulting type distribution is shown in Table 2a.

Type	Frequency	Fraction (%)
1	6	15.0
2	20	50.0
3	5	12.5
B	3	7.5
NA	6	15.0
Total	40	100.00

(a) Classification by lowest significant SSD

Type	Frequency	Fraction (%)
1	6.42	16.0
2	18.61	46.5
3	4.07	10.2
B	4.90	12.3
NA	6.00	15.0
Total	40	100.00

(b) Mixture classification

Table 2: Type overview for 40 subjects in ÇK

The table illustrates a pronounced heterogeneity of reasoning across the 40 subjects, with a distribution that has some similarities to commonly observed distributions in the level-k literature. The levels 1 and 2 are present in significant shares and also level-3 and Bayesian behavior is detected. The level-k distribution is hump-shaped, similar to estimated distributions in strategic interaction. By a large margin, the modal behavior turns out to be level-2 play, supporting the hypothesis that naïve inference is an important element in social learning. The results from my implementation of this framework in the experiment with intra-team communication are very similar (section 4.2).

[11] Table 9 on page 32 reports the data for each subject and type.

[12] In the following, I will use the p-value not only in its original meaning in the hypothesis tests, but with slight abuse also as my best type-independent measure of closeness between fingerprints.

Like subject 8 in Figure 1b, 7 players choose thresholds $\theta \in \{-10, -5, 0, 5, 10\}$ in all 15 situations. Players classified as level-2 account for 103 of the 136 instances (75.7%) of cascading ($\theta \in \{-10, 10\}$). The model thus proposes a behavioral explanation for the information cascading detected in ÇK.

The experiment was not designed to discriminate between types in the way I do. The differences between types' fingerprints in the occurring sequence of situations are discriminating in the sense that the ideal fingerprints are never congruent and the SSDs are always distinct. Still, I need to account for cases in which subjects' fingerprints have a small SSD for more than one ideal type. To obtain a structure similar to a mixture model I classify subjects that are significantly close to, say, two types as partly one type and partly the other type. The share attributed to each type is chosen anti-proportionally to the p-value. Table 2b shows that the main results remain unchanged when taking into account how well types can be discriminated.

Costa-Gomes and Crawford (2006) check for unknown types by testing whether subject's fingerprints are close to each other without being close to an ideal type. Such "clusters" of multiple subjects would suggest a common, but not predicted way of thinking about the situation. Unfortunately, such a specification test is not feasible with the original ÇK-data since each subject decides in different 15 situations. My implementation of this game with intra-team communication (section 4.2) confronts all subjects with the same situations. Further insights into the reasoning will thus be provided from the messages as well as the specification test.

4. Data from team communication

In order to put the results of the previous section under scrutiny and get a better insight into the reasoning process, I conducted two experiments with an intra-team communication protocol that yields incentivized written accounts of subjects' reasoning (see Burchardi and Penczynski, 2012).[13] In these experiments, teams of two players are confronted with 6 decisions that arose in the original studies by AH and ÇK, respectively. In both cases, I chose decision situations that occurred later in the sequence in order for a substantial history of public information to exist.

The communication protocol incentivizes the individuals' messages within the

[13]Instructions are reprinted in sections B and D of the supplementary material.

team as follows. Participants are randomly assigned into teams of two players. The two members are connected through the modified chat module of the experiment software.[14] Once the situation is observed, each team member can state a so-called "suggested decision" and justify this in a written message. As soon as this is done for the first three situations, the suggestions and messages get exchanged simultaneously. In a next step, both team members individually state their "final decision" for the three first decisions. It is known to them that for each situation one of the two members' final decisions will be chosen randomly by the computer to count as the "team's action". This construction provides incentives to state the full reasoning underlying the suggested decision in a clear and convincing way.

In contrast to the original design in Burchardi and Penczynski (2012), subjects first write suggested decisions for *three* situations before both the suggested decisions and the messages are simultaneously exchanged. This ensures that the first three messages are written prior to any communication with the team partner and reflect only the individual's reasoning. The same procedure is repeated for the second three decisions.[15]

The message is entered freely without space or time limitations. In the AH experiment, apart from the suggested decision, another structured part of the communication consists of quantifying the confidence in the suggested decision. Subjects can put percentages between 50 and 100, indicating the subjective probability that their choice coincides with the true state of the world. As part of the incentivized communication stage, this is similar to eliciting individual beliefs.

The messages are classified independently by two research assistants and yield one main piece of information. For each individual message the RAs indicate the level of reasoning that the message corresponds to most closely. For this task, the RAs are introduced to the level-k model and receive detailed instructions about characteristics of the individual types.[16]

The following features of reasoning are derived from the model and should be present for the message to be classified as a certain level. Random level-0 play results from misunderstanding the nature of the game or from putting arguments that are orthogonal to any reasonable inference from private signals and public

[14]The experiments were programmed and conducted with the software z-Tree (Fischbacher, 2007).

[15]Having messages exchanged after three rather than one suggested decision has the advantage that all three messages are not "contaminated" by the team partner's arguments. In the original setup with an immediate exchange, this would only be true for the very first message.

[16]Instructions for the RAs are reprinted in sections C and E of the supplementary material.

actions. Level-1 play features an open disregard of others' actions or a strong emphasis on the unambiguity of the own signal in contrast to others' decisions. Level-2 play requires that others' actions are taken at face value as a signal and predecessors are not differentiated. In contrast to this, level-3 play distinguishes individual predecessors and evaluates the information content of the action for each of them.

Both RAs first provide independent sets of classifications. After this, both are anonymously informed about all classifications of the other RA and have the possibility to revise their own classification. This iteration serves to reconsider diverging classifications and to screen errors or misperceptions. The information is merged by using only those classifications that coincide between the two RAs. The RAs agree in a large majority of classifications (AH: 567 (541) out of 636, 89.2% (85.1%); ÇK: 218 (194) out of 252, 86.5% (77.0%); pre-revision in brackets). Burchardi and Penczynski (2012) provide further evidence for the robustness and replicability of this kind of classification. Overall, 17 out of 1272 individual RA level classifications imply a different decision than observed in the suggested decision. 10 of those come from 5 decisions in which the two RAs agree, suggesting an irregularity on the subject side.

4.1. The Anderson and Holt framework

The experiment was conducted in 7 experimental sessions at the LEEDR Laboratory of Cornell University. 106 subjects in teams of two were taking decisions in 6 social learning situations that arose in the original experiment of AH.[17] In the experiment, the binary setting described in 1.2 has equally likely states $\omega \in \{A, B\}$ reflected by two urns A and B. The signals $s \in \{\text{White}, \text{Black}\}$ are informative draws from the urns, where urn A contains 2 white and 1 black ball and urn B contains 1 white and 2 black balls. Possible actions are $a \in \{A, B\}$.

Results. Table 3 gives an overview of the 6 urn choices. The first line gives the information (history and signal) available to subjects. The next two lines give means of the suggested decisions and confidences observed in the experiment. The suggested decisions take values A or B and the confidences take values between 50 and 100. Urn choices 1 and 3 were least controversial with a large

[17] In particular, subjects are told at the beginning that they will experience situations that arose in a previous experiment which was played by individuals. Subsequently, the game is explained in the instructions and finally they see the information about the specific situation on the screen.

majority choosing B and A, respectively, and high average confidences of nearly 70. Urn choice 4 was the most controversial with 53.8% choosing B. The lowest confidence, however, was indicated for urn choice 2. The last four rows of the table indicate the expected behavior of the types proposed in the model.

		Urn choice					
		1	2	3	4	5	6
Information*		$BBAb$	$BAABAb$	$BAAa$	$BBABa$	$ABAAAb$	$BAAAb$
Suggested decision**		0.953	0.689	0.057	0.538	0.415	0.547
Confidence		69.76	60.75	69.14	64.27	67.17	64.68
Probability	0	0.5	0.5	0.5	0.5	0.5	0.5
of choice B	1	1	1	0	0	1	1
	2/3/B	1	0.5	0	1	0	0

Notes: For information, A and B denote observed actions, a and b denote private signals favoring A or B, respectively. Suggested decisions are reported in fractions of choice B.

Table 3: Suggested decisions and confidences as well as choice probabilities for different ideal types in the 6 urn choices ($N = 106$)

The messages from the intra-team communication give direct information about individuals' reasoning. In the upcoming analysis, I focus on the first three decisions that were taken prior to any contact with the team partner.[18] I pool the information for the first three decisions by forming the union of the level classifications. Table 4 summarizes the data by giving the subjects' lowest and highest level over these three decisions.[19] 29 subjects (27%) are not classified, mostly because they did not write any message. 79% of the remaining 77 subjects are attributed a constant level.

The overall picture of this table reflects a large heterogeneity of reasoning. Focusing on the numbers of subjects on the diagonal with coinciding lower and upper bound, there is a pronounced heterogeneity with a hump-shaped distribution like in the previous analysis. Also, the marginal distributions are hump-shaped

[18] The classification for the last 3 decisions is reported in Table 10 on page 33. A Wilcoxon signed-rank test does not reject the hypothesis that the level lower bounds or the level upper bounds are different between decisions 1-3 (Table 4) and decisions 4-6 (Table 10, $p_{lower} = 0.515$, $p_{upper} = 0.470$). To control for order effects, the order of choices was reversed in the last three sessions. If the choices are not numbered, the analysis presents the temporal order of the decisions, i. e. the "first three decisions" will always refer to the decisions taken without prior conversation with the team partner.

[19] The union excludes any outcome of a non-classification (NA). If all three decisions are classified differently or as NA by both RAs, the subject appears under NA. For more details, Table 11 on page 33 reports this data by decision. Tables 12 and 13 on pages 34 and 35 report the data of both RAs on an individual level.

		Highest level (∅ 1.66)					
		0	1	2	3	NA	Total
	0	8	3	3	1	0	15
Lowest level	1		12	6	0	0	18
(∅ 1.38)	2			40	3	0	43
	3				1	0	1
	NA					29	29
	Total	8	15	49	5	29	106

Table 4: Level classification in first three decisions

and feature modes at level-2. The overall modal interval is {2} with 40 subjects classified consistently as level-2.

Messages. With the classification being in accordance with the results in the ÇK-data, I want to illustrate each type's qualitative features in detail with the help of the messages. This will show how the subjects' views on the inference situation coincide with the kinds of reasoning that the model postulates. Tables 14 to 19 (pages 36 to 41) reproduce the messages of chosen subjects.

Level-0. The model is based on uninformative level-0 reasoning, which is modeled by random play. The messages in Table 14 show that subjects misunderstand the game or self-report to have no understanding of the game. Accordingly, they predominantly report the lowest possible confidence in their suggested decision of 50.

Level-1. The messages that are categorized as level-1 reasoning illustrate an increment of understanding beyond level-0. Table 15 shows that subjects mostly explain their understanding as to why the private draw is informative. Often, a confidence of 66 or 67 is reported, which stems from the isolated understanding of the one private draw that indicates the true urn with a probability of $\frac{2}{3}$. Subjects follow their signal without exception. The reasoning frequently does not include the other players' actions. If it does, then it expresses confusion or ignorance as to how to make use of this information. In level-1 players' minds, a model of other players is not existing, but the understanding suffices to process the information from the own private draw.

Level-2. Moving on to messages that were classified to reflect level-2 reasoning, Tables 16 and 17 show how the subjects take into account the history of others' actions in addition to their private signal. However, this perception of other players is not very differentiated as all observed actions A and B are given the same weight. Often, the observed history is summarized in statistics reporting, for example, that 3 out of 5 predecessors selected A, without taking into account the predecessors' choice situation. This information is mostly combined with the odds of the private draw being correct. Importantly, this kind of reasoning is exhibited irrespective of facing similar or mixed public information. As a result, subjects can become very confident in their own choice.

Level-3. Finally, Tables 18 and 19 present communication that was classified as reflecting level-3 reasoning. The messages show clearly how players differentiate between predecessors. They do so by putting themselves in the shoes of their predecessors, making use of the knowledge they have about their history and trying to back out their private draws. It becomes apparent that the population belief is updated using the public information available.

The model of their predecessors that these players have in mind takes into account that the signal is only one ingredient in the decision process. The other ingredient that influences the actions is the observed history. Thus, these players are sophisticated enough to model processes in other players' minds which – as shown in earlier messages – some of these other players exhibit and some do not achieve.

4.2. The Çelen and Kariv framework

The experiment was conducted in 4 experimental sessions at the Laboratory of the University of Mannheim. 42 subjects in teams of two were taking decisions in 6 social learning situations that arose in the original experiment of ÇK. In the experiment, the setting described in 1.3 has equally likely states $\omega \in \{A, B\}$ determined by the sign of the sum of 8 continuous signals that are uniformly distributed on $[-10, 10]$. Subjects take decisions by entering a threshold in the interval $[-10, 10]$. For each subject, the computer uses the threshold and the signal to determine the binary state that the subject chooses.

Results. Table 5 gives an overview of the 6 scenarios which were chosen to discriminate between the different types. The first line gives the information about the history available to subjects. The following lines give mean and standard deviation of the choices in $[-10, 10]$ as well as predicted play by level-k and Bayesian players.

		Round					
		1	2	3	4	5	6
	History	AAB	$BBAA$	$AAABA$	BBB	$AABA$	$BABA$
Sugg. decision (mean)		0.17	0.18	-3.52	4.29	-2.50	-0.48
Sugg. decision (s.d.)		4.20	3.95	5.26	5.43	4.84	3.02
	Level-1	0	0	0	0	0	0
Predicted	Level-2	-5	0	-10	10	-10	0
choice	Level-3	-2.5	-5	-5	7.5	-5	-5
	Bayesian	1.25	-5.625	-4.6875	8.75	-4.375	-3.125

Table 5: Empirical and theoretical decisions in the 6 rounds ($N = 42$)

Like in section 3, the 6 suggested decision per subject can be used to identify their type as illustrated in Figure 2. The analysis of the SSD to the ideal type of level-k and Bayesian reasoning yields the results illustrated in Table 6.[20] The individual weights in the mixture-type specification of Table 6b are reported in Table 23 on page 44. Similar to the results in sections 3 and 4.1, pronounced type heterogeneity and a mode behavior of level-2 can be observed.

