When groups should not imitate their most successful members

Jan K. Woike (Woike@mpib-berlin.mpg.de)
Max Planck Institute for Human Development, Center for Adaptive Rationality (ARC), Lentzeallee 94,
14195 Berlin, Germany

Jean-Philippe Bonardi
Université de Lausanne, Quartier UNIL-Dorigny, Bâtiment Internef
1015 Lausanne, Switzerland

Rocio Garcia-Retamero
Department of Experimental Psychology, University of Granada,
Granada, Spain

Abstract

The imitation of successful peers is often heralded as an intelligent shortcut to reduce individual learning costs. Using computer simulations, we demonstrate that this advice can be ill-founded and harmful in a cognitive inference task involving continuous learning. In particular, success-based imitators perform worse than both learners who integrate the learning experience of all group members and isolated learners. We report on sensitivity analyses for this phenomenon and offer explanatory mechanisms.

Keywords: group decision making; imitation; social learning; computer simulation

Introduction

The results of a recent social learning tournament (Rendell et al., 2010) suggest that it always pays to copy successful others when faced with a choice between individual learning and group learning. Yet, imitation learning is not as prevalent in the biological world as could be expected (Rieucau & Giraldeau, 2011). Humans, at least, often orient themselves towards successful peers to shorten periods of individual exploration and try to imitate the best group member (Garcia-Retamero, Takezawa, & Gigerenzer, 2009; Garcia-Retamero, Takezawa, Woike, & Gigerenzer, 2013). Yet, it can be argued that there are situations in which this strategy does not pay off (Denrell, 2005). In this simulation study we want to illustrate one such situation and compare individual learning with several social learning strategies in a sequential cognitive inference task based on the framework used in Garcia-Retamero, Takezawa, and Gigerenzer (2006).

The learning task and learning strategies

The learning task

We investigated the behavior of virtual decision makers using a paired-comparison task. This is an inference task, in which agents have to decide which of two objects has the higher criterion value. The basis for this inference are the values of a set of dichotomous variables (henceforth called cues) that characterize the two objects. An environment in this study consists of a set of N objects that are associated with criterion values and k binary cue values. The agents follow a strictly non-compensatory inference strategy: for each pair of objects an agent i looks up cues in an agent-specific order $O_i = (o_{1,i}, o_{2,i}, \ldots, o_{k,i})$ until a cue discriminates between the two objects (i.e., until the cue value is different for the two compared objects). In this case the object with the higher cue value (i.e., the value that indicates a higher criterion value) is selected. If none of the cues discriminates, a random decision is made. All agents are given the directions of all cues in each environment. For each cue, this direction is determined to maximize the number of correct decisions assuming that all possible pairs of objects in an environment are known and considered. In this setup, the choice of the cue order alone determines the success or failure of the agent in a given environment. The problem of finding the best cue order has been proven to be computationally intractable when the whole data set is available (Martignon & Schmitt, 1999). For the case of off-line learning (i.e., for situations in which all cue values for all objects in a decently sized sample are known), a strategy called take-the-best (TTB), which determines the order of cues O_i by ranking them according to their ecological validity, performs well across a variety of problems, especially for generalization tasks (Czerlinski, Gigerenzer, & Goldstein, 1999; Gigerenzer & Brighton, 2009). Our study focuses on on-line learning instead: For a learner in an unfamiliar decision setting and without any prior knowledge, learning has to be based on experience on a trial-by-trial basis (Hertwig, Barron, Weber, & Erev, 2004). The learner has no prior access to information about objects or environment, and full information might be costly and time-consuming if not impossible to obtain.

