Attractivity Weighting: Take-the-Best's Foolproof Sibling

Paul D. Thorn (thorn@phil.hhu.de)

Gerhard Schurz (schurz@phil.hhu.de)

DCLPS, Theoretical Philosophy, University of Duesseldorf,
Universitaetsstr. 1, Duesseldorf 40225, Germany

Abstract

We describe a prediction method called "Attractivity Weighting" (AW). In the case of cue-based paired comparison tasks, AW predicts a weighted average of the cue values of the most successful cues. In many situations, AW's prediction is based on the cue value of the most successful cue, resulting in behavior similar to Take-the-Best (TTB). Unlike TTB, AW has a desirable characteristic called "access optimality": Its long-run success is guaranteed to be at least as great as the most successful cue. While access optimality is a desirable characteristic, concerns may be raised about the short-term performance of AW. To evaluate such concerns, we here present a study of AW's short-term performance. The results suggest that there is little reason to worry about the short-run performance of AW. Our study also shows that, in random sequences of paired comparison tasks, the behavior of AW and TTB is nearly indiscernible.

Keywords: Bounded Rationality; Ecological Rationality; Attractivity Weighting; Take-the-Best; Meta-induction.

Prediction Games

The object of study within the present paper is a prediction method known as "Attractivity Weighting" (AW). AW was introduced (under the name "weighted meta-induction") as a possible response to Hume's problem of induction (Schurz, 2008) based on findings in mathematical learning theory (cf. Cesa-Bianchi & Lugosi, 2006). The formal properties that make AW an attractive prediction method were demonstrated in the context of so called "prediction games." Prediction games will serve as a framework for evaluating prediction strategies within the paper. The main aim of the paper will be to assess the short-term performance of AW. As a secondary result, we will observe the near indiscernability of the behavior of AW and another method, called "Take-the-Best," in random sequences of paired comparison tasks.

Informally, a prediction game consists of a sequence of events, along with a collection of participating prediction methods. In a series of *rounds*, the participating methods deliver predictions about the character of successive elements of the event sequence, predicting the character of the first event in the first round, the second in the second, etc. At the end of each round, the actual value of the just considered event is revealed.

Formally, a prediction game is defined as a pair, (E, M), where E = (e_1, e_2, ...) is an infinite sequence of events, and M is a finite set of methods. For example, the elements of E may be the weather conditions (rainy or not rainy) for a sequence of days. For convenience, the elements of E are assumed to be real numbers in the interval [0, 1]. The

elements of M are of two sorts. Methods of the first sort, called "object methods," make their predictions independently of the other methods. Methods of the second sort, called "meta-methods," make their predictions based on the predictions of the object methods. As a basis for making their predictions, it is assumed that each meta-method has 'access' to the present and past predictions of each object method.

Within a prediction game, the *prediction of a method*, m, of the value of the nth event, e_n, is denoted $P_n(m)$. The normalized *loss* for an individual prediction, $P_n(m)$, is a function of the distance between the prediction and the event's value, and takes a value in [0, 1]. The *score* for a method, m, for event n is denoted $S_n(m)$, and is defined as 1 *minus* the loss for the prediction. By default, we assume that losses are measured by *linear* distance, that is: $S_n(m) = 1 - |e_n - P_n(m)|$.[1]

The results of the following section also depend on the assumption that each method makes a prediction concerning each event in the event sequence. In order to accommodate naturally occurring situations where some methods do not make a prediction concerning some events, we treat non-predictions as *ersatz* predictions, as distinguished from *genuine* predictions. Ersatz predictions are recorded as a prediction of 0.5, and scored accordingly (cf. Martignon & Hoffrage, 1999).

Prediction games represent a relatively general framework. For example, the framework is apt for representing cue-based paired comparison tasks. In that case, (i) cues are treated as object methods, and (ii) methods, such as Take-the-Best (TTB), that make their predictions based on cues are treated as meta-methods: TTB proceeds by ordering the available cues (object methods) according to their 'ecological' validity, and imitates the prediction of the first cue (object method) in the ordering that delivers a genuine prediction. The *ecological validity* of a method m, as of round n, is equated with the average score of the method for the genuine predictions that it made within the first n rounds.