Type	Frequency	Fraction (%)
1	10	23.8
2	15	35.7
3	3	7.1
B	4	9.5
NA	10	23.8
Total	42	100.00

(a) Classification by lowest significant SSD

Type	Frequency	Fraction (%)
1	10.40	24.8
2	13.22	31.5
3	4.41	10.5
B	3.98	9.5
NA	10.00	23.8
Total	42	100.00

(b) Mixture classification

Table 6: Type overview for 42 subjects

Messages. Qualitatively, the messages in this experiment are very similar to the ones reported previously. The level classification yields the results shown in

[20]The analysis of the first three decisions is less powerful and yields significant non-random play only for slightly more than 50% of the players. Results for the first 3 and second 3 decisions are reported in Tables 20 and 21. Like in the AH experiment (section 4.1), the communication has no detectable effect on the sophistication since the results are not significantly different from each other (Wilcoxon rank-sum, $p = 0.703$.)

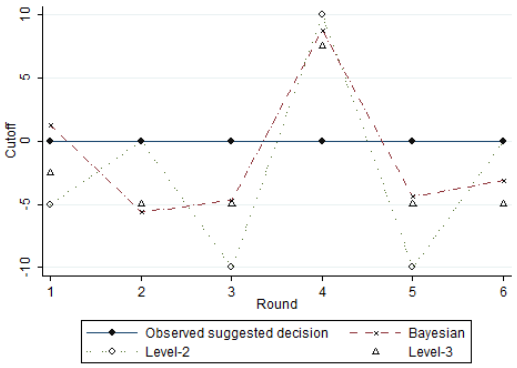

(a) Level-1 play (cutoff = 0, subject 36).

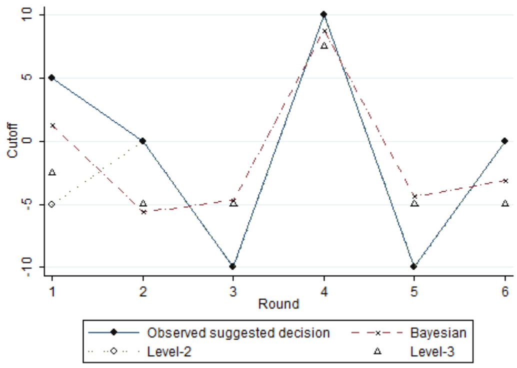

(b) Level-2 play (subject 5).

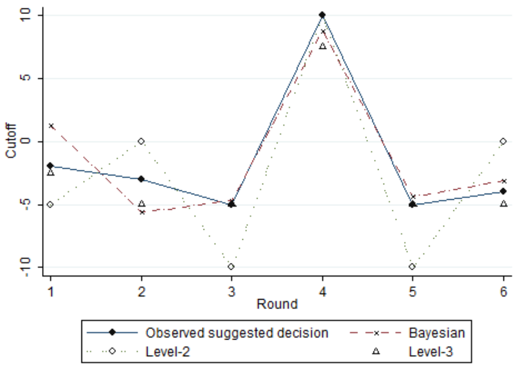

(c) Level-3 play (subject 26).

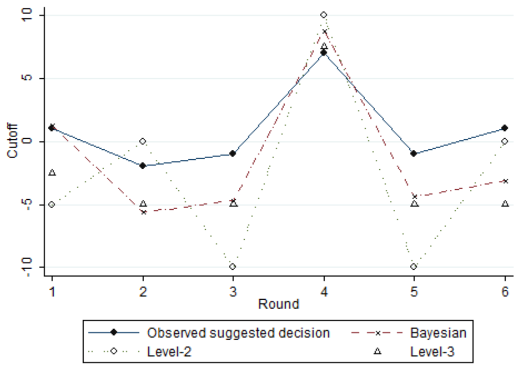

(d) Bayesian play (subject 12).

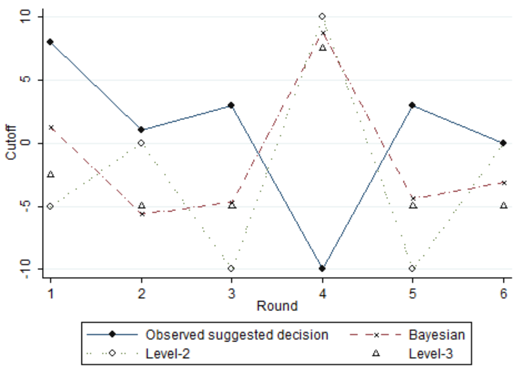

(e) Level-0 play (mistaken sign, subject 35).

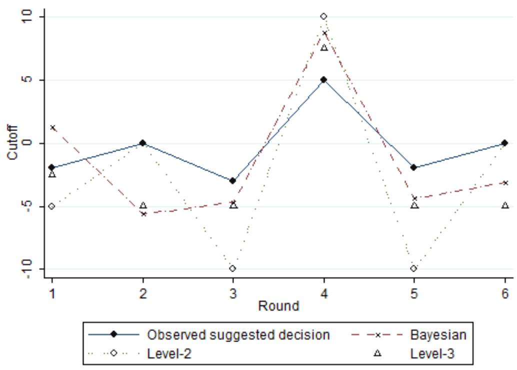

(f) Level-1 play (subject 9).

Figure 2
Subjects' and types' "fingerprints" in the team experiments à la ÇK

Table 7. The reasoning is found to be heterogeneous and the mode behavior is level-2 to a similar extent as before.

		Highest level (⌀ 2.03)					
		0	1	2	3	NA	Total
	0	1	0	1	0	0	2
Lowest level	1		3	0	0	0	3
(⌀ 1.86)	2			18	3	0	21
	3				3	0	3
	NA					13	13
	Total	1	3	19	6	13	42

Table 7: Level classification in first three decisions

Specification Test. It is now possible to investigate whether types of reasoning are common which have not been considered in the theory. In contrast to section 3, all subjects in this experiment face the same set of decisions. In a specification test as in Costa-Gomes and Crawford (2006), all fingerprints in the data are used as "pseudotypes" and compared to other subjects' fingerprints. This potentially identifies clusters of players that play similarly and therefore probably have a similarly structured reasoning.

The clusters are formed according to the conditions given in Costa-Gomes and Crawford (2006), who define a cluster "as a group of two or more subjects such that: (a) each subject's original estimated type has smaller likelihood than the pseudotypes of all other subjects in the group; and (b) all subjects in the group make 'sufficiently similar' guesses" (pp. 1761-1762). In my case, condition (a) translates into the pseudotypes having a lower p-value in the comparison with random play than the originally estimated ideal type. Condition (b) is implemented by defining "sufficiently similar" as significantly closer than random play ($p < 0.05$).

The analysis yields six clusters as shown in Table 8 (Table 23 reports cluster affiliation by subject). It will be seen that it is mostly some kind of decision error which is common among players in a cluster. Cluster D features four players who would be categorized based on the SSD if they had not persistently put the wrong sign on their thresholds, thus exhibiting probably the most simple and extreme case of decision error (see Figure 2e for an example). By the messages, they are indeed classified partially as level-1 or level-2. Clusters B, C, and F feature players who are significantly closer than random play only to level-2 play. In each cluster, players' thresholds have a common feature which makes them closer to each other than to the ideal level-2 type. Players in cluster B and C start off with higher numbers, in cluster B they play a moderate fourth action. Players in cluster F start off with lower numbers.

Cluster E features four players who have mostly been classified as level-2 or level-3 by the fingerprints and the messages and who share a pronounced tendency towards negative numbers which is not featured by other player of the same reasoning.

Although the decisions were chosen to discriminate as much as possible between level-2, level-3 and Bayesian, the specification test yields one large set A of overlapping and non-nested clusters of 19 players. Most of these players are classified by the fingerprints and messages as level-2, level-3 or Bayesian. This suggests that while the data is theoretically able to discriminate between types, due to decision noise this is only partially possible with the experimental data. This is a clear limitation of the data in the ÇK framework, which is somewhat alleviated in the original data due to a larger set of 15 decisions. In the setup with team communication, the messages allow to take a closer look at the reasoning and thus to investigate the validity of the fingerprint results.

Cluster	Subjects	Characteristic
A	2, 6, 7, 9, 12, 13, 14, 16, 21, 22, 23, 27, 29, 33, 34, 37, 38, 39, 42	Level-2, Level-3 or Bayesian
B	4, 17	Level-2 with higher round 1, moderate round 4
C	5, 28, 31	Level-2 with higher round 1
D	11, 19, 35, 40	Mistaken sign
E	15, 18, 32, 41	Negative tendency
F	20, 30	Level-2 with lower round 1

Notes: A is not a proper cluster, it is the union of overlapping, non-nested clusters.

Table 8: Clusters in the specification test

Level-3 and Bayesian types. Level-3 and Bayesians can be distinguished more easily in the ÇK than in the AH framework. While the classification of the messages was only investigating level reasoning, the level-3 messages can be checked for beliefs of rational play by others. To illustrate the functioning of the fingerprint analysis and the detailed insights of the messages, I want to discuss two players who could exhibit Bayesian reasoning. Table 24 reports their messages in the first three decisions translated from German.

The first example, subject 26, starts out in the first message clearly stating the assumption that predecessors played rationally. He is calculating correctly the expected signals following the first two A's in the history, but makes a rough guess for the B in third position, which turns out too small a correction. Since the sequences in decisions 2 and 3 are even longer, the subject contents himself with guesses taking into account that later players in the sequence have a strong signal when they play against the majority. The fingerprint analysis identifies him as 80% level-3 and 20% Bayesian (Table 23). This follows from his guesses being closer to level-3 guesses, despite his assumption of rational preceding play (see Figure 1c).

The second example, subject 9, starts out relatively similar as subject 26. However, the forming of the expectation is biased towards 0 since a signal between 0 and 10 results in an expected threshold of 2. In the third message, this is exhibited even stronger. In the second message, the reasoning is not differentiating predecessors at all and appears to be level-2. The very hesitant movements away from 0 make her being classified as 53% level-1, with the remaining weight distributed across level-2, level-3 and Bayesian (see Figure 2f).

5. Discussion

The fact that the level-k structure known from strategic thinking successfully describes reasoning in social learning situations suggests that the model captures general features of decision making processes. The insights from this study help to flesh out some of these processes.

The results highlight how level-0, level-1, and level-2 are fundamentally different: A level-0 player does not understand the game enough to even process her own signal. Level-1 players understand the rules of the game and can use the signal, but do not understand how to learn from others since no model of others is conceived that tells what the history implies. Level-2 players then are able to model others in probably the least sophisticated way by assuming they play simply according to their own signal.

From level-2 onwards, the reasoning is characterized by the subjects' ability to model other peoples' reasoning in their head. The view that the step from level-1 to level-2 is a cognitively very important one has been put forward by Coricelli and Nagel (2009), who show in a neuroeconomic study that the level-2 reasoning involves activity in other parts of the brain than level-1 reasoning does. In particular, the additional activity is in the medial prefrontal cortex, an area that is connected with mentalizing, i. e. thinking about the mental activity of others. In the level-k framework, the model of others is reflected by the population belief, which in the literature has been specified as degenerate on level $k-1$ or non-degenerate on levels lower than k.

The psychology literature suggests that a plausible starting point or anchor for a model of others' knowledge or reasoning is one's own knowledge or reasoning, which is then adjusted to the specific "other" (Tversky and Kahneman, 1974; Gordon, 1986; Nickerson, 1999). In this light, a degenerate population belief can be seen to reflect the first conceptualization of the other as someone that thinks in the same way as oneself thought until now. Best responding to this view of the other turns a level-$k-1$ player into a level-k player who likens all others to his past level-$k-1$ self. This best response iteration might go on to level-$k+1$, etc. Alternatively, further deliberation might result in a more refined model of others in the sense that they might be less sophisticated than the player just was, making the (adjusted) population belief non-degenerate. This gives a cognitive perspective on the two competing assumptions of degenerate and non-degenerate population beliefs.

In any case, both the non-degenerate population belief and the additional best response can be viewed as a refinement of the initial model of others. Such refinements should occur in players of level-2 and higher. A possible reason for the rare observation of levels beyond 3 (Arad and Rubinstein, 2012) might be that eventually the correct specification of the population belief is deemed more fruitful than mechanically thinking through further best responses.

6. Conclusion

This paper proposes a level-k model of social learning which introduces heterogeneous, hierarchical types in the context of inference. In contrast to much of the social learning literature, the experimental investigation is able to differentiate these various types thanks to two innovations. On the one hand, within the informative framework of social learning introduced by Çelen and Kariv (2004) the predicted behavior differs by type, enabling an analysis of individual reasoning on the basis of their data. On the other hand, independently of the framework, the nature of reasoning can be investigated with the help of written accounts obtained from incentivized intra-team communication as introduced by Burchardi and Penczynski (2012).

The analyses find strong new evidence for predictions of the model, in particular the heterogeneity of reasoning across individuals and the presence of the particular types of reasoning. Level-2 forms of reasoning are found to be the prevalent nature of social learning, implying that subjects have simplified models of others in their head and do not take into account others' specific situation. Rational expectations would imply that people hold the correct model of others in their minds. The model and the results explain why rational expectations fail to hold in a social learning context (Weizsäcker, 2010).

The model captures subjects' behavior well and gives a new interpretation of the hierarchical structure of reasoning. In particular, it reflects the gradual better understanding of the situation (level-0, 1, and 2) as well as the stepwise refinement of the model held of others' reasoning (level-2, 3, etc.). Applying the concept of levels of reasoning to social learning is not only a useful extension of this theory's scope, it also highlights the importance and usefulness of iterated best responses as a framework to think about reasoning in social settings.