Individual learning with the validity algorithm

In on-line learning situations TTB cannot be easily applied since an agent generally has no basis for accessing or estimating the actual cue validities Todd and Dieckmann (2012) propose a learning mechanism that can be used in this setting, the *validity algorithm*. The validity algorithm starts out with a random cue order for the first trial and stores the values of two variables for each cue i: the number of observed discriminations d_i (i.e., the number of observed object pairs with different cue values for the two objects), and the number of correct predictions $c_i \leq d_i$ (i.e., the number of observed object pairs with different cue values, for which the object with positive cue value has a higher criterion value than the object with negative cue value). Both variables are set to zero for all cues at the start. In each learning trial a decision is made using the trial's cue order, and feedback is received on the correctness of this decision. For the first cue in the order that discriminates d_i is incremented by 1 and if the prediction turns out to be correct, c_i is incremented as well. After each trial, cue validities for all cues are estimated as:

$$\hat{v}_i = \begin{cases} \frac{c_i}{d_i} & d_i > 0 \\ 0.5 & d_i = 0 \end{cases} \tag{1}$$

A new cue order is established for the next trial by ranking the validity estimates and adapting the cue order accordingly. Dieckmann and Todd (2004) observe that by using the validity algorithm individual performance can approach the performance that is obtained by using the ecological cue validities from the start (i.e., the performance of individuals that calculate cue validities based on the full sample and order cues according to these validities). Yet, the on-line learning process is likely to be slow and convergence cannot be guaranteed. Group learning has been proposed as a way to speed up this learning process. In this paper, we investigated whether this is in fact the case.

Group Learning Strategies

While an isolated individual is often condemned to learn on a personal trial and error basis, humans often find themselves situated in groups that offer ways to overcome this predicament. In the current study we compare the individual performance in the sequential learning task with the performance of members in learning groups using various group learning strategies.

Group members alternate between blocks of individual trials and group exchange phases. After each individual trial, each individual's cue order is updated using the validity algorithm. In the group exchange phase individuals exchange information according to the social rule they have been assigned to. A social cue order is determined, and all individuals adjust their individual cue order. Afterwards each individual's memory is altered in accordance with the new cue order (see below) and the next individual trial block begins.

We implement the following group learning algorithms: 1) *imitation,* 2) the *plurality* rule, and 3) the *averaging* rule.

Imitation An easy way to learn from others is achieved by simply imitating their behavior. If the relevant aspects can be observed or communicated, some individuals can avoid undergoing a longer learning process by copying the result of others. Since it is highly likely that not everyone who is observed is suited to be an adequate model, a degree of specificity is well-advised.

The *imitate-the-best* rule proceeds by first identifying the individual in the group who achieved the best performance in the preceding trial block. In case of ties, this individual is randomly chosen among those with the highest number of successes. This individual's cue order is then chosen to be the resulting group cue order and every other individual copies this cue order.

One parameter for this strategy is the number of past observations considered for determining the most successful individual. This parameter has been set to the size of the individual trial block. A trade-off has to be considered here: The longer the time frame, the more observations can be evaluated and the performance measurement might well be more reliable. On the other hand, the more observations are considered the higher the chance that an individual changes the cue order used between the trials, so that older observations may be less relevant or even misleading in regard to evaluating this individual's present cue order.

Plurality The *plurality rule* in standard choice contexts proceeds by letting individuals vote for their preferred option and the option with the most votes (the plurality of votes) is chosen. A variant of this rule is the "'majority rule"' that implies strictly speaking that there could not be a decision without an absolute majority of votes. So if three alternatives receive 40%, 35%, and 25% of the votes, respectively, the plurality rule consistently chooses the first alternative, even if none of the alternatives has obtained more than 50% of the votes.

The transfer to ordering cues is straight-forward: For the first and each subsequent rank position (but the last) a vote will be held, where cues whose rank has already been established cannot be voted for. Each individual selects the non-ranked cue that comes first in the individual's cue order and the cue with the plurality of votes (or a random cue among those cues tied for the plurality of votes) is ranked at the position that is voted for, until the complete cue order is established. This implies that an individual can vote for the same cue more than once and that the *plurality rule* uses $k-1$ voting steps for k cues.

Averaging One of the principles that underlies evidence-based approaches to management, medicine and education, is the systematic collection and analysis of empirical evidence that can inform practice. Observations are collected in databases and the effectiveness of a treatment is determined via meta-analysis across studies.

