Limited Search and Environmental Regularities: Evaluating the Process Coherence of Take-the-best

Michael D. Lee and Shunan Zhang
Department of Cognitive Sciences
University of California, Irvine

Abstract

Heuristic decision-making models, like take-the-best, rely on environmental regularities. They conduct a limited search, and ignore available information, by assuming there is structure in the decision-making environment. Take-the-best relies on at least two regularities: *diminishing returns*, which says that information found earlier in search is more important than information found later; and *correlated information*, which says that information found early in search is predictive of information found later. We develop new approaches to determining search orders, and to measuring cue discriminability, that make the reliance of take-the-best on these regularities clear, and open to manipulation. We then demonstrate, in the well-studied German cities environment, when and how these regularities provide grounds for the limited search used by take-the-best. In particular, we show that when the two environmental regularities are present, the decisions made on limited search are unlikely to have changed after exhaustive search, but that both regularities are necessary. We discuss how this measure of process coherence provide justification for, and bounds on, models of non-compensatory decision-making like take-the-best.

Introduction

The 1992 Olympics was the first time professionals from the US NBA league were allowed to play in the men's basketball competition. The US 'Dream Team', filled with stars like Michael Jordan, Magic Johnson, Larry Bird, Scottie Pippen and Charles Barkley was one of the most dominant teams ever assembled for any sporting competition. Their closest game was a 117–85 victory in the final over Croatia, and head coach Chuck Daly never felt the need to call a timeout during the tournament.

Making predictions about the outcomes of sporting contests is notoriously difficult, but the Dream Team made some predictions easy. Imagine trying to predict whether or not the US would beat its first opponent in the tournament, Angola, and examining the players in each team in sequence. At some point early in the US list—maybe after Jordan, Johnson, Bird, Pippen, and Barkley—there would be no need to look further. No matter who else was on the US roster, or the Angolan roster, the outcome is already clear. The Dream Team also made predictions easy during the course of games. With about 5 minutes to play in the first half against Angola, the US led 45 to 8. It was clear the US would win by a large margin, without watching the rest of the game.

Both of these decisions about a US victory are *non-compensatory*, because not all of the available information is used. The remaining player rosters are not examined, and the rest of the game is not watched. Yet the decision to forego further information seems sensible in both cases. It is not a reaction to limited time or cognitive resources, but a recognition of the nature of the environment in which decision-making is taking place. The first few US players are so good that there are no other players who could lead Angola to victory, and the first half performances of the teams are highly predictive of second half performances.

Fast and Frugal Heuristics

The adaptive value of non-compensatory decisions lies at the heart of the 'fast and frugal' approach to modeling cognition (Gigerenzer & Goldstein, 1996; Gigerenzer & Todd, 1999). This approach has developed simple and effective heuristic models of human judgment and choice, built around two compelling ideas. One is that decision-making needs to be fast, because the world is competitive. The other is that good decision-making mechanisms should be frugal, in the sense of being simple, because the world is changeable place. Simplicity makes them robust, just as a machine with few moving parts is unlikely to break.

As Gigerenzer and Todd (1999) emphasize, neither of these motivating principles are about cognitive limitations. Fast and frugal heuristics are effective to the extent they seize on the opportunities presented by environmental regularities. They are rooted in properties of environments, not limits of memory, bounds on cognitive processing capabilities, or other internal constraints. What is important is that the heuristics mesh with their environment, allowing limited capacity cognitive processes to function effectively (Simon, 1956, 1990; Todd & Gigerenzer, 2003).

Two of the most important environmental regularities are highlighted in making predictions about the Dream Team. First, if the environment is searched so additional information provides *diminishing returns*, with less useful information being found as search progresses, it may be sensible to make an early decision. Second, if there is *correlated information*, so that what is found early in search is predictive of what will be found later, it can also be sensible to make an early decision.

The basic idea of fast and frugal heuristics is that human decision-making exploits these sorts of possibilities. By assuming the environment has structure, non-compensatory decision processes can be used that are fast, robust and accurate.

Overview

In this paper, we study one of the most prominent fast and frugal heuristics—the take-the-best model of decision-making—in terms of the diminishing returns and correlated information environmental regularities. To do this, we analyze the take-the-best model in a way that allows its behavior under different assumptions about the environment to be studied. We also extend the way heuristics like take-the-best have previously been justified, focusing not on the outcome or correspondence measure of their decision accuracy, but on a process or coherence measure of whether more exhaustive search changes the decisions suggested by limited search. Through a case study with the widely-considered German cities environment, we show that there are good grounds for the limited search assumed by take-the-best, but only when both the environmental regularities are satisfied.

Environmental Regularities and Take-The-Best

Perhaps the best developed fast and frugal heuristic is take-the-best (Gigerenzer & Goldstein, 1996; Gigerenzer & Todd, 1999), which chooses between two alternatives, each represented by a set of binary cues. Take-the-best says that people search these cues in a specific order, determined by a measure called cue validity, and terminate search as soon as a discriminating cue is found.

