PRIVATE ENTERPRISE RESEARCH CENTER
TEXAS A&M UNIVERSITY

PERCspectives on Research

September 2009

Eureka Learning

When learning to play a game, what factors determine an individual's ability to improve her performance as she plays similar games? Specifically, what techniques lead to an increase in the probability of winning a two outcome perfect information game without chance moves (PI-game)? Such games can be solved by backwards induction in principle, but in practice the complexity of PI-games can quickly overwhelm most people's computational capacity, as several studies have shown. What does it take for a player to solve these games when the rank of the game exceeds the player's ability to find the solution via backwards induction?

In ***PERC*** Working Paper #0909, ***PERC Rex Grey Professor*** John Van Huyck and coauthor C. Nicholas McKinney Jr. explore how individuals can experience "eureka" moments while playing a game, and how this affects strategies and outcomes. Effective heuristics, or rules-of-thumb, allow a player to rapidly find a solution that is close to the optimal solution. They are not von Neumann and Morgenstern strategies that assign an action to every information set controlled by the player. Instead players are taught to recognize winning positions and develop heuristics that increase the likelihood of reaching those winning positions. This is a common method covered in the books designed to teach chess. Both novice and chess masters are taught to recognize winning positions. The more winning positions that a player learns to recognize, the greater her chances of winning the game, as noted elsewhere in the literature. This logic can be applied to any PI-game.

When winning positions are not transparent, players often develop heuristics that increase the probability of success. For example, in the game of tic-tac-toe, people are often taught to put the first mark in the center square. While putting the first mark in the center square will not guarantee a win (or tie), it does dramatically decrease the chance of losing.

As people learn to play PI-games, they may learn to see and apply more and more heuristics. The authors explore how this learning occurs by investigating the possibility that players experience a eureka moment in which they adopt a heuristic that improves their chances of winning the PI-game.

The authors address three specific questions: When, given a number of PI-games in a short period of time, does player performance improve? If performance does improve, can it be attributed to learning new heuristics or are the players actually increasing their reasoning depth? And finally, is the improvement in performance slow and gradual, or does it occur in discrete jumps that correspond to Eureka Learning?

To answer these questions, the authors design an experiment in which subjects play the game of Nim against a procedurally rational algorithm. Nim is a two-player game in which players alternate taking one or more stones from one of *m* rows. The player who takes the last stone wins. Nim games are ideal for Eureka Learning because the investigator can vary complexity by changing the rank of the game.

A person with no prior experience with Nim is unlikely to use a procedurally rational algorithm. However, by mastering a few heuristics, a person could learn to play more complicated Nim games and improve her chances of winning. One of the most effective heuristics is *Strategy Copying*, which occurs in a two row game (m=2), when one's opponent must choose from two equal rows. The strategy consists of taking the same number of stones from the unchosen row. A player who understands the strategy copying heuristic will look for opportunities to transform the game into a subgame with two equal rows, which she will then win by copying her opponent's strategy.

The subject pool for the experiment is composed of 71 individuals recruited from the Texas A&M University student body. Subjects play two sets of 47 games against a computer opponent. The games are presented in a different order in the second treatment, and the rows are scrambled to make it difficult to recognize the similarities between the two treatments. The experiment is designed to test if the subjects' performance increases over time. If performance does improve, the individual games are examined to see what the subjects are learning.

The results indicate that subject performance does improve over time. In many games the difference is small, but it is statistically significant. Some subjects behave as if they have learned to recognize winning positions that improve their overall performance, but subjects do not develop procedurally rational strategies that lead to substantively rational play in games of rank 5 or more.

Although there is evidence of learning, the results are not perfect. In Eureka Learning, a player's probability of winning in a series of Nim games would be negatively correlated with rank up to the point in which she learned a new heuristic. From then on, the player's new strategy would lead to wins in all games in which the heuristic was applicable. Nine of the subjects tested fit the profile of a Eureka Learner for two row strategy copying games, but there is little evidence of Eureka Learning of more complicated heuristics.

Probabilistic models of performance indicate that player performance improves as the game is played. Response time when interacted with strategy copying provides some evidence that subjects play strategy copying games faster. The role of time in the logistic model of play provides evidence that subjects develop heuristics that increase their performance.

The broader impact of these findings should not be ignored. We face complicated decisions in which the eventual outcome is well beyond our ability to backwards induct on a daily basis. We place our retirement savings in mutual funds, not because we know what they will be worth 30 years from now. We do it because, as a general rule, the market will outperform most other forms of investment in the long run. We choose to attend college because, as a general rule, college graduates live longer lives and have higher lifetime earnings. Virtually all of life's long term decisions are made using heuristics designed to simplify the complicated.

