Proceedings of the 2010 Winter Simulation Conference
B. Johansson, S. Jain, J. Montoya-Torres, J. Hugan, and E. Yücesan, eds.

NAÏVE LEARNING ALGORITHMS UTILIZED FOR THE PREDICTION OF STOCK PRICES TO COMPARE ECONOMIC MODELS OF DECISION MAKING

Caleb Krell

The University of Oklahoma
3117 SW 103rd St.
Oklahoma City, Oklahoma, 73159, USA

Hank Grant

The University of Oklahoma
202 W. Boyd St., Room 124
Norman, Oklahoma 73019-0631, USA

ABSTRACT

An advance in economic thought is in the area of behavioral economics where traditional models of rational decision-making are challenged by newer models of behavior such as Prospect Theory. This is coupled with a world where algorithms have abilities to learn, remember and evolve over time to make better decisions. The advances on these two fronts are forcing the world of markets to be analyzed from a different angle. This work is a look at markets to compare traditional expected utility theory of economic decision-making to the newer idea of Prospect Theory. Two learning algorithms, based on traditional expected utility and Prospect Theory, are designed and then compared under several scenarios designed to replicate various market conditions faced by investors. Deviations were analyzed to measure the effectiveness of the two algorithms and also the two models of economic decision making, where it was found that risk averseness described by Prospect Theory will lead to greater deviations in expected prices than more traditional models of economic decision making. This is for several reasons, including risk aversion can, in most situations, lead to suboptimal economic decisions.

1 INTRODUCTION

This research started off with research on financial derivatives in conjunction with simulation and modeling, yet it has evolved in to something more. It is just the first step into the intersection of research on financial derivatives, behavioral economics, simulation, modeling and learning algorithms. These various areas tie in together through the modeling of the market place. How does one model decision making and how does or can this affect the market and prices, is the over arching question. More specifically the question at hand is which would be better models of economic decision making, traditional rational choice models or newer behavioral economics. This studied via the creation of two naïve learning algorithms, one based in expected utility and the other based on Prospect Theory. These were then simulated to make predictions of a price process of a fictitious stock, as an investor would do. Thusly, there are three main thrusts of this work. First, is the analysis of neoclassical theories of economic decision making and Prospect Theory. Then there is the analysis of learning algorithms and specifically the construction of the naïve learning algorithms. Finally, simulation and modeling are addressed. The performance of the algorithms was compared by analyzing the deviations between the predictions and simulated price process.

2 LITERATURE REVIEW

Despite the fact that people have used financial derivatives for centuries, especially options, the fair pricing of these instruments has always proved difficult. This prevented wider use of these instruments, for the lack of some generally accepted pricing method prevented the creation of a secondary market. In fact, the Chicago Board of Trade established the Chicago Board Options Exchange in April of 1973 (Hull 2008). There have many option pricing models, yet the one that has brought about the greatest change is the Black-Scholes. While this was earth shattering, it was only able to have a closed form solution due numerous assumptions, and it is these assumptions that create problems as the world doesn't normally fit these assumptions. The ideas of constant interest rates, the ability to only price European style options, the delta neutral replicating portfolio, lack of arbitrage and lognormal price process of the underlier were all challenged.

One method of dealing with relaxing the assumptions of the Black-Scholes is to utilize various numerical methods. Three methods are Monte Carlo simulation, lattice structures and various execution algorithms. Binomial trees and lattices can be fast when the dimensions of the state vector are low (1 to sometimes 3 dimensions) (Broadie et al 2000). Simulation space requirements grow linearly with the dimensions of the state vector which makes this approach attractive for very complex multiple factor, exotic options, where as binomial lattices grow exponentially with respect to the dimensions of the state variable (Broadie et al 2000). Monte Carlo simulation has been gaining popularity as computing power has been increasing. The first example of using Monte Carlo simulation to price options was published by Boyle in 1977 (Charnes 2000). The use of Monte Carlo simulation is an attractive numerical method because of the flexibility that is implicit in the method. However, issues of the distribution of asset prices and returns and the resulting Brownian motion are unresolved. This is because the same stochastic processes are utilized to create the various sample paths for the underlier. For the case of exotic, path dependent, Bermuda options, which are a mixture of American and European options, Andersen and Broadie have taken a different approach. To deal with the issue of path dependency, an algorithm that looks for the best exercising strategy is implemented. Here the underlier is modeled as a martingale through a Markov process, with the underlier price is the state and the algorithm looks to maximize the value of the option by finding the best exercising date. Another critique of the standard Monte Carlo simulation method, from Duan and Simonato, 1995, is regarding the martingale property of the estimate of the underlier. The simulated sample paths of the underlier many times fail to posses the martingale property because simulation can only approximate the theoretical properties due to finite repetitions and the qualities of the random number generator (Duan and Simonato 1995).

Like the modeling of a production facility and specifically the machine repair model, liquidity holes can appear seemingly out of the blue and without memory of last event. This property fits exceeding well within the framework of Brownian motion because Brownian motion also has a memoryless aspect in that the next drift or step in the randomwalk has no relation to the steps before. Previous steps do not affect the direction of the current step; they only give the current starting point. However, markets, during a liquidity hole type of event, have exhibited some path dependency (Taleb 1997). This market memory can also be seen in non-liquidity hole events such as a new market high or low through the activity surrounding those price levels despite the volatility that lead to such events (Taleb 1997). This partial memory raises issues regarding modeling. To be completely accurate a model most likely should contain aspects of both the memoryless random walk and the slight memory during the market events it is relevant. However, this is extremely difficult from a mathematical, theoretical standpoint and also from a numerical, simulation modeling perspective.