References

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

Anderson, Lisa R., "Payoff Effects in Information Cascade Experiments," *Economic Inquiry*, October 2001, *39* (4), 609–15.

__ **and Charles A. Holt**, "Information Cascades in the Laboratory," *American Economic Review*, December 1997, *87* (5), 847–62.

Arad, Ayala and Ariel Rubinstein, "The 11-20 Money Request Game: Evaluating the Upper Bound of k-Level Reasoning," *American Economic Review*, 2012, *102* (7).

Banerjee, Abhijit V, "A Simple Model of Herd Behavior," *The Quarterly Journal of Economics*, August 1992, *107* (3), 797–817.

Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch, "A Theory of Fads, Fashion, Custom, and Cultural Change in Informational Cascades," *Journal of Political Economy*, October 1992, *100* (5), 992–1026.

Brocas, Isabelle, Juan D. Carillo, Stephanie W. Wang, and Colin F. Camerer, "Imperfect choice or imperfect attention? Understanding strategic thinking in private information games," Technical Report, University of Southern California 2010.

Burchardi, Konrad B. and Stefan P. Penczynski, "Out Of Your Mind: Eliciting Individual Reasoning in One Shot Games," Working Paper, Stockholm University June 2012.

Camerer, Colin F., Teck-Hua Ho, and Juin-Kuan Chong, "A Cognitive Hierarchy Model of Games," *The Quarterly Journal of Economics*, August 2004, *119* (3), 861–898.

Çelen, Boğaçhan and Shachar Kariv, "Distinguishing Informational Cascades from Herd Behavior in the Laboratory," *American Economic Review*, June 2004, *94* (3), 484–498.

Cipriani, Marco and Antonio Guarino, "Herd Behavior in a Laboratory Financial Market," *American Economic Review*, 2005, *95* (5), pp. 1427–1443.

Coricelli, Giorgio and Rosemarie Nagel, "Neural correlates of depth of strategic reasoning in medial prefrontal cortex," *Proceedings of the National Academy of Sciences*, 2009, *106* (23), 9163–9168.

Costa-Gomes, Miguel A. and Vincent P. Crawford, "Cognition and Behavior in Two-Person Guessing Games: An Experimental Study," *American Economic Review*, December 2006, *96* (5), 1737–1768.

__ , __ , **and Bruno Broseta**, "Cognition and Behavior in Normal-Form Games: An Experimental Study," *Econometrica*, September 2001, *69* (5), 1193–1235.

Crawford, Vincent P. and Nagore Iriberri, "Level-k Auctions: Can a Nonequilibrium Model of Strategic Thinking Explain the Winner's Curse and Overbidding in Private-Value Auctions?," *Econometrica*, November 2007, *75* (6), 1721–1770.

Dominitz, Jeff and Angela A. Hung, "Homogeneous Actions and Hetergeneous Beliefs: Experimental Evidence on the Formation of Information Cascades," Econometric Society 2004 North American Winter Meetings 64, Econometric Society August 2004.

Drehmann, Mathias, Jorg Oechssler, and Andreas Roider, "Herding and Contrarian Behavior in Financial Markets: An Internet Experiment," *American Economic Review*, December 2005, *95* (5), 1403–1426.

Eyster, Erik and Matthew Rabin, "Cursed Equilibrium," *Econometrica*, September 2005, *73* (5), 1623–1672.

__ **and** __ , "Naïve Herding in Rich-Information Settings," *American Economic Journal: Microeconomics*, November 2010, *2* (4), 221–43.

Fischbacher, Urs, "z-Tree: Zurich toolbox for ready-made economic experiments," *Experimental Economics*, June 2007, *10* (2), 171–178.

Goeree, Jacob K., Thomas R. Palfrey, Brian W. Rogers, and Richard D. McKelvey, "Self-Correcting Information Cascades," *Review of Economic Studies*, 07 2007, *74* (3), 733–762.

Gordon, Robert M., "Folk Psychology as Simulation," *Mind and Language*, 1986, *1* (2), 158–171.

Hung, Angela A. and Charles R. Plott, "Information Cascades: Replication and an Extension to Majority Rule and Conformity-Rewarding Institutions," *American Economic Review*, December 2001, *91* (5), 1508–1520.

Kübler, Dorothea and Georg Weizsäcker, "Limited depth of reasoning and failure of cascade formation in the laboratory," *Review of economic studies*, 2004, *71*, 424–441.

Nagel, Rosemarie, "Unraveling in Guessing Games: An Experimental Study," *American Economic Review*, December 1995, *85* (5), 1313–1326.

Nickerson, Raymond S., "How We Know–and Sometimes Misjudge–What Others Know: Imputing One's Own Knowledge to Others," *Psychological*

Bulletin, November 1999, *125* (6), 737–759.

Nöth, Markus and Martin Weber, "Information Aggregation with Random Ordering: Cascades and Overconfidence*," *The Economic Journal*, 2003, *113* (484), 166–189.

Oberhammer, Clemens and Andreas Stiehler, "Does Cascade Behavior in Information Cascades Reflect Bayesian Updating?," Technical Report, Max Planck Institute of Economics 2003.

Smith, Lones and Peter Sørensen, "Pathological Outcomes of Observational Learning," *Econometrica*, March 2000, *68* (2), 371–398.

Stahl, Dale O. and Paul W. Wilson, "On Players' Models of Other Players: Theory and Experimental Evidence," *Games and Economic Behavior*, July 1995, *10* (1), 218–254.

Tversky, Amos and Daniel Kahneman, "Judgment under Uncertainty: Heuristics and Biases.," *Science*, 1974, *185* (4157), 1124–1131.

Weizsäcker, Georg, "Do We Follow Others When We Should? A Simple Test of Rational Expectations," *American Economic Review*, December 2010, *100* (5), 2340–60.

Willinger, Marc and Anthony Ziegelmeyer, "Are More Informed Agents Able to Shatter Information Cascades in the Lab?," in Patrick Cohendet, Patrick Llerena, Hubert Stahn, and Gisele Umbhauer, eds., *The Economics of Networks: Interaction and Behaviours*, New York: Springer, 1998, pp. 291–305.

Ziegelmeyer, Anthony, Frédéric Koessler, Juergen Bracht, and Eyal Winter, "Fragility of Information Cascades: An Experimental Study using Elicited Beliefs," Technical Report, Max Planck Institute for Research into Economic Systems 2008.

A. Additional tables

Subject	Level 1		Level 2		Level 3		Bayesian	
ID	SSD	p	SSD	p	SSD	p	SSD	p
1	549.13	0.671	1475.83	0.801	1143.33	0.750	1255.61	0.697
2	522.00	0.585	1022.00	0.183	870.75	0.199	850.50	0.202
3	714.61	0.964	946.61	0.095	777.61	0.224	823.28	0.140
4	331.00	0.068	131.00	0.000	134.75	0.000	99.08	0.000
5	1500.00	1.000	375.00	0.000	656.25	0.107	610.19	0.040
6	289.86	0.029	1020.36	0.237	796.35	0.306	721.46	0.178
7	725.00	0.970	200.00	0.000	368.75	0.004	441.56	0.012
8	700.00	0.954	25.00	0.000	268.75	0.003	170.86	0.000
9	255.00	0.012	320.00	0.000	203.75	0.000	298.23	0.001
10	627.00	0.862	97.00	0.000	245.75	0.001	177.84	0.000
11	687.17	0.943	275.17	0.000	156.42	0.000	168.43	0.000
12	423.00	0.260	818.00	0.027	426.75	0.009	595.65	0.016
13	661.25	0.914	51.25	0.000	117.50	0.000	97.09	0.000
14	400.01	0.197	450.01	0.002	363.01	0.003	494.64	0.006
15	838.61	0.997	394.61	0.001	310.11	0.002	318.56	0.001
16	883.00	0.999	158.00	0.000	289.25	0.002	305.77	0.001
17	0.02	0.000	723.65	0.059	536.24	0.034	529.73	0.033
18	100.00	0.000	475.00	0.014	243.75	0.001	269.21	0.002
19	186.39	0.001	548.09	0.023	555.54	0.041	575.62	0.055
20	224.69	0.005	485.09	0.004	316.89	0.001	328.71	0.002
21	524.25	0.593	120.25	0.000	140.75	0.000	112.44	0.000
22	391.00	0.175	656.00	0.035	528.50	0.042	587.05	0.042
23	421.00	0.254	141.00	0.000	137.25	0.000	124.54	0.000
24	242.25	0.008	52.25	0.000	162.25	0.000	115.79	0.000
25	308.25	0.043	743.25	0.012	415.75	0.007	607.40	0.013
26	1386.52	1.000	783.02	0.062	929.27	0.552	936.15	0.324
27	1171.15	1.000	1182.65	0.163	845.15	0.150	947.88	0.159
28	444.25	0.324	194.25	0.000	209.25	0.000	214.01	0.000
29	1028.33	1.000	258.33	0.000	490.58	0.006	389.41	0.001
30	514.66	0.561	602.66	0.019	455.91	0.019	456.18	0.010
31	893.26	0.999	537.26	0.003	538.26	0.030	522.84	0.008
32	619.25	0.847	939.25	0.043	684.25	0.073	834.31	0.064
33	5.00	0.000	885.00	0.093	652.50	0.054	740.49	0.071
34	396.98	0.190	1331.88	0.579	1110.53	0.705	1001.84	0.453
35	899.00	0.999	389.00	0.001	521.50	0.045	494.42	0.016
36	500.00	0.511	175.00	0.000	243.75	0.001	168.24	0.000
37	197.13	0.002	659.63	0.011	268.38	0.001	475.27	0.005
38	103.00	0.000	768.00	0.156	545.50	0.127	723.30	0.202
39	1500.00	1.000	575.00	0.013	881.25	0.444	669.35	0.066
40	807.11	0.995	505.11	0.023	727.31	0.309	665.61	0.174

Table 9: Type analysis in data of Çelen and Kariv (2004)

		Highest level (⌀ 1.76)					
		0	1	2	3	NA	Total
	0	5	0	6	0	0	11
Lowest level	1		13	5	0	0	18
(⌀ 1.50)	2			42	3	0	45
	3				2	0	2
	NA					30	30
	Total	5	13	53	5	30	106

Table 10: Level classification in second 3 decisions

	Urn choice					
	1	2	3	4	5	6
Information	$BBAb$	$BAABAb$	$BAAa$	$BBABa$	$ABAAAb$	$BAAAb$
Level-0	5	7	7	8	5	8
Level-1	5	10	3	12	16	17
Level-2	40	29	44	39	40	36
Level-3	3	2	3	1	4	4
NA	53	58	49	46	41	41
Total	106	106	106	106	106	106