Attractivity Weighting

The main object of our discussion will be the meta-method AW. AW's predictions are formed by taking a weighted average of the predictions of the (accessible) object methods. The weights that AW assigns to the object

---

[1] More generally, the results of the following section depend on the assumption that the loss function, L, is *convex*, i.e., for all *r*, d_i (distance *i*), and d_j (distance *j*): $L((1-r)d_i + rd_j) \leq (1-r)L(d_i) + rL(d_j)$.

methods are called "attractivities." The attractivity of an object method, at a given round, is determined by comparing the average score of the object method with the average score of AW.[2] If the average score of the object method is less than or equal to AW's score, then the attractivity of the method is *zero*. If the average score of the object method is greater than that of AW, then the attractivity of the method is equal to the difference between the two averages. Thus the attractivity of a method m after the nth round is defined as follows, where $\bar{S}_n(m)$ denotes the average score of a method m for the first n rounds (i.e., $\bar{S}_n(m) = \sum_{i=1}^{n}(S_i(m)/n)$):

$$at_n(m) = \begin{cases} \bar{S}_n(m) - \bar{S}_n(AW) & \text{if } \bar{S}_n(m) - \bar{S}_n(AW) > 0, \\ 0, & \text{else.} \end{cases}$$

AW's prediction in round n is based on the attractivities assigned in round n−1. In the case where all object methods are assigned zero attractivity in round n−1, it is stipulated that AW imitates the prediction of the object method whose ecological validity is the greatest as of round n−1 (with ties broken by a randomized tie-breaker). The method imitated by AW in such cases is denoted "maxVal$_{n-1}$."[3] In round one, AW predicts 0.5. Formally, the predictions of AW are defined as follows, where m ranges over the set of accessible object methods, and $n > 1$:

$$P_n(AW) = \begin{cases} \dfrac{\sum_m P_n(m)\, at_{n-1}(m)}{\sum_m at_{n-1}(m)} & \text{if for some m: } at_{n-1}(m) > 0, \\ P_n(maxVal_{n-1}), & \text{else.} \end{cases}$$

In the following section, we explain the result that AW has a desirable characteristic called "access optimality."

The Virtues of AW

An important characteristic of AW is that it is access optimal: In the long-run, the mean score of AW is guaranteed to converge to the mean score of the best scoring object method to which it has access. In other words, where "maxSuc$_n$" denotes the average score of the best scoring object method as of round n, $\bar{S}_n(AW)$ goes to maxSuc$_n$, as n goes to infinity. In the short-term, the difference between the average score of AW and of the best scoring object method is bounded by $\sqrt{k/n}$, where k is the number of object methods (i.e., for all prediction games, and all n: $\bar{S}_n(AW) + \sqrt{k/n} \geq$ maxSuc$_n$) (Schurz, 2008). If the number of accessible object methods, k, is large, then the *worst case* short-term

performance of AW need not be very good. We will return to this point below.

Access optimality is an important characteristic, in the context of Hume's problem of induction. Hume's problem is easily illustrated within the framework of prediction games. Within prediction games, the character of the event sequence is unconstrained. This means that, regardless of the character of the preceding n−1 events, event n may take any value whatsoever (within [0, 1]). This implies that there is no sure-fire way to exploit the observation of past events in order to make accurate predictions about future unobserved events. Access optimality offers a means to mitigate this problem: Granted that there is no way to ensure good performance within a prediction game, applying AW ensures that one does no worse than the best scoring object method to which one has access.

To date, the only prediction methods that are known to be access optimal are variants of AW (e.g., variants that employ exponentially weighted attractivities). Beyond, this, it is demonstrable that all one-favorite meta-methods, including TTB, are not access optimal. That is, any method that forms its prediction, for each event, by imitating the prediction of a single object method (or cue) is not access optimal. It is also demonstrable that well-known weighting methods such as multiple linear regression and Franklin's Rule (see below) are not access optimal.