A somewhat similar strategy that can be employed by groups in the setting of the simulation is the pooling of evidence across all individuals within the group. To find a cue order, cues are evaluated by using the collected experience of all group members, there is no voting or evaluation of individual solutions. The rule is called averaging rule, because validities of cues are calculated based on average cue information. The average validities v_i^a in this case are calculated as

$$v_i^a = \frac{\frac{\sum_{j=1}^{n_g} c_j}{n_g}}{\frac{\sum_{j=1}^{n_g} d_j}{n_g}} = \frac{\sum_{j=1}^{n_g} c_j}{\sum_{i=j}^{n_g} d_j}. \tag{2}$$

Table 1: Data sets used to create the simulation task.

Data Set (Source)	Number of Cases	Number of Cues	Validity Average (Range)	Criterion	Cues (Selection)
1) Forbes 500 (StatLib)	79	5	.80 (.67–1.00)	Profit (in million $)	Market value, assets, sales, number of employees, profits, cash flow, sector
2) Ice Cream (StatLib)	29	5	.71 (.52–.97)	Ice cream consumption (4 weeks)	Temperature information, lagged temperature, family income, price, year
3) Minimum Wage (UCLA)	301	11	.54 (.46–.75)	Change in full time employees	Information about state and company, changes in employees, times, registers, salaries , etc.
4) Wildcat Strikes (Simonoff, 2003)	163	4	.60 (.35–.74)	Number of strikes in a company	Number of grievances, union status, rotation status, workforce size
5) CPU Performance (UCI MLR)	209	6	.87 (.79–.95)	Relative CPU Performance	Machine Cycle Time, Cache Memory, Main Memory, Number of channels
6) Land Rent (Weisberg, 1985)	58	4	.71 (.56–.96)	Rent paid per acre	Average rent, cow density, proportion of pasture land, liming requirement
7) Professors' Salary (Rice, 1995)	51	5	.79 (.55–.98)	Salary	Rank, number of years in current rank, highest degree earned, number of years since degree
8) Software Development (JSE)	104	4	.71 (.53–.86)	Total Work Hours	Function point count, operating system, database management system, language
9) Home Prices (StatLib)	117	5	.71 (.51–.95)	Home Price	Square feet, age, taxes, city area, city location, home features
10) Stock Market (UCLA)	368	9	.58 (.50–.73)	Percentage of price change at 26 weeks	Average volatility, price/sales, price/cash, debt/equity, profit margin, ROI, etc.

The group-based cue order is then constructed by ordering cues in descending order of average validity. The number of individual cue discriminations d_i for all group members is adjusted to the average number of discriminations for this cue in the group (which may be a non-integer number). Note that the same cue order would result from using the ratio of the sum of all discriminations and successful predictions (both denominator and numerator are divided by the same number), but the averaging scales the information back to the level of group members. Averaging individual validities though, would ignore the number of observations that each value is based on and would lead to different results in the general case.

Memory alteration

In Dieckmann and Todd (2004), it is assumed that c_i and d_i, the number of discriminations and successful discrimination by each cue, are recorded and recalled accurately by individuals. In this simulation we employ a variant that reflects a more realistic, imperfect memory: in fixed intervals, both c_i and d_i are randomly mutated with the constraint that the resulting order of cues has to remain constant. Our variant is therefore likely to perform worse than the original algorithm.

While this memory update is a handicap for individual learning, it is actually a vital step for the group learning algorithms: after each group phase, an individual that replaces his cue order by the newly constructed group-based order faces a problem otherwise: if he retains his old memories unchanged, there will be a high probability that his cue order will change back to the original order when the validity algorithm is employed in individual learning. Only if the first trial after the group learning phase leads to a change in the order of estimated cue validities will the group phase have any effect on the individual.