Seeing the limitations of the models using Brownian motion, many have utilized more complex stochastic processes, either in the form of multi-dimensional Brownian motion or processes like the Orstein-Uhlenbeck stochastic processes. Riberio and Hodges (2004) used an Orstein-Uhlenbeck stochastic process in conjunction with Brownian motion to modeling the prices of storable commodities. This is because Brownian motion doesn't have a mean reversion property that appears in commodity prices (Riberio and Hodges 2004) due to fact that future prices must become the spot prices at maturity.

While methodologies of pricing are of the highest importance, the methods utilized by markets to actually determine a spot price to clear the market is also important, as it forms the basis of the modeling of the pricing methodologies. Economists have usually turned to auctions to develop a baseline for prices. What makes these auctions intriguing is what the various bids say about the bidders' risk taking and their expectations. Contreras et al, developed an iterative clearing algorithm to attempt to optimally clear the day ahead electricity market which was simulated under both single and multi-round auctions. However, the algorithms utilized input costs to minimize generation costs and maximize profits (Contreras et al 2001). Another take on the electricity markets is a simulation of market participants instead of the simulation of a single market clearing algorithm. In order to understand the various market constructions, a market was simulated with various bidders or agents that utilized an agent based simulation (Bower and Dunn 1999). Here the agents made bids based on a naïve learning algorithm.

The topic of auctions and agents leads to game theory and behavioral economics as well as more general decision theory. This is because of the overlap between markets and market participants and their decision making processes. One way to look at markets or auctions is as a game. As viewed in terms of a game, there are players (agents or bidders) with strategies and various payoffs for the various strategies utilized. Given this case, auctions can be modeled as a non-cooperative game that may have a Nash Equilibrium (Harrison 1989). Here, in multi-round auctions it is the first bid that matters the most as the agent attempts to get as close to their expected utility and yet beat the other agents (Harrison 1989). With the traditional ideas of rational decision making being challenged by research in game theory and studies showing agents making sub-optimal decisions, then the next step is to go deeper in to the alternative decision criteria, such as risk aversion, and why does this normally sub-optimal basis for decisions occur. The first alternative to the traditional rationality was Prospect Theory (Chateauneuf and Wakker 1999). In 1979 Daniel Kahneman and Amos Tversky published their analysis of decision under risk. Here they challenged expected utility theory and rational choice as the model for economic behavior because of the various anomalies that have been observed. The utilities of payoffs or outcomes, under expected utility, are weighted by their respective probabilities (Kahneman and Tversky 1979). Via numerous experiments, it was found that there is a certainty effect, where agents over weigh payoffs that are perceived to be certain relative to those that are just probable. This was found when people were given choices of payoffs along with their respective probabilities, and a majority of people failed to make decisions based on the weighted expected value (Kahneman and Tversky 1979). Additionally, it was found that in terms of negative payoffs, people were more willing to take a highly probable loss that was greater than a smaller certain loss, even though this choice had a smaller expected value. This lead to the observation that most people are risk averse in the positive domain and risk seeking in the negative domain of payoffs (Kahneman and Tversky 1979). In another experiment it was found that wealth starting points are of great importance. This implies that the change plays an important role in decision making under uncertainty, which is contrary to expected utility as that implies that the change in wealth is not important but the final asset position (Kahneman and Tversky 1979). This also implies that there is some reference point for the evaluations of various payoffs. Another explanation of the anomalies observed by Kahneman and Tversky is mental accounting. Rabin and Thaler argue that mental accounting is a key component in the explanation of these anomalies. It was observed that people track transactions and choices in isolation (Rabin and Thaler 2001). Additionally, these risk averse agents can be viewed as "myopic loss-averters" and are susceptible to a piecemeal extraction of assets. The so called "disposition effect" is where investors hold on losing stocks for too long and sell off winning stocks far too quickly (Grinblatt and Han 2005). Additionally, this can create a spread between the fundamental value of a stock and the equilibrium price.

While Prospect Theory has given one explanation of certain observed anomalies, it has not been taken by all to be the paradigm used to explain economic theory. There are those who disagree. The main critisim of Prospect Theory here is that the scenarios that were contrived for the original experiments did not represent reality and thus lead to a distorted idea of how people make choices under uncertainty (Levy and Levy 2002). Not only the value function has been challenged but also the whole idea of participants in the market place operating under these new rules has been challenged. In the face of this evidence, a competing idea is that those that make decisions that fall in line with prospect theory are inexperienced and can learn, and as they learn their actions will approach the traditional neoclassical models (List 2004). Experiments were conducted to look at all aspects where expected utility and Prospect Theory differ. Added to this, subjects were either educated in the markets and experienced traders or they were complete novices (List 2004). What was discovered is that the results were mixed, where novices were acting as predicted by Prospect Theory and those who were experienced were acting as predicted by expected utility (List 2004). What was also very interesting was that as scenarios were played and explained, the subject's actions would morph over time to approach expected utility (List 2004).