It is easy to see why one-favorite meta-methods are not access optimal. Consider a prediction game with ten object methods, and a one-favorite method called "Mono." Suppose that the predictions of the object methods are highly accurate when they are not imitated by Mono, and highly inaccurate when they are imitated by Mono (so that there is a negative correlation between the score of an object method and its being imitated). In such circumstances, the predictions of Mono will be highly inaccurate, while, for each event, the predictions of nine of the ten object methods will be highly accurate. Beyond the theoretical possibility of situations in which one-favorite methods fail to perform well (as illustrated by a simulation presented in Schurz and Thorn, 2016, fig. 3), there is a wide range of naturally occurring situations where one-favorite methods, such as TTB, perform poorly. The problem arises in situations where the payoff for performing a given action is an inverse function of the number of individuals who perform the action. Such cases may arise when the task is to determine where one should go in attempting to gather a seasonal resource (e.g., fish or berries). In such cases, widespread adoption of TTB applied to cues concerning the past productivity of given locations will drive each member of a population of TTBers to attempt to gather resources at the same location, resulting in a poor mean payoff for the TTBers.[4] Similar dynamics may be observed in a wide range of tasks, including market entry problems, career

---

[2] An alternative variant of AW, called "intermittent AW" in (Schurz & Thorn, 2016), may be formulated by identifying the attractivity of an object method with the difference between its ecological validity and the ecological validity of (intermittent) AW. We here investigate AW, since it is more frugal (cf. Newell, Rakow, Weston, & Shanks, 2004).

[3] The access optimality of AW holds regardless of how AW makes its predictions in cases where all object methods have zero attractivity.

[4] In action games, where actions take the place of predictions, a weighted average of several actions is interpreted as a probabilistic mixture, i.e., the weights are interpreted as the probabilities of emulating respective actions.

choice, cuing problems, route selection, departure time selection, etc. While TTB performs poorly in such tasks, the access optimality of AW ensures good performance. For empirical studies of human performance in such tasks, see Rapoport, Seale, Erev, and Sundali (1998), Rapoport, Stein, Parco, and Seale (2004), and Rapoport, Gisches, Daniel, and Lindsey (2014).

Despite its access optimality, it is possible to raise concerns about the short-term performance of AW, as we did in Schurz and Thorn (2016). In order to form a clearer picture of the concern, consider 'typical' environments where (i) the observed ecological validities of the accessible cues quickly approach their actual (long-run) ecological validities, as the number of observed items increases, and (ii) the average score of cues does not vary according to their use by meta-methods (contrary to the environments described in the preceding paragraph). We call environments meeting conditions (i) and (ii) "non-elusive."

It is possible to distinguish two sorts of non-elusive environment: *compensatory* (where there are methods of linear weighting that outperform TTB), and *non-compensatory* (where no method of linear weighting outperforms TTB) (cf. Martignon & Hoffrage, 1999). Concerns may be raised about the performance of AW in both sorts of environment. *First*, AW will perform worse than TTB, in the short-run, in non-compensatory environments. In such environments, TTB's one-favorite approach is appropriate, and while it is demonstrable that AW will adopt the behavior of a one-favorite method in the long-run (proceeding until only one object method has an attractivity greater than zero), this will take some time. The extent to which TTB will outperform AW, in the short-run, in typical non-compensatory environments is an open question. Using the simulation studies reported below, we attempted to answer this question. *Second*, AW (like TTB) will typically perform worse than *some* other weighting strategies (both in the long and short-run) in compensatory environments. For example, Schurz and Thorn (2016) show that Franklin's Rule (described below) outperforms AW in some compensatory environments. One possible solution to the present problem would be to include a variety of meta-methods that are known to perform well in compensatory environments among the set of methods to which AW has access, thereby permitting AW to emulate another weighting method, different from AW, in appropriate situations. We call this refined version of AW "vAW." (TTB could be adapted in a similar manner.) A residual worry may be raised regarding this proposal: Regardless of the capacity of vAW to emulate the behavior of a well-adapted weighting method in the long-run, the performance of the weighting method will probably exceed the performance of vAW in the short-run. This worry is analogous to the one that arose in the comparison of AW to TTB in typical non-compensatory environments: AW will lose out, in the short-run, in the midst of learning which method it should emulate in the long-run. Once again, we will use the simulation studies reported below to address the magnitude of this problem.