Although these heuristics are important, learning them by experience may be difficult. In Nim games, spending five minutes with a teacher is probably worth more than playing 94 games, given that only about 10 percent of the subjects were able to discover the strategy copying heuristic in the first treatment and successfully repeat it thereafter.

The Adequacy of Educational Cost Functions: Lessons from Texas

Policymakers, educators, and increasingly, litigators, demand reliable estimates of the resources needed to provide an adequate education. This demand has led to studies using a variety of techniques to generate estimates of the resources needed for an education. Two popular methods of generating these estimates are the professional judgment method and econometric cost functions. Recently these methods have been put to the test in Texas, in a lawsuit that confounded the issue of the school funding mechanism with the issue of the resources needed to achieve an adequate education.

PERC Research Fellow Timothy Gronberg, ***PERC Jordan Professor of Economics*** Dennis Jansen, and coauthor Lori Taylor use the Texas litigation studies as a lens through which to explore best practices in the estimation of educational cost functions. Their analysis in ***PERC*** Working Paper #0906 highlights five key decisions that researchers must make when using the cost function methodology in an educational setting, and explores the implications of the various possible choices using recent data on public schools in Texas. As the analysis demonstrates, some common practices in cost function analyses of education are not best practices, and these deviations from best practice can have a significant impact on the estimated cost of an adequate education.

In 2004, 46 school districts brought suit against the state of Texas, arguing that the Texas school finance formula was unconstitutional. The plaintiff districts argued that the formula provided inadequate funding, that cost adjustments for student need were too low, and that as a result, school districts were forced to raise their local property tax rates to the maximum allowed under the formula, thereby removing local discretion over property tax rates.

Three studies of educational adequacy were used as evidence in the trial and made very different predictions about the cost of achieving adequacy in the plaintiff districts. The first, conducted by Timothy Gronberg, Dennis Jansen, Lori Taylor, and Kevin Booker (GJTB) at the request of the Texas legislature, predicted that it would cost in total an additional $861 thousand to raise all of the plaintiff districts up to the threshold level.

Using the same performance thresholds, a second cost function study, performed at the request of the plaintiffs, predicted that it would cost at least an additional $457 million, an amount over 500 times greater than the GJTB estimate. The third study, a professional judgment analysis, predicted that providing an adequate education in the 46 plaintiff districts would cost even more, an additional $683-$830 million.

The authors use the striking differences in the predictions from these cost function studies to motivate a discussion of best-practice techniques in cost function analyses. While cost function analysis is a particularly attractive tool for estimating the cost of an adequate education, implementing the methodology is far from simple. Researchers must make a number of decisions that can have considerable influence on the results of their analyses. Some of the most challenging decisions arise from the potential for school district inefficiency, the choice of functional form, the need for appropriate measures of educational outputs, the treatment of heteroskedasticity, and the problem of endogeneity.

The authors explore the implications of the various possible choices in these five areas using recent data on public schools in Texas. They generate baseline estimates for school district expenditures per pupil and use the model to illustrate how researchers' methodology decisions can greatly affect their final cost predictions.

One such decision is whether or not to use pupil-weighting in the estimation. Many researchers transform data at the school or district level into individual-level data by weighting each school or district by the number of students served. While pupil-weighting may be a justifiable way to correct for heteroskedasticity, this practice undervalues the informational contribution of small schools and overvalues the contribution of very large schools. It also significantly increases the estimated cost of an adequate education. The authors predict the cost of an average-quality education using their baseline estimates and compare it to the predicted cost using pupil-weighted data, and find a difference of $535 million per year solely attributable to pupil-weighting.

Other practices used to deal with challenges such as school district inefficiency and endogeneity of explanatory variables yield cost predictions that, in some cases, are substantially higher than the baseline estimates. The authors illustrate how a researcher's choice between alternative techniques applied to the same data can produce results that are surprisingly different.

Education cost function estimation is a challenging enterprise. The appropriate conclusion, however, is not to abandon cost function estimation as either hopelessly flawed in general or as flawed and inapplicable to public schools. As noted in other studies, such a move would cede the field to other approaches — such as the professional judgment approach — that can be subjected to equally vociferous criticism. Rather, the appropriate conclusion is that it is critically important to demand best-practice techniques from any analyst of educational costs.