To make clear the reliance of take-the-best on the environmental regularity of diminishing returns, we develop a new approach to determining cue search orders. To make clear the reliance of take-the-best on the environmental regularity of correlation information, we develop a refined measure of cue discriminability. Each of these theoretical developments is best described in concrete terms, and so we introduce them in terms of the German cities environment used in our later evaluation.

The German Cities Environment

The German cities environment is probably the most widely-used environment for studying fast and frugal heuristics. It represents 83 cities using 9 cues, and lists the population of each city, with the goal of using the cues to decide which of two cities has the larger population (Gigerenzer & Goldstein, 1996). For example, Hamburg is defined as having a soccer team in the Bundesliga, being the state capital, an exposition site, having an intercity train line, and having a university. Dortmund is defined similarly, except it is not a state capital, but is in the industrial belt. The goal of take-the-best is to decide which of two cities, like Hamburg and Dortmund has the larger population, based on their cues.

Two key measures associated with cues in environments are *discriminability* and *validity*. Discriminability measures how often a cue discriminates between stimuli, as state capital and industrial belt discriminate between Hamburg and Dortmund. Validity measures how often a cue, given that it discriminates, belongs to the stimulus with the higher criterion value. Since Hamburg has more people than Dortmund, state capital is a valid cue for the comparison, but industrial belt is not.

Figure 1 shows the discriminability and validity of all 9 cues in the German cities environment. Discriminability is naturally bounded between 0 and 0.5, while validity is naturally bounded between 0.5 and 1. The national capital cue has low discriminability, because only Berlin has the cue, and so it only discriminates when Berlin is one of the cities in the paired comparison. But it has perfect validity because, when it does discriminate, it always gives the correct answer, since Berlin is Germany's most populous city. Meanwhile, the university cue has high discriminability, because about half the cities have universities and half do not. But it has a validity of about

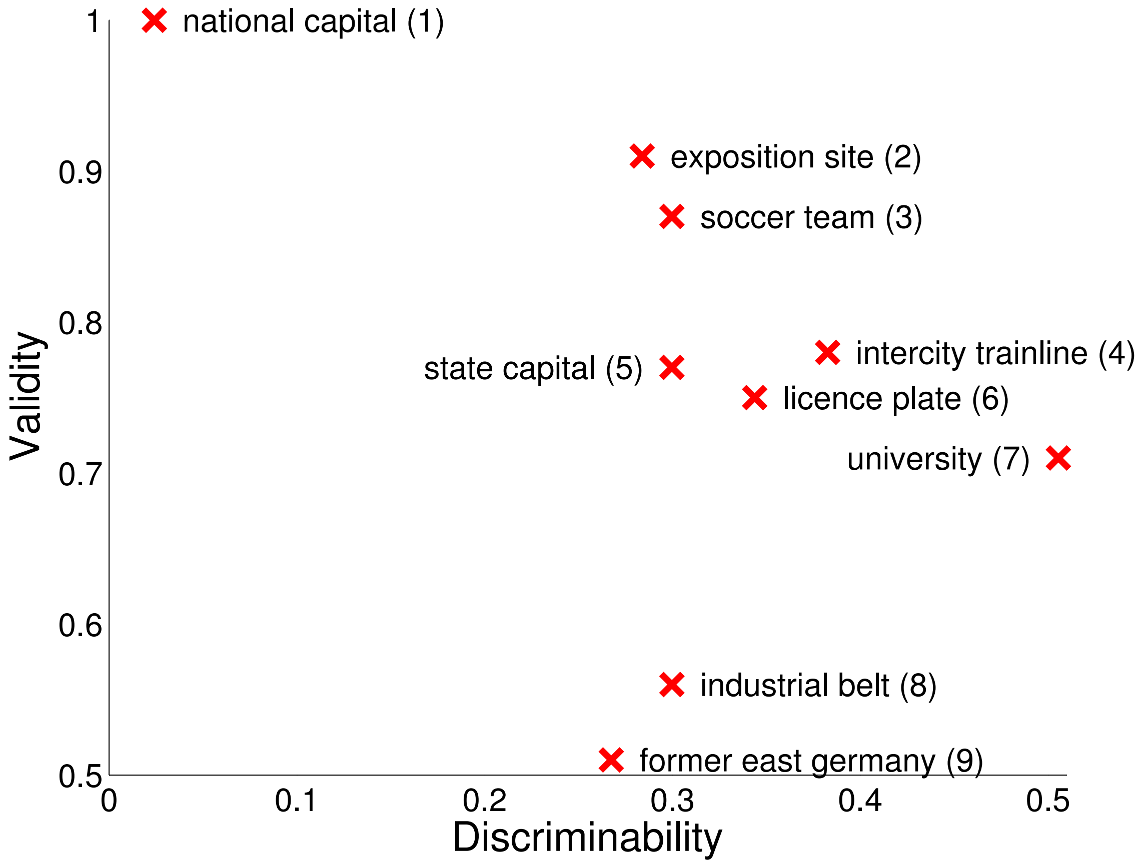

Figure 1. Discriminability and validity of the 9 cues in the German Cities environment.