In the simulations we therefore use the following mechanism for altering an individual's memory: The cue order and the number of discriminations per cue are retained. Only in the case that any d_i is zero, it will be changed to one. The c_i on the other hand are based on percentages drawn from the uniform distribution $U(0.5;1)$. A minimum of 0.5 is chosen, as cue directions are known a priori, and the minimum cue validity under this constraint is 0.5. A set of k numbers r_i is drawn from this distribution and sorted in descending order $(r_1, r_2, \ldots, r_k)$. The number of successes in memory are subsequently calculated as $c_i = r_i \cdot d_i$. This procedure guarantees that the ordering of cues performed by the validity algorithm will result in the specified cue order.

The retention of all d_i leads to a gradually decreasing probability of switches between cue positions, as more and more observations are needed to change the ratio c_i/d_i substantially and the probability to bridge the randomly determined gaps between successive cues is affected accordingly. This property is shared with the original validity algorithm. This memory alteration can be equally applied to group learners and isolated learners, and as a rule it is applied following each group exchange phase (isolated learners are yoked to the randomization schedule of social learners).

Simulation Setup

Environments Agents in this simulation have to solve the paired comparison task in environments constructed from data sets. None of the data sets was artificially created or hypothetically derived. All were collected in ecologically meaningful economic contexts. In each environment, agents were confronted with object pairs whose members were randomly selected from the cases in the data sets.

Using a range of empirical environments allows to increase the generalizability of our conclusions. The data sets were chosen for variability, with semantic variance and different types and numbers of cues and cases. Table 1 summarizes the set of ten environments used to generate the inference task. The data sets *Forbes 500*, *Ice Cream* (Kadiyala, 1970), and *Home Prices* were taken from the Statlib collection of data sets[1]. *Minimum Wage* (Card & Krueger, 1994) and *Stock Prices* were taken from the UCLA collection of statistical

[1]http://lib.stat.cmu.edu/

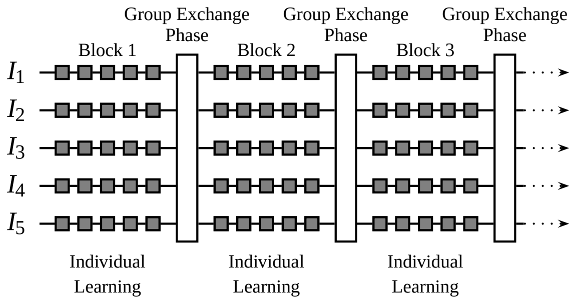

Figure 1: Simulation structure for a group of 5 simulated individuals

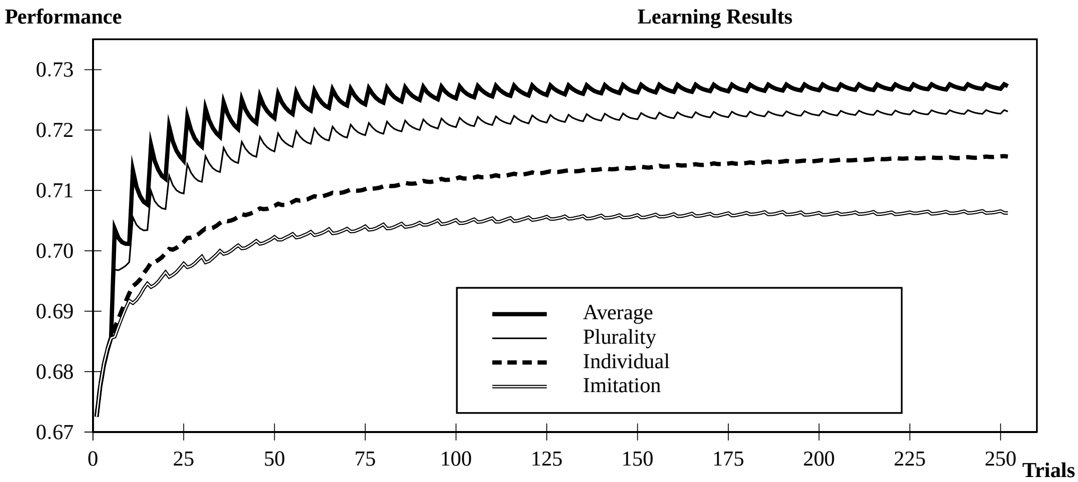

Figure 2: Performance in the base condition: lines depict the average expected accuracy of individual cue orders across individuals and environments for the four simulated learning strategies