Table 11: Level classification by urn choice

		Decision 1 (BBAb)					Decision 2 (BAABAb)					Decision 3 (BAAa)					Decision 4 (BBABa)					Decision 5 (ABAAAb)					Decision 6 (BAAAb)				
Session	Subject ID	Sugg. dec.	Confidence	Fingerprint	RA 1	RA 2	Sugg. dec.	Confidence	Fingerprint	RA 1	RA 2	Sugg. dec.	Confidence	Fingerprint	RA 1	RA 2	Sugg. dec.	Confidence	Fingerprint	RA 1	RA 2	Sugg. dec.	Confidence	Fingerprint	RA 1	RA 2	Sugg. dec.	Confidence	Fingerprint	RA 1	RA 2
1	1	B	75	{0, 1, 2, 3}	2	2	B	60	{0, 1, 2, 3}	2	2	A	75	{0, 1, 2, 3}	2	2	A	60	{0, 1}	2	1	B	75	{0, 1}	1	1	B	60	{0, 1}	1	1
1	2	B	0	{0, 1, 2, 3}	2	2	A	50	{0, 2, 3}	.	0	A	50	{0, 1, 2, 3}	.	.	B	50	{0, 2, 3}	2	2	A	50	{0, 2, 3}	.	.	A	0	{0, 2, 3}	.	.
1	3	B	75	{0, 1, 2, 3}	2	2	B	75	{0, 1, 2, 3}	2	2	A	75	{0, 1, 2, 3}	2	2	B	80	{0, 2, 3}	2	2	A	60	{0, 2, 3}	0	2	A	50	{0, 2, 3}	.	.
1	4	B	70	{0, 1, 2, 3}	.	.	B	70	{0, 1, 2, 3}	.	.	A	70	{0, 1, 2, 3}	2	1	B	50	{0, 2, 3}	.	.	B	70	{0, 1}	1	1	A	55	{0, 2, 3}	2	2
1	5	B	80	{0, 1, 2, 3}	2	1	B	80	{0, 1, 2, 3}	1	1	A	80	{0, 1, 2, 3}	2	1	A	80	{0, 1}	1	1	B	80	{0, 1}	1	1	B	80	{0, 1}	1	1
1	6	B	75	{0, 1, 2, 3}	0	0	B	75	{0, 1, 2, 3}	.	1	A	75	{0, 1, 2, 3}	2	2	B	65	{0, 2, 3}	2	2	A	65	{0, 2, 3}	2	2	A	65	{0, 2, 3}	2	2
1	7	B	95	{0, 1, 2, 3}	2	2	A	50	{0, 2, 3}	0	2	A	90	{0, 1, 2, 3}	2	2	B	75	{0, 2, 3}	2	2	A	70	{0, 2, 3}	2	2	A	70	{0, 2, 3}	2	2
1	8	B	70	{0, 1, 2, 3}	2	2	B	55	{0, 1, 2, 3}	2	2	A	70	{0, 1, 2, 3}	2	2	B	50	{0, 2, 3}	2	2	A	55	{0, 2, 3}	2	2	A	60	{0, 2, 3}	2	2
1	9	B	67	{0, 1, 2, 3}	2	1	B	67	{0, 1, 2, 3}	.	1	A	67	{0, 1, 2, 3}	.	.	A	67	{0, 1}	1	1	B	67	{0, 1}	1	1	B	67	{0, 1}	1	1
1	10	B	0	{0, 1, 2, 3}	.	.	A	0	{0, 2, 3}	0	2	A	0	{0, 1, 2, 3}	.	.	B	0	{0, 2, 3}	2	2	A	0	{0, 2, 3}	2	2	A	0	{0, 2, 3}	2	2
1	11	B	70	{0, 1, 2, 3}	2	1	B	50	{0, 1, 2, 3}	.	0	A	60	{0, 1, 2, 3}	2	2	B	60	{0, 2, 3}	2	2	A	60	{0, 2, 3}	2	2	A	60	{0, 2, 3}	2	2
1	12	B	50	{0, 1, 2, 3}	.	.	A	50	{0, 2, 3}	.	.	B	60	{0}	.	.	A	50	{0, 1}	.	0	B	60	{0, 1}	1	1	B	50	{0, 1}	.	.
2	13	B	0	{0, 1, 2, 3}	2	2	B	0	{0, 1, 2, 3}	2	1	A	100	{0, 1, 2, 3}	2	2	B	55	{0, 2, 3}	2	2	A	100	{0, 2, 3}	2	2	A	60	{0, 2, 3}	2	2
2	14	B	70	{0, 1, 2, 3}	2	1	B	70	{0, 1, 2, 3}	2	1	A	70	{0, 1, 2, 3}	2	1	A	70	{0, 1}	.	.	B	60	{0, 1}	.	.	B	60	{0, 1}	.	.
2	15	B	100	{0, 1, 2, 3}	3	3	B	70	{0, 1, 2, 3}	2	3	A	100	{0, 1, 2, 3}	2	2	B	60	{0, 2, 3}	2	2	A	100	{0, 2, 3}	2	2	A	70	{0, 2, 3}	2	2
2	16	B	0	{0, 1, 2, 3}	2	1	B	65	{0, 1, 2, 3}	0	0	A	60	{0, 1, 2, 3}	0	0	A	67	{0, 1}	1	0	B	67	{0, 1}	.	.	B	67	{0, 1}	0	.
2	17	B	0	{0, 1, 2, 3}	.	.	A	50	{0, 2, 3}	.	.	A	50	{0, 1, 2, 3}	0	0	B	0	{0, 2, 3}	.	.	B	50	{0, 1}	.	0	B	50	{0, 1}	.	.
2	18	B	75	{0, 1, 2, 3}	2	2	B	50	{0, 1, 2, 3}	.	0	A	75	{0, 1, 2, 3}	.	.	B	50	{0, 2, 3}	2	2	A	50	{0, 2, 3}	2	2	B	50	{0, 1}	3	3
2	19	B	80	{0, 1, 2, 3}	2	2	B	50	{0, 1, 2, 3}	2	2	A	80	{0, 1, 2, 3}	2	2	B	60	{0, 2, 3}	2	2	A	75	{0, 2, 3}	2	2	A	75	{0, 2, 3}	2	2
2	20	B	85	{0, 1, 2, 3}	2	2	B	50	{0, 1, 2, 3}	.	2	A	80	{0, 1, 2, 3}	2	2	B	55	{0, 2, 3}	2	2	A	60	{0, 2, 3}	2	2	A	60	{0, 2, 3}	2	2
2	21	B	69	{0, 1, 2, 3}	.	.	A	69	{0, 2, 3}	.	0	A	69	{0, 1, 2, 3}	.	.	B	69	{0, 2, 3}	.	.	A	99	{0, 2, 3}	.	.	A	0	{0, 2, 3}	.	.
2	22	B	80	{0, 1, 2, 3}	.	.	A	50	{0, 2, 3}	.	.	A	60	{0, 1, 2, 3}	.	.	B	70	{0, 2, 3}	.	.	A	60	{0, 2, 3}	.	.	A	60	{0, 2, 3}	.	.
2	23	B	67	{0, 1, 2, 3}	.	.	B	67	{0, 1, 2, 3}	.	.	A	67	{0, 1, 2, 3}	.	.	A	75	{0, 1}	.	.	B	70	{0, 1}	.	.	B	67	{0, 1}	.	.
2	24	B	70	{0, 1, 2, 3}	2	2	A	70	{0, 2, 3}	2	0	A	70	{0, 1, 2, 3}	2	2	B	60	{0, 2, 3}	2	2	A	80	{0, 2, 3}	2	2	A	70	{0, 2, 3}	2	2
2	25	B	75	{0, 1, 2, 3}	2	2	A	65	{0, 2, 3}	2	2	A	80	{0, 1, 2, 3}	2	2	B	65	{0, 2, 3}	2	2	A	85	{0, 2, 3}	2	2	A	75	{0, 2, 3}	2	2
2	26	B	66	{0, 1, 2, 3}	1	1	A	50	{0, 2, 3}	2	2	A	66	{0, 1, 2, 3}	2	2	B	60	{0, 2, 3}	2	2	A	60	{0, 2, 3}	2	2	A	60	{0, 2, 3}	2	2
2	27	B	50	{0, 1, 2, 3}	0	.	B	50	{0, 1, 2, 3}	0	0	A	50	{0, 1, 2, 3}	0	0	A	50	{0, 1}	0	0	B	50	{0, 1}	.	.	B	50	{0, 1}	0	0
2	28	B	50	{0, 1, 2, 3}	2	2	B	60	{0, 1, 2, 3}	.	.	A	50	{0, 1, 2, 3}	.	.	A	70	{0, 1}	.	.	B	70	{0, 1}	.	.	B	70	{0, 1}	.	.
2	29	B	75	{0, 1, 2, 3}	2	2	A	60	{0, 2, 3}	2	2	A	75	{0, 1, 2, 3}	2	2	B	70	{0, 2, 3}	2	2	A	70	{0, 2, 3}	2	2	A	50	{0, 2, 3}	0	0
2	30	B	70	{0, 1, 2, 3}	.	.	B	60	{0, 1, 2, 3}	.	.	A	65	{0, 1, 2, 3}	.	.	A	60	{0, 1}	.	.	A	70	{0, 2, 3}	2	2	A	60	{0, 2, 3}	2	2
3	31	B	70	{0, 1, 2, 3}	.	.	B	70	{0, 1, 2, 3}	.	.	B	0	{0}	.	.	B	80	{0, 2, 3}	2	2	A	90	{0, 2, 3}	2	2	A	80	{0, 2, 3}	.	.
3	32	B	0	{0, 1, 2, 3}	.	.	B	66	{0, 1, 2, 3}	.	.	A	66	{0, 1, 2, 3}	.	.	A	66	{0, 1}	.	.	B	67	{0, 1}	.	.	B	66	{0, 1}	.	.
3	33	B	50	{0, 1, 2, 3}	0	0	B	80	{0, 1, 2, 3}	.	.	A	55	{0, 1, 2, 3}	0	0	A	65	{0, 1}	.	.	A	88	{0, 2, 3}	.	.	A	75	{0, 2, 3}	.	.
3	34	B	50	{0, 1, 2, 3}	.	.	A	60	{0, 2, 3}	.	.	A	50	{0, 1, 2, 3}	.	.	A	50	{0, 1}	.	.	B	50	{0, 1}	.	.	B	50	{0, 1}	.	.
3	35	B	67	{0, 1, 2, 3}	.	1	B	67	{0, 1, 2, 3}	.	.	A	67	{0, 1, 2, 3}	.	.	A	67	{0, 1}	.	.	B	67	{0, 1}	.	.	B	67	{0, 1}	2	.
3	36	B	90	{0, 1, 2, 3}	2	2	B	50	{0, 1, 2, 3}	.	2	A	90	{0, 1, 2, 3}	2	2	B	51	{0, 2, 3}	2	2	A	60	{0, 2, 3}	2	2	A	55	{0, 2, 3}	2	2
3	37	B	50	{0, 1, 2, 3}	.	.	A	50	{0, 2, 3}	.	.	A	50	{0, 1, 2, 3}	.	.	B	50	{0, 2, 3}	2	2	A	50	{0, 2, 3}	2	2	A	50	{0, 2, 3}	.	2
3	38	B	66	{0, 1, 2, 3}	2	2	B	50	{0, 1, 2, 3}	2	2	A	75	{0, 1, 2, 3}	2	2	B	60	{0, 2, 3}	2	2	A	70	{0, 2, 3}	2	2	A	70	{0, 2, 3}	2	2
3	39	B	0	{0, 1, 2, 3}	.	.	A	0	{0, 2, 3}	.	2	A	80	{0, 1, 2, 3}	2	2	A	50	{0, 1}	.	.	A	50	{0, 2, 3}	.	.	B	50	{0, 1}	0	0
3	40	B	75	{0, 1, 2, 3}	2	2	B	50	{0, 1, 2, 3}	2	2	A	90	{0, 1, 2, 3}	2	1	B	75	{0, 2, 3}	2	2	A	90	{0, 2, 3}	2	2	A	60	{0, 2, 3}	2	2
3	41	B	75	{0, 1, 2, 3}	2	2	A	60	{0, 2, 3}	2	2	A	60	{0, 1, 2, 3}	2	1	B	60	{0, 2, 3}	2	2	A	65	{0, 2, 3}	2	2	B	50	{0, 1}	0	0
3	42	B	75	{0, 1, 2, 3}	0	0	B	75	{0, 1, 2, 3}	.	.	A	75	{0, 1, 2, 3}	.	.	A	75	{0, 1}	1	1	B	75	{0, 1}	2	1	B	75	{0, 1}	.	.
3	43	B	70	{0, 1, 2, 3}	2	2	B	60	{0, 1, 2, 3}	2	2	A	75	{0, 1, 2, 3}	2	2	B	0	{0, 2, 3}	2	2	A	70	{0, 2, 3}	2	2	B	50	{0, 1}	3	2
3	44	B	75	{0, 1, 2, 3}	2	2	B	75	{0, 1, 2, 3}	1	1	A	75	{0, 1, 2, 3}	2	1	A	50	{0, 1}	1	2	B	50	{0, 1}	1	2	B	50	{0, 1}	1	1
3	45	B	65	{0, 1, 2, 3}	.	.	A	55	{0, 2, 3}	2	2	A	70	{0, 1, 2, 3}	2	2	B	70	{0, 2, 3}	2	2	A	70	{0, 2, 3}	2	2	A	70	{0, 2, 3}	3	3
3	46	B	75	{0, 1, 2, 3}	2	2	A	50	{0, 2, 3}	2	3	A	75	{0, 1, 2, 3}	.	.	B	50	{0, 2, 3}	2	3	A	50	{0, 2, 3}	2	.	A	50	{0, 2, 3}	.	.
3	47	B	50	{0, 1, 2, 3}	2	0	B	50	{0, 1, 2, 3}	.	.	B	100	{0}	.	.	A	50	{0, 1}	.	1	A	50	{0, 2, 3}	.	.	B	0	{0, 1}	.	.
3	48	B	75	{0, 1, 2, 3}	2	2	B	75	{0, 1, 2, 3}	.	.	A	75	{0, 1, 2, 3}	2	2	A	75	{0, 1}	0	0	B	75	{0, 1}	0	0	B	75	{0, 1}	1	0
4	49	B	50	{0, 1, 2, 3}	.	.	B	100	{0, 1, 2, 3}	.	.	A	60	{0, 1, 2, 3}	.	.	A	50	{0, 1}	.	.	A	50	{0, 2, 3}	.	.	B	75	{0, 1}	.	.
4	50	B	50	{0, 1, 2, 3}	.	.	B	50	{0, 1, 2, 3}	1	1	A	50	{0, 1, 2, 3}	2	1	A	75	{0, 1}	.	1	B	75	{0, 1}	1	1	B	75	{0, 1}	1	1
4	51	B	50	{0, 1, 2, 3}	0	0	B	50	{0, 1, 2, 3}	0	0	B	50	{0}	0	0	B	0	{0, 2, 3}	0	0	B	50	{0, 1}	0	0	B	50	{0, 1}	0	0
4	52	B	0	{0, 1, 2, 3}	2	2	A	65	{0, 2, 3}	.	2	A	80	{0, 1, 2, 3}	2	2	B	80	{0, 2, 3}	2	2	A	100	{0, 2, 3}	2	2	A	100	{0, 2, 3}	.	2
4	53	B	75	{0, 1, 2, 3}	2	2	A	60	{0, 2, 3}	0	0	A	75	{0, 1, 2, 3}	2	2	A	50	{0, 1}	1	1	A	50	{0, 2, 3}	2	2	B	50	{0, 1}	1	1
4	54	B	80	{0, 1, 2, 3}	1	1	B	60	{0, 1, 2, 3}	1	1	A	70	{0, 1, 2, 3}	2	2	A	66	{0, 1}	1	1	B	60	{0, 1}	1	1	B	70	{0, 1}	1	1
4	55	B	75	{0, 1, 2, 3}	2	2	B	75	{0, 1, 2, 3}	2	2	A	100	{0, 1, 2, 3}	2	2	A	75	{0, 1}	1	1	A	75	{0, 2, 3}	2	2	A	75	{0, 2, 3}	2	2
4	56	B	75	{0, 1, 2, 3}	.	.	A	50	{0, 2, 3}	.	.	A	75	{0, 1, 2, 3}	2	2	B	50	{0, 2, 3}	2	2	A	50	{0, 2, 3}	2	2	A	50	{0, 2, 3}	2	2
4	57	B	75	{0, 1, 2, 3}	.	1	B	60	{0, 1, 2, 3}	1	1	A	85	{0, 1, 2, 3}	2	2	B	60	{0, 2, 3}	2	2	A	60	{0, 2, 3}	2	2	A	50	{0, 2, 3}	2	2
4	58	B	50	{0, 1, 2, 3}	1	1	B	50	{0, 1, 2, 3}	.	.	A	50	{0, 1, 2, 3}	.	.	B	50	{0, 2, 3}	.	.	A	50	{0, 2, 3}	.	.	A	0	{0, 2, 3}	.	.
4	59	B	75	{0, 1, 2, 3}	2	2	A	50	{0, 2, 3}	2	2	A	85	{0, 1, 2, 3}	2	2	B	60	{0, 2, 3}	2	2	A	67	{0, 2, 3}	2	2	A	60	{0, 2, 3}	2	2
4	60	B	50	{0, 1, 2, 3}	2	1	B	50	{0, 1, 2, 3}	1	1	A	50	{0, 1, 2, 3}	.	.	A	0	{0, 1}	.	.	B	0	{0, 1}	.	.	B	0	{0, 1}	1	1
4	61	B	75	{0, 1, 2, 3}	.	.	A	50	{0, 2, 3}	.	.	A	75	{0, 1, 2, 3}	.	.	B	60	{0, 2, 3}	.	.	A	75	{0, 2, 3}	.	.	A	55	{0, 2, 3}	.	.
4	62	B	66	{0, 1, 2, 3}	2	2	A	67	{0, 2, 3}	2	2	A	78	{0, 1, 2, 3}	2	2	B	53	{0, 2, 3}	.	.	A	50	{0, 2, 3}	.	.	A	50	{0, 2, 3}	2	2
4	63	B	67	{0, 1, 2, 3}	2	2	B	50	{0, 1, 2, 3}	2	2	A	50	{0, 1, 2, 3}	2	2	A	75	{0, 1}	0	0	A	50	{0, 2, 3}	2	2	B	60	{0, 1}	1	1
4	64	A	50	{0}	.	0	B	100	{0, 1, 2, 3}	1	1	A	100	{0, 1, 2, 3}	1	1	A	100	{0, 1}	1	1	B	100	{0, 1}	0	1	B	100	{0, 1}	1	1