0.7, because cities without universities are sometimes more populous than cities with universities.

Balancing Discriminability and Validity

Take-the-best assumes that cues are searched in decreasing order of their validity. This corresponds to a strong assumption about the environmental regularity of diminishing returns, because it guarantees that those discriminating cues with the most information are found earliest.

It might be, however, that this assumption is too strong. Discriminability focuses on finding information. Validity focuses on information being highly important if it is found. In this sense, discriminability and validity represent a natural trade-off in determining the order of cue search, and that trade-off is visible in Figure 1. Evidence from experimental investigations of search rules (e.g., Newell, Rakow, Weston, & Shanks, 2004; Rakow, Newell, Fayers, & Hersby, 2005) finds both discriminability and validity can be relevant to search, with individual differences and task constraints influencing how they combine to determine the order of search (e.g., Martignon & Hoffrage, 1999).

A simple way to formalize the combination was introduced recently by Lee and Newell (in press), and is pursued here. Rather than following just validity, the cue search order is determined by giving a weight w to the validity of each cue, the remaining weight $1-w$ to discriminability. The search order is then based on ordering cues according to the sum of these two weighted components. Setting $w = 1$ therefore gives a validity based search, as for the original take-the-best heuristic. Setting $w = 0$ gives a discriminability based search, and intermediate values balance both measures in determining search order.

Extending Discriminability to Capture Correlation

Take-the-best is consistent with the environmental regularity of correlated information. By employing a one-reason approach to search termination, it assumes that information from future cues will not present strongly contradictory evidence. This is again a strong assumption. There is empirical evidence that people sometimes search in non-compensatory ways, but extend their consideration beyond the lone first discriminating cue. People sometimes look for two or three or more reasons to make a decision, even if they do not search exhaustively (e.g., Dhami, 2003).

A less extreme way to capture the idea of correlated information in a heuristic like take-the-best is to refine the cue discriminability measure. Correlated environments are those where the information provided by one cue is consistent with information provided by other cues. One way to measure this correlation is to break the traditional measure of discriminability into two parts. These parts measure whether or not a cue is consistent with other cues, in terms of the stimulus it favors when it discriminates. We call the consistent part the positive discrimination rate d^{+}, and the inconsistent part the negative discrimination rate d^{-}.

Because they are new theoretical measures, we need to explain how d^{+} and d^{-} are calculated, and this requires a bit more formality. If the kth cue discriminates in favor of stimulus A over stimulus B, the definition of cue validity means that the log-odds evidence it provides in favor of stimulus A is $\log\left(v_k/\left(1-v_k\right)\right)$ (Bergert &

Table 1: The cue validity v, discriminability d, positive discriminability d^+, and negative discriminability d^- for the nine cues in the German cities environment.

	National capital	Exposition site	Soccer team	Intercity trainline	State capital	Licence plate	University	Industrial belt	Former East Germany
v	1.00	0.91	0.87	0.78	0.77	0.75	0.71	0.56	0.51
d	0.024	0.284	0.300	0.383	0.300	0.344	0.505	0.300	0.267
d^+	0.022	0.260	0.217	0.298	0.225	0.268	0.370	0.110	0.143
d^-	0.002	0.024	0.083	0.084	0.075	0.076	0.136	0.189	0.125

Nosofsky, 2007; Lee & Cummins, 2004). It is natural to express evidence on the log-odds scale, because this allows the evidence contributed by different cues to be summed (e.g., Cover & Thomas, 2006).[1]

The evidence that all of the cues *except* the kth cue provide in favor of stimulus A or stimulus B can be formed by such a sum. We write all of the cues except the kth that discriminate in favor of stimulus A as $k' \in A$, and similarly write all of the cues except the kth that discriminate in favor of stimulus B as $k' \in B$. Then the total evidence is

$$t_{k'} = \sum_{k' \in A} \log \frac{v_{k'}}{1 - v_{k'}} - \sum_{k' \in B} \log \frac{v_{k'}}{1 - v_{k'}}.$$

This total $t_{k'}$ will be positive if the other cues overall provide evidence in favor of stimulus A, and negative if they favor stimulus B. The positive discrimination rate d^+ for the kth is then the proportion of times, over all stimulus pairs, when the kth cue discriminates, that it favors the same alternative as that favored by the remaining cues, as calculated from the sign of $t_{k'}$. The negative discrimination rate d^- is the proportion of times the alternative favored by the kth cue is different from that favored by the other cues over all stimulus pairs. The positive and negative discrimination rates partition[2] the traditional discrimination rate d, with $d = d^+ + d^-$. Intuitively, d_k^+ measure how often the kth cue discriminates in favor of the same alternative as that favored by the other cues.