The Simulations

In order to evaluate the short-term performance of AW, we simulated prediction games using data sets characterizing natural environments. In particular, we tested the performance of AW, along with a number of other meta-methods, using the twenty data sets used by Czerlinski, Gigerenzer, and Goldstein (1999) in evaluating the performance of TTB. These data sets are heterogeneous, and representative of a wide range of environments, involving the prediction of city population, attractiveness of persons, high school dropout rates, homelessness rates, mortality rates, house prices, professor salaries, automobile fuel consumption, body fat, fish fertility, mammal sleep duration, biodiversity, rainfall, and atmospheric conditions.[5] As with Czerlinski, Gigerenzer, and Goldstein (1999), we used the data to formulate paired comparison tasks, i.e., tasks where a method must judge which of two objects (e.g., two German cities) has a greater *criterion value* (e.g., population) on the basis of a number of cues (e.g., whether a city has a university).

For each data set, we generated 1,000 prediction games (with finite event sequences).[6] Each prediction game, for a given data set, was generated by forming a random sequence of the set of all pairs of objects that the data set concerned. For example, for the data set that concerned the population of German cities, we generated a random sequence of the pairs of German cities (excluding repeats, e.g., if (Munich, Hamburg) was included (Hamburg, Munich) was not). The sequence of object pairs formed the basis for a sequence of paired comparison tasks. The value, e_i, for the ith element of event sequence was read off of the sequence of object pairs: If the first element of the pair had a greater criterion value, then $e_i = 1$, otherwise $e_i = 0$. The participating object methods were simply the cues for the corresponding data set. Cue values were determined as follows: (i) if the first element of the pair had a higher cue value, then the cue predicted that the first element has a higher criterion value (i.e., $P_i(cue) = 1$), (ii) if the second element of the pair had a higher cue value, then $P_i(cue) = 0$, and (iii) if the cue values were identical, then the result was an ersatz prediction, with $P_i(cue) = 0.5$. Notice that the preceding manner of defining cue predictions leaves open the possibility that any given cue's validity, for the set of all events, is less than 0.5. But if the validity of a cue is very low (and below 0.5), then a meta-method such as TTB may wish to form its predictions by predicting counter to the cue. To accommodate this possibility, we allowed meta-methods to consider the 'counter-cue', corresponding to any given

---

[5] We used the data sets with dichotomized cue values that are available at: http://www-abc.mpib-berlin.mpg.de/sim/Heuristica/

[6] Given the standard error of the mean (not reported) for each of the mean values reported below, samples of 1,000 were needlessly large, according to reasonable standards for reporting simulation results (Bindel & Goodman, 2009).

cue, that is: (i) if a cue predicts 0, its counter-cue predicts 1, (ii) if a cue predicts 1, its counter-cue predicts 0, and (iii) if a cue predicts 0.5, its counter-cue also predicts 0.5.

The participating meta-methods included AW and TTB, along with an array of other methods, including Franklin's Rule, the Minimalist, multiple linear regression (MLR), and Dawes' Rule. The latter three methods, along with TTB, were considered by Czerlinski, Gigerenzer, and Goldstein (1999).[7] Finally, we considered a variant of AW, called "vAW," that treats Franklin's Rule, MLR, and Dawes' Rule as if they were accessible object methods. The methods made their predictions as follows:

AW. The data sets employed in our simulations result in many cases where some of the relevant cues do not deliver genuine predictions. We introduced a slight modification of AW, in order to improve its performance in such situations: In determining AW's prediction (i.e., $P_n(AW)$), according to the equation described above, we assumed that, in any given round, m ranges only over those cues (and counter-cues) that delivered a genuine prediction in that round. In the case where no cue delivers a prediction AW predicts 0.5. AW remains access optimal with these modifications, so long as we restrict ourselves to prediction games where, for each cue, the average score of AW's predictions in cases where the cue was attractive but made no genuine prediction exceeds 0.5, in the long-run.

vAW. As AW, save that Franklin's Rule, MLR, and Dawes' Rule are also treated as accessible object methods

TTB. In each round n, TTB forms its prediction by (i) ordering the accessible cues and counter-cues by their observed ecological validity as of round n−1, and (ii) emulates the prediction of the first cue (or counter cue) in the ordering that delivers a genuine prediction. If no cue makes a prediction or if all cues have undefined ecological validities (as in round one), TTB predicts 0.5.

