

Does Imitation Benefit Cue Order Learning?

Rocio Garcia-Retamero,^{1,2} Masanori Takezawa,^{1,3} and Gerd Gigerenzer¹

¹Max Planck Institute for Human Development, Berlin, Germany

²Universidad de Granada, Spain

³Tilburg University, The Netherlands

Abstract. Inferences are often based on uncertain cues, and the accuracy of such inferences depends on the order in which the cues are searched. Previous research has shown that people and computers progress only slowly in individual learning of cue orderings through feedback. A clue to how people (as opposed to computers) solve this problem is social learning: By exchanging information with others, people can learn which cues are relevant and the order in which they should be considered. By means of simulation, we demonstrate that imitate-the-best and imitate-the-majority speed up individual learning, whereas a third social rule, the Borda rule, does not. Imitate-the-best also leads to a steep increase in learning after a single social exchange, to cue orders that are more accurate than ecological validity, and to faster learning than when individuals gain the learning experience of all other group members but learn without social exchange. In two experiments, we find that people speed up cue learning in a similar way when provided with social information, both when they obtain the information from the experimenter or in free discussions with others.

Keywords: cue learning, imitate-the-best, imitate-the-majority, heuristics, social learning

In daily life, we frequently make inferences about current and future states of the world. For instance, when New York city taxi drivers pick up customers in the Bronx, Brooklyn, or other dangerous neighborhoods, they need to screen potential passengers to decide whether they are trustworthy. Taxi drivers have to make up their minds quickly. Many report that in cases of doubt, they drive past people to assess them before pickup. An error in judgment means losing a fare, being robbed, or even murdered. Refusing a harmless client, however, means losing money. Inferences concerning trustworthiness are based on uncertain cues, and there are some cues for trust that are shared by virtually all drivers, including older over younger, and female over male (Gambetta & Hamill, 2005).

In this article, we consider the problem of which of the two alternatives, varying on several dichotomous cues, has a higher value on a quantitative criterion (i.e., a two-alternative forced-choice task), such as which of the two passengers is the more trustworthy one based on cues like age or sex. Research indicates that people often base such inferences on a sequential analysis of a few cues, rather than weighting and adding several of them (e.g., Bröder & Gaissmaier, 2007; Garcia-Retamero & Dhami, 2009; Garcia-Retamero, Hoffrage, & Dieckmann, 2007; Lee & Cummins, 2004; Rieskamp & Hoffrage, 2008; Rieskamp & Otto, 2006). Sequential processing is effective if cues are ordered according to their usefulness rather than randomly (Garcia-Retamero & Dieckmann, 2006). One formal model for such a process is take-the-best (Gigerenzer & Goldstein, 1996), which orders cues according to validity, that is, the relative frequency of

making a correct inference given that the cue discriminates between the two alternatives. For instance, Bergert and Nosofsky (2007) compared a generalized version of take-the-best with a “rational” weighted additive model, and concluded that the vast majority of participants adopted take-the-best. Nosofsky and Bergert (2007) found similar results when comparing take-the-best with exemplar models.

The Problem of Learning Good Cue Orderings

Laboratory experiments using two-alternative forced-choice tasks showed that without explicit knowledge about cue validities, people seem to be rather slow in learning which cues are good and producing orderings that approximate ecological validity, even within the confined range of the 4–6 cues usually available in these experiments (Rakow, Hinvest, Jackson, & Palmer, 2004; see also Garcia-Retamero & Hoffrage, 2006). Yet the reason for this is not necessarily to be found in people’s minds alone but also in the task. For instance, Newell and Shanks (2003) had four cues with ecological validities 0.80, 0.75, 0.70, and 0.69, and gave participants 60 trials to learn a cue ordering. To learn the order of the validities of the last two cues, however, one would need 100 trials in which the third cue discriminates and another 100 trials for the fourth cue in order to experience that one results in 70 and the other in 69 correct inferences, assuming perfect memory. The same

problem holds for Newell, Rakow, Weston, and Shanks' (2004) experiment. Thus, learning of ecological cue orderings can be a lengthy process, depending on the number of cues and the differences between cue validities. Moreover, research in multiple cue probability learning suggests that interference effects exist when multiple cues are available and their validities have to be learned concurrently. For instance, if irrelevant cues are present, the utilization of valid cues is reduced and the accuracy of judgments is lower than that in a condition in which the irrelevant cues are not included (Brehmer, 1973; Edgell & Hennessey, 1980; see Newell, Lagnado, & Shanks, 2007, for a review). How can individuals learn efficient cue orderings if cue validities are not available beforehand?

One could assume that people can use take-the-best and update a cue ordering by using only the cues that they had searched. That is, a cue ordering could be acquired by learning-while-doing. Note that learning while one makes inferences with take-the-best or another lexicographic heuristic (exploitation) demands even more trials than when all cue values for the two objects in each trial are provided, as in Newell and Shanks' exploration phase (see above). The reason it requires more trials is that in each trial, the validity of only one cue can be updated, since the heuristic stops search immediately. This "worst-case" scenario is what is addressed in the present article. We first analyze by means of computer simulation which social exchange rules can speed up learning. Then in two experiments we test whether humans take advantage of social learning in ways similar to the simulation.

The question of how a cue ordering is learned has been considered only recently by Todd and Dieckmann (in press; see also Todd & Dieckmann, 2005). In a series of computer simulations, the authors evaluated the performance of the validity algorithm, a simple learning rule for forced-choice paired comparison tasks that updates cue orderings by validity on a trial-by-trial basis. The accuracy of the cue orderings resulting from the application of the validity algorithm was tested using the German cities data set, which consists of the 83 German cities having more than 100,000 inhabitants at the time (Fischer Welt Almanach, 1993). These cities were described by nine cues such as whether a city has a university (see Table 1).

The validity algorithm starts with a random cue ordering and searches one cue at a time until it finds a cue that discriminates between the alternatives, which is used to make the decision (i.e., the algorithm chooses the alternative favored by the first discriminating cue). After each decision, feedback is provided and the cue ordering is updated. The validity algorithm retains two pieces of information for each cue: how often a cue stopped the search (and led to a decision), and a count of the correct decisions. The validity of each cue is computed by dividing the number of current correct decisions by the number of current discriminations.¹

Todd and Dieckmann (in press) tested the performance of the validity algorithm in 100 trials, that is, in 100 paired comparisons that were randomly selected from the set of German cities. There were nine cues. The differences between their validities ranged between 0.01 and 0.15. The authors tested predictive accuracy, that is, how well the cue ordering learned in each trial would do if it were applied to the entire set of the paired comparisons (i.e., 3,403 pairs of 83 cities). The simulation results showed that the validity algorithm performance soon rose above that achieved by a random ordering. Simulated individuals, however, made slow progress in learning efficient cue orderings: Even after updating cue validities through feedback for 100 trials, accuracy remained well behind that when the ecological cue validities were used. This is not just a problem for the validity algorithm, but also applies to other ordering principles that require learning conditional weights such as β weights. The delta rule (used in neural networks) and the validity algorithm were the fastest learning rules, but both needed more than 2,900 trials to match the ecological validity ordering.

In a corresponding experimental setting, Todd and Dieckmann (in press) also showed that, just like computers, after 100 trials, participants were slow in learning cue orderings by validity when they could update such orderings through feedback. This result is in line with the previous findings from two-alternative forced-choice tasks and multiple cue probability learning mentioned above, and shows that learning good cue orderings can take a long time.