data sets[2], *Wildcat Strikes* accompanies (Simonoff, 2003)[3]. The Machine Learning Repository at UCI (Asuncion & Newman, 2007) contributed the *CPU Performance* data (Ein-Dor & Feldmesser, 1987), Weisberg (1985) references the *Land Rent* data. The *Professors' Salary* data were taken from Rice (1995)[4], the *Software Development* data are based on Matson and Huguenard (2007).

Each data set was transformed for use in the simulation. We dichotomized all non-binary cue variables (0/1) using the median (for the *Land Rent* and the *Professor Salary* data) or the mean of each variable (for the rest of the environments). In some cases, only a subset of the original variables was included as some cue-criterion relationships could not be sensibly interpreted. In two cases the number of original variables was reduced to make their inclusion in the simulation feasible. All transformations and selections were applied before running any of our simulations.

Simulation parameters In our base condition, groups consist of five individuals each. All individuals start with random cue orders. Blocks of individual learning trials are interspaced with phases of group exchange (see Figure 1). There are five trials per individual learning block, and in each trial an individual samples one pair of objects, makes an inference and obtains feedback. Each individual j stores three pieces of information: the number of successful predictions in the current trial block (s_j), the number of discriminations that each of the cues in the data set made ($d_{i,j}$ for each cue i, i.e., the number of decisions based on this cue) and the number of successful predictions that were based on each cue i ($c_{i,j}$). Only the number of successes is reset at the beginning of each trial block, the other variables are changed by individual learning and the memory alteration procedure described above.

The simulation proceeds for 50 trial blocks and group exchange phases (i.e., for 250 rounds of individual trials). For each of the three social rules – *imitation*, *plurality* and *averaging* – 2,000 groups (10,000 individuals) were simulated for each data set in the base condition, while we simulated 10,000 isolated learners for comparison.

[2] http://www.stat.ucla.edu/data/
[3] http://www.stern.nyu.edu/j̃simonof/AnalCatData
[4] http://www.amstat.org/publications/jse/jse/data/archive.html

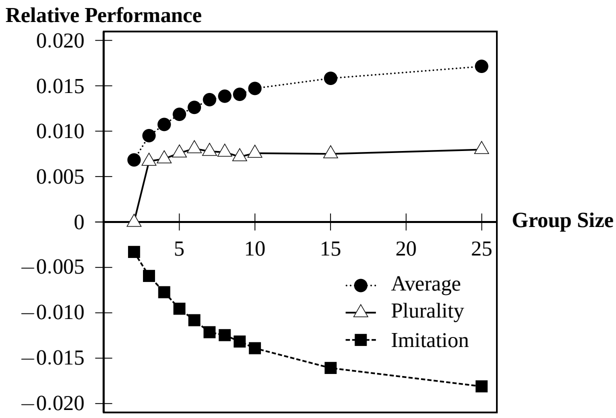

Figure 3: Effect of group size on performance relative to isolated learners after 50 group-exchange phases

Results

Base condition

Round-wise results for the base condition are presented in Fig. 2. Groups that implement the averaging rule and the plurality rule perform better than isolated learners immediately after the first social learning phase with the averaging rule being the best learning algorithm. On the other hand, groups relying on imitation perform even worse than isolated learners and while their performance increases over time, the distance between imitators and isolated learners actually widens.

The results point out a faulty component of the imitation strategy in this context: groups implementing imitation learning are unable to pick out the truly successful strategies based on sample information. What drives these differences between learning strategies and how robust are these findings? A closer analysis of single environments reveals a stable ordering of the algorithms (for each paired comparison between algorithms: $p = .002$, $N = 10$, two-sided binomial test), the results do not seem to be due to particular properties of specific data sets. In the following, we examine the sensitivity of the observed pattern regarding group-size.