Table 1 shows the cue validity v, discriminability d, positive discriminability

[1]Unfortunately, some previous research has added the cue validities themselves to combine evidence, which has no information theoretic justification, and has likely led to errors in the experimental design and the interpretation of results in earlier work.

[2]This means that the extremely rare case where the other cues favor neither alternative is distributed equally to the positive and negative discrimination rates.

d^+, and negative discriminability d^- for the nine cues in the German cities environment. The validity and discriminability measures are those already presented in Figure 1. The refinement of the traditional unitary discriminability measure into positive and negative discriminabilities shows that most cues have much greater positive discriminability. This means that they tend to discriminate in favor of the alternative favored by the other cues. The notable exception is the industrial belt cue, which has a higher negative discriminability. This means that, when one city is in the industrial belt of Germany, but the other is not, other cues will tend to provide evidence for the non-industrial belt city being larger. But, generally, the greater positive than negative discriminabilities show that there is correlated information in the German cities environment.

Evaluating Take-the-best

Previous evaluations of take-the-best have focused on the accuracy of the decisions it makes (e.g., Gigerenzer & Goldstein, 1996; Martignon & Hoffrage, 2002). The general finding has been that, despite its very limited search, take-the-best often matches, and sometimes exceeds, the accuracy of benchmark statistical methods that use all of the available cue information. This is a reasonable approach to evaluation, focused on the assessment of external outcomes. A natural complementary approach to assessment is to consider processes, and justify the internal coherence of a decision process, regardless of its outcome. The distinction between the assessment of process and outcome is an important one in the decision sciences (e.g., Simon, 1976), sometimes presented a distinction between correspondence and coherence (Dunwoody, 2009; Hammond, 2007). While people cannot always make accurate decisions, they can always follow effective decision processes, and it seems important to assess both aspects of decision-making (e.g., Lee, 2006).

For non-compensatory search, one measure of a reasonable decision process is that it terminates once the current decision is unlikely to change. If it becomes clear that further search is unlikely to change the current decision, it is sensible to stop searching. Regardless of the accuracy of that decision, further search will not change the outcome and, in that sense, the decision is internally coherent. So, our assessment of take-the-best focuses on whether it stops when further search is unlikely to change the current decision.

Evaluation Approach

The probability of a decision maker changing their mind can be calculated for any search order over cues, given the discriminability and validity measures for those cues. The four examples in Figure 2 demonstrate the evaluation approach we used, using the comparison of Stuttgart and Paderborn as a concrete example.

In the top left panel of Figure 2, a search order based entirely on validity is used. The first cue, national capital, does not discriminate the cities, but the second cue, exposition site, does, and favors Stuttgart. The gray lines show the possible evidence paths from the second cue onwards, together with thicker black short horizontal lines showing the relative probability of each evidence tally after each subsequent cue has been examined. These paths and distributions come from knowing the discriminability of each cue, which gives the probability that additional evidence will be supplied by the cue, and the validity of each cue, which gives the magnitude of the potential evidence.

The analysis now asks "how likely it is that exhaustive search could change the decision from Stuttgart to Paderborn?" This is determined by the distribution of final evidence shown by the histogram after the ninth cue. Those evidence tallies that favor Paderborn are highlighted in black, and give the probability that exhaustive search will reverse the decision implied by the second cue. If this probability is very small, there are grounds for non-compensatory decision making, because even exhaustive search will likely not lead to a different decision.

The top right panel of Figure 2 shows the analysis for the same search order, but using positive and negative discrimination rates. Now, if the current evidence favors one alternative, the probabilities of future discrimination in favor or against that alternative follow the positive and negative rates, rather than being evenly split between the two possibilities. Because positive discrimination rates are usually larger the distribution of final evidence tallies gives greater support the currently favored alternative. This is clear in the top right panel, through an upward shift in the evidence path, and in the distribution of final tallies. There is now only a very small probability of exhaustive search changing the decision from Stuttgart to Paderborn.

The bottom two panels of Figure 2 show the same two analyses, but for a different search order. This is a search order based only on discriminability, starting with the most discriminating university cue and moving to the least discriminating national capital cue. As before, the first cue does not discriminate between Stuttgart