Does Social Learning Improve the Performance of Cue Order Learning? A Simulation Study

In contrast to many laboratory tasks, in real-world environments people often exchange information with other individuals before making a judgment. Consider once again the question of which passenger might be more trustworthy. Rather than collecting information only individually, taxi drivers might also discuss with other drivers which cues are more useful. Thus, people could learn to order cues both individually and socially by exchanging information.

In group decision making, several authors have analyzed whether individuals are able to gain valuable knowledge from other group members to improve inferential accuracy. Empirical results show that groups are often more accurate than average individuals (Garcia-Retamero, Takezawa, & Gigerenzer, 2008; Hastie & Kameda, 2005; Hill, 1982; Kameda & Nakanishi, 2002; Kerr & Tindale, 2004). For instance, even when individuals' quantitative judgments are systematically biased, the performance of the average

¹ The resulting (sample-based) cue ordering when using the validity algorithm should not be confused with the ecological validity ordering, as Dougherty, Franco-Watkins, and Thomas (2008) did. The latter is based on the complete knowledge of the environment, whereas the former models people's subjective updating of the ordering of validities based on samples (Gigerenzer, Hoffrage, & Goldstein, 2008). It is this subjective, sample-based ordering, that is the input to take-the-best or to other cognitive strategies.

Table 1. Description of the eight inference problems used in the simulation

Description of the real-world problems	Average VAL (DR)	Standard deviation of VAL (DR)	Range of VAL (DR)	Skewness of VAL (DR)	Kurtosis of VAL (DR)
<i>Rainfall</i> : Predicting the amount of rainfall after cloud seeding for 24 weather observations (Woodley, Simpson, Biondini, & Berkeley, 1977), described by six cues (e.g., suitability for seeding or percent cloud cover)	0.62 (0.49)	0.08 (0.07)	0.76–0.51 (0.52–0.34)	0.80 (–2.40)	2.38 (6.00)
<i>Homelessness</i> : Predicting the rate of homelessness of 50 US cities (Tucker, 1987), described by six cues (e.g., unemployment rate or public housing)	0.58 (0.51)	0.06 (0.00)	0.68–0.52 (0.51–0.51)	1.60 (–2.40)	3.05 (6.00)
<i>Lifespan</i> : Predicting the lifespan of 58 mammals (Allison & Cicchetti, 1976), described by eight cues (e.g., brain weight or gestation time)	0.74 (0.51)	0.12 (0.00)	0.93–0.53 (0.51–0.50)	–0.20 (–2.41)	–0.15 (4.50)
<i>Population</i> : Predicting the population of the 83 German cities with at least 100,000 inhabitants (Fischer Welt Almanach, 1993), described by nine cues (e.g., whether the city has a soccer team or university)	0.76 (0.30)	0.16 (0.13)	1.00–0.51 (0.51–0.02)	–0.30 (–0.80)	–0.39 (3.1)
<i>Obesity</i> : Predicting obesity at age 18 of 58 children (Weisberg, 1985), described by 11 cues (e.g., height at age 9 or leg circumference at age 9)	0.67 (0.51)	0.10 (0.01)	0.87–0.56 (0.51–0.49)	0.80 (–3.00)	0.18 (9.1)
<i>Fuel consumption</i> : Predicting the average motor fuel consumption per person for each of the 48 contiguous US (Weisberg, 1985), described by seven cues (e.g., population of the state or number of licensed drivers)	0.73 (0.51)	0.09 (0.00)	0.82–0.57 (0.51–0.51)	–1.40 (–2.60)	1.28 (7.0)
<i>Car accidents</i> : Predicting the accident rate per million vehicle miles for 39 segments of highway (Weisberg, 1985), described on 13 cues (e.g., segment's length and average traffic count)	0.65 (0.48)	0.07 (0.08)	0.76–0.52 (0.51–0.23)	–0.10 (–2.90)	–0.97 (8.9)
<i>Mortality</i> : Predicting the mortality rate in 60 US cities (McDonald & Schwing, 1973), described by 15 cues (e.g., average January temperature and percentage of population aged 65 or older)	0.67 (0.51)	0.09 (0.00)	0.81–0.54 (0.51–0.51)	0.02 (–3.20)	–0.64 (11.1)

Note. VAL, validity; DR, discrimination rate.

estimates comes close to models that weigh and add cues (see Einhorn, Hogarth, & Klemperer, 1977; Hogarth, 1978). Hastie and Kameda (2005) theoretically demonstrated that following the decision of the majority has the same error-reduction function in discrete decision-making tasks. In a series of computer simulations, they also showed that decisions supported by the majority or by the most capable person in the group achieve higher accuracy than those made by average individuals.

All in all, research in group decision making has identified situations in which group-based judgments are superior to individual judgments. Although previous research has focused on group consensus processes, the findings may be relevant to our research question. Our hypothesis is that the exchange of information can help individuals to solve the problem of learning good cue orderings in less time. To test this hypothesis, we conducted a series of computer simulations in which we evaluated the success of several social rules when they were implemented in the validity algorithm for ordering cues. In two subsequent experiments, we analyze whether people's behavior conforms to the results of the simulations.

Method

In our simulations, we used data for a variety of inference problems that have been studied in different disciplines, including psychology, economics, and computer science. The task was to infer, on the basis of several cues, which of the two objects had a higher criterion value. Altogether we considered eight inference problems (see Table 1). The problems differed by the number of objects considered and the number of cues provided for making a decision.

In the simulations, we included three conditions of social learning and two of individual learning. In the three social learning conditions, a group of five simulated individuals went through a trial block of five paired comparisons. In these trials, the simulants individually updated a cue ordering with feedback by applying the validity algorithm while making inferences with take-the-best. In the first trial, simulants started with random cue orderings. For each trial, the corresponding cue ordering was used to look up cues until a discriminating cue was found, which was in turn used to make the decision. After each decision, feedback was provided and the cue ordering was updated. The group as a whole received the same set of paired comparisons, whereas different simulated individuals within a group received different paired comparisons. Therefore, groups in the three social conditions came up with the same cue orderings after the first trial block, whereas each simulant within a group typically came up with a different cue ordering.

After the trial block, simulants exchanged information about the cue orderings that they learned individually with the other group members. All simulants within the group used a social rule, which differed among the three conditions, to arrive at a single social cue ordering (Hastie & Kameda, 2005). The first social rule was imitate-the-majority: Individuals vote for the cue with the highest

validity, and the one with the most votes becomes the top cue in the social cue ordering. Everyone then votes for the next best cue, and this process is repeated for all the cues. In cases of ties, one cue is randomly selected. The second social rule was imitate-the-best: Individuals imitate the cue ordering of the most successful group member. Unlike imitate-the-majority, imitate-the-best requires ordinal information about success. Finally, according to the Borda rule (named after the French mathematician Jean-Charles de Borda, who designed a voting system for the French Academy of Science to elect its members), voters rank all cues, and the sum of ranks is taken as the social ordering. Among all three social rules, imitate-the-best involves the least computation.

After exchanging information socially, simulants worked individually on the next block of five trials, in which they looked cues up using the social cue ordering instead of the cue orderings that they had learned on their own. The social cue ordering was then updated with feedback using the validity algorithm on an individual basis. The process of exchanging social information, computing a social cue ordering, and updating the social cue ordering individually was repeated for each of the 20 trial blocks (or 100 trials in total).