Table 12: Individual data in the team communication experiment à la Anderson and Holt (1997), sessions 1-4

		Decision 1 (BBAb)					Decision 2 (BAABAb)					Decision 3 (BAAa)					Decision 4 (BBABa)					Decision 5 (ABAAAb)					Decision 6 (BAAAb)				
Session	Subject ID	Sugg. dec.	Confidence	Fingerprint	RA 1	RA 2	Sugg. dec.	Confidence	Fingerprint	RA 1	RA 2	Sugg. dec.	Confidence	Fingerprint	RA 1	RA 2	Sugg. dec.	Confidence	Fingerprint	RA 1	RA 2	Sugg. dec.	Confidence	Fingerprint	RA 1	RA 2	Sugg. dec.	Confidence	Fingerprint	RA 1	RA 2
5	65	B	75	{0, 1}	.	.	B	67	{0, 1}	.	.	A	67	{0, 1}	.	.	A	67	{0, 1, 2, 3}	2	.	B	67	{0, 1, 2, 3}	.	.	B	67	{0, 1, 2, 3}	.	.
5	66	A	65	{0, 2, 3}	2	2	A	80	{0, 2, 3}	2	2	B	60	{0, 2, 3}	2	2	A	100	{0, 1, 2, 3}	.	.	B	100	{0, 1, 2, 3}	2	2	B	100	{0, 1, 2, 3}	2	2
5	67	B	70	{0, 1}	.	.	B	70	{0, 1}	.	.	A	70	{0, 1}	.	.	A	60	{0, 1, 2, 3}	.	.	B	50	{0, 1, 2, 3}	.	.	B	60	{0, 1, 2, 3}	.	.
5	68	B	75	{0, 1}	0	0	B	75	{0, 1}	0	0	A	75	{0, 1}	1	1	A	75	{0, 1, 2, 3}	2	1	B	75	{0, 1, 2, 3}	2	1	B	75	{0, 1, 2, 3}	.	.
5	69	B	75	{0, 1}	.	.	B	75	{0, 1}	.	.	A	75	{0, 1}	.	.	A	50	{0, 1, 2, 3}	.	.	B	63	{0, 1, 2, 3}	1	1	B	63	{0, 1, 2, 3}	1	1
5	70	B	60	{0, 1}	1	1	B	65	{0, 1}	1	1	A	75	{0, 1}	0	0	A	60	{0, 1, 2, 3}	2	1	B	55	{0, 1, 2, 3}	2	2	A	60	{0}	2	0
5	71	B	67	{0, 1}	1	1	B	66	{0, 1}	1	1	A	67	{0, 1}	1	1	A	50	{0, 1, 2, 3}	0	0	B	67	{0, 1, 2, 3}	.	.	B	80	{0, 1, 2, 3}	2	2
5	72	B	67	{0, 1}	1	1	B	67	{0, 1}	.	.	A	67	{0, 1}	.	.	A	67	{0, 1, 2, 3}	.	.	B	67	{0, 1, 2, 3}	.	.	B	67	{0, 1, 2, 3}	.	.
5	73	A	60	{0, 2, 3}	2	2	A	60	{0, 2, 3}	2	2	B	60	{0, 2, 3}	2	2	A	60	{0, 1, 2, 3}	2	2	B	50	{0, 1, 2, 3}	2	2	B	60	{0, 1, 2, 3}	2	2
5	74	A	50	{0, 2, 3}	.	.	A	50	{0, 2, 3}	3	3	A	50	{0, 1}	0	0	A	70	{0, 1, 2, 3}	3	3	A	0	{0, 2, 3}	2	2	B	70	{0, 1, 2, 3}	3	3
5	75	B	50	{0, 1}	.	.	B	60	{0, 1}	.	.	B	50	{0, 2, 3}	.	.	A	50	{0, 1, 2, 3}	.	.	B	50	{0, 1, 2, 3}	.	.	B	60	{0, 1, 2, 3}	.	.
5	76	B	70	{0, 1}	3	3	B	70	{0, 1}	3	3	B	50	{0, 2, 3}	3	3	A	66	{0, 1, 2, 3}	3	3	B	66	{0, 1, 2, 3}	3	3	B	66	{0, 1, 2, 3}	.	.
5	77	B	80	{0, 1}	0	0	B	85	{0, 1}	.	.	A	80	{0, 1}	.	.	A	65	{0, 1, 2, 3}	2	2	A	65	{0, 2, 3}	.	.	B	65	{0, 1, 2, 3}	2	2
5	78	B	50	{0, 1}	1	1	A	50	{0, 2, 3}	2	2	B	50	{0, 2, 3}	.	.	A	75	{0, 1, 2, 3}	2	2	B	50	{0, 1, 2, 3}	2	2	B	75	{0, 1, 2, 3}	2	2
6	79	B	50	{0, 1}	0	0	B	50	{0, 1}	0	0	A	50	{0, 1}	0	0	A	50	{0, 1, 2, 3}	0	0	B	50	{0, 1, 2, 3}	0	0	B	50	{0, 1, 2, 3}	0	0
6	80	A	90	{0, 2, 3}	2	2	A	98	{0, 2, 3}	2	2	B	85	{0, 2, 3}	2	2	A	70	{0, 1, 2, 3}	2	2	B	50	{0, 1, 2, 3}	1	1	B	95	{0, 1, 2, 3}	2	2
6	81	B	66	{0, 1}	1	1	B	66	{0, 1}	1	1	A	66	{0, 1}	1	1	A	50	{0, 1, 2, 3}	2	2	A	50	{0, 2, 3}	.	.	B	50	{0, 1, 2, 3}	2	2
6	82	A	75	{0, 2, 3}	2	2	A	60	{0, 2, 3}	2	2	B	55	{0, 2, 3}	2	2	A	75	{0, 1, 2, 3}	.	.	B	50	{0, 1, 2, 3}	2	2	B	75	{0, 1, 2, 3}	.	.
6	83	B	50	{0, 1}	.	.	B	50	{0, 1}	.	.	B	50	{0, 2, 3}	.	.	A	50	{0, 1, 2, 3}	.	.	B	50	{0, 1, 2, 3}	.	.	B	50	{0, 1, 2, 3}	.	.
6	84	B	67	{0, 1}	1	1	B	67	{0, 1}	1	1	A	67	{0, 1}	.	.	A	75	{0, 1, 2, 3}	2	2	B	67	{0, 1, 2, 3}	2	2	B	70	{0, 1, 2, 3}	2	1
6	85	B	66	{0, 1}	1	1	B	66	{0, 1}	1	1	A	66	{0, 1}	1	1	A	66	{0, 1, 2, 3}	1	1	B	66	{0, 1, 2, 3}	1	1	B	66	{0, 1, 2, 3}	1	1
6	86	A	0	{0, 2, 3}	2	2	A	60	{0, 2, 3}	2	2	A	50	{0, 1}	.	.	A	50	{0, 1, 2, 3}	.	.	A	50	{0, 2, 3}	.	.	B	50	{0, 1, 2, 3}	.	.
6	87	B	50	{0, 1}	.	.	A	50	{0, 2, 3}	.	.	B	50	{0, 2, 3}	.	.	A	70	{0, 1, 2, 3}	.	.	B	50	{0, 1, 2, 3}	.	.	B	70	{0, 1, 2, 3}	.	.
6	88	A	75	{0, 2, 3}	2	2	A	75	{0, 2, 3}	2	2	B	70	{0, 2, 3}	2	2	A	70	{0, 1, 2, 3}	2	2	B	50	{0, 1, 2, 3}	0	0	B	75	{0, 1, 2, 3}	.	1
6	89	B	67	{0, 1}	.	.	B	67	{0, 1}	.	.	A	50	{0, 1}	.	.	B	50	{0}	.	.	B	50	{0, 1, 2, 3}	.	.	A	60	{0}	.	.
6	90	A	50	{0, 2, 3}	2	2	A	50	{0, 2, 3}	0	0	B	50	{0, 2, 3}	2	2	A	70	{0, 1, 2, 3}	2	2	B	50	{0, 1, 2, 3}	0	0	B	60	{0, 1, 2, 3}	2	2
6	91	B	60	{0, 1}	.	.	B	60	{0, 1}	3	3	B	70	{0, 2, 3}	2	2	A	80	{0, 1, 2, 3}	3	3	B	70	{0, 1, 2, 3}	3	3	B	90	{0, 1, 2, 3}	3	3
6	92	B	50	{0, 1}	.	0	A	50	{0, 2, 3}	.	.	B	50	{0, 2, 3}	.	.	A	70	{0, 1, 2, 3}	.	.	A	65	{0, 2, 3}	2	2	B	75	{0, 1, 2, 3}	.	.
7	93	B	50	{0, 1}	3	3	A	50	{0, 2, 3}	3	3	B	70	{0, 2, 3}	2	2	A	50	{0, 1, 2, 3}	2	2	B	50	{0, 1, 2, 3}	2	2	B	80	{0, 1, 2, 3}	2	2
7	94	B	80	{0, 1}	1	1	A	75	{0, 2, 3}	2	2	B	50	{0, 2, 3}	.	.	A	90	{0, 1, 2, 3}	2	2	A	80	{0, 2, 3}	2	2	B	95	{0, 1, 2, 3}	2	2
7	95	B	75	{0, 1}	.	1	B	75	{0, 1}	.	.	A	75	{0, 1}	.	.	B	50	{0}	.	.	A	50	{0, 2, 3}	.	.	A	50	{0}	.	.
7	96	B	100	{0, 1}	1	1	B	75	{0, 1}	1	1	A	75	{0, 1}	0	0	A	55	{0, 1, 2, 3}	.	.	B	50	{0, 1, 2, 3}	.	.	B	57	{0, 1, 2, 3}	.	.
7	97	B	0	{0, 1}	1	1	B	50	{0, 1}	.	1	A	100	{0, 1}	1	1	A	50	{0, 1, 2, 3}	1	1	B	100	{0, 1, 2, 3}	2	1	A	80	{0}	.	.
7	98	B	60	{0, 1}	.	1	B	60	{0, 1}	1	1	A	60	{0, 1}	.	.	A	60	{0, 1, 2, 3}	2	2	A	50	{0, 2, 3}	2	2	B	50	{0, 1, 2, 3}	.	.
7	99	A	80	{0, 2, 3}	2	2	A	80	{0, 2, 3}	2	2	B	70	{0, 2, 3}	2	2	A	80	{0, 1, 2, 3}	2	2	B	50	{0, 1, 2, 3}	2	2	B	80	{0, 1, 2, 3}	2	2
7	100	A	80	{0, 2, 3}	2	2	A	80	{0, 2, 3}	2	2	B	80	{0, 2, 3}	2	2	A	85	{0, 1, 2, 3}	2	2	A	50	{0, 2, 3}	2	3	B	85	{0, 1, 2, 3}	2	2
7	101	A	75	{0, 2, 3}	2	2	A	90	{0, 2, 3}	2	2	B	80	{0, 2, 3}	2	2	A	95	{0, 1, 2, 3}	2	2	B	50	{0, 1, 2, 3}	2	2	B	80	{0, 1, 2, 3}	2	2
7	102	B	75	{0, 1}	.	.	B	75	{0, 1}	.	.	A	90	{0, 1}	.	.	A	0	{0, 1, 2, 3}	.	.	B	65	{0, 1, 2, 3}	.	.	B	100	{0, 1, 2, 3}	.	.
7	103	B	50	{0, 1}	.	.	A	80	{0, 2, 3}	.	.	A	80	{0, 1}	.	.	A	50	{0, 1, 2, 3}	.	.	A	75	{0, 2, 3}	.	.	B	75	{0, 1, 2, 3}	.	.
7	104	B	80	{0, 1}	.	.	A	80	{0, 2, 3}	.	.	A	75	{0, 1}	.	.	A	70	{0, 1, 2, 3}	.	.	B	85	{0, 1, 2, 3}	.	.	B	80	{0, 1, 2, 3}	.	.
7	105	A	80	{0, 2, 3}	2	2	A	80	{0, 2, 3}	.	.	B	80	{0, 2, 3}	.	.	A	80	{0, 1, 2, 3}	.	.	B	80	{0, 1, 2, 3}	.	.	B	80	{0, 1, 2, 3}	.	.
7	106	A	90	{0, 2, 3}	2	2	A	90	{0, 2, 3}	.	.	B	80	{0, 2, 3}	.	.	A	95	{0, 1, 2, 3}	2	2	B	50	{0, 1, 2, 3}	.	0	B	90	{0, 1, 2, 3}	2	1