Learning and learning algorithms is the final aspect explored as it weasels its way into the modeling of decision making and games. There are numerous types of learning algorithms that utilize a wide range of techniques and are used in an extremely wide variety of situations. Yet they all perform the same operations, they recall, learn and change over time. They are dynamic. Research conducted by Dow in 1989 showed that it is possible to derive organizational structure through agents who start off poorly informed but learn as they play iterative strategies in non-cooperative games. The agents and the system that emerges converges to a unique organizational structure (Dow 1989). What is interesting here is how scenarios become modeled as games which are then modeled as a network, for instance, the prisoners' dilemma becomes an adaptive network (Dow 1989). An interesting application is in the utilization of neural network learning algorithms to develop solutions to problems that have no closed form solutions. Hutchinson, Lo and Poggio's research used neural networks to achieve what others have used Monte Carlo simulation to do (Hutchinson et al 1994). The network takes the basic inputs of an option's price as inputs and uses previous data to train the network, thus creating a network that can capture the dynamics of the underlier's price process (Hutchinson et al 1994). Yet this requires a large training set of data. However, this approach can be computationally expensive and complex and this has been dealt with over time. New sorting techniques have been established and developed so as to reduce this computational cost with great success (Deb, Pratap, Agarwal and Meyarivan 2002). One of the most exciting areas is support vector machines. This class of learning includes methods of support vector classification and support vector regression. The initial idea is that it utilizes techniques grounded in statistical learning theory and characterizes the properties of learning machines so that they might generalize well to unseen data (Smola and Scholkopf 2004). One unique aspect of support vector machines is how it utilizes kernels to map the data into higher dimensions in order to create a hyperplane that separates the data in to its various classes. Once this operation has been done, the data can then be mapped back to the original sample space in addition with the hyperplane (which at this point can either be linear or nonlinear) separating the data (Smola and Scholkopf 2004). One of the most intriguing applications has been in the arena of quantitative finance to predict stock prices. Fundamental variables and technical indicators have been used as the training data for support vector regression algorithms to make predictions of a stock's time series (Ince and Trafalis 2004). Another advantage of support vector machines is that they can handle incomplete data and high-dimensionality (Gavrishchaka and Ganguli 2003).

3 METHODOLOGY

After much research, it was decided to look at neoclassical rationality versus Prospect Theory in terms of a basis of trading and analysis of the movements of securities and their derivatives. Ideas from learning algorithms and agent based decision making was decided to be included. Simulation was to be a key ingredient to tie it all together. With this tall order and amount of entropy (if one will), the goal became to find an idea that would incorporate all these aspects of research that has spanned a year and a half. Thus, every major topic that was covered in the research has found a home here. The idea of simulating a naïve learning algorithm that utilized both of the concepts found in expected utility theory of neoclassical rationality and those of Prospect Theory in order to make predictions of a fictitious stock price that followed a geometric Brownian stochastic process.

While there are various simulation software packages available to modelers and analysts, this project required the inclusion of user written C. The AweSim simulation environment was a straight forward choice because of the ability to include C files and customize the simulation to fit the needs of the model as well as data collection. Another advantage is in computing time. As the AweSim environment is straightforward and uncluttered it allows the user to keep everything as simple or as complicated as needed, thereby reducing computing time and the required amount of computing power needed.

Additional software to be utilized was Microsoft's Visual Studio 2008 and Excel for Mac 2008. As AweSim accepts user inputs in both visual basic and C languages, writing the code for the learning algorithm naturally followed that it would be written in C. This is due to the inherent flexibility of C and a working familiarity with the language.

The aspect of this project that brings in all the aspect of the research is the algorithms themselves. It is here that the rubber meets the road. In deciding in which direction to head with the algorithm to best represent the research and the goals of the project, the conclusion was reached that it must incorporate learning and be a base for which the concepts of expected utility and Prospect Theory could be demonstrated. The algorithm makes use of naïve learning as it can recall previous data, decisions and strategies utilized.

The goal or more aptly, the purpose, of the algorithm was utilize several strategies to make predictions of a fictional stock's price process. A fictitious stock price process was generated for the simulation and the algorithm attempted to predict each following movement of the price process. Two different algorithms were created, one each for the concepts of expected utility and Prospect Theory. While they were two different algorithms, many aspects were shared. First, they utilized the same method of the naïve learning process. Also, they both chose strategies in similar manners. The expected utility algorithm analyzed the previous data, calculated weighted expected values and made a final prediction. This is the same process the Prospect Theory algorithm follows the same logic, but with a twist. Depending on the value calculated, it would instead make an adjustment that follows the ideas of the s-shaped value function. Deviations between the price process and the predicted values were recorded by the AweSim software along with the strategies that were utilized.

As discussed in the literature, there are several options for modeling a security's price process. One option is the Orstein-Uhlenbeck process. But this process has a much stronger mean-reversion aspect. Moreover, the fictitious stock is assumed not to have any arbitrage, dividends, short interest and is log-normally distributed. While there are many issues with theses assumptions that can affect the price process in the real world, this simplification allows for a straightforward analysis of the algorithms in their first experiment. Due to these assumptions, it was settled that the price process would be modeled for the simulation with geometric Brownian motion of the form:

$$S_t = S_{t-1}e^{((\mu-\sigma 2/2) + \sigma N(0,1))}$$

Where σ is variance, μ is the drift and N(0,1) is a sample from the standard Gaussian normal distribution. Additionally, several scenarios in the experiments were created by adding other terms on to this price process. No drift, positive drift and negative drift scenarios were created. Also, to model seasonal trending, a sinusoidal function was added on to the geometric Brownian motion. The scenarios with and without drifts followed the above formula with only the μ term either equaling 0, +0.001 or -0.001. The seasonal trend equation was:

$$S_t = S_{t-1}e^{((\mu-\sigma 2/2) + \sigma N(0,1))} + 2*\sin(0.5 * t)$$

While this seasonality would lead to a highly improbable price process, it did achieve the goal of creating a price process that would test the algorithms with cyclical data.