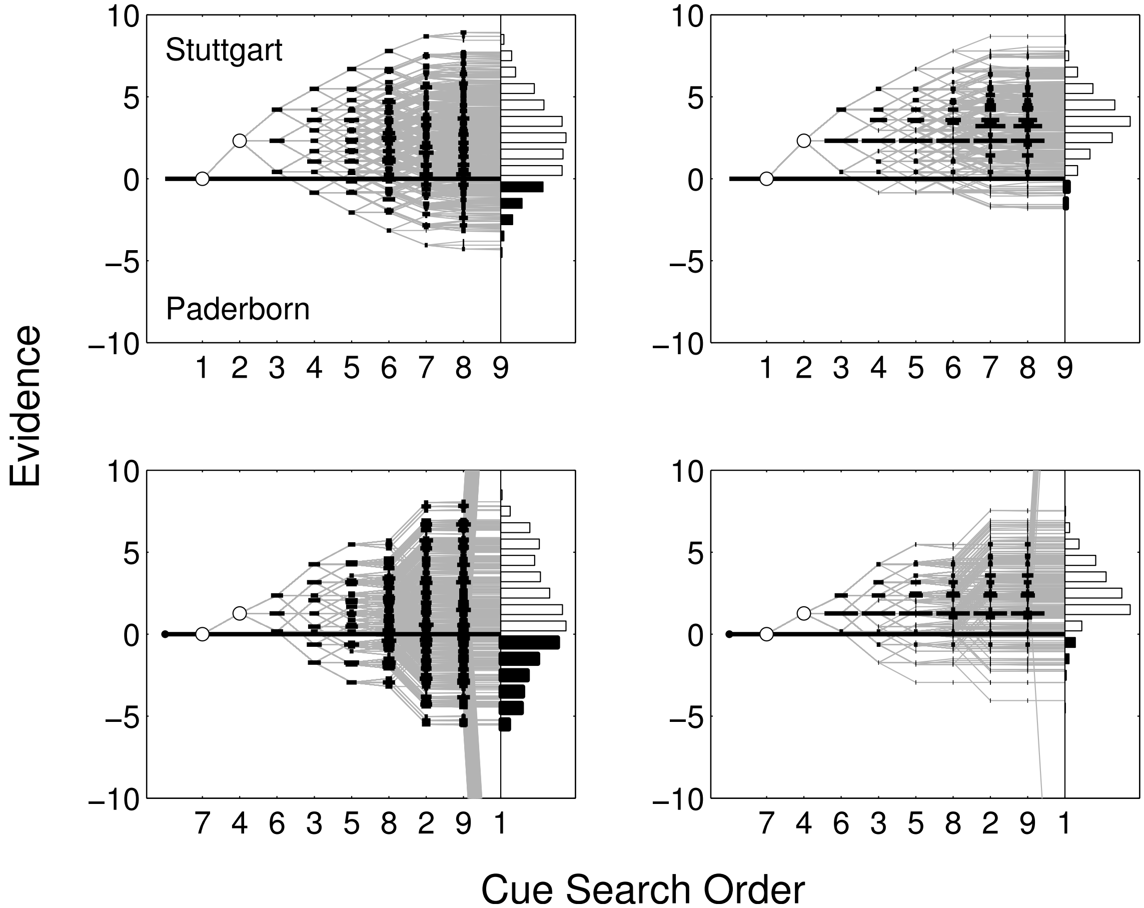

Figure 2. Evidence paths, and distributions of final tallies, for a comparison of Stuttgart and Paderborn where the first discriminating cue favors Stuttgart. Each panel shows by gray lines the possible evidence paths for future cues, culminating in a distribution of final evidence tallies. The final tallies agreeing with the current decision are shown in white, while those corresponding to the alternative decision are shown in black. All four panels consider the case where two cues have been searched, and the current evidence favors choosing Stuttgart. Panels in the top row corresponds to validity-based search, while those in the bottom row corresponds to discriminability-based search. Panels in the left column corresponds to using traditional discriminability to assess future evidence, while the right column corresponds to using positive and negative discriminabilities.

and Paderborn, but the second intercity trainline cue does, and favors Stuttgart. The bottom left panel of Figure 2 shows the evidence paths that follow, based on this search order, using the traditional single measure of discriminability. The distribution of final tallies shows there is a strong probability of the decision being changed to Paderborn after exhaustive search.

The bottom right panel of Figure 2 considers the same search order, but uses the positive and negative discrimination rates. The distribution of final tallies now shows a much lower probability of Paderborn being favored by exhaustive search.

The four examples in Figure 2 show two key trends, corresponding to two environmental regularities. One trend is made by the top versus bottom panels. The top panels show the results of search based on validity, which builds into decision-making an assumption of diminishing returns. These diminishing returns are visually clear from the convex nature of the evidence paths, as later cues provide successively less evidence. This means that early decisions are less likely to be over-turned by later evidence, because the later evidence is less compelling.

The bottom row of panels shows the change when a different search strategy, based on discriminability in this case, is used. Now later cues can provide strong evidence, the evidence paths can move quickly toward one alternative or another late in search, and the distribution of final evidence tallies covers both choices. Basing search on cue validity corresponds to assuming diminishing returns, and provides grounds for non-compensatory decision-making. When this assumption is not made, early decisions can be over-turned by later evidence.

The other key trend is shown by the left versus right panels in Figure 2. The left panels show final evidence tallies based on the traditional measure of discriminability. Those on the right show the tallies coming from using positive and negative discrimination rates. These measures allow the environmental regularity of correlated information to be incorporated in the analysis. In the panels on the right, once evidence favors a decision, future discriminating cues are generally expected also to favor that decision, because positive discriminability is greater than negative discriminability for most cues in the German cities environment. What is seen early in search is generally predictive of what is seen later, and so early decisions are unlikely to be over-turned. The presence of correlated information provides grounds for non-compensatory decision-making.