Table 13: Individual data in the team communication experiment à la Anderson and Holt (1997), sessions 5-7

Subject ID	(H_t, s_t)	Message	Sugg. Decision	Confidence	Classification
27	BBAb	I don't know this is all probability so this is a guess	B	50	$\emptyset$
	BAABAb	I have no clue.	B	50	$\{0\}$
	BAAa	I have no clue.	A	50	$\{0\}$
51	BBAb	randomly choose my selection. i will choose B, u select A and computer will take decison randomly on behalf of us	B	50	$\{0\}$
	BAABAb	randomly choose my selection. i will choose B, u select A and computer will take decison randomly on behalf of us	B	50	$\{0\}$
	BAAa	randomly choose my selection. i will choose B, u select A and computer will take decison randomly on behalf of us	B	50	$\{0\}$
79	BAAAb	idk, just a guess	B	50	$\{0\}$
	ABAAAb	just a guess	B	50	$\{0\}$
	BBABa	just a guess	A	50	$\{0\}$

Table 14: Examples of messages classified as level-0

Subject ID	(H_t, s_t)	Message	Sugg. Decision	Confidence	Classification
58	BBAb	since the private draw is black, urn b is more likely to be the true urn.	B	50	$\{1\}$
71	BAAAb	The ball picked was black, right? So, there's a greater chance that it came from urn B than urn A	B	67	$\{1\}$
	ABAAAb	The ball picked was black. There are two black balls in urn B and only one black ball in urn A. Therefore, there's a 2/3 chance the ball came from urn B and a 1/3 chance it came from urn A.	B	66	$\{1\}$
	BBABa	Urn A contains two white balls and urn B contains one. Therefore, there's a 2/3 chance the ball came from urn A.	A	67	$\{1\}$
72	BAAAb	It would be foolish not to choose the urn with the greater proportion of black balls.	B	67	$\{1\}$
81	BAAAb	it doesnt matter what the other teams found there is still a higher probability that a black ball is chosen from the urb B therefore i would choose b	B	66	$\{1\}$
	ABAAAb	there is still a higher chance that the black ball is chosen (2/3 versus 1/3) therefore i would still choose B it doesnt matter what the other teams chose	B	66	$\{1\}$
	BBABa	it doesnt matter the order of sequence of the other teams i would still choose A since the probability of choosing a white ball is twice that of a black ball (2/3 vs. 1/3) therefore, i would still choose A	A	66	$\{1\}$
84	BAAAb	I don't see how the previous draws affect anything. We got black, it's 2/3 chance from B, 1/3 chance from A unless I'm completely confused.	B	67	$\{1\}$
	ABAAAb	2/3 chance that black ball is from B, 1/3 chance from A	B	67	$\{1\}$
85	BAAAb	there is a 66.7% chance that the ball is black. i'm not sure how the previous draws would affect our guess	B	66	$\{1\}$
	ABAAAb	I think it would be the best to stick with urn B since the probability is in favor of that.	B	66	$\{1\}$
	BBABa	I went with A because the chances it came from A are more likely that B.	A	66	$\{1\}$
97	BAAAb	If the other teams are mostly going with A, and we have a black ball, I think our best choice is B.	B	–	$\{1\}$
	ABAAAb	I think we have a better chance selecting urn B.	B	50	$\emptyset$
	BBABa	I think since we are drawing a white ball, our best chance is to select from urn A.	A	100	$\{1\}$
98	BAAAb	Urn B has the most black balls	B	60	$\emptyset$
	ABAAAb	Even though th teams before mostly picked A, I don't think that should have a bearing on the choice here; B still has the higher statistical probability of having a black ball	B	60	$\{1\}$

Table 15: Examples of messages classified as level-1

Subject ID	(H_t, s_t)	Message	Sugg. Decision	Confidence	Classification
1	BBAb	It is more likely that a black ball came from Urn B, and 2/3 of the other teams have made this same choice.	B	75	{2}
	BAABAb	A black ball is more likely to have come from Urn B. Not enough teams have gone before us to be affected by the fact that 3/5 of them selected A	B	60	{2}
	BAAa	A white ball is more likely to be from Urn A, and 2/3 of the previous teams have selected this urn as well	A	75	{2}
8	BBAb	Since 2 out of 3 of the previous teams chose urn B, and our draw corresponds to that choice, I'd say that urn b is a good bet.	B	70	{2}
	BAABAb	this is a tough one, but based on our draw, B is a good choice. Also, 2 out of 5 other teams chose B	B	55	{2}
	BAAa	2 out of 3 previous teams chose A and our draw corresponds to A as well.	A	70	{2}
56	BAAa	I'm choosing this based on the other teams and our draw. Also, I think I forgot to press enter for the last message... So for the last one, I made that decision based on the other teams draws not ours.	A	75	{2}
80	BAAAb	3 out of 4 previous teams are picking A so that is more likely they probably decide to pick A because see a white ball drawn	A	90	{2}
	ABAAAb	there are 5 teams before us and 4 of them choose A, so i am confident that it is A most of them must see a white ball from A	A	98	{2}
	BBABa	4 teams before us and 3 of them picked B so I would pick B	B	85	{2}
82	BAAAb	Most people must have seen a white draw for there to be predominatley A being chosen.	A	75	{2}
	ABAAAb	It seems that most people saw a white ball in their private draw.	A	60	{2}
	BBABa	Hard to say, still though most people seem to have seen a black ball.	B	55	{2}
86	BAAAb	Only because 3 teams chose A	A	–	{2}
	ABAAAb	4 out of 5 say A	A	60	{2}
88	BAAAb	3 white balls in a row from the black urn is only a (1/3)times(1/3) times (1/3) chance - therefore, very low. I suspect that people chose urn A if they got a white ball	A	75	{2}
	ABAAAb	I assume teams chose urn A if they drew white. The chance of getting 4 whites from Urn B is very low therefore I believe it is urn A. Even though we drew a black, it's still greater probability that it is A	A	75	{2}
	BBABa	It appears from previous teams choices that 3 of the 4 balls were black. Therefore, even though we drew a white ball, it's more likely that our urn is B	B	70	{2}

Table 16: Examples of messages classified as level-2 (continued next table)

Subject ID	(H_t, s_t)	Message	Sugg. Decision	Confidence	Classification
100	BAAAb	since 3 of the teams before us chose A, we might assume that they got a white ball. we may have just gotten a black ball by chance from urn A	A	80	{2}
	ABAAAb	since most teams before us have chosen A, we may assume that most of them got the white ball from their drawing. using probability, it's more likely that we got a black by with 1/3 chance than the previous teams getting a white ball from urn B with much lower probability	A	80	{2}
	BBABa	we can assume that team 1 got the black ball, team 2 probably also got the black ball, and team 3 probably got the white ball. probability wise, it's more likely that its from urn B	B	80	{2}
101	BAAAb	Okay so the teams before us chose A more frequently, which means they private draw must have been white, with one black. We got a black one, but since its a 50% chance that both urns could be the true urn, we can attribute ours to simple probability Thus, we have a 75% chance of it being A More frequency-better chance it comes from one urn than the other	A	75	{2}
	ABAAAb	Frequency of a is greater than B, thus, it's a good chance its urn A because although we picked black 4 whites to two blacks shows a 2 to 1 ratio meaning that for everty 3 balls, 2 are white and one is black this reflects the probabilty of urn A	A	90	{2}
	BBABa	The first team must have seen a black ball and chose B second team probably got black as well so chose B third team probably got white and decided to go with A 4th team got black again and went with B Thus if we assume all the choices before are rational then the choices are more black than white thus B MUST be the rationally right answer as in if you don't pick B, it would be quite irrational	B	80	{2}
106	BAAAb	Given that we were shown a black ball, there is a 2/3 chance that it is in B. However, the other teams chose A, which means that they were given a white ball, meaning that there is a greater chance of A	A	90	{2}

Table 17: Examples of messages classified as level-2

Subject ID	(H_t, s_t)	Message	Sugg. Decision	Confidence	Classification
15	BBAb	The first team definitely saw a black ball, they went in blind and chose B The second team must have also saw a black ball, and corroborated with what team 1 saw, the 3rd team saw a white ball, which contradicted what the previous team saw It is possible that the 2nd team saw white, but just went along with team 1 so, I would say, out of 4 draws, there's certainly 1 black, 1 maybe black, and 1 white our draw is a black one I am pretty confident that we should vote on urn B	B	100	$\{3\}$
	BAABAb	team 1 saw black team 2 saw white, for sure team 3 saw white for sure, since it could be influenced by either team 1 or 2, but chose to go with white team 4 saw black team 5 saw white team 5 has 2 guaranteed black before it, and 2 guaranteed white before it, but chose white, so probably got white our draw is black, there's 3 white, 2 black before us I'd say its 50 50, let's go with black, urn B since we're sure about what we saw	B	70	$\{2, 3\}$
	BAAa	same logic, lets go with A since majority came out as white	A	100	$\{2\}$
46	BBAb	Teams 1 and 2 probably saw a black ball, and team 3 a white ball 3 blacks and 1 white is much more likely for urn b	B	75	$\{2\}$
	BAABAb	Team 1 probably drew a black ball, as they have no information Team 2 must have drawn white, if they chose A despite team 1's decision to choose black Team 3 probably drew white as well, as the previous teams decisions are split so they decide on their own team 4 must have drawn black, as them drawing white would be totally agains ttheir choice to choose B It's 50-50	A	50	$\{2, 3\}$
	BAAa	draws are almost purely based on the individual ball draw, as the disparity of information in the teams previous theirs indicates 3 white balls to 1 black ball, basically	A	75	$\emptyset$
74	ABAAAb	It seems the other teams saw more white balls though it is possible the third and the fourth team actuall saw black balls but they had to make a random choince.	A	50	$\{3\}$
	BBABa	team 1 and team 4 drew black balls and team 3 and team 5 drew white. therefore it is a 50/50 chance.	A	50	$\{0\}$

Table 18: Examples of messages classified as level-3 (continued next table)

Subject ID	(H_t, s_t)	Message	Sugg. Decision	Confidence	Classification
93	BAAAb	Well, from our observations it seems that we would be choosing Urn B. However, if we look at previous guesses, a desire may be made to choose Urn A from belieivng that other people must be drawing white balls. We have to consider the fact that other teams are also being influenced by previous decisions. For example, team 4 may have drawn a black ball and wanted to choose Urn B, but didn't in order to match the previous people. Therefore, I still think we should choose Urn B.	B	50	{3}
	ABAAAb	Even though we drew a black ball, feel that for 80% of people to have chosen A, this may indicate that white is the majority. However, I'm not sure at all, due to previous people also being influenced by previous decisions. We know that the first decision is not based off of previous decisions, therefore we can assume that the first person drew a white ball. The second person contrasted the first person so we can assume they went by their ball and chosen B. The third person was going off of a balanced group where half of the people had chosen whtie and half of the people had chosen black, therefore their decision is valid too. After that it was based on how people analyzed it somewhat. I believe the fourth person would still have stayed with what they got so I think A is better.	A	50	{3}
	BBABa	We know that the first person would stick to what ball they drew, therefore it can be decided to be B. The second person would have no reason to follow the first person since its just one statistic so they are probably accurate too. The third person needed a reason to deviate so they were definitely A. The last person probably would have stayed with what they got so I think it was B, regardless of our drawing.	B	70	{2}

Table 19: Examples of messages classified as level-3

Type	Frequency	Fraction (%)
(0)	17	40.5
1	12	28.6
2	10	23.8
3	2	4.8
B	1	2.4
Total	42	100.00

(a) Classification by lowest significant SSD

Type	Frequency	Fraction (%)
(0)	17.00	40.5
1	12.00	28.6
2	10.00	23.8
3	1.89	4.5
B	1.11	2.6
Total	42	100.00

(b) Mixture classification

Table 20: Type overview for 42 subjects in decisions 1-3

Type	Frequency	Fraction (%)
(0)	18	42.9
1	9	21.4
2	8	19.0
3	0	0.0
B	7	16.7
Total	42	100.00

(a) Classification by lowest significant SSD

Type	Frequency	Fraction (%)
(0)	18.00	42.9
1	9.00	21.4
2	8.55	20.4
3	0.33	0.8
B	6.11	14.6
Total	42	100.00

(b) Mixture classification

Table 21: Type overview for 42 subjects in decisions 4-6

		Highest level (⌀ 2.11)					
		0	1	2	3	NA	Total
	0	1	0	1	2	0	4
Lowest level	1		2	2	0	0	4
(⌀ 1.68)	2			15	2	0	17
	3				3	0	3
	NA					14	14
	Total	1	2	18	7	14	42