For control, we introduced two individual learning conditions. In the Individual-1 condition, each simulant received the same set of paired comparisons as those in the corresponding social learning conditions but did not exchange any information socially. In the Individual-5 condition, each simulant within the group went through the entire set of paired comparisons that the five simulants in the social learning conditions received (i.e., 25 trials in total). The Individual-1 condition is the baseline for evaluating the benefits of the three social learning rules, and the Individual-5 condition enables evaluating whether gains by social rules can be explained solely by the larger number of paired comparisons that a group experiences over an individual.

Predictive accuracy of the validity algorithm was examined for each of the eight inference problems (see Table 1) in the five learning conditions (i.e., imitate-the-majority, imitate-the-best, Borda, Individual-1, and Individual-5), providing 40 conditions altogether. For each of these conditions, we simulated 1,000 groups of five simulants, who made 100 decisions between randomly selected paired comparisons (i.e., 20,000,000 observations in total).

Results and Discussion

Figure 1 shows the performance of imitate-the-majority, imitate-the-best, and the Borda rule while learning cue orderings compared to that of Individual-1 (thin solid jagged line) and the more competitive version of individual learning (i.e., Individual-5; thick solid jagged line). The performance of the rules was averaged across the 1,000 groups and the eight inference problems. As benchmarks, the performance of a random ordering of cues (lower straight line) and the ecological ordering (upper straight line) are shown.

Both imitate-the-majority and imitate-the-best improved the performance of individual cue order learning substan-

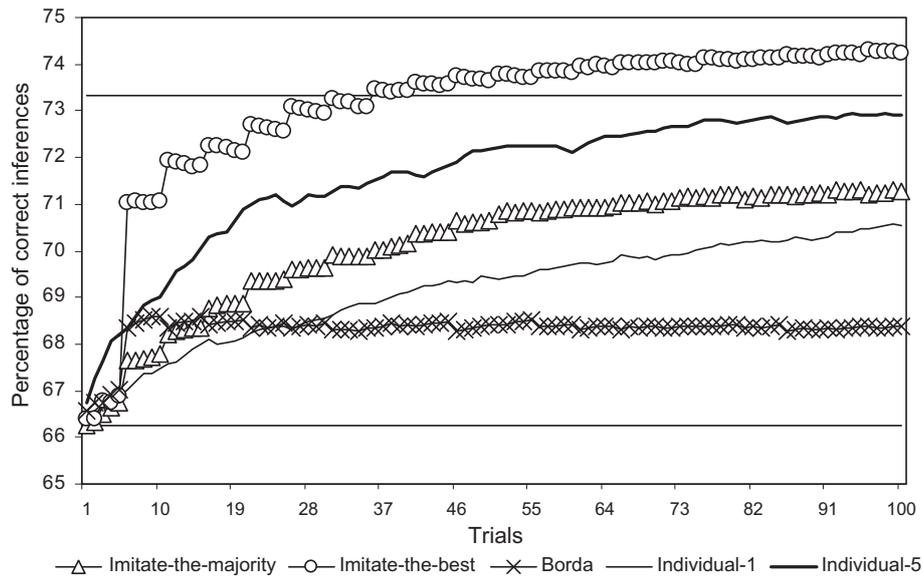


Figure 1. Predictive accuracy of imitate-the-majority, imitate-the-best, and Borda rule when implemented in the validity algorithm for 100 trials between randomly selected paired comparisons. For comparison, the performance of individual learning according to validity (Individual-1; thin solid jagged line), a more competitive version of individual learning (Individual-5; thick solid jagged line), ecological ordering (upper straight line), and random ordering (i.e., lower straight line) were added.

tially, with a particular boost by imitate-the-best after the first social exchange. The Borda rule, however, was less successful and did not even reach the performance level of individual learning. Imitate-the-best led to a steep increase in accuracy after the first exchange and eventually made even more correct decisions than can be obtained with the exact ecological cue validities. Most important, imitate-the-best alone led to more accurate decisions than the Individual-5 benchmark did. Since Individual-5 had exactly the same information as the groups who used imitate-the-best, the performance of imitate-the-best cannot be reduced to that of individual learning.

As a second performance measure, we also computed the tau correlation between the subjective cue orderings and the ecological cue validities for each condition. Consistent with the previous analyses, the correlation increased from .25 to .65 and from .23 to .47 when using imitate-the-best or imitate-the-majority, respectively, whereas increments were barely visible when using the Borda rule or the less competitive version of individual learning (i.e., Individual-1). In such cases, correlations only increased from .24 to .27 and from .25 to .30, respectively. When using the more competitive version of individual learning (i.e., Individual-5), the correlation increased from .23 to .55. These results, therefore, are consistent with those presented above and suggest that imitate-the-best is an efficient strategy for improving individual cue order learning, whereas the Borda rule that adds individuals' ranks fails.

Note that the combination between individual learning and imitate-the-best generated cue orderings with higher accuracy than that achieved by ecological validity

(Figure 1). An individual analysis shows that at Trial 100, 45% of the cue ordering led to more accurate inferences, whereas 34% arrived at orderings that matched ecological validity, and only the other cue orderings performed below. Note that ordering cues by validity is a heuristic rather than an optimizing principle; for instance, ordering by validity ignores the dependencies between cues (Brighton & Gigerenzer, 2008; Martignon & Hoffrage, 2002). The results in Figure 1 show that imitate-the-best can not only perform better than the validity orderings computed from the sample that an individual has experienced so far, but also better than the ecological validity ordering (which no individual can know unless they have complete information about the entire environment). This suggests that the effect of imitate-the-best cannot be reduced to a convenient speeding up of learning the validity order by indirectly increasing the sample size (see the Individual-5 learning condition), but actually results, on average, in cue-orders that outperform subjective (sample-based) validity and even ecological validity.

How robust are the results of this simulation? First, the average result shown in Figure 1 holds for each of the eight inference problems in Table 1. Second, we modified two parameters of the basic simulation: the group size (2, 10, 25, and 100 individuals) and the number of trials after which social information is exchanged (5, 25, and 50 trials). The results showed that the higher the group size, the higher the group's accuracy, regardless of the social rule they used. Group size showed diminishing returns; for instance, the increase in accuracy from 5 to 25 individuals was larger than that from 25 to 100. Furthermore, the higher the frequency

of exchanging social information, the higher the performance was in the long run. Again, imitate-the-best proved to be the most effective way of resolving the problem of slow individual learning. For instance, when 10 simulants exchanged information on just one occasion (after 50 trials), this single social exchange was sufficient to achieve the accuracy of the ecological validity ordering. In summary, the superiority of imitate-the-best, as observed in the basic simulation, remained stable with an increasing number of simulants in the group, even when the number of social exchange opportunities was reduced.

Does Social Learning Solve the Limitations of Individual Learning? Empirical Studies

The simulation results provide a theoretical context for empirical studies. Would real people profit from social exchange like the simulated ones did? To our knowledge, there is no experimental research that tests whether social learning improves individual cue order learning. It is also unclear which strategy people use to process social information in cue order learning. To answer these questions, we conducted two experiments using a similar procedure to that used in the simulation study. On the basis of several cues, participants had to infer which of the two objects had a higher criterion value. In both experiments, one group in which individuals learned cue orderings by feedback was compared with groups where individuals received information after each trial block about the cue orderings of all group members. In Experiment 1, the social information was provided by the experimenter, whereas in Experiment 2, no such information was given, but participants were allowed to talk to each other without any constraints.