Results

The results of our analyses of the German cities environment are summarized by the 'wheel' in Figure 3. The wheel rotates from the bottom left clockwise to the bottom right, increasing the value of w to move from pure discriminability based search when $w = 0$ to pure validity based search when $w = 1$. The inner rim of the wheel shows w increasing. The outer rim denotes by circular markers each time the increase in w leads to a change in the search order, and the orders themselves are listed, using the cue numbering from Figure 1.

The histograms outside the search orders in Figure 3 show the extent of search, using the criterion that search should continue while there is a greater than 5% chance of exhaustive search leading to a different decision.[3]

What is being measured is the proportion of cues, beyond the first discriminating cue, searched before making a decision. This dependent variable is called the Proportion of Extra Cues (PEC), and ranges between 0 and 1 (Newell & Lee, 2009). For example, if two cities are being compared, and the third cue is the first one that discriminates between them, but the fourth and fifth cues are also searched, then the PEC measure is 2/7, since 2 extra cues were searched, out of the 7 possible. When PEC=0, the first discriminating cue terminates search, as in take-the-best, and when PEC=1, search exhausts all of the cues. The gray histograms correspond to the distribution of the PEC measure over all city pairs using the traditional discriminability measure,. The white histograms correspond to the distribution of the PEC measure using the positive and negative discriminability measures.

Two key results are clear. One is that, as the emphasis in search order shifts to validity, search becomes less extensive. In other words, as the assumption of diminishing returns is built into the search order, there are stronger grounds for stopping sooner. The other key result is that, when positive and negative discriminability are used, search becomes much less extensive. In other words, when the correlation of cues in the environment is considered, it often becomes clearer earlier that the likelihood of a decision being reversed is small, and the grounds for stopping earlier are again stronger.

[3]The choice of 5% is obviously not principled, but is often a default choice in statistical inference. Using values near 5% lead to qualitatively similar results, and behave in the ways one would expect. As the criterion becomes more stringent, search is more extensive. Our goal is not to study or justify this criterion, but to show the impact of the environmental regularities for a reasonable criterion setting.