Table 22: Level classification in second 3 decisions

Subject ID	Decision 1 (AAB)			Decision 2 (BBAA)			Decision 3 (AAABA)			Decision 4 (BBB)			Decision 5 (AABA)			Decision 6 (BABA)			Fingerprint Analysis Decisions 1-6				Cluster
	Sugg. dec.	RA 1	RA 2	Sugg. dec.	RA 1	RA 2	Sugg. dec.	RA 1	RA 2	Sugg. dec.	RA 1	RA 2	Sugg. dec.	RA 1	RA 2	Sugg. dec.	RA 1	RA 2	Level-1	Level-2	Level-3	Bayesian	
1	2	.	.	5	.	.	0	.	.	-1	.	.	0	1	1	-1	.	.	1.00	0.00	0.00	0.00	
2	8	.	.	2	.	.	2	.	.	4.32	.	.	7.31	.	.	-1.11	.	.	0.00	0.00	0.00	0.00	A
3	0	1	1	0	.	.	0	.	.	0	.	.	0	.	.	0	.	.	1.00	0.00	0.00	0.00	
4	4	0	2	0	2	1	-8	2	2	2	2	0	-10	2	2	0	2	2	0.00	1.00	0.00	0.00	B
5	5	.	.	0	.	.	-10	.	.	10	2	2	-10	.	.	0	.	.	0.00	1.00	0.00	0.00	C
6	0	.	.	0	.	.	-5	2	2	7	.	.	0	.	.	3	.	.	0.27	0.32	0.00	0.41	A
7	-2	.	.	-2	3	3	-4	.	.	7	.	.	3	3	3	0	.	.	0.27	0.00	0.32	0.41	A
8	-4	2	2	-3	.	.	-10	.	.	10	.	.	0	.	.	-10	.	.	0.00	0.00	1.00	0.00	
9	-2	2	3	0	2	2	-3	3	3	5	2	2	-2	2	2	0	2	2	0.53	0.12	0.18	0.17	A
10	5	0	0	0	.	.	0	.	.	0	0	0	1.11	.	.	-2	.	.	1.00	0.00	0.00	0.00	
11	0	0	1	0	1	1	0	1	1	-10	.	.	0	.	.	0	.	.	0.00	0.00	0.00	0.00	D
12	1	2	1	-2	3	3	-1	.	.	7	2	2	-1	3	3	1	0	0	0.33	0.00	0.16	0.51	A
13	-5	2	2	0	.	.	-8	.	.	4	2	2	-4	.	.	0	.	.	0.00	0.84	0.16	0.00	A
14	4	0	0	4	0	0	2	2	2	9	2	2	2	.	0	4	0	0	0.00	0.00	0.00	0.00	A
15	-8	.	0	-8	.	.	-8	.	.	0	2	2	-9	2	0	-1	.	.	0.00	1.00	0.00	0.00	E
16	-2	2	0	0	2	2	-10	2	2	8	2	0	-1	2	2	0	2	2	0.00	0.67	0.16	0.17	A
17	-1	.	.	0	2	2	-8	2	2	0	2	1	-8	2	2	0	2	2	0.00	1.00	0.00	0.00	B
18	0	.	.	-10	.	.	-5	.	.	0	.	.	-6	.	.	-6	.	.	0.00	0.00	0.59	0.41	E
19	3	.	.	5	.	.	6	.	.	-2	.	.	3	.	.	2	.	.	0.00	0.00	0.00	0.00	D
20	-10	.	.	0	2	2	-10	2	2	8	2	2	-10	2	2	-1	2	0	0.00	1.00	0.00	0.00	F
21	4	2	2	3	.	.	2	2	2	9	2	2	3	2	2	7	.	.	0.00	0.00	0.00	0.00	A
22	-5	2	2	0	2	2	-8	2	2	9	2	2	-2	2	2	-2	3	3	0.00	0.70	0.21	0.09	A
23	1	2	2	0	2	2	-4	2	2	8	2	2	-1	2	2	0	2	2	0.10	0.14	0.19	0.57	A
24	-1	.	.	1	.	.	-1	.	.	-3	.	.	1	.	.	-1	.	.	1.00	0.00	0.00	0.00	
25	7	0	2	7	0	2	7	.	.	10	.	.	0	.	.	-9	.	.	0.00	0.00	0.00	0.00	
26	-2	3	3	-3	3	3	-5	3	3	10	.	.	-5	3	3	-4	2	3	0.00	0.02	0.80	0.19	
27	0	1	1	0	1	1	-9	.	.	8	2	2	-0.5	2	1	0	.	.	0.00	0.44	0.21	0.35	A
28	4	.	.	6	.	.	-10	2	2	8	2	2	-10	2	2	5	.	0	0.00	1.00	0.00	0.00	C
29	3	2	0	2	2	2	-1	2	2	3.5	3	3	-1	2	2	2.5	0	0	1.00	0.00	0.00	0.00	A
30	-8	2	2	0	2	2	-10	2	2	10	2	2	-8	2	2	0	2	2	0.00	1.00	0.00	0.00	F
31	0	.	.	0	2	1	-10	2	2	10	2	2	-10	2	2	0	1	1	0.00	1.00	0.00	0.00	C
32	0	.	.	-10	.	.	-10	3	2	5	.	.	-10	.	.	0	.	.	0.00	0.59	0.00	0.41	E
33	3.1	.	.	0	.	.	1	.	.	4.25	.	.	3.1	.	.	0	.	.	1.00	0.00	0.00	0.00	A
34	0	.	.	8	.	.	7	.	.	10	.	.	-8	.	.	0	.	.	0.00	0.00	0.00	0.00	A
35	8	0	2	1	2	0	3	2	2	-10	2	2	3	0	.	0	1	1	0.00	0.00	0.00	0.00	D
36	0	0	1	0	.	1	0	.	.	0	1	1	0	.	.	0	.	.	1.00	0.00	0.00	0.00	
37	-2	3	3	0	.	2	-4	2	2	2	2	.	-1.5	2	0	1	.	.	0.91	0.03	0.03	0.03	A
38	2	2	0	4	.	.	0	2	2	8	2	2	-4	2	2	0	2	2	0.00	0.00	0.00	0.00	A
39	0	2	2	-1	3	3	-2	.	0	2	2	3	-1.5	3	3	-0.5	.	.	0.99	0.00	0.00	0.00	A
40	4	.	.	6	.	.	3	2	0	-7	2	2	5	.	.	-3	.	.	0.00	0.00	0.00	0.00	D
41	-8	2	2	-8	.	.	-10	.	.	6	.	.	-10	.	.	-4	.	.	0.00	0.76	0.24	0.00	E
42	-1	2	2	0.5	2	0	-7	2	2	9	2	2	-3	3	3	0	.	.	0.00	0.61	0.15	0.24	A

Table 23: Individual data in the team communication experiment à la ÇK

Subject ID	(H_t)	Message	Sugg. Decision	Classification
26	AAB	Hello. I propose that we assume that all participants before decided rationally. The sequence suggests that at action 1 $y > 0$. Expected value would so far be 5. Was this considered in action 2, the decision for all numbers greater than -5 would be A too. Expected value (5+2,5)/2=3,75. Was this considered in action 3, b was chosen only for strongly negative numbers. This is why I would propose b as well, that is to say a negative number.	-2	$\{3\}$
	BBAA	Decider 3 and 4 had more information, expected value consequently positive. Would therefore tend to A. At each $x > -3$ jump to A.	-3	$\{3\}$
	AAABA	Last participant has again most information. Apparently participant 4 had a strongly negative number. Still, I would be to chose A at negative numbers up to -5.	-5	$\{3\}$
9	AAB	Purely mathematically assumed: ;-) I am fourth and have information A,A,B. The first should have set 0 as marginal point to have a 50 50 chance. The second can thus be sure that the first had a positive number. He thus sets his point a bit higher at 2. With two A decisions the third can assume that both previous numbers were positive, so he can courageously bet on A with -2. This might be too courageous, he has maybe a negative number. Long speech, short meaning A with still high risk -2	-2	$\{2, 3\}$
	BBAA	I am fifth with B,B,A,A as information. This give a 50:50 chance so I would guess a 0.	0	$\{2\}$
	AAABA	6. with information A,A,A,B,A assume: 1. chooses 0 $\rightarrow$ positive number, 2. chooses 0 as well $\rightarrow$ positive number, 3. becomes more courageous chooses -2 $\rightarrow$ possibly small negative number but in sum still everything positive, 4. still more courageous -3, now we would be back in the area around zero, so 5 chooses 0 again $\rightarrow$ has a positive number. Means we should be more courageous :D hehe I think this makes little sense	-3	$\{3\}$

Table 24: Examples of messages potentially reflecting Bayesian reasoning

A. The model in the Eyster and Rabin framework

The social learning environment used in Eyster and Rabin (2010) has continuous signals $s_t \in [0,1]$ and actions $a_t \in [0,1]$. The signal obeys $s = Pr[\omega = 1|s]$. Given an assumed payoff function $g_t(a_t;\omega) = -(a_t - \omega)^2$, player t maximizes her expected payoff by setting $a_t = E[\omega|s_t, H_t]$.

Level-0 play is random and not connected to the private signal. The level-0 strategy is hence

$$\sigma^0(s_t, H_t) \sim U[0,1].$$

In the notation of Eyster and Rabin (2010), random play over $[0,1]$ leads to an expectation of the log odds of $\ln\left(\frac{a_t}{1-a_t}\right) = 0$.

The best response to uninformative preceding play is to follow the private information. A level-1 player t thus follows strategy

$$\sigma^1(s_t, H_t) = s_t.$$

Due to the equivalence of level-1 actions and signals, a level-2 player combines the private information with predecessors' actions. The strategy $\sigma^2(s_t, H_t)$ is implicitly defined as

$$\ln\left(\frac{\sigma^2}{1-\sigma^2}\right) = \left[\sum_{i=1}^{t-1} \ln\left(\frac{a_i}{1-a_i}\right)\right] + \ln\left(\frac{s_t}{1-s_t}\right).$$

Given $\sigma^1(s_t, H_t)$, the actions of predecessors a_i are taken at face value and processed as if they reflected the private information. This level-2 specification is Eyster and Rabin's (2010, p. 230) specification of BRTNI play. If all players t in a sequence follow this strategy and take previous actions at face value, the early signals will be overweighted as follows

$$\ln\left(\frac{\sigma^2}{1-\sigma^2}\right) = \left[\sum_{i=1}^{t-1} 2^{t-1-i} \ln\left(\frac{s_i}{1-s_i}\right)\right] + \ln\left(\frac{s_t}{1-s_t}\right).$$

The best response to such behavior is to undo the overweighting, back out the true signals of predecessors, and combine them with the own signal. Given the public history and the informative action space, the individual signal s_i, $\forall i < t$ can be backed out perfectly. Level-3 ends up playing like a rational Bayesian player that combines all information in strategy $\sigma^3(s_t, H_t)$ so that

$$\ln\left(\frac{\sigma^3}{1-\sigma^3}\right) = \sum_{i=1}^{t} \ln\left(\frac{s_i}{1-s_i}\right)$$

Of course, this hinges on the assumption that the population belief of this level-3 player is indeed correct and would deliver different results if the predecessors were actually level-0 or level-1 players. In this rich-information setup it is always possible to back out the private signals of non-level-0 players. The prerequisite for doing this correctly is to have a correct population belief of player i, that is, to know how player i used her signal when choosing her action. Then, the signal can be backed out by simply observing her action and all prior actions. This framework features the possibility to back out any signal that was used in a deterministic strategy. In other words, any level-k play, $k > 0$, is fully informative given the correct population belief about it.

Like in the Anderson and Holt setup, higher levels of reasoning are conceivable but are only interesting because of their population belief. In particular, level-4 might believe that players' actions simply summarize the previous signals in an optimal fashion (level-3). Beyond this self-awareness of level-3, no further insights are obtained when considering higher levels.

B. Experiment instructions (AH framework)

Welcome to the experiment!

Introduction

You are about to participate in an experiment in team decision making. The experiment is funded by Cornell University. Please follow the instructions carefully.

In addition to the participation fee of $5, you may earn a considerable additional amount of money. Your decisions determine the additional amount. You will be instructed in detail how your earnings depend on your decisions. All that you earn is yours to keep, and will be paid to you in private, in cash, after today's session.

It is important to us that you remain silent and do not look at other people's screens. If you have any questions or need assistance of any kind, please raise your hand, and an experimenter will come to you. If you talk, shout out loud, etc., you will be asked to leave. Thank you.

The experiment consists of a test round and 2 parts. Part I is designed as a warm-up with 3 trivia questions. Part II consists of 2 rounds of 3 decisions each. In all 6 decisions your task will be identical and success is identically rewarded.

Since this is a team experiment, you will be randomly matched with another participant in this room, to form a team that plays as one entity. Your teammate will change every round, so please do not assume content of previous communication to be known by your new partner. The way you interact as a team to take decisions will be the same throughout all rounds.

Now, let me explain how your **Team's Action** is determined. In fact, both your teammate and you will enter a **Final Decision** individually and the computer will choose randomly which one of your two final decisions counts as your team's action. The probability that your teammate's final decision is chosen is equal to the probability that your final decision will be chosen (i. e. your chances are 50:50). However, you have the possibility to influence your partner's final decision in the following way: Before you enter your final decision, you can propose to your partner a **Suggested Decision** and send him one and only one text **Message**. *Note that this message is your only chance to convince your partner of the reasoning behind your suggested decision. Therefore, use the message to explain your suggested decision to your teammate.* After you finish entering your suggested decision and your message, these will be shown to your teammate. Simultaneously, you will receive your partner's suggested decision and message. Both of you will then make your final decision. As outlined above, once you both enter your final decision, the computer chooses randomly one of your final decisions as your team's action.

If you have any questions at this point, please raise your hand. In order for you to get familiar with the messaging system, you will now try it out in a **Test Period**. Please turn the page for further instructions.

Test period

A participant in this room is now randomly chosen to be your teammate. The **Test Period** has two rounds, with communication in each round. Since this is only a test, your earnings will not depend on anything that happens now. In both test rounds you will need to send and receive pieces of information. The information consists of the answer to a question and one given phrase. After the successful exchange, you will enter the number again. This way, the communication structure is identical to the one in the experiment rounds.

The messenger allows **Messages** of any size. However, you have to enter the message line by line since the input space is only one line. Within this line you can delete by using the usual "Backspace" button of your keyboard. By pressing "Enter" on the keyboard, you add the written sentence to the message. Please note that only added sentences will be sent and seen by your partner. *The words in the blue input line will* **not** *be sent*. You can always delete previously added sentences by clicking the "Clear Input" button. The number of lines you send is not limited. You can therefore send messages of any length. You finally send the message to your partner by clicking the "Send Message" button.

If you have any questions at this point, please raise your hand. When you are ready, please click the "Ready" button to start the **Test Period**.