If people's behavior conforms to the results of the simulations, those who can exchange information will learn more efficient cue orderings than those who only learn individually if they use imitate-the-majority or imitate-the-best. As a consequence, they will make more correct inferences and focus on high validity cues more often.

Experiment 1

Method

Participants

A total of 240 students (131 women and 109 men, average age 24 years, range 18–35) at the Free University of Berlin participated in the experiment. Participants were randomly assigned to one of four equally sized learning conditions ($n = 60$). The computerized task was conducted on groups of five individuals and lasted 1 h. Participants received a show-up fee of €8 plus half of the amount they earned in

the task, with an average payment of €11 (ranging from €8 to €17).

Stimuli and Design

Participants had to infer which of two job candidates (displayed column-wise) for an open position would be more productive in the future. To make these inferences, they could search for information about six cues describing the candidates (i.e., whether they had organizational skills, social skills, positive letters of recommendation, computer skills, whether they spoke foreign languages, or were reliable). These cues are common for assessing job candidates (see Garcia-Retamero & Rieskamp, 2008; Garcia-Retamero et al., 2008). The cues had a positive or negative value for each candidate, represented as “+” and “–”, respectively. The order in which the cues were presented on the screen was fixed for each group of participants but varied randomly between groups of participants. Likewise, the position of the two candidates (left or right on the screen) varied randomly.

We created an inference problem that consisted of 30 pairs of candidates described by six cues with the cue validities 0.83, 0.75, 0.69, 0.64, 0.58, and 0.50, and cue discrimination rates 0.40, 0.40, 0.43, 0.37, 0.40, and 0.40. The discrimination rate of a cue is the number of pairs in which cue values differ between alternatives. Neither cue validities nor cue discrimination rates were given to the participants. For each group of participants, cue labels were randomly assigned to the different cues. Participants made 210 inferences, broken down into seven trial blocks comprising 30 paired comparisons each. The same set of paired comparisons was presented within each block but in random order. After each trial block, participants ranked cues according to the subjective cue validities. Specifically, participants were told: “In the following you should rank the cues in the order of their validity. What does validity mean? Suppose one candidate has a positive evaluation and the other candidate has a negative evaluation. The validity tells you how probable it is that the candidate with the positive evaluation is also the candidate who will be more productive.”

The experiment included four conditions. In the social-information condition, after ranking cues according to the subjective cue validities, each participant received information about the cue orderings of all the members of his/her group in that trial block. The imitate-the-majority condition was identical to the previous condition, except that participants were also informed about the cue ordering computed according to the majority rule. In the imitate-the-best condition, participants received information about the final payoff of all group members in the trial block in addition to information about their cue orderings. In short, we refer to these three conditions as “social learners”. Finally, in the individual condition, participants did not receive any information after ordering cues according to the subjective validities.

In sum, the experimental design had two factors: the learning condition (social-information, imitate-the-majority, imitate-the-best, and individual; between-subjects) and trial block, with seven repetitions of the 30 pairs of objects (within-subjects).

Procedure

The experimental instructions told the participants to imagine that they worked for a company that had grown substantially. It was their job to make recommendations for a new personnel. They had to make choices between pairs of candidates and select the one that would be most productive on the basis of six descriptive cues. Participants could search for information about these cues by clicking boxes on the computer screen. Once a box with information on a cue for both candidates was opened, the cue values remained visible until a decision was made. For each cue that was looked up, 1 cent was deducted from a participant's overall payoff. After a decision was made by clicking on a button, outcome feedback was displayed. For each correct decision, participants earned 7 cents; for incorrect decisions, no money was deducted. The current balance of their account was always visible on the computer screen.

Five dependent variables were recorded in each trial: (1) the number of cues searched for, (2) which cues were searched for, (3) the order in which the cues were searched for, (4) whether the decision was correct, and (5) the time that the participants invested in searching for cues.² Additionally, after each trial block, participants' cue orderings were recorded according to the subjective validity of the cues.

Results

Do Social Learners Learn More Efficient Cue Orderings Than Individual Learners Do?

Figure 2 shows that in all three social conditions, the tau correlation between participants' subjective cue orderings and the ecological validity ordering increased much faster than in the individual learning condition. In fact, individual learning was quite slow and barely visible, and from the second to the last trial block, social learners achieved higher correlations than individual learners did. Can the benefit of social exchange be explained solely by the larger number of paired comparisons a group indirectly experiences over an individual? To answer this question, participants in the individual learning conditions were randomly assigned to groups of five individuals, and social cue orderings in the artificially generated groups were computed by using imitate-the-best or imitate-the-majority. These groups received the same amount of information than those in the social conditions but did not have the opportunity for social exchange. The correlation between these cue orderings and the ecological validity ordering was computed. As Figure 2 shows, in the artificially generated individual conditions, correlations do not reach the social learning groups. For imitate-the-majority the results are even indistinguishable from

individual learning. Overall, these results show that people – like simulants – do indeed learn more efficient cue ordering through social exchange.

Do Social Learners Search for the Most Valid Cue More Often Than Individual Learners Do?

Figure 3 shows the percentage of trials in which each cue was looked up for each trial block. For instance, the most valid cue was looked up by the participants in all four conditions about half of the time before the first social exchange (i.e., after the first block of 30 trials). For each of the social conditions, this percentage increased to about 71–82% after the first social exchange, and continued to increase but at a less accelerated pace up to 80–90%. Individual learners, in contrast, learned to look up the best cue more often over trials, but their average did not reach 70%. A similar but less pronounced effect can be seen for the second-best cue, except for the social-information condition, which had the least social information. On the other end, the least valid cue was looked up before the first social exchange about as often as or more than the most valid one, but this frequency decreased from trial block to trial block in the three social conditions, not in the individual one. The intermediate cues showed no difference. Figure 4 shows a similar pattern for the reaction times, that is, the time between trial onset and the mouse click on the cue. Participants in the social conditions were faster than individual learners at learning to look for the most valid cue and did not look up the least valid cue as quickly. An analysis of variance is consistent with both of these results, showing an interaction between learning condition, the cue ranking, and the trial block, $F(46.7, 3,676.2) = 1.52$, $p = .013$, $\eta_p^2 = .10$, and $F(46.7, 3,676.2) = 2.54$, $p = .001$, $\eta_p^2 = .10$, for frequency and response time, respectively.³

Do Social Learners Learn to Make Correct Inferences More Often Than Those Who Just learn Individually?

Inferential accuracy increased from 65% ($SEM = 0.25$) to 74% ($SEM = 0.25$) in the social-information condition; from 65% ($SEM = 0.28$) to 74% ($SEM = 0.25$) in the imitate-the-majority condition, and from 66% ($SEM = 0.25$) to 73% ($SEM = 0.31$) in the imitate-the-best condition. In contrast, inferential accuracy only increased from 64% ($SEM = 0.26$) to 67% ($SEM = 0.30$) in the individual condition. From the third to the last trial block, social learners achieved higher accuracy than individual learners did. Consistent with these results, an analysis of variance with percentage correct as the dependent variable shows an interaction between the learning conditions and the trial block, $F(14.6, 1,144.4) = 1.97$, $p = .015$, $\eta_p^2 = .10$.

² Time measures were corrected excluding observations outside the ± 2 SD interval around the individual mean.

³ Degrees of freedom for the analyses containing repeated-measures factors were corrected by using the Greenhouse and Geisser (1959) technique.