As the algorithms were naïve learning algorithms, the strategies that were used were already created and then were chosen by the algorithm. This is opposed to learning algorithms that learn and create new strategies. This was chosen due to time constraints. Moreover, there is more room for research and growth. Nevertheless, six different strategies were chosen as a starting point. Four of the six strategies are finding trends in the data. The fifth makes an estimate based off a sample from the standard normal distribution, and the sixth makes predictions based on Fibonacci support levels and the standard normal distribution. The algorithms saved the values of the price process in various variables for up to ten periods previous and updated these values every tenth of a time unit.

The first strategy is a simple five period moving average, so it predicts the next movement will be the five period moving average. The second strategy uses a ten period moving average. A ten period geometric average was utilized for the third strategy. For the fourth strategy, a ten period Bayesian average with C = 5 was used. The fifth strategy simply takes the previous value for the price process and then adds on a sample from the standard normal distribution. Finally, the sixth strategy is a little more complicated. Many traders that used technical analysis or sometimes also referred to as charting techniques use Fibonacci support levels. Here the previous two values of the price process are compared with various values from the Fibonacci sequence. If the values of the price process exceed the percentages of either 61.53%, 38.19% or 23.52% then it is thought that the price will move up and otherwise it will be heading down. 61.53% is the percentage difference between any two consecutive numbers in the Fibonacci sequence and so on. Once it determines which way it believes the price process to be heading it then adds or subtracts an absolute value taken from the standard normal distribution.

The difference between the expected utility algorithm and the Prospect Theory algorithm in the actual predictions and strategies is that in situations where price falls below a reference point, the Prospect Theory algorithm makes an adjustment of:

$$P_t = (2.25 * P_t) - (1.25 * P_{ref})$$

This represents the value function's risk aversion levels in the prediction. The expected utility algorithm simply makes a prediction with the strategies. Additionally, this adjustment was given a bottom limit as it could lead to negative predictions which fail to make sense regarding prices as they have zero as a bottom limit.

In all cases the algorithms are always calculating the values from each strategy along with the deviations of each strategy and a five period cumulative deviation. Additionally, the algorithms are always calculating weights, which are used when a new strategy is to be chosen by the algorithm. The deviations of each prediction are calculated by:

$$\text{Deviation} = | S_t - P_t |$$

And this is done for every strategy. When the cumulative deviations becomes greater than 3.5 the algorithm will proceed to switch strategies. Which strategy that will be chosen is based off of a comparison of the various five predictions, deviations, cumulative deviations and calculated weights of all six strategies. The strategy with the smallest of all these aspects is then chosen to be the new current strategy. This point is also where the reference point for the Prospect Theory algorithm is set.

4 EXPERIMENTS

With the algorithms ready and scenarios decided upon, the experiments fell into place. There were issues of choosing simulation run length, number of runs. The price process starting point and how the data was to be collected; not to mention exactly which data was to be collected. Additionally, there was the issue of what statistically analysis was to be done on the collected data. However, this all came together in a straightforward manner. The simulations were all run on a Apple MacBook Air with a 1.6 GHz Intel Duo Core processor with 2 Gigabytes of RAM running Apple's OS X and Parallels 2 running a virtual machine of Microsoft's Windows XP.

First the starting point of the price process started with a value of 20. Yet, this value became problematic in a couple of the scenarios. While it proved to be fine for the scenarios where μ was either 0 or some positive number (representing no drift and positive drift), it would lead to bad data in the other two scenarios. This is because it would force the price process to close to zero extremely fast. While a crashing stock can be an interesting case, this failed to serve the purposes of the investigation of the experiment for it never gave the algorithms a chance to run long enough to make enough predictions to analyze. Nevertheless, a starting value of 40 was selected, allowing the algorithms to learn and make predictions.

The price process' starting point was also kept the same for ever scenario for continuity across all scenarios and experiments.

The four scenarios represent four very common trends that appear with stocks. First, the scenario without any drift is a situation that is very similar to a stock of a highly mature company which has already gone through periods of growth. It could also represent a value stock that could begin a new period of growth. A positive drift is analogous to a growth company or even a bubble that is building. The negative drift would be similar to a situation that was faced in the financial sector of the stock market in September 2008 where investors had lost faith in the investment banks and their share prices steadily fell. Finally, the sinusoidal scenario is analogous to the share price of a company whose profits follow with some sort of seasonality like refining stocks or agriculture and certain types of retailers. However, the seasonality had a greater frequency than that is normally observed, which allowed for the algorithms to work harder to make accurate predictions.

In the first scenario, the price process meandered up and down without a real discernable trend. This is in contrast to the simulated Brownian motion of the other scenarios. Scenario two clearly had a strong positive trend. Despite the clear upward trend there was some variability and uncertainty as to how it plays out. The negative drift appeared to be the inverse of scenario two's positive drift, and this is exactly the case. Also, while there was a strong negative trend in each run, none ever actually reach zero, which allowed for the two algorithms to make meaningful predictions. The fourth scenario demonstrates Brownian motion without a drift but with oscillations of a high frequency. It is these oscillations that gave both algorithms more difficult time making accurate predictions.

The run length was also an issue, along with the number of runs. It was decided that the length of each run would be 1950 minutes. This represents the total number of trading minutes during a week on the New York Stock Exchange. The simulations were run for 195 time units but observations of the deviations and all the variables were taken at each tenth of a time unit. The algorithms, also, made updates every tenth of a time unit, thus representing a minute of trading. It was also decided that each scenario would have 100 runs thereby yielding 100 observations of the average deviations and 100 cumulative deviations. Also recorded was each strategy that was used to make each prediction. However, this produced massive amounts of data. The four scenarios were run 100 times for each of the two algorithms which generated over five million data points. This was a far too large of a data set to work with so it was reduced to just 100 observations each of average deviation and the cumulative deviation for each of the eight groups of runs. This generated a reasonable amount of data that would work with the analysis. This data was collected through an AweSim network and saved into a .dat file. This was then imported into Microsoft Excel for the calculation of the averages and editing for the statistical analyses. Once the data was properly formatted in Excel, the analysis was able to be preformed. Eight one way analysis of variance (ANOVA) were calculated. One ANOVA was calculated for each of the four scenarios and the average and cumulative deviations.