Experiment - Part II

You are about to start **Part II** of the experiment.[21] This part consists of 2 periods of 3 decisions each.[22] In each period you will team up with a different, randomly chosen participant. Therefore you will make 3 urn choices with a given partner.

The 6 scenarios for your urn choices are taken from another experiment. They are randomly chosen and simply allow me to confront you with more scenarios than if we played it out ourselves. I will now explain the setup to you in the same fashion the participants of the original experiment were instructed. The original experiment was not done by teams, but I will explain it with teams for simplicity.

In each urn choice your task is the following:

Your team is asked to **predict** from which **randomly chosen urn** a ball was drawn. It is equally likely that urn A or urn B will be the true urn. That is, there is a 50 percent chance that Urn A is the true urn, and a 50 percent chance that Urn B is the true urn. **Urn A** contains **2 white** balls, and **1 black** ball. **Urn B** contains **1 white** ball and **2 black** balls. Therefore, there is a $\frac{2}{3}$ chance that a white ball comes from urn A and a $\frac{2}{3}$ chance that a black ball comes from urn B.

To help determine which urn is the right one in a given scenario, you will see one ball, drawn at random, from the urn. **Only** your team will see the outcome of this private draw. And your team will only see this **one draw**, you will get the same information as your teammate. After each draw, the ball is returned to the container before making the next draw. Therefore, each team will have **one private draw**, with the **ball being replaced** after each draw. This way each team draws from an urn that contains 3 balls in total.

[21] The experiment consists of two parts, part I being independent of part II.

[22] Table 25 was shown to give an overview.

When it is your turn to decide, bold text in the top of your screen will give you information about your draw. If the ball the computer has randomly drawn for you is white, your window will read, "**Your team's private draw from the urn is WHITE.**" Your window will read, "**Your team's private draw from the urn is BLACK.**", if the ball the computer has randomly drawn for you is black.

In each scenario, **decisions are made sequentially**, i.e. one team after the other. The order in which teams decide in a given scenario is determined randomly. Once the first team has agreed on its action based on its private draw, the second team will be asked to make its decision. The members of team 2 will see their private draw, and **also** which urn was chosen by the first team. Until the end of the round, all following teams will be **informed about the chosen urns of all earlier teams**, but not about the other teams' draw from the urn.

In each round, you see the three scenarios and will write down your **Suggested Decision** and **Message** for each of them. After that you see your partner's **Suggested Decision** and **Message** for each of the three scenarios and will make your **Final Decisions**. Finally, you will be informed about your **Team's Action** and the true urn. The order of events is illustrated in the table below. For each of the 6 scenarios, if your team's action and the chosen urn coincide, you will individually earn \$1.00 (your teammate will get \$1.00 as well).

As described earlier, you will send your teammate a **Suggested Decision** and a **Message**. Remember to explain in the message your reasoning behind your suggested decision. *(And note again that the words in the blue input line will* **not** *be sent. Press "Enter" to add them to the message.)* After this information is exchanged, both of you enter your **Final Decision**, from which the computer randomly chooses the **Team's Action**.

As part of the communication, you can quantify your **Confidence** for each choice you make. You can put numbers between 50% and 100%, where 50% implies that you think both urns are equally likely and higher numbers reflect your higher confidence in your choice, up to certainty (100%).

Let me summarize the main points: (1) In each scenario, it is equally likely that urn A or urn B is the true urn. (2) There is a $\frac{2}{3}$ chance that a white ball came from urn A and a $\frac{2}{3}$ chance that a black ball came from urn B. (3) In the knowledge of the previous teams' actions and your draw, choose either urn A or urn B. (4) Like you, other teams saw one draw from the true urn and their predecessors' urn choice. (5) If your team chooses the true urn, you will earn \$1.00.

If you have any questions at this point, please raise your hand. When you click the "Ready" button, you will start the first round of the experiment.

C. Classification instructions (AH framework)

In the following I will describe the classification process for the analysis of an experiment. Subjects play a game of incomplete information. Their reasoning will be classified along the lines of a model of level of reasoning, which will be laid out in the first section. It is set up in analogy to the level-k model in complete information settings as introduced by Nagel (1995) and Camerer et al. (2004).

Follow the instructions of this booklet. Read them entirely to get an overview and then start the classification. The game is the social learning experiment as implemented by Anderson and Holt (1997), which I assume you are fully familiar with. The subjects were put in 6 different situations that occurred in the original study by Anderson and Holt (1997) and decided in a team about their action, which would be remunerated according to the true urn in Anderson and Holt (1997).

The model will introduce four different ways of reasoning in the context of this game. The aim of the classification is to take a look at the 6 messages per player and connect him/her with one or

	Round	Urn choice	Action
↓	1	1	Suggested Decision
↓		2	
↓		3	
↓		1	Final Decision
↓		2	
↓		3	
↓		1	Get result
↓		2	
↓		3	
↓	2	4	Suggested Decision
↓		5	
↓		6	
↓		4	Final Decision
↓		5	
↓		6	
↓		4	Get result
↓		5	
↓		6	

Table 25: Order of events in Part II.

more types. Details will be explained below. Please limit yourself to making inferences only from what can clearly be derived from the message stated, i.e. do not try to think about what the player *might have thought*.

Reasoning types in the model

In the context of social learning, reasoning types differ in the ways they process private information (in form of draws from an urn) and public information (in form of actions of predecessors). A natural model of level of reasoning starts out with a random level-0 type and builds a hierarchy based on the number of iterated best responses. In this game, this hierarchy turns out as follows:

Level-0 Randomizing between A and B, irrespective of own private signal.

Level-1 If others' signal is not informative, like level-0 random play, then the best response is to get informed by the own signal. Therefore, level-1 players are following their own private signal.

Level-2 Since level-1 play is fully informative about the private signals, the best response in light of this is to follow the majority if it is ahead by more than one signal difference. A level-2 player will thus only follow his signal if the previous actions were split equally between the two urns or one urn was chosen just one time more often.

Level-3 Since level-2 play is as informative as Bayesian play, level-3 play, a best-response, is like original Bayesian play. This implies that the play and beliefs of level-3 players are identical to play and beliefs in the Bayesian equilibrium if everybody is level-3.

Starting from this model, I will now explain some expected message contents for the individual levels of reasoning.

Level-0 **Random play** should result from messages that clearly do not understand the nature of the game or put reasons for play that are orthogonal to any reasonable inference using private signal and previous actions.

Level-1 Level-1 is characterized by **disregarding others' signal**. In the message, this can be an open disregard of others' actions or an emphasis on the unambiguity of the own signal.

Level-2 A level-2 players with a degenerate population belief (as introduced in the model above) implies that all **others' actions are taken at face value** like a signal. It follows that the ratio of As vs. Bs in the previous actions might weighted against the 2/3 chance of the own signal to be correct. In any case, a level-2 player **never engages in a differentiation of individual predecessors**. If everybody is assumed to play their signal, this implies that others are regarded as a homogeneous crowd that simply differs in the signals they received not in their position in the sequence. Alternatively to a degenerate population belief, a level-2 player might think that some players played random, like level-0. Then, the signals are not taken at full face value and probably the own signal is valued more than the observed actions.

Level-3 heterogeneity Under a degenerate population belief, only level-3 players **distinguish individual predecessors** in the extent to which their actions reflect their private signals. It is therefore a characteristic of level-3 players to differentiate informative and uninformative observed actions depending on the position of predecessors in the sequence. One example is that a level-3 player will note in a history of AAAA that the later players might just have followed the majority, he therefore does not infer those players signal with certaint. Put starkly, observing BBBA implies that a level-3 player rules out the fourth player to be level-2, an observation only level-3 will make.

The data

You will see different situations in which the players decided to choose urn A or urn B.

Each situation is explained in brackets by the information available to the subject. The capital letters indicate previous players' actions. The small letters indicate the private signal received by the player at hand. In addition to the sent message, you see the suggested decision, which is 1 for urn A and 2 for urn B. The communication was structured in the sense that it gave the players the possibility to indicate the confidence in the own urn choice. This number is given to you as well.

The classification

I would like you to classify messages into **one of the 4 levels** described above. Indicate the closest type under 'level' in the according space.

Indicate whether the population belief is **degenerate** on the next lower level or **non-degenerate**. If it is non-degenerate, denote which types (in terms of levels) are assumed to be present.

In addition to this classification, I ask you to indicate **any difference** between the observed reasoning and the type given. This might include differences in the belief about when others start imitating others' actions, the weight the public information receives compared to the private information, the way the population updating is done by a level-3 player, etc. Also, if you think the potential level could be more than one, indicate your considerations here.

D. Experiment instructions (ÇK framework)

The experiment instructions in German combine the communication instructions (see Introduction and Test period above) with instructions from Çelen and Kariv (2004) translated to German and adapted to the team setting. A translation is available upon request.

E. Classification instructions (ÇK framework)

In the following I will describe the classification process for the analysis of an experiment. Subjects play a game of incomplete information. Their reasoning will be classified along the lines of a model of level of reasoning, which will be laid out in the first section. It is set up in analogy to the level-k model in complete information settings as introduced by Nagel (1995) and Camerer et al. (2004).

Follow the instructions of this booklet. Read them entirely to get an overview and then start the classification. The game is the social learning experiment as implemented by Çelen and Kariv (2004), which I assume you are fully familiar with. The subjects in my experiment were put in 6 different situations that occurred in the original study by Çelen and Kariv (2004) and decided in a team about their action, which would be remunerated according to the true signal in Çelen and Kariv (2004).

The model will introduce five different ways of reasoning in the context of this game. The aim of the classification is to take a look at the 6 messages per player and connect him/her with one or more types. Details will be explained below. Please limit yourself to making inferences only from what can clearly be derived from the message stated, i.e. do not try to think about what the player *might have thought*.

Reasoning types in the model

In the context of social learning, reasoning types differ in the ways they process private information (in form of draws from an urn) and public information (in form of actions of predecessors). Çelen and Kariv (2004) analyse Bayesian behavior, which is fully rational and believes others to be fully rational. Please consult the paper to get a thorough understanding of this kind of reasoning in this context.

A natural model of level of reasoning starts out with a random level-0 type and builds a hierarchy based on the number of iterated best responses. In this game, this hierarchy turns out as follows:

Level-0 Playing a random threshold on the action space.

Level-1 If others' signals are not informative, like level-0 random play, then the best response is to get informed by the own signal. Therefore, level-1 players are following their own private signal and set a threshold of 0.

Level-2 With a threshold of 0, the observed action of a predecessor indicates whether the signal was in $[-10, 0)$ or $[0, 10]$. With this information, the best response is to add 5 or subtract 5 for each previous action A or B in the expected value of the sum of the signals. A level-2 player will thus expect a sum of signals of 0 if the history has equally many A's as B's and, e.g., 5 if one more B's than A's have been observed and 10 if two more B's than A's have been observed. The threshold will be set as the negative of the expected sum.

Level-3 Thresholds of -10 or 10 lead to uninformative play since the action is decided irrespective of the signal. Level-3 players hence only use information from level-2 players that are not cascading, i.e. play -5, 0, or 5.

Level-4 Level-4 players only use information from level-3 players that are not cascading. Although level-4 play is distinct from level-3, we will not pursue this or higher levels in our classification.

Starting from this model, I will now explain some expected message contents for the individual levels of reasoning.

Level-0 **Random play** should result from messages that clearly do not understand the nature of the game or put reasons for play that are orthogonal to any reasonable inference using private signal and previous actions.

Level-1 Level-1 is characterized by **disregarding others' action**. In the message, this can be an open disregard of others' actions or an emphasis on the *un*ambiguity of the own signal.

Level-2 A level-2 player with a degenerate population belief (as introduced in the model above) implies that all **others' actions are equally informative** and indicative of the sign of the signal and thus each one changes the expected sum of signals by $+5$ or -5. It follows that the relative number of As vs. Bs in the previous actions determine the threshold. In any case, a level-2 player **never engages in a differentiation of individual predecessors**. If everybody is assumed to play their signal, this implies that others are regarded as a homogeneous crowd that simply differs in the signals they received not in their position in the sequence.

Level-3 heterogeneity Level-3 players are the first in the hierarchy to **distinguish individual predecessors** in the extent to which their actions are informative about their private signals. In particular, level-2 actions might be completely uninformative if they result from a threshold of 10 or -10. At the same time, actions that result from a signal below a threshold of -5 or above a threshold of 5 induce a strong change in the expected sum of signals. Therefore, sum actions might have a stronger influence on the level-3's threshold than others. It is therefore a characteristic of level-3 players to differentiate more or less informative observed actions depending on the position of predecessors in the sequence and the likely threshold they had set. One example is that a level-3 player will note in a history of AB that the second player's signal must have been very low since a threshold of -5 was still undercut. This informativeness changes to uninformativeness when the threshold is expected to be 10 or -10. Then, the action is uninformative.

Bayesian More than level-3 players, Bayesian players realize that *every* predecessor has most likely set a different threshold. This way, particular later actions in the sequence can have a strong influence on the expected sum of signals and thus the set threshold.

The data

You will see different situations in which the players observed histories of A's and B's and set a Suggested Decision between -10 and 10.

Each situation is described by the information available to the subject. In addition to the sent message, you see the suggested decision.

The classification

I would like you to classify messages into **one of the 5 levels** described above. Indicate the closest type under 'level' in the according space.

Indicate whether the population belief is **degenerate** on the next lower level or **non-degenerate**. If it is non-degenerate, denote which types (in terms of levels) are assumed to be present.

In addition to this classification, I ask you to indicate **any difference** between the observed reasoning and the type given. This might include differences in the belief about what threshold is assumed to be behind others actions, the way the population updating is done by a level-3 player, etc. Also, if you think the potential level could be more than one, indicate your considerations here.

Please also note that some messages contain considerations regarding the threshold to set etc. which are orthogonal to the inferential considerations. Please make sure to only classify according to the strategic content.