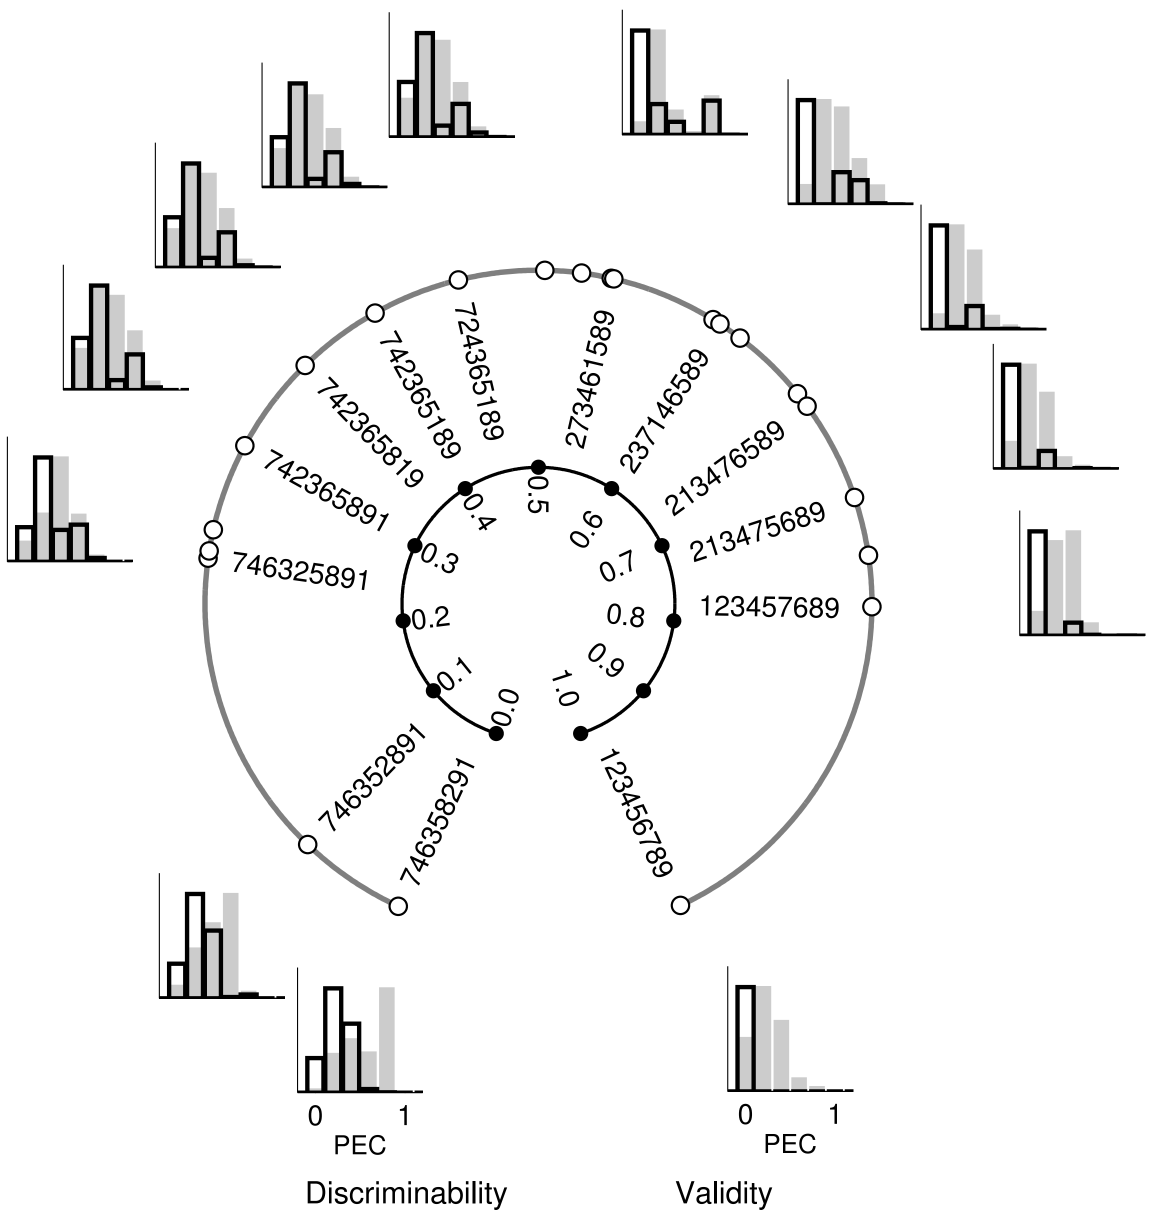

Figure 3. Results of manipulating the search order by emphasizing validity or discriminability, and manipulating the measure of discriminability, on the proportion of cues beyond the first discriminating one that must be searched to reduce the probability of a changed decision below 5%. The inner rim shows the change in the w parameter that weights validity in determining the cue search order, ranging from strictly discriminability-based search at the bottom-left, to strictly validity-based search at the bottom right. The outer rim shows the change in patterns of the actual search orders by circular markers, and provides the details for a selected representative subset of these orders. The histograms for these selected order shows the distribution of the Proportion of Extra Cues (PEC) measure, over all possible questions, assuming both traditional discriminability (shaded gray) and positive and negative discriminability (unshaded).

In this sense, the results in Figure 3 provide support for take-the-best as an effective decision-making heuristic. But the results also show the bounds of the argument. When one or other or both of the environmental assumptions are not met, search is more extensive than the first discriminating cue, and often much more so. The gray histograms, when a correlated environment is not incorporated, always show search extending well beyond the first discriminating cues for many of the questions. The histograms not in the bottom-right corner, considering discriminability as part of the search order, and so not relying on diminishing returns, also show more extensive search, however discriminability is measured.

It is worth noting that it is not easy to measure the effect of environmental assumptions on performance, using an outcome measure of accuracy. The benchmark accuracy, gained from always considering all of the cues, is that 77.4% of city pairs are correctly answered. The 44 heuristics presented in Figure 3, comprised of 22 search orders by 2 types of discriminability measurement, all lie between 76.9% and 77.7%. Thus, it is clear that the heuristics match the accuracy of extensive search, and that there is very little variation in the final outcome. The interesting variation, in terms of search, and assumptions about environmental correlation, are captured by the process-oriented coherence analysis presented in Figure 3.

Discussion

In this paper, we have considered two key theoretical assumptions underpinning the take-the-best heuristic for decision-making. These are both assumptions about the information structure of the environment: that it is has diminishing returns, so that evidence found later in search is less important than evidence found earlier; and that it has correlated information, so that information found early in search is likely to be consistent with information found later.

Demonstrating that the grounds for take-the-best rely on these assumptions requires the ability to manipulate the assumptions. We proposed a richer set of search orders, generalizing the validity search of take-the-best, so that the assumption of diminishing returns could be manipulated. We also proposed dividing discriminability into positive and negative components, to capture the assumption of correlated information.

Our results, for the German cities environment, show that combining diminishing returns and correlated information provide grounds for non-compensatory decision-making. When these assumptions are met, the probabilities that exhaustive search would change an initial decision are very small. The first discriminating cue favors the same alternative as exhaustive search favors. The detailed results in Figure 3 show that both environmental regularities are important, and quantify their effect. The results also show, however, the bounds on the justification for limited search. When the basis of search moves away from validity, and so does not emphasize diminishing returns, more extensive search is needed. Similarly, if the correlational structure of an environment is not assumed, many more cues than the first discriminating one need to be examined to make it likely a decision will not change.