The Minimalist (Min). As TTB, save that the cue order is determined at random.

Franklin's Rule (FR). In each round n, Franklin's Rule predicts a weighted average of the genuine predictions of cues (or counter-cues) whose observed ecological validity is at least 0.5. The weight for each cue/counter-cue, in round n, is proportional to its ecological validity as of round n−1. As with TTB, if no cue makes a prediction or if all cues have undefined ecological validities (as in round one), Franklin's Rule predicts 0.5.

Dawes' Rule. As Franklin's Rule, save that each cue is assigned equal weight.

MLR. For each round n, the predictions of MLR were determined by finding the ordinary least squares multiple regression model for predicting the criterion values from the

cue values, based on the objects that had been observed prior to round n. MLR's prediction about which of two objects has a higher criterion value was determined by which object had a higher predicted criterion value according to the regression model. For rounds where there is insufficient data to find a regression model, MLR predicts 0.5.

AW's Short-term Performance

Within our simulations, there was no considerable lag in the performance of AW in comparison to TTB, or in the performance of vAW in comparison to the best performing alternative weighting strategy (which was generally MLR). Figure 1 gives a sense of these results, showing the average scores of the considered meta-methods at various rounds of the simulated prediction games (i.e., the averages for the 20,000 prediction games, based on 1,000 simulations for each of the 20 data sets). For the sake of readability, we do not plot the values for Franklin's Rule, whose mean performance was similar to TTB and AW.
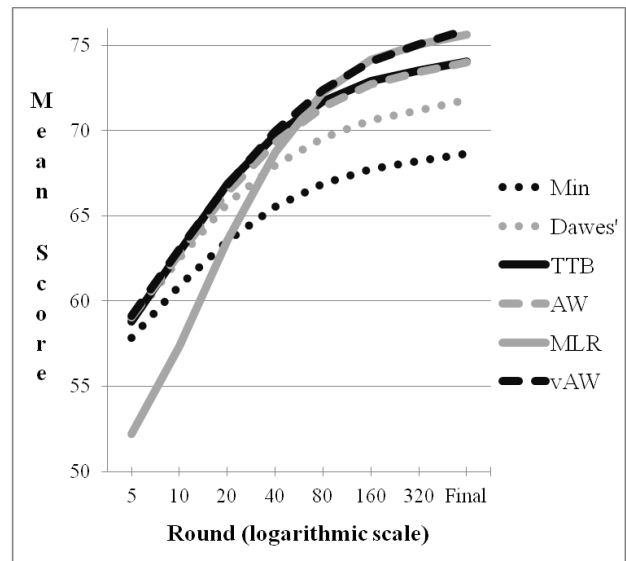

Figure 1: Average scores for the competing meta-methods, at various rounds.

In assessing the gravity of the concerns raised in the preceding section, it is important to note that the performance of AW closely matches that of TTB. While TTB usually performed slightly better than AW, there were only four data sets where its mean performance exceeded that of AW by more than 0.01, at any stage. This indicates that the access optimality of AW does *not* come at the expense of short-term performance, in comparison with TTB. Similarly, on average, vAW closely approximated the performance of the best performing method, from the early through the later stages of respective games. This indicates that a simple modification of AW enables performance that quickly matches the performance of alternative weighting methods, when appropriate.

---

[7] We followed Czerlinski, Gigerenzer, and Goldstein (1999) in including MLR as a meta-method. In an extended study, we plan to also include logistic regression.

For purposes of comparison with the results of Czerlinski, Gigerenzer, and Goldstein (1999), we recorded the average frugality of the competing meta-methods (i.e., the number of cue value pairs that the meta-methods needed to access before making a prediction). Table 1 shows the results.

Subsequent to observing the similarity of AW and TTB in terms of mean scores and mean frugality, we formed the hypothesis that, in non-elusive environments, the behavior of AW was very similar to that of TTB. In the following section, we report some results that support this hypothesis.

Table 1: Mean frugality of the competing meta-methods.

