2014 Russell Ackoff Doctoral Student Fellowship Proposal

Berkeley J. Dietvorst, 3rd Year Doctoral Student in OPIM

Advisor: Joseph Simmons, Associate Professor in OPIM

Understanding Algorithm Aversion

Project Background:

Imagine that as an admissions officer for a university, it is your job to decide which student applicants to admit to your institution. Because your goal is to admit the applicants who will be most likely to succeed, this decision requires you to forecast students' success based on the information in their applications. There are at least two ways to make these forecasts. The more traditional way is for you to review each application yourself and make a forecast about each one. This is called the *human method*. Alternatively, you could rely on an evidence-based algorithm to make these forecasts. This may involve building a statistical model using the data of past students; this model would tell you how to optimally combine each piece of information in the students' applications in order to make your forecasts. This is called the *algorithm method*.

Dawes (1979), Meehl (1954), and many others (Grove et al, 2000) have demonstrated that algorithmic forecasts are more accurate than human forecasts in many domains, including admissions decisions, parole decisions, and medical diagnoses. And, of course, when algorithmic forecasts are better, people should prefer to rely on algorithms rather than humans for forecasting. But do they?

There is empirical (e.g. Onkal et al., 2009; Promberger & Baron, 2006) and anecdotal evidence suggesting that forecasters prefer humans to superior algorithms. But the reasons for this preference are unknown. In this project, we investigate three questions about people's use of algorithmic forecasts: 1) When are people likely to use inferior human forecasters in favor of superior algorithms? 2) Why do people resist using algorithms for forecasting? 3) How can we increase the likelihood that people will use an algorithm?

The Present Research:

<u>1) When are people likely to use inferior human forecasters in favor of superior algorithms?</u>

In this line of research, we investigated how seeing an algorithm or a human forecaster perform (and therefore err) affects people's likelihood of choosing the algorithm instead of the human forecaster. We hypothesized that people are especially unlikely to use a superior algorithm after they have seen it err. Following this hypothesis, we conducted three studies showing that people who are exposed to predictions (and consequently errors) from an algorithm are less likely to use the algorithm, and that they make less accurate predictions as a result.

In Study 1, we tasked participants with predicting the percentiles of graduated MBA students based on information from their applications. Participants were randomly assigned to four conditions. Participants either: practiced making forecasts for 15 students with feedback on their accuracy, received estimates from a statistical model for 15 students with feedback on its accuracy, both of these, or neither of these. Then, participants decided whether to tether their monetary incentives to the model or themselves for 10 future predictions.

We found that participants who were exposed to the statistical model were less likely to tether their incentives to the model (24%) than participants who only made their own estimates (63%) ($p<.001$). Even participants who saw the model's estimates and were outperformed by the statistical model were less likely to rely on the model's estimates (31%) than participants who only made their own estimates ($p<.001$). Further, watching the model perform had a negative effect on performance, precisely because participants who saw the model perform were less likely to choose it.

Studies 2 and 3 demonstrated the robustness of these findings using the same paradigm. In Study 2, we used a different prediction task, payment rule, and an even better performing model. In Study 3, we used the same task as Study 1; however, participants were yoked to a past participant from Study 1, viewed the past participant's estimates, and chose between the model and the past participant. We found the same pattern of results in both studies: participants who were exposed to the model were less likely to use it than those who only saw the human perform, even if they witnessed the model outperform the human.

This portion of our project is approaching completion. We are currently editing a manuscript detailing this work, and we plan on submitting it for publication within the next month. I am seeking funding so that I am able to travel and present this work at conferences. I have already presented this work at the Society for Judgment and Decision making, and I am scheduled to present it at the Society for Consumer Psychology. I also hope to present it at the Behavioral Decision Research in Management Conference, the Academy of Management, and the Association for Consumer Research Conference.