More general conclusions about the relative impact and usefulness of each assumption, and their interaction, requires a much more extensive study of a broad range of environments. Unfortunately, many of the other environments studied in the fast and frugal literature (e.g., Czerlinski, Gigerenzer, & Goldstein, 1999) formed binary cues by taking median splits of continuous variables. By construction, this means all of the cues have discriminabilities of 0.5, and so create impoverished environments from the perspective of the current analyses. With appropriate environments, however, our findings give some theoretical tools and initial results to motivate that broader exploration.

Another line of future work is to extend take-the-best as a decision-making heuristic, and evaluate these extensions. Our results present a detached analysis of when and why take-the-best works in structured environments. But, the mechanism we used for manipulating search orders immediately gives a set of possible new decision-making heuristics (see Lee & Newell, in press). The mechanism we developed for breaking discriminability into positive and negative components, to assess correlated information, does not immediately give rise to new heuristics. It does, however, give a theoretical opening for their development. For example, one sensible stopping rule would require a discriminating cue with a positive discrimination rate larger than the negative discrimination rate. This would correspond to stopping only when the discriminating cue was more likely than not in agreement with the information that would be provided by further search. Possibilities like these seem worth exploring.

What our results do show is when and why non-compensatory heuristics like take-the-best are justified in stopping their search. When environments have diminishing returns, and when environments have correlated information, exhaustive search does not change decisions, and there are good grounds for limited search.

Asked before the Angola game what he knew about the Angolan team, US player Charles Barkley replied: "All I know about Angola is Angola's in trouble." That was a non-compensatory opinion, unlikely to have been changed on the basis of further reflection. It proved to be accurate.

Acknowledgments

We thank Ben Newell for very helpful comments on an earlier version. Support from the Air Force Office of Scientific Research Award FA9550-11, and Australian Research Council Grant DP110100797, are gratefully acknowledged.

References

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 & Cognition*, *33*(1), 107–129.

Cover, T. M., & Thomas, J. A. (2006). *Elements of information theory.* New York, NY: Wiley.

Czerlinski, J., Gigerenzer, G., & Goldstein, D. G. (1999). How good are simple heuristics? In G. Gigerenzer, P. M. Todd, & The ABC Research Group (Eds.), (p. 97-118). Oxford University Press.

Dhami, M. K. (2003). Psychological model of professional decision-making. *Psychological Science*, *14*, 175–180.

Dunwoody, P. T. (2009). Introduction to the special issue: Coherence and correspondence in judgment and decision making. *Judgment and Decision Making*, *4*, 113–115.

Gigerenzer, G., & Goldstein, D. G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, *103*(4), 650–669.

Gigerenzer, G., & Todd, P. M. (1999). *Simple heuristics that make us smart.* New York: Oxford University Press.

Hammond, K. R. (2007). *Beyond rationality: The search for wisdom in a troubled time.* Oxford University Press.

Lee, M. D. (2006). A hierarchical Bayesian model of human decision-making on an optimal stopping problem. *Cognitive Science*, *30*, 555–580.

Lee, M. D., & Cummins, T. D. R. (2004). Evidence accumulation in decision making: Unifying the "take the best"and "rational" models. *Psychonomic Bulletin & Review*, *11*(2), 343–352.

Lee, M. D., & Newell, B. R. (in press). Using hierarchical Bayesian methods to examine the tools of decision-making. *Judgment and Decision Making*.

Martignon, L., & Hoffrage, U. (1999). Why does one-reason decision making work? a case study in ecological rationality. In G. Gigerenzer, P. M. Todd, & The ABC Research Group (Eds.), (p. 119-140). Oxford University Press.

Martignon, L., & Hoffrage, U. (2002). Fast, frugal and fit: Simple heuristics for paired comparison. *Theory and Decision*, *52*, 29–71.

Newell, B. R., & Lee, M. D. (2009). Learning to adapt evidence thresholds in decision making. In N. Taatgen, H. van Rijn, J. Nerbonne, & L. Shönmaker (Eds.), *Proceedings of the 31st Annual Conference of the Cognitive Science Society.* Austin, TX: Cognitive Science Society.

Newell, B. R., Rakow, T., Weston, N. J., & Shanks, D. R. (2004). Search strategies for decision making: The success of 'success'. *Journal of Behavioral Decision Making*, *17*, 117–130.

Rakow, T., Newell, B. R., Fayers, K., & Hersby, M. (2005). Evaluating three criteria for establishing cue-search hierarchies in inferentialjudgment. *Journal of Experimental Psychology: Learning, Memory & Cognition*, *31*, 1088–1104.

Simon, H. A. (1956). Rational choice and the structure of environments. *Psychological Review*, *63*, 129–138.

Simon, H. A. (1976). From substantive to procedural rationality. In S. J. Latsis (Ed.), *Method and appraisal in economics* (pp. 129–148). London: Cambridge University Press.

Simon, H. A. (1990). Invariants of human behavior. *Annual Review of Psychology*, *41*, 1–19.

Todd, P. M., & Gigerenzer, G. (2003). Bounding rationality to the world. *Journal of Economic Psychology*, *24*, 143–165.