Method	Frugality
Min	2.2
TTB	2.3
AW	2.5
vAW	7.1
MLR	7.3
Dawes'	7.4
FR	7.4

The Near Indiscernibility of AW and TTB in Non-elusive Environments

In attempting to determine what judgment strategies humans use in paired comparison tasks, two sorts of data have been most important: *outcome patterns* and *process tracing* (cf. Bröder, 2012). Outcome patterns consist of data concerning the responses that subjects provide in the face of given paired comparison tasks (Bröder, 2003; Bröder & Schiffer, 2006; Rieskamp & Otto, 2006; Rieskamp, 2008). Rather than considering responses, process tracing monitors information acquisition patterns (Newell & Shanks 2003; Newell, Weston, & Shanks 2003). Information concerning response times (Bröder & Gaissmaier, 2007), and subject self-reporting (Walsh & Gluck, 2016), may also be relevant in determining what judgment strategies humans use.

By the analysis of outcome patterns, and (to a lesser extent) process tracing, some psychological studies are thought to corroborate the claim that TTB plays *some* role in human reasoning. Based largely on studies of Bröder (2003), Rieskamp and Otto (2006), and Rieskamp (2008) (that investigated human behavior in non-elusive environments), the received view is that human beings are adaptive in the strategies they use in making predictions, and that, in appropriate environments, subjects are disposed to (learn to) use TTB in making predictions. Note, however, that the theoretical conclusions of the above mentioned studies are based on maximum likelihood techniques (Bröder, 2003) and model fitting (Rieskamp & Otto, 2006; Rieskamp, 2008), and that these methods are based on comparative evaluations of candidate hypotheses. This means (as is acknowledged by the authors of the aforementioned studies) that data that strongly supports one hypothesis (e.g., that subjects are using TTB to accomplish

a given prediction task) among a given pool of hypotheses, doesn't necessarily support the hypothesis over others that were not considered (e.g., the hypothesis that subjects are using AW to accomplish the task). Our general point is certainly not a new one, and it is not our intent to disparage the use of comparative methods of evaluation. Our only intent is to suggest that existing data does not tell in favor of the adaptive use of TTB over that of AW. In order to support the present claim, we report results bearing on the degree of confluence between TTB and AW, both in terms of their predictions (outcome patterns), and informational demands (process tracing).

The data reported in Tables 2 and 3 is from the simulations described above. Table 2 reports the percentage of trials in which the predictions of respective meta-methods agreed with the predictions of TTB. The averages reported here are the averages of the averages for 1,000 simulations for the 20 environments. In other words, we first collected the averages for each of the 20 environments. We then took the averages of those averages. Table 2 reports the average percentage of trials in which the informational demands of respective meta-methods (i.e., the profile of cues accessed) were identical to the informational demands of TTB. The averages reported are, again, the averages of the averages for 1,000 simulations for the 20 environments.

Table 2: Mean percentage of predictions identical with TTB.[8]

Method	Agreement w/ TTB
AW	98%
FR	93%
MLR	85%
Dawes'	83%
Min	83%

Table 3: Mean percentage of cases where cues accessed were identical with TTB.

Method	Agreement w/ TTB
AW	93%
FR	13%
MLR	12%
Dawes'	12%
Min	19%

As is evident from the data, the behavior of AW is very similar to that of TTB in non-elusive environments. It is of interest to note, for example, that the degree of overlap between AW and TTB is far greater than the degree of predictive fit between subject behavior and the best fitting models that have been offered in the literature, such as that of Rieskamp and Otto (2006). Absent the explicit intention

---

[8] If we require that all of AW's predictions are rounded to the nearest integer, then the percentage agreement with TTB is 99.5.

to create a prediction task in which TTB and AW produce different predictions, it is unlikely that one would produce a task that could be used as evidence in favor of the adaptive use of TTB over AW, and vice versa. Matters are complicated by the fact that human subjects appear to be adaptive in the methods that they deploy, employing weighting methods by default (Bröder, 2003; Rieskamp & Otto, 2006; Rieskamp, 2008).

Conclusions

The results of our simulations suggest that there is no reason to worry about the short-term performance of AW, in comparison to TTB. Similarly, a variant of AW can be formulated that performs well in both compensatory and non-compensatory environments.