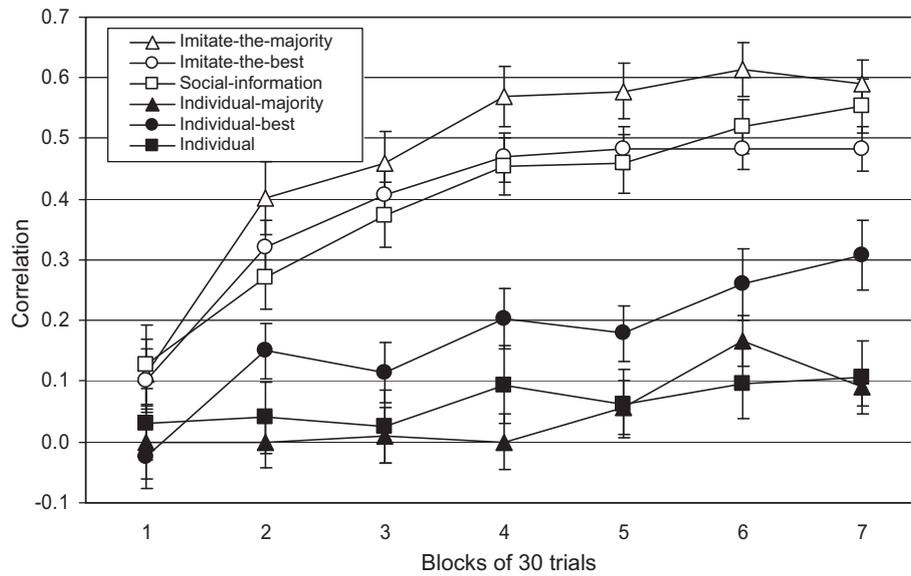


Figure 2. Correlation between participants' subjective cue rankings (after each trial block) and the ranking according to ecological validity across the seven trial blocks in Experiment 1. Error bars represent one standard error.

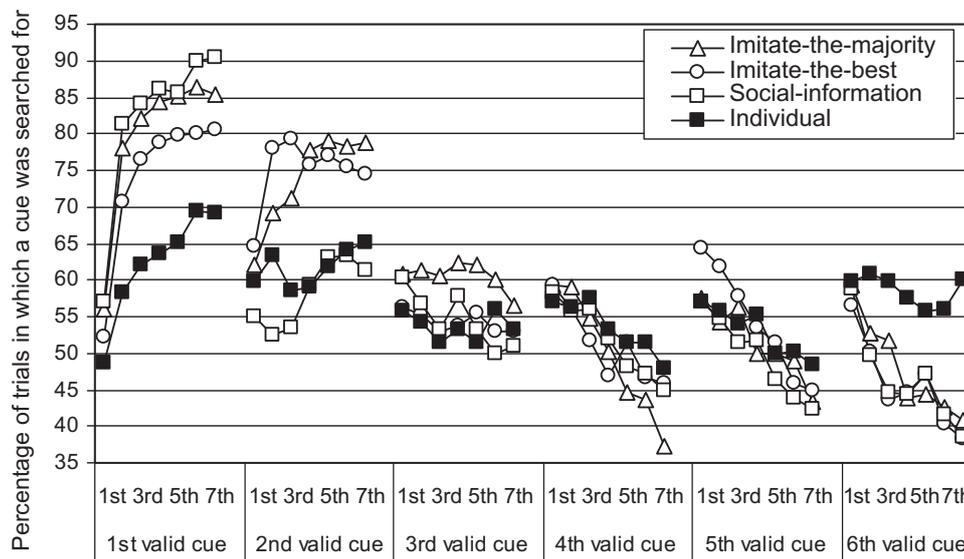


Figure 3. Percentage of trials in which a cue was searched for depending on cue ranking (most to less valid cue) and the trial block in Experiment 1.

Discussion

Experiment 1 showed that social exchange during individual learning speeds up the learning of efficient cue orderings and increases the frequency of looking up the most valid cues as well as the accuracy of inferences. This effect is not due to differences in the amount of information that participants in different conditions searched for, as both social and

individual learners searched for the same number of cues (i.e., about 3.5 of the six cues per trial). Experiment 1 could not, however, answer the question of which strategy social learners use, because the social information was provided by the experimenter rather than searched for by the participants. In the real world, we rarely receive a summary of other individuals' cue orderings but would have to approach people to obtain information. To test whether the results general-

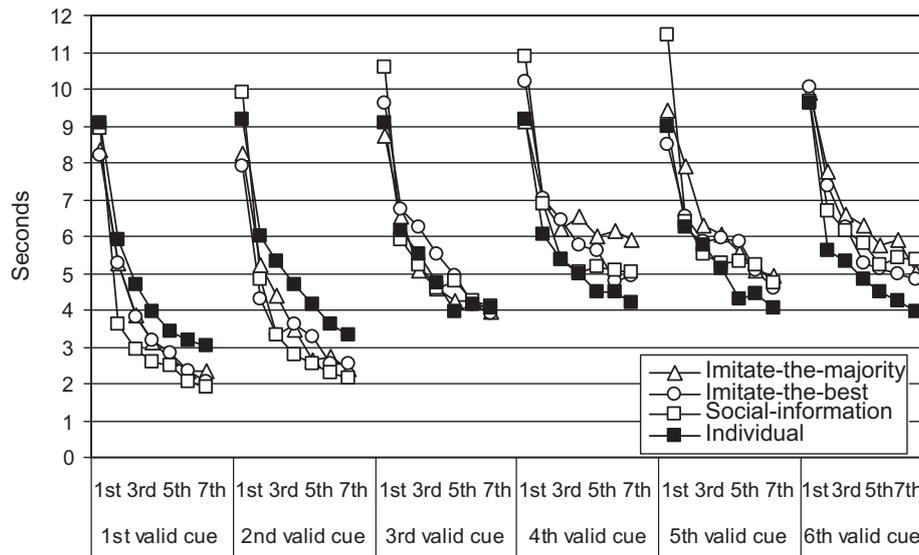


Figure 4. Time between trial onset and the mouse click on the cue depending on cue ranking (most to less valid cue) and the trial block in Experiment 1.

ize to situations in which people can freely communicate, we used a group discussion setting in Experiment 2.

Experiment 2

In this experiment, participants had no constraints on the information they could talk about. Would participants naturally engage in exchanging information about relevant cues in this situation? If so, would social learners prefer to imitate the most competent group member or to aggregate social information by using the majority rule?

Method

Participants

A total of 120 students (55 men and 65 women), whose average age was 26 (range 19–39), participated in the experiment. Participants were randomly assigned to one of two equally sized learning conditions ($n = 60$). The computerized task was conducted in groups of five individuals and lasted 1 h. Participants received a show-up fee of €8 plus half of the amount they earned in the task, with an average payment of €10.40 (ranging from €8 to €15).

Design and Procedure

The instructions and the procedure of Experiment 2 were identical to those of Experiment 1, except that the three

social conditions were replaced by a discussion condition. Specifically, after each trial block, individuals ranked cues according to subjective cue validity and then met in groups of five individuals to converse freely for 10 min. Participants were not instructed on what to discuss. Afterwards, the group members had to agree on a cue ranking according to cue validities. In contrast, in the individual condition, participants did not meet or exchange any information with other individuals. In sum, the experimental design had two factors: the learning condition (discussion and individual; between-subjects) and trial block, with seven repetitions of the 30 pairs of objects (within-subjects).

Results

Do Participants in the Discussion Condition Learn More Efficient Cue Orderings Than Those Who Learn Individually?