5 RESULTS

The expectation from the literature was that the Prospect Theory based algorithm would consistently under perform. This is due to the severe risk aversion. Additionally, it has been previously shown that inexperienced participants in the markets follow closely to the value function of the Prospect Theory whereas more experienced participants' actions are more in line with expected utility. The experiments and their simulations have soundly supported these expectations. In each experiment, the expected utility based algorithm had the least deviations, and thusly had superior performance of the Prospect Theory based algorithm. All the differences were found to be statistically significant.

The first ANOVA calculation is an analysis of the average deviations. The second is an analysis of the cumulative deviations over the course of each run. Both have 100 observations. The average deviations was calculated by finding the average of the deviations of each tenth of a time unit for each run, thusly yielding an average for each run for a total of 100 observations. Here, the expected utility algorithm had an average deviation of 0.65358 for all the runs where as the Prospect Theory had an average of

1.18868. The analysis found the F statistic had a value of 117.758 while the F critical value was 3.8889. This means that we must reject and there is a statistical difference, which implies that the expected utility algorithm has better performance regarding the no drift scenario. The analysis of the cumulative deviations leads to the same conclusion as here the F statistic was 125.9577, and the F critical value was also 3.8889. This serves to strengthen the results.

The third and fourth ANOVA calculations were similar. However, the difference was the scenario, as here there was a positive drift in the price process. The average of the average deviations for the expected utility algorithm was 0.8112 and 2.41096 for the Prospect Theory algorithm. This difference was also found to be statistically significant as the F statistic was 203.6875 and the F critical value was again 3.8889. The second part is the same as above in that it was an analysis of the cumulative deviations over the runs and also because it gave the same results.

The negative drift scenario analysis is featured in calculations five and six. Here the calculations yields similar results to the previous two. The only striking difference is how much closer the average deviations are. The expected utility algorithm had an average of the average deviations of 0.3715 and the Prospect Theory algorithm had an average of 0.6691. While this was still statistically significant (the F statistic was 175.577 and the F critical value was 3.8889), it shows a large improvement over all other scenarios. Yet again, the expected utility algorithm outperformed. This is also seconded by the analysis of the cumulative deviations, where the F statistic was 175.4019 and the F critical value was 3.8889.

Finally the last ANOVA's were the analysis of the sinusoidal adjustment. Here one can see the most dramatic results. The expected utility algorithm's average of average deviations was 0.9732 and the Prospect Theory algorithm's average was 3.3676. This scenario gave the most challenging price process to predict, as shown by the increase in the average deviations. Another interesting calculation is the variance of both algorithms. The variance of the Prospect Theory algorithm was nearly eight times greater than the expected utility algorithm (0.8005 versus 0.1051). Nevertheless, even in this case, there is a statistical difference between the two with the expected utility algorithm outperforming again.

The first scenario with no drift, the two algorithms demonstrated there are two striking differences. The first is the scale of the average deviations, with the expected utility contained in a much narrower band. Also, there is a greater variance in the Prospect Theory's averages. Again, here we can actually see the increase in the variance between the two algorithms, only this is from the positive drift scenario. The striking thing here is how the expected utility algorithm hovers around 0.8. Additionally, the Prospect Theory algorithm has some wild spikes in addition to the significantly greater average deviations. The average deviations from the negative drift scenario the price process steadily decline. This is also the scenario in which the Prospect Theory algorithm best performed. However, as these charts show, the expected utility algorithm had a lower average of average deviations. Although, here, the variance of the expected utility algorithm was greater than previous. The sinusoidal price process was much more challenging for both algorithms. The expected utility algorithm had more variance than before in addition to a greater average. However, the Prospect Theory algorithm had a slightly higher average than previous.

6 CONCLUSIONS AND FUTURE WORK

Several conclusions are able to be drawn from this work. This is due to the fact that there are numerous facets that touch a number of areas. There is also much more work that can be done, for this is just a beginning. The convergence of machine learning, behavioral economics, game theory and modeling can be taken much further. Nevertheless, there are some key conclusions that are more direct.

The first conclusion is that the data demonstrated that the expected utility based algorithm made better predictions on a consistent basis. In all four scenarios, it made predictions with much smaller deviations. Even in the negative drift scenario, where it could be thought that risk aversion would be beneficial, there are several reasons for this. First, it has been shown that the more experienced market participants follow closer to the traditional expected utility model of economic behavior, and despite the larger numbers of novice market participants following more closely to Prospect Theory, these more experienced participants would be able to take advantage and drive the overall market to behave more in line with expected utility theory. This does imply that using Prospect Theory in this generalized manner may be an inappropriate methodology for modeling the market, and may be better suited for a more agent based modeling approach. This also implies that behavioral economics might have limitations regarding the modeling of markets. While it has given us many advantages in the realm of understanding how people make choices, it may not necessarily be good for describing how groups make choices or how the aggregate make choices.