Our simulations also suggest that the behavior of AW is nearly indiscernible from that of TTB, in non-elusive environments. It seems, then, that the possibility cannot be excluded that AW plays some role in human reasoning, inasmuch as there is data consistent with the hypothesis that subjects are adaptive users of TTB, in non-elusive environments. Considerations of 'prior intuitive plausibility' may favor the hypothesis that subjects use TTB rather than AW. On the other hand, the foolproof nature of AW, underwritten by its access optimality, suggests that AW is more adaptive than TTB: Inasmuch as we expect human cognition to be adapted to its environment (an environment in which the application of one-favorite methods can lead to catastrophic failure), there is some reason to expect that something like AW plays some role in human cognition. Such musings are, of course, an impetuous (rather than a substitute for) proper empirical studies, which have yet to be conducted.

Acknowledgments

References

Bindel, D., and Goodman, J. (2009). *Principles of Scientific Computing*. Retrieved from http://www.cs.nyu.edu/courses/spring09/G22.2112-001/book/book.pdf

Bröder, A. (2003). Decision making with the "adaptive toolbox": Influence of environmental structure, intelligence, and working memory load. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 29,* 611-625.

Bröder, A. (2012). The quest for take the best - Insights and outlooks from experimental research. In Todd, P., & Gigerenzer, G. (Eds.), *Ecological rationality: Intelligence in the world*, New York: Oxford University Press.

Bröder, A., & Schiffer, S. (2006). Adaptive flexibility and maladapative routines in selecting fast and frugal decision strategies. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *32,* 904-918.

Cesa-Bianchi, N., & Lugosi, G. (2006). *Prediction, learning, and games.* Cambridge: Cambridge Uniersity. Press.

Czerlinski, J., Gigerenzer, G., & Goldstein, D.G. (1999). How good are simple heuristics? In: Gigerenzer et al. (Eds.) *Simple heuristics that make us smart*, (pp. 97-118).

Gigerenzer, G., Todd, P.M., & The ABC Research Group (Eds.). (1999). *Simple heuristics that make us smart*. Oxford: Oxford University Press.

Martignon, L., & Hoffrage, U. (1999). Why does one-reason decision making work? In: Gigerenzer et al. (Eds.). *Name of the book* (pp. 119-140).

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,* 53-65.

Newell, B.R., Weston, N.J., & Shanks, D.R. (2003). Empirical tests of a fast-and-frugal heuristic: Not everyone "takes-the-best." *Organizational Behavior & Human Decision Processes*, *91*, 82-96.

Rapoport, A., Gisches, E., Daniel, T., & Lindsey R. (2014). Pre-trip information and route choice decisions with stochastic travel conditions: Experiment. *Transportation Research Part B*, 68, 154-172.

Rapoport, A. Seale, D.A., Erev, I., & Sundali, J.A. (1998). Equilibrium play in large market entry games. *Management Science*, 44, 119-141.

Rapoport, A., Stein, W.E., Parco, J.E., & Seale, D.A. (2004). Equilibrium play in single-server queues with endogenously determined arrival times. *Journal of Economic Behavior and Organization*, 55, 67-91.

Rieskamp, J., & Otto, P. (2006). SSL: A theory of how people learn to select strategies. Journal of Experimental Psychology: General, 135(2), 207-236.

Rieskamp, J. (2008). The importance of learning when making inferences. *Judgment and Decision Making*, 3(3), 261-277.

Schurz, G. (2008). The meta-inductivist's winning strategy in the prediction game: A new approach to Hume's problem. *Philosophy of Science, 75,* 278-305.

Schurz, G., & Thorn, P. (2016). The Revenge of Ecological Rationality: Strategy-Selection by Meta-Induction Within Changing Environments. *Minds and Machines*, 26(1-2), 31-59.

Todd, P. M., & Gigerenzer, G. (Eds.). (2012). *Ecological rationality: intelligence in the world.* New York: Oxford University Press.

Walsh, M.M., & Gluck, K.A. (2016) Verbalization of Decision Strategies in Multiple-Cue Probabilistic Inference. *Journal of Behavioral Decision Making*, *29,* 78-91.