The effect of larger groups

To test the effect of group size we simulated 1000 groups for each group size across all data sets and algorithms. We chose group sizes of two to ten members and in addition 15, 20 and 25 members and averaged results after fifty group exchange phases. Fig. 3 demonstrates the effect of this variation on groups implementing the different learning strategies in terms of the relative difference to the average accuracy of isolated learners: The averaging strategy profits the most from larger group sizes with diminishing marginal returns. For the plurality rule, the addition of new members is beneficial up to about five members, while additional members do neither improve nor decrease the performance. Most strikingly though, the performance of imitators decreases with group size, showing the worst result for 25 members.

Discussion and Conclusion

For the plurality rule a group size of two is equivalent to a random choice when the two members disagree, therefore the increase in performance for the first few members is easily explained. The effect of further members has to do with inertia: After each group exchange phase, all individuals start with the social cue order established in that phase. To change the social order in the following social learning phase, a plurality of individuals must have changed their order in the same way. In a larger group, the plurality choice can be considered to be more valid, but it becomes more difficult to change the social order, as each individual learns from different paired comparisons resulting in potentially different changes in individual cue orders. As a plurality of votes is necessary for a change this can lead to inertia.

Imitation learning is yet a special case: With larger group sizes it is more probable to find orders with a successful past track record. Based on the observation that even a random cue order achieves a decent rate of successful comparisons,there is a good chance of finding an individual with perfect or near-to-perfect track record for a block of trials. What creates a first problem, is the fact that the chosen individuals will have learned based on mostly or exclusively successful comparisons, as any failure lowers the success score. The success of individuals could be due to encountering simpler problems in the environment (that can be solved using many cue orders). As candidate solutions are developed by individual learning only, the group will retain strategies that are based on biased samples, on an ”under-sampling of failure” (Denrell, 2003). A second problem for imitation strategies is that a suboptimal change in a single individual order combined with an unfortunate distribution of encountered object pairs can be transferred to all group members in one group exchange phase. It is much more difficult to retain learning increments using this strategy.

In this study, we have identified one class of situations, in which imitation learning based on observed success is not only worse than often feasible social learning alternatives, but it is even worse than isolated individual learning. The problems encountered by the learning algorithm are likely to be endemic in other relevant decision contexts. The inferiority of the learning principle may be difficult to notice, as the performance of imitators increases over trials as for all other strategies, and there might be no other strategies for comparison. Stereotypical causal explanation patterns will assume that something special and unique to the individual must have been responsible for success. This has to be seen in sharp contrast with situations in which the quality of a strategy can be judged using external or logical criteria. The performance of imitation based on mere past success constitutes a lower bound for imitation performance. Concrete evidence for decision makers falling into the trap of taking past success as a proxy for future performance and ignoring sampling and selection processes is given by Offerman and Schotter (2009). In their experiments participants tend to copy the behavior of decision makers that took great risks and were lucky to get the best possible outcome using a risky strategy with a lower expected value than alternative strategies. In an ongoing learning context these problems might well be attenuated and decision makers should eschew the choice of naive benchmarking procedures and mindless imitation of the most successful (Denrell & Liu, 2012; Pfeffer & Sutton, 2006), if they want to channel the power of social learning into obtaining the best possible results.

References

Asuncion, A., & Newman, D. J. (2007). *UCI Machine Learning Repository* (http://www.ics.uci.edu/∼mlearn/MLRepository.html). University of California, Irvine, School of Information and Computer Sciences.

Card, D., & Krueger, A. (1994). Minimum wages and employment: A case study of the fast-food industry in new jersey and pennsylvania. *The American Economic Review*, *84*(4), 772–793.