<u>2) Why do people resist using algorithms for forecasting?</u>

We are just starting this line of research. We have found preliminary evidence suggesting that people believe that humans are better than algorithms at: detecting targets that are exceptions to the rule, learning from mistakes, treating targets as individuals, and improving with practice. Additionally, we have found that people think that the performance of a human forecaster is more acceptable than the performance of an equally accurate algorithm across varying levels of accuracy.

I am seeking funding to run follow-up studies in order to better understand which, if any, of the beliefs stated above contribute to people's avoidance of algorithms. In these experiments, we will manipulate participants' beliefs about an algorithm and test to see if this manipulation has increased the probability that they will use the algorithm. For example, we would assign half of participants in one experiment to a condition where it is explained that an algorithm adjusts its coefficients as it collects more data, and therefore, improves with practice. Then we would test to see if participants who have received this information are more likely to use the algorithm for incentivized forecasts.

<u>3) How can we increase the likelihood that people will use an algorithm?</u>

In this line of research, we are aiming to find ways of presenting the choice between a human forecaster and an algorithm that will increase people's chances of using the algorithm. So far, we have found that people are more likely to choose an algorithm (instead of a human forecaster) to complete many forecasts (100 or more) than they are to complete one forecast. This effect is not due to concerns about efficiency; participants believe that algorithms have a larger advantage in terms of accuracy when both forecasters are completing a larger number of forecasts.

I am seeking funding to run follow-up studies investigating why people are more likely to use an algorithm for larger numbers of forecasts. Further, I plan to run additional studies searching for other ways to present the choice between a human forecaster and an algorithm that will increase people's likelihood of using the algorithm.

Explanation of Budget:

Participant Payment. We plan to run at least 6 studies on Amazon Mechanical Turk investigating why people resist using algorithms for forecasting and why people are more likely to use an algorithm when making many forecasts. On average, these studies will cost $0.50 per participant. Each of the experiments will have 2 conditions and 200 participants per condition. Therefore, the total cost of these studies will be $1200.

Conference Travel. I am requesting funding for travel to two academic conferences where I plan to present this work. The two conferences that I hope to attend with money from this fellowship are the Behavioral Decision Research in Management Conference, which will be hosted by the London business school, and the Association for Consumer Research Conference, which will be held in Baltimore, Maryland. Details of the expenses associated with these conferences are listed below.

Behavioral Decision Research in Management; London, July 2014: $2,200

- Flights: $1600
- 3 nights at conference hotel: $400
- Registration fees: $200

Association for Consumer Research; Baltimore, MD, October 2014: $600

- Train Tickets: $100
- 3 nights at conference hotel: $300
- Registration fees: $200

Budget summary:

Expense	**Amount Requested**
Participant payment	$ 1,200
Conference travel	$ 2,800
Total	**$ 4,000**

Current Resources:

There is no grant or departmental funding for this project. The Operations and Information Management Department gives its students $800 each academic year for conference travel, but I have already used this money to attend the Society for Judgment and Decision Making Conference in November 2013.

Advisor Signature:

Joseph Simmons, Associate Professor in OPIM

References:

Dawes, R. M. (1979). The robust beauty of improper linear models in decision making. *American Psychologist*, 34(7), 571–582.

Grove, W. M., Zald, D. H., Lebow, B. S., Snitz, B. E., & Nelson, C. (2000). Clinical versus mechanical prediction: A meta-analysis. *Psychological Assessment*, 12(1), 19–30.

Meehl, P. E. (1954). Clinical versus statistical prediction: A theoretical analysis and review of the literature. Minneapolis: University of Minnesota Press.

Önkal, D., Goodwin, P., Thomson, M., Gönül, S., & Pollock, A. (2009). The relative influence of advice from human experts and statistical methods on forecast adjustments. *Journal of Behavioral Decision Making*, 22(4), 390-409.

Promberger, M., & Baron, J. (2006). Do patients trust computers?. *Journal of Behavioral Decision Making*, 19(5), 455-468.