The second conclusion is the limitations of using past trend data for future decisions and limitations of the naïve learning methodology. This is why for future work there are several directions to head. First would be to apply the algorithms to real stock data to observe how they would perform. Another direction would be to explore more complex methods of machine learning such as either neural networks and/or support vector regression. As these are much more complex forms of learning algorithms there is much knowledge to gain. Additionally, there is combination of genetic games, genetic learning/programming. This could have a keen ability for modeling of markets over time, especially when used in conjunction with agent based simulations.

REFERENCES

Altonji, J.G. and Pierret, C.R. 2001. Employer Learning and Statistical Discrimination. *The Quarterly Journal of Economics*, 313-350.

Anderson, L. and Broadie, M. 2004. Primal-Dual Simulation Algorithm for Pricing Multidimensional American Options. *Management Science*, 50, 9, 1222-1234.

Baldi, P., Caramellino, L. and Iovino, M.G. 1999. Pricing General Barrier Options: A Numerical Approach Using Sharp Large Deviations. *Mathematical Finance*, 9, 4, 293-322.

Battiti, R. 1992. First- and Second –Order Methods for Learning: Between Steepest Descent and Newton's Method. *Neural Computation*, 4, 141-166.

Bernstein, P.L. 1998. *Against the Gods: The Remarkable Story of Risk*, John Wiley & Sons, Inc., New York, NY.

Bessembinder, H. and Lemmon, M.L. 2002. Equilibrium Pricing and Optimal Hedging in Electricity Forward Markets. *The Journal of Finance*, 57, 3, 1347-1382.

Black, F. and Scholes, M. 1973. The Pricing of Options and Corporate Liabilities. *The Journal of Political Economy*, 81, 3, 637-654.

Bleichrodt, H., Pinto, J.L. and Wakker, P.P. 2001. Making Descriptive Use of Prospect Theory to Improve the Prescriptive Use of Expected Utility. *Management Science*, 47, 11, 1498-1514.

Blix, M. 2003. A Gas Futures Volatility Model. *Stockholm School of Economics*, Stockholm, Sweden.

Bower, J. and Bunn, D. 2000. A Model-Based Comparison of Pool and Bilateral Market Mechanisms for Electricity Trading. *Energy Journal*, 21, 1-29.

Broadie, M. and Glasserman, P. 1996. Estimating Security Price Derivatives Using Simulation. *Management Science*, 42, 2, 269-285.

Broadie, M. and Glasserman, P. 1997. Pricing American-Style Securities Using Simulation. *Journal of Economic Dynamics and Control*, 21, 1323-1352.

Broadie, M., Glasserman, P. and Ha, Z. 2000. Pricing American Options by Simulation Using a Stochastic Mesh with Optimized Weights. *Probabilistic Constrained Optimization: Methodology and Applications*, pp 32-50.

Carr, P.P. and Bandyopadhyay, A. 2000. How to Derive the Black-Scholes Equation Correctly? http://www.staff.uiuc.edu/~bandyopa.

Chanteauneuf, A. and Wakker, P. 1999. An Axiomatization of Cumulative Prospect Theory for Decision Under Risk. *Journal of Risk and Uncertainty*, 18, 2, 137-145.

Chang, F.R. 2004. *Stochastic Optimization in Continuous Time*, Cambridge University Press, New York, NY.

Chao, H. and Peck, S. 1996. A Market Mechanism For Electric Power Transmission. *Journal of Regulatory Economics*, 10, 25-59.

Charnes, J.M. 2000. Using Simulation for Option Pricing. *Presented at 2000 Winter Simulation Conference*, Orlando, Fl.

Chidambaran, N.K., Lee, C.W.J. and Triguerors, J.R. 1998. An Adaptive Evolutionary Approach to Option Pricing via Genetic Programming. *NYU Working Paper No. FIN-98-086*, New York University, New York, NY.

Contreras, J., Candiles, O., Ignacio de la Fuente, J. and Gomez, T. 2001. Auction Design in Day-Ahead Electricity Markets. *IEEE Transactions on Power Systems*, 16, 1, 88-96.

Cortazar, G. 2000. Simulation and Numerical Methods in Real Options Valuation. *Working Paper*, Pontificia Universidad Catolica de Chile-General.

Crosby, J. 2005. Pricing Commodity Options in a Multi-Factor Jump-Diffusion Model Using Fourier Transforms. *LloydsTSB,* johnc2205@yahoo.com.

Deb, K., Pratap, A., Agarwal, S. and Meyarivan, T. 2002. A Fast and Elitist Multiobjective Genetic Algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, 6, 2, 182-197.

Dow, G.K. 1990. The Organization as an Adaptive Network. *Journal of Economic Behavior and Organization*, 14, 159-185.

Duan, J.C. and Simonato, J.G. 1998. Empirical Martingale Simulation for Asset Prices. *Management Science*, 44, 9, 1218-1233.

Dubey, P., Neyman, A. and Weber, R.J. 1981. Value Theory without Efficiency. *Mathematics of Operations Research*, 6, 1, 122-128.

Dufresne, D. 2001. The Integral of Geometric Brownian Motion. *Advances in Applied Probability*, 33, 1, 223-241.

Dumais, S., Platt, J., Heckerman, D. and Sahami, M. 1998. Inductive Learning Algorithms and Representations for Text Categorization. *Conference on Information and Knowledge Management*, 148-155. Bethesda, MD.

Emmer, S., Kluppelberg, C. and Korn, R. 2001. Optimal Portfolios with Bounded Capital At Risk. *Mathematical Finance*, 11, 4, 365-384.

Fiegenbaum, A. and Thomas, H. 1988. Attitudes Toward Risk and the Risk-Return Paradox: Prospect Theory Explanations. *The Academy of Management Journal*, 31, 1, 85-106.