Figure 5 shows that all conditions (including the artificially generated ones) started out with about the same tau correlation between participants' subjective cue orderings and the ecological validity ordering, but thereafter, the correlation increased much faster in the discussion condition than in the rest. As in Experiment 1 (Figure 2), individual learning was quite slow, and from the second to the last trial block, social learners achieved higher correlations than the individual learners did. This result, which cannot be explained solely by the larger number of paired comparisons a group indirectly experiences over an individual, shows that participants can learn by means of free discussion as efficiently as simulated individuals or real individuals who receive the relevant information prepackaged by the experimenter.

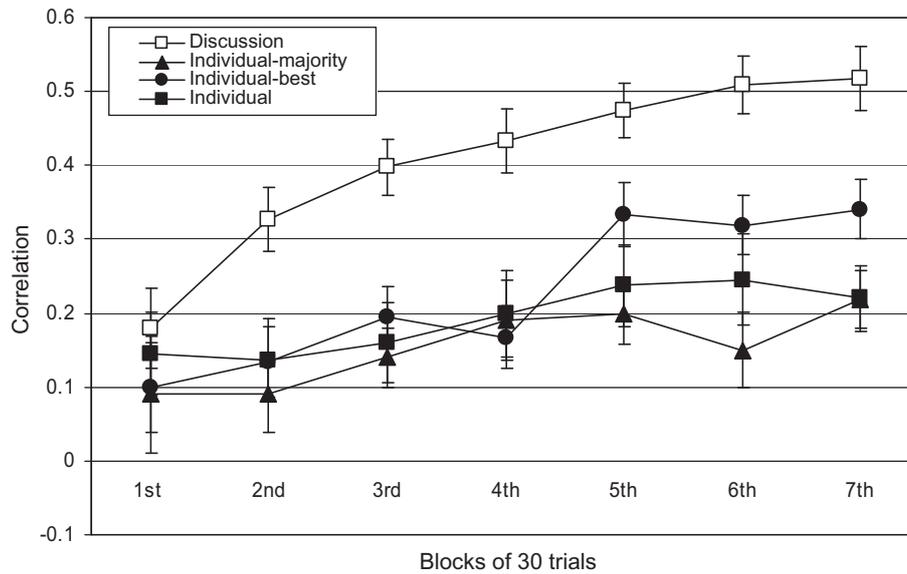


Figure 5. Correlation between participants' subjective cue rankings (after each block) and the ranking according to ecological validity across the seven trial blocks in Experiment 2. Error bars represent one standard error.

Which Strategy Did the Participants Use When They Exchanged Information?

Results showed that the correlation between subjective cue orderings and the cue ordering of the most competent group member increased from .26 ($SEM = 0.07$) to .58 ($SEM = 0.07$) from the first to the last trial block. In contrast, the correlation with the cue ordering of the majority only increased from .26 ($SEM = 0.06$) to .35 ($SEM = 0.06$). Furthermore, the analyses of the recorded conversations that participants had in the discussion sessions showed that in 65.7% of the sessions, at least one participant mentioned information about payoffs. These results are consistent with the hypothesis that participants exchanged information about their payoffs in the group discussion sessions and imitated the cue ordering of the best group member.

Do People in the Discussion Condition Search for the Most Valid Cue More Often Than Individual Learners?

Figure 6 shows, for each trial block, the percentage of trials in which each cue was looked up. The most valid cue was looked up by the participants in both conditions equally often before the first social exchange (i.e., after the first block of 30 trials with feedback). For the discussion condition, this percentage increased to about 80%, whereas this number remained around 65% for individual learners. For the less valid cues, participants in the discussion condition learned to avoid these faster than those in Experiment 1 (Figure 3). Figure 7 shows a similar pattern for the reaction times, that is, the time between trial onset and the mouse click on the cue. Participants in the discussion condition were faster than individual learners in learning to look for the most valid cue

and did not look up the least valid cue as quickly, replicating the corresponding result in Experiment 1. An analysis of variance is consistent with both these results showing an interaction between learning condition, the cue ranking, and the trial block, $F(13.4, 1,575.0) = 3.98, p = .001, \eta_p^2 = .10$, and $F(13.4, 1,575.0) = 1.55, p = .04, \eta_p^2 = .04$, for frequency and response time, respectively. These results replicate those in Experiment 1 in a situation where participants could gather information by talking freely with other group members.

Do People Who Exchange Information About the Task Make Correct Inferences More Often Than Those Who Learn Individually?

Inferential accuracy increased from 64% ($SEM = 0.20$) to 73% ($SEM = 0.31$) in the discussion condition. In contrast, inferential accuracy only increased from 65% ($SEM = 0.20$) to 67% ($SEM = 0.35$) in the individual condition. From the third to the last trial block, participants in the discussion condition achieved greater accuracy than participants in the individual condition. Consistent with these results, an analysis of variance with percentage correct as the dependent variable showed an interaction between the learning conditions and the trial block, $F(4.6, 543.5) = 3.27, p = .008, \eta_p^2 = .05$.

When social learners were able to discuss the task freely with other group members, they also became more frugal (i.e., they searched for less of the available information) than individual learners in making inferences. This result, along with the increase in inferential accuracy, explains why participants who exchanged information obtained a higher overall payoff than those who learned individually (€9.50 vs. €8.90).

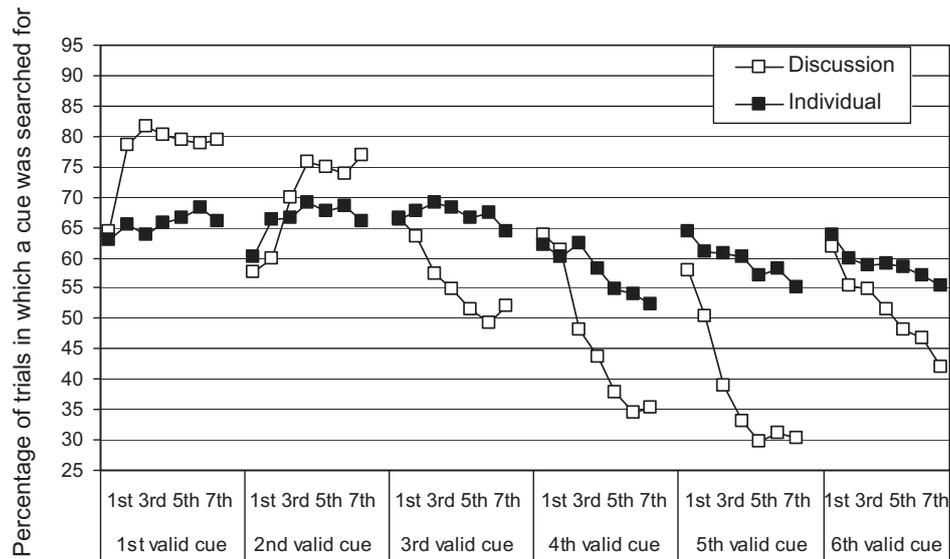


Figure 6. Percentage of trials in which a cue was searched for depending on cue ranking (most to less valid cue) and the trial block in Experiment 2.