Czerlinski, J., Gigerenzer, G., & Goldstein, D. G. (1999). How good are simple heuristics. In G. Gigerenzer, P. M. Todd, & the ABC Reseach Group (Eds.), *Simple heuristics that make us smart* (p. 97-118). Oxford University Press, USA.

Denrell, J. (2003). Vicarious learning, undersampling of failure, and the myths of management. *Organization Science*, 227–243.

Denrell, J. (2005). Should we be impressed with high performance? *Journal of Management Inquiry*, *14*(3), 292–298.

Denrell, J., & Liu, C. (2012). Top performers are not the most impressive when extreme performance indicates unreliability. *Proceedings of the National Academy of Sciences*, *109*(24), 9331–9336.

Dieckmann, A., & Todd, P. M. (2004). Simple ways to construct search orders. In K. Forbus, D. Gentner, & R. Regier (Eds.), *Proceedings of the 26th annual conference of the cognitive science society* (pp. 309–314). Mahwah, NJ: Erlbaum.

Ein-Dor, P., & Feldmesser, J. (1987). Attributes of the performance of central processing units: a relative performance prediction model. *Commun. ACM*, *30*(4), 308–317.

Garcia-Retamero, R., Takezawa, M., & Gigerenzer, G. (2006). How to learn good cue orders: When social learning benefits simple heuristics. In R. Sun & N. Miyake (Eds.), *Proceedings of the 28th annual conference of the cognitive science society* (pp. 1352–1358). Mahwah, NJ: Erlbaum.

Garcia-Retamero, R., Takezawa, M., & Gigerenzer, G. (2009). Does imitation benefit cue order learning? *Experimental Psychology*, *56*, 307–320.

Garcia-Retamero, R., Takezawa, M., Woike, J. K., & Gigerenzer, G. (2013). Social learning: a route to good cue orders. In R. Hertwig, U. Hoffrage, & the ABC Research Group (Eds.), *Simple heuristics in a social world* (pp. 343–354). New York: Oxford University Press.

Gigerenzer, G., & Brighton, H. (2009). Homo heuristicus: Why biased minds make better inferences. *Topics in Cognitive Science*, *1*(1), 107-143.

Hertwig, R., Barron, G., Weber, E., & Erev, I. (2004). Decisions from experience and the effect of rare events in risky choice. *Psychological Science*, *15*(8), 534–539.

Kadiyala, K. R. (1970). Testing for the independence of regression disturbances. *Econometrica*, *38*, 97–117.

Martignon, L., & Schmitt, M. (1999). Simplicity and robustness of fast and frugal heuristics. *Mind Mach*, *9*(4), 565-593.

Matson, J. E., & Huguenard, B. R. (2007). Evaluating aptness of a regression model. *Journal of Statistics Education*, *15*(2).

Offerman, T., & Schotter, A. (2009). Imitation and luck: an experimental study on social sampling. *Games and Economic Behavior*, *65*(2), 461–502.

Pfeffer, J., & Sutton, R. I. (2006). Evidence-based Management. *Harvard Business Review*, *84*(1), 62–74.

Rendell, L., Boyd, R., Cownden, D., Enquist, M., Eriksson, K., Feldman, M., et al. (2010). Why copy others? insights from the social learning strategies tournament. *Science*, *328*(5975), 208–213.

Rice, J. A. (1995). *Mathematical statistics and data analysis* (2nd ed.; D. Press, Ed.). Belmont, California.

Rieucau, G., & Giraldeau, L. (2011). Exploring the costs and benefits of social information use: an appraisal of current experimental evidence. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *366*(1567), 949–957.

Simonoff, J. S. (2003). *Analyzing Categorical Data*. New York: Springer.

Todd, P. M., & Dieckmann, A. (2012). Simple rules for ordering cues in one-reason decision making. In P. M. Todd, G. Gigerenzer, & the ABC Research Group (Eds.), *Ecological rationality: Intelligence in the world* (pp. 1393–1400). New York: Oxford University Press.

Weisberg, S. (1985). *Applied linear regression* (Wiley, Ed.). New York.