Fiegenbaum, A. 1990. Prospect Theory and the Risk-Return Association: An Empirical Examination in 85 Industries. *Journal of Economic Behavior and Organization*, 14, 187-203.

Fonesca, C.M. and Fleming, P.J. 1993. Genetic Algorithms for Multiobjective Optimization: Formulation, Discussion and Generalization. *Genetic Algorithms: Proceedings of the Fifth International Conference*, San Mateo, CA, 416-423.

Fu, M.C. and Hu, J.Q. 1995. Sensitivity Analysis for Monte Carlo Simulation of Option Pricing. *Probabilty in the Engineering and Informational Sciences*, 9, 3, 417-446.

Fu, M.C., Laprise, S.B., Madan, D.B., Su, Y. and Wu, R. 2001. Pricing American Options: A Comparison of Monte Carlo Simulation Approaches. *Journal of Computational Finance*, 4, 3, 39-88.

Gavrishchaka, V.V. and Ganguli, S.B. 2003. Volatility forecasting from multiscale and high-dimensional market data. *Neurocomputing*, 55, 285-305.

Ghahramani, Z. and Beal, M.J. 2000. Propagation Algorithms for Variational Bayesian Learning. *Advances in Neural Information Processing Systems*, 13.

Glasserman, P., Heidelberger, P. and Shahabuddin, P. 2002. Portfolio Value-At-Risk with Heavy-Tailed Risk Factors. *Mathematical Finance*, 12, 3, 239-269.

Grinblatt, M. and Han, B. 2005. Prospect theory, mental accounting, and momentum. *Journal of Financial Economics*, 78, 311-339.

Harris, L. 1989. S&P 500 Cash Stock Price Volatilities. *The Journal of Finance*, 44, 5, 1155-1175.

Harrison, G.W. 1989. Theory and Misbehavior of First-Price Auctions. *The American Economic Review*, 79, 4, 749-762.

Harrison, G.W. 1992. Theory and Misbehavior of First-Price Auctions: Reply. *The American Economic Review*, 82, 5, 1426-1443.

Herkommer, D. 2006. Correlation Effects in Credit Risk Models with Incomplete Accounting Information. *Goethe University*, Frankfurt, Germany.

Horn, J., Nafpliotis, N. and Goldberg, D.E. 1994. A Niched Pareto Genetic Algorithm for Multiobjective Optimization, in *Proceedings of the First IEEE Conference on Evolutionary Computation*, 1, 82-87.

Hull, J.C. 2008. *Fundamentals of Futures and Options Markets*, Pearson Prentice Hall, Upper Saddle River, NJ.

Hutchinson, J.M., Lo, A.W. and Poggio, T. 1994. A Nonparametric Approach to Pricing and Hedging Derivative Securities Via Learning Networks. *The Journal of Finance*, 49, 3, 851-889.

Ince, H. and Trafalis, T.B. 2007. Kernel principle component analysis and support vector machines for stock price prediction. *IIE Transactions*, 39, 629-637.

Jonsson, M. and Sircar, K.R. 2002. Partial Hedging in a Stochastic Volatility Environment. *Mathematical Finance*, 12, 4, 375-409.

Kahneman, D. 2003. Maps of Bounded Rationality: Psychology for Behavioral Economics. *The American Economic Review*, 93, 5, 1449-1475.

Kahneman, D. and Thaler, R. 1991. Economic Analysis and the Psychology of Utility: Applications to Compensation Policy. *The American Economic Review*, 81, 2, 341-346.

Karesen, K.F. and Husby, E. 2000. A Joint State-Space Model for Electricity Spot and Futures Prices. *Norsk Regnesentral*, Norwegian Computing Center, Oslo, Norway.

Law, A.M. 2007. *Simulation Modeling & Analysis*. New York: McGraw Hill, Inc.

Levy, M. and Levy, H. 2002. Prospect Theory: Much Ado about Nothing? *Management Science*, 48, 10, 1334-1349.

List, J.A. 2004. Neoclassical Theory versus Prospect Theory: Evidence from the Marketplace. *Econometrica*, 72, 2, 615-625.

Mansouri, H. and Trafalis, T.B. 2007. Reinforcement Learning: An On-Line Framework Using Suport Vectors. *Proceedings of the 2007 International Joint Conference on Neural Networks*, Piscataway, NJ, 1278-1282.

Matsuda, K. 2005. *Introduction to Brownian Motion*, Department of Economics, The City University of New York.

Medasani, S. and Krishnapuram, R. 1997. Determination of the Number of Components in Gaussian Mixtures Using Agglomerative Clustering. Proceedings of the ICNN, Houston, TX, 1412-1417.

Miller, D.A. and Zurada, J.M. 1997. Dynamics of Structural Learning with an Adaptive Forgetting Rate. *International Conference on Neural Networks*, 3, 1827-1832.

Myagkov, M. and Plott, C.R. 1997. Exchange Economies and Loss Exposure: Experiments Exploring Prospect Theory and Competitive Equilibria in Market Environments. *The American Economic Review*, 87, 5, 801-828.

Pritsker, A.A.B. and O'Reilly, J.J. 1999. *Simulation with Visual SLAM and AweSim.* New York: John Wiley & Sons, Inc.

Rabin, M. and Thaler, R.H. 2001. Anomalies: Risk Aversion. *Journal of Economic Perspectives*, 15, 1, 219-232.

Rabin, M. and Thaler, R.H. 2002. Response from Mathew Rabin and Richard H. Thaler. *The Journal of Economic Perspectives*, 16, 2, 229-230.