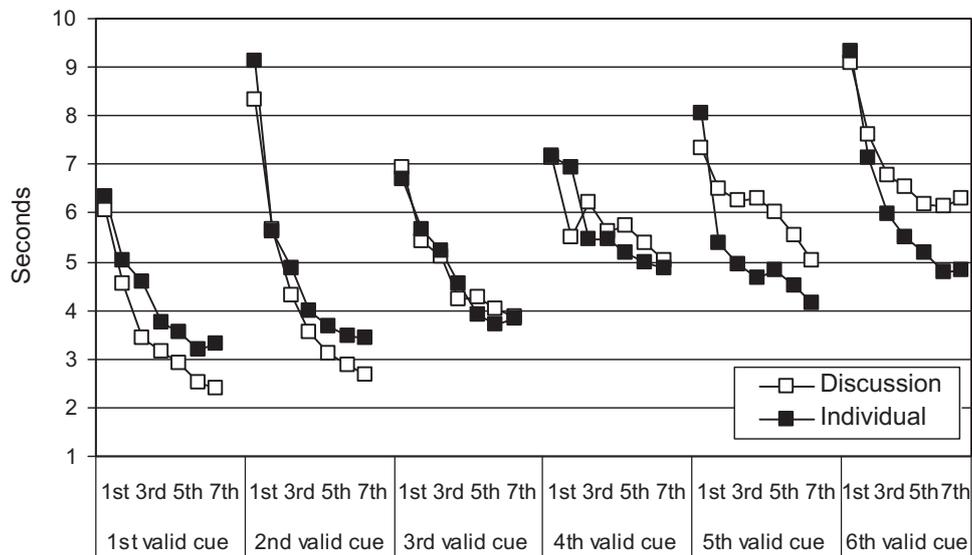


Figure 7. Time between trial onset and the mouse click on the cue depending on cue ranking (most to less valid cue) and the trial block in Experiment 2.

General Discussion

In this paper, we addressed the problem that individual learning to order cues by feedback is relatively slow, as demonstrated by the previous experiments and computer simulations. By means of simulation, we showed that individual learning can be expedited by two social rules, imitate-the-majority and imitate-the-best, but not with the Borda rule. Moreover, imitate-the-best, unlike imitate-the-majority, was more accurate than ecological validity and was also better than Individual-5, that is, an individual with

all the learning experience to which all group members are exposed. In two experiments, the improvement observed in the simulations was replicated for real participants, both when the relevant social information was provided by the experimenter and when it was not, but the participants were able to engage in free discussion. Overall, these findings suggest that imitate-the-best and imitate-the-majority, but not the Borda rule, are efficient strategies for improving individual cue order learning. Results in the group experiments further indicate that imitate-the-best is a prime strategy when people exchange information about cue orderings.

Social and Individual Learning

Figure 6 provides a hypothesis about a key difference between social and individual learning of cue orderings. Previous theories on individual learning assumed a simultaneous updating of the entire order of N cues, either by updating validity or by other means such as tally and swap (e.g., Todd & Dieckmann, 2005). That would suggest that the speed of learning is about equal for all cues, regardless of whether the cues are relevant. However, Figure 6 shows that when people can freely exchange information, they learn to search for the best cue more frequently after the first exchange, and there is little further learning. In contrast, learning for the second cue occurs mainly after the second social exchange. In the course of learning, all the cues with lower validity are searched for less often, apparently in keeping with a more continuous decrease in the percentage of trials that are searched. The pattern suggests that in situations with unconstrained social exchange, people try to determine the best cue first and the second-best only later, and so on.

This observation provides a different view on cue learning than does the picture of simultaneous updating of validities, or more generally, weights, as in Bayesian probability updating. Learning cue ordering is assumed to be sequential: First, try to establish what the best cue is out of N cues. Then, using the remaining $N - 1$ cues, determine what the second-best cue might be. Such a sequential procedure reduces the problem space quickly, and allows for a stable ordering that does not require constant updating. This corresponds well with the observation that in laboratory tasks, participants tend to settle into a routine at some point, that is, move from an exploration-oriented phase to an exploitation phase (Bröder & Schiffer, 2006; see also Betsch & Haberstroh, 2005, for a summary).

Open Questions

To the best of our knowledge, the present paper is the first to propose that social learning can help individual cue order learning. At the same time, it leaves several questions open for future research. Although our results were robust across the inference problems in Table 1 and are also applicable to various learning algorithms reported by Todd and Dieckmann (in press), they may depend on the nature of the task. Other authors (e.g., Hastie & Kameda, 2005) investigated how the members of a group who made decisions individually arrived at a collective decision. That is, they focused on social learning of decisions. In contrast, our focus was on social learning of the information that individuals use to make decisions. Despite these differences, we all found that cognitively simple social rules perform better than individual learning. Which social rules improve individual learning in other tasks, where learning is slow, still needs to be studied. Other open questions relate to the fact that our experiments embodied an ideal situation where feedback was given after each individual trial and was always correct, and external search rather than internal search in memory was used. If feedback is scarce and partly

erroneous, as in many aspects of our lives, or if inferences are made from memory, individual learning can be expected to slow down even further than shown in Figures 2 and 4. Is the benefit of social learning here comparable to what we found, or is it equally slowed down? In situations where feedback is scarce or when inferences are made from memory, will people still tend to imitate the successful member, or rather the majority? Furthermore, participants in the experiments had to learn about an environment in which the cue validities varied, but not the discrimination rates. If discrimination rates also vary, then a combination between validity and discrimination rate, such as success (Martignon & Hoffrage, 2002; Newell et al., 2004), is a viable alternative principle for cue ordering. In this case, one could learn through social exchange about which cues allow decision making more often, that is, have higher discrimination rates, or greater success. We do not know of any research that addresses these questions.

We began this research with the question: What social learning rules can speed up individual learning so that cue orders approximating the ecological validity order can be found faster? The results indicate that there is an interesting follow-up question: Which social rules can improve orders beyond ecological validity, and why?

Social Learning as Part of the Study of Bounded Rationality

According to Gigerenzer and Selten (2001), bounded rationality asks two questions: What are the heuristics in the adaptive toolbox? And, in which environments is a given heuristic ecologically rational? The study of heuristics, both for preferences (Payne, Bettman, & Johnson, 1993) and for inferences (Gigerenzer, 2008), has focused to a large degree on individual decision making, and initially paid little attention to the question of how heuristics and their building blocks, such as cue orderings, are learned. Recently, the question of learning has been addressed by a number of researchers (e.g., Brighton & Gigerenzer, 2008; Rieskamp & Otto, 2006), but predominantly in the form of individual reinforcement learning. Yet increasingly evident are the limitations of mere individual learning, including the lack of reliable feedback in many situations, the potential dangerousness of individual learning by feedback in matters of food and health, and the relative slowness of individual learning even if feedback is ideal and safe. These environmental factors that limit the efficacy of individual learning call for special attention to forms of social learning in the study of bounded rationality. The research presented in this article does not abandon individual learning, but tries to open up a perspective on how people rely on both individual and social learning to improve their heuristics over time. It also shows an area of intersection between the research on group decisions and the human tendency to conform to the majority's behavior. By integrating these various theoretical strands, we might eventually better understand how people mix social learning with individual learning, and how this mix depends on the structure of the social environment.

References

- Allison, T., & Cicchetti, D. (1976). Sleep in mammals: Ecological and constitutional correlates. *Science*, *174*, 732–734.
- Bergert, F. B., & Nosofsky, R. M. (2007). A response-time approach to comparing generalized rational and take-the-best models of decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *33*, 107–129.
- Betsch, T., & Haberstroh, S. (2005). Current research on routine decision making: Advances and prospects. In T. Betsch, & S. Haberstroh (Eds.), *The routines of decision making* (pp. 359–376). Mahwah, NJ: Erlbaum.
- Brehmer, B. (1973). Note on the relation between single-cue probability learning and multiple-cue probability learning. *Organizational Behavior and Human Performance*, *9*, 246–252.
- Brighton, H., & Gigerenzer, G. (2008). Bayesian brains and cognitive mechanisms: Harmony or dissonance? In N. Chater, & M. Oaksford (Eds.), *The probabilistic mind: Prospects for Bayesian cognitive science* (pp. 189–208). New York: Oxford University Press.
- Bröder, A., & Gaissmaier, W. (2007). Sequential processing of cues in memory-based multiattribute decisions. *Psychonomic Bulletin & Review*, *14*, 895–900.
- Bröder, A., & Schiffer, S. (2006). Adaptive flexibility and maladaptive routines in selecting fast and frugal decision strategies. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *32*, 904–918.
- Dougherty, M. R., Franco-Watkins, A. M., & Thomas, R. (2008). Psychological plausibility of the theory of probabilistic mental models and the fast and frugal heuristics. *Psychological Review*, *115*, 199–213.
- Edgell, S. E., & Hennessey, J. E. (1980). Irrelevant information and utilization of event base rates in nonmetric multiple-cue probability learning. *Organizational Behavior and Human Performance*, *26*, 1–6.
- Einhorn, H. J., Hogarth, R. M., & Klempner, E. (1977). Quality of group judgment. *Psychological Bulletin*, *84*, 158–172.
- Fischer Welt Almanach* [Fischer World Almanac]. (1993). Frankfurt, Germany: Fischer.
- Gambetta, D., & Hamill, H. (2005). *Streetwise. How taxi drivers establish their customers' trustworthiness*. New York: Russel Sage.
- Garcia-Retamero, R., & Dhami, M. K. (2009). Take-the-best in expert-novice decision strategies for residential burglary. *Psychonomic Bulletin and Review*, *16*, 163–169.
- Garcia-Retamero, R., & Dieckmann, A. (2006). Una revisión crítica del enfoque de los heurísticos rápidos y frugales. *Revista Latinoamericana de Psicología*, *38*, 509–522.
- Garcia-Retamero, R., & Hoffrage, U. (2006). How causal knowledge simplifies decision making. *Minds & Machines. Special Volume on Causality, Uncertainty and Ignorance*, *16*, 365–380.
- Garcia-Retamero, R., Hoffrage, U., & Dieckmann, A. (2007). When one cue is not enough: Combining fast and frugal heuristics with compound cue processing. *Quarterly Journal of Experimental Psychology*, *60*, 1197–1215.
- Garcia-Retamero, R., & Rieskamp, J. (2008). Do people treat missing information adaptively when making inferences? *The Psychological Record*, *58*, 547–568.
- Garcia-Retamero, R., Takezawa, M., & Gigerenzer, G. (2008). Comunicación grupal y estrategias de toma de decisiones. *Psicothema*, *20*, 753–759.
- Gigerenzer, G. (2008). *Rationality for mortals*. New York: Oxford University Press.
- Gigerenzer, G., & Goldstein, D. G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, *103*, 650–669.
- Gigerenzer, G., Hoffrage, U., & Goldstein, D. G. (2008). Fast and frugal heuristics are plausible models of cognition: Reply to Dougherty, Franco-Watkins & Thomas (2008). *Psychological Review*, *115*, 230–239.
- Gigerenzer, G., & Selten, R. (2001). *Bounded rationality: The adaptive toolbox*. Cambridge, MA: MIT Press.
- Greenhouse, S. W., & Geisser, S. (1959). On methods in the analysis of profile data. *Psychometrika*, *24*, 95–112.
- Hastie, R., & Kameda, T. (2005). The robust beauty of majority rules in group decisions. *Psychological Review*, *112*, 494–508.
- Hill, G. W. (1982). Group versus individual performance: Are $N + 1$ heads better than one? *Psychological Bulletin*, *91*, 517–539.
- Hogarth, R. M. (1978). A note on aggregating opinions. *Organizational Behavior and Human Performance*, *21*, 40–46.
- Kameda, T., & Nakanishi, D. (2002). Cost-benefit analysis of social/cultural learning in a nonstationary uncertain environment: An evolutionary simulation and an experiment with human subjects. *Evolution and Human Behavior*, *23*, 373–393.
- Kerr, N. L., & Tindale, R. S. (2004). Group performance and decision making. *Annual Review of Psychology*, *55*, 623–655.
- Lee, M. D., & Cummins, T. D. R. (2004). Evidence accumulation in decision making: Unifying the “take the best” and the “rational” models. *Psychonomic Bulletin & Review*, *11*, 343–352.
- Martignon, L., & Hoffrage, U. (2002). Fast, frugal, and fit: Simple heuristics for paired comparison. *Theory & Decision*, *52*, 29–71.
- McDonald, G. C., & Schwing, R. C. (1973). Instabilities of regression estimated relating air pollution to mortality. *Technometrics*, *15*, 463–482.
- Newell, B. R., Lagnado, D. A., & Shanks, D. R. (2007). *Straight choices: The psychology of decision making*. Hove, UK: Psychology Press.
- Newell, B. R., Rakow, T., Weston, N. J., & Shanks, D. R. (2004). Search strategies in decision making: The success of “success”. *Journal of Behavioral Decision Making*, *17*, 117–137.
- Newell, B. R., & Shanks, D. R. (2003). Take the best or look at the rest? Factors influencing “one-reason” decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *29*, 53–63.
- Nosofsky, R. M., & Bergert, F. B. (2007). Limitations of exemplar models of multi-attribute probabilistic inference. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *33*, 999–1019.
- Payne, J. W., Bettman, J. R., & Johnson, E. J. (1993). *The adaptive decision maker*. New York: Cambridge University Press.
- Rakow, T., Hinest, N., Jackson, E., & Palmer, M. (2004). Simple heuristics from the adaptive toolbox: Can we perform the requisite learning? *Thinking and Reasoning*, *10*, 1–29.
- Rieskamp, J., & Hoffrage, U. (2008). Inferences under time pressure: How opportunity cost affect strategy selection. *Acta Psychologica*, *127*, 258–276.
- Rieskamp, J., & Otto, P. E. (2006). SSL: A theory of how people learn to select strategies. *Journal of Experimental Psychology: General*, *135*, 207–236.
- Todd, P. M., & Dieckmann, A. (1987). Heuristics for ordering cue search in decision making. In L. K. Saul, Y. Weiss, & L. Bottou (Eds.), *Advances in neural information processing systems* (17, pp. 1393–1400) Cambridge, MA: MIT Press.

- Todd, P. M., & Dieckmann, A. (in press). In Simple rules for ordering cues in one-reason decision making. In Todd, P. M., Gigerenzer, G., & the ABC Research Group (Eds.), *Ecological rationality: Intelligence in the world*. New York: Oxford University Press.
- Tucker, W. (1987). Where do the homeless come from? *National Review*, 25, 34–44.
- Weisberg, S. (1985). *Applied linear regression*. New York: Wiley.
- Woodley, W. L., Simpson, J., Biondini, R., & Berkeley, J. (1977). Rainfall results 1970–75: Florida area cumulus experiment. *Science*, 195, 735–742.

Rocio Garcia-Retamero

Facultad de Psicología
Universidad de Granada
Campus Universitario de Cartuja
18071 Granada
Spain
E-mail rretamer@ugr.es

Received May 26, 2008

Revision received August 30, 2008

Accepted September 7, 2008