Ribeiro, D.R. and Hodges S.D. 2004. A Two-Factor Model for Commodity Prices and Futures Valuation. *EFMA 2004 Basel Meetings Paper*, Working Paper Series, University of Warwick, UK.

Ribeiro, D.R. and Hodges, S.D. 2005. A Contago-Constrained Model for Storable Commodities. *Journal of Futures Markets*, 25, 11, 1025-1044.

Roelfsema, M.R. 2000. *Non-linear Index Arbitrage: Exploiting the Dependencies between Index and Stock Options*. CD Delft, Netherlands: Delft University Press.

Ross, S.M. 1985. *Introduction to Probability Models.* New York: Academic Press.

Roth, A.E. and Rothblum, U.G. 1982. Risk Aversion and Nash's Solution for Bargaining Games with Risky Outcomes. *Econometrica*, 50, 3, 639-648.

Samuelson, L. 2002. Evolution and Game Theory. *The Journal of Economic Perspectives*, 16, 2, 47-66.

Schultz, W. 2004. Neural coding of basic reward terms of animal learning theory, game theory, microeconomics and behavioural ecology. *Current Opinion in Neurobiology*, 14, 139-147.

Smola, A.J. and Scholkopf, B. 2004. A Tutorial on Support Vector Regression. *Statistics and Computing*, 14, 3, 199-222.

Stein, M.L. 2002. Fast and Exact Simulation of Fractional Brownian Surfaces. *Journal of Computational and Graphical Statistics*, 11, 3, 587-599.

Takriti, S., Krasenbrink, B. and Wu, L.S.Y. 2000. Incorporating Fuel Constraints and Electricity Spot Prices into the Stochastic Unit Commitment Problem. *Operations Research*, 48, 2, 268-280.

Taleb, N. 1997. *Dynamic Hedging: Managing Vanilla and Exotic Options*. New York: John Wiley & Sons, Inc.

Trafalis, T.B., Couellan, N.P. and Bertrand, S.C. 1997. Training of Supervised Neural Networks via a Nonlinear Primal-Dual Interior-Point Method. *International Conference on Neural Networks*, 3, 2017-2021.

Truman, A. 2006. *Ito Calculus for Pedestrians*, University if Wales Swansea.

Tversky, A. and Kahneman, D. 1981. The Framing of Decisions and the Psychology of Choice. *Science*, 211, 4481, 453-458.

Tversky, A. and Kahneman, D. 1979. Prospect Theory: An Analysis of Decision Under Risk. *Econometrica*, 47, 2, 263-291.

Tversky, A. and Kahneman, D. 1992. Advances in Prospect Theory: Cumulative Representation of Uncertainty. *Journal of Risk and Uncertainty*, 5, 297-323.

Tzikas, D.G., Wei, L., Likas, A., Yang, Y., and Galatsanos, N.P. 2006. A Tutorial on Relevance Vector Machines for Regression and Classification with Applications. *EURASIP News Letter, 17, 2, 4-23.*

Watt, R. 2002. Defending Expected Utility Theory. *The Journal of Economic Perspectives*, 16, 2, 227-229.

Whitcher, B. 2001. Simulating Gaussian Stationary Processes with Unbounded Spectra. *Journal of Computational and Graphical Statistics*, 10, 1, 112-134.

Woo, C.K., Horowitz, I. and Hoang, K. 2001. Cross Hedging and Forward-Contract Pricing of Electricity. *Energy Economics*, 23, 1, 1-15.

Verhoeff, T. 1998. The Trader's Dilemma: A Continuous Version of the Prisoner's Dilemma. *Eindhoven University of Technology*, Eindhoven, Netherlands.

AUTHOR BIOGRAPHIES

CALEB KRELL has earned his Master of Science and Bachelor of Science from the University of Oklahoma. He currently resides in Oklahoma City. His research interests are in the application of simulation and operations research to the areas of financial derivatives, markets and general economic issues. His email address is <caleb.d.krell-2@ou.edu>.

HANK GRANT, Ph.D. joined the faculty at the University of Oklahoma in December of 1993 as Director of the School of Industrial Engineering and Southwestern Bell Professor. Since his arrival, he has been actively involved in all aspects of the growth of School as well as in the startup of the Center for the Study of Wireless Electromagnetic Compatibility. He is currently Director of the Center and Dugan Professor of Industrial Engineering and was recently the S. R. Noble Presidential Professor. Dr. Grant was recently award the Outstanding Industrial Engineer Award for 2001 from Purdue University. Prior to joining the University of Oklahoma, Dr. Grant was with the National Science Foundation in Washington, DC, where he directed programs in Production Systems, Engineering Design, and Operations Research. While at the NSF, he also worked on the design and implementation of a new science station at the South Pole.

Before that, he was Director of the Measurement and Manufacturing Systems Laboratory with Hewlett-Packard in HP Labs. In that capacity, he was responsible for developing the five to ten year vision of HP's manufacturing requirements as well as a new instrument system architecture for the company. Before joining HP, Dr. Grant was involved in the startup and development of two Industrial Engineering software companies: Pritsker Corporation and FACTROL. Dr. Grant received the 1988 Entrepreneur of the Year in High Technology Award from the Institute of American Entrepreneurs (Arthur Young/Venture) for his role in starting FACTROL, which provided factory floor scheduling tools to Fortune 500 companies throughout the US and also in Europe. Dr. Grant is a Fellow of the Institute of Industrial Engineers and is a member of the following societies: INFORMS, the Institute of Industrial Engineers, Tau Beta Pi and The Institute of American Entrepreneurs. Dr. Grant received his Ph.D. from Purdue University in Industrial Engineering in 1980. His email address is <hgrant@ou.edu>.