Modeling and Evaluating Electricity Options Markets with Intelligent Agents

Derek W. Lane[∇], Charles W. Richter, Jr., Gerald B. Sheblé
[∇]Systems Engineering Program
Electrical and Computer Engineering Department
Iowa State University
Ames, Iowa 50011
Phone: (515) 294-3046, Facsimile: (515) 294-4263, E-mail: gsheble@iastate.edu

Abstract: Under deregulation, the formation of electricity markets is a topic of great interest in the power industry and in financial institutions worldwide. Using derivative financial instruments (including options) becomes important for hedging against uncertainty and managing risk--limiting exposure to adverse market conditions. Black and Scholes' equation is often used to value options, but its validity is questionable due to assumptions that may not hold for electricity, most notably the assumption of log-normally distributed prices for the underlying commodity. In this research, a put options market for electricity is modeled. Adaptive agents trade in this market to maximize profit. They are not forced to use an explicit economic or financial model (e.g., Black-Scholes) in their valuation. A genetic algorithm (GA) is used to find alternate valuations that are used to generate buy and sell signals. The results show that it is possible to evolve profitable valuations for use with buying and selling options in this simple model. Reasons for and implications of this finding (e.g., that Black-Scholes may not be a good method for pricing electricity derivatives) are discussed.

Keywords: Black-Scholes, options pricing, adaptive agents, agent-based economics, risk management, power system deregulation.

I. INTRODUCTION

Throughout the world, the electric industry is in the midst of major changes designed to promote competition. No longer vertically integrated with guaranteed customers and suppliers, electric generators and distributors will have to compete to sell and buy electricity. The traditional electric utilities of the past will find themselves in a highly competitive environment. Some countries and regions of the US (e.g., California, PJM) are already operating in a restructured environment. There does not yet appear to be a standardized final market structure that works for all areas, but each market that springs up adds to our experience and helps us make the next market implementation work a little better and more competitively. The authors believe that, to some degree depending on the market implementation, regional commodity exchanges will play a key role in buying and selling electricity.

This research assumes a framework which has been described in detail by Sheblé [10]. Companies presently having both generation and distribution facilities would be divided into separate profit and loss centers. Power is generated by generation companies (GENCOs), transported via transmission companies (TRANSCOs) and distrbution companies (DISTCOs), and is sold to energy service companies (ESCOs) representing the end-consumers. The North American Electricity Reliability Council (NERC) sets the reliability and security standards. Energy mercantile associations (EMAs) will emerge in this competitive electric industry. The EMAs will promote liquidity and as an intermediate partner to all multilateral trades, they will provide assurance to traders, that they need not worry trading because about a defaulting contract partner.

This framework, described by Sheblé [10, 11], allows for cash (consists of spot and forward markets), futures and planning markets. See Figure 1. The spot market allows for trading power each hour (or other duration, e.g., 30 minutes) in the next 30 days. Forward contracts allow energy traders to buy or sell firm electricity contracts as specified in the contract from 1 to 18 months into the future. The futures market allows traders to purchase a non-firm electricity contract for a given month in the future (e.g., 1 to 18 months). Futures contracts provide a means for electricity traders to manage their risk. The planning market is a longer-term market used to develop capital for building large items like new plants and transmission lines.

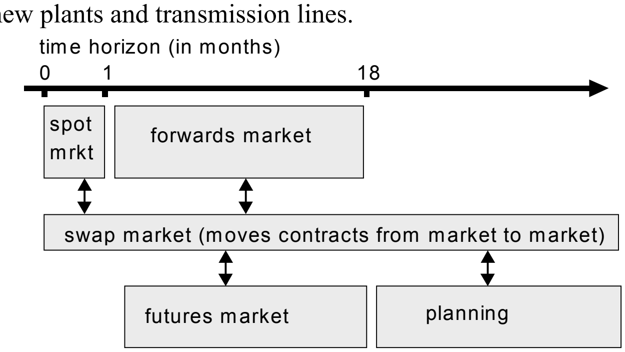

Figure 1 Interconnection between the markets

Options markets (for both futures and physicals) for electric energy are expected to be common and will be an important means of mitigating risk. An option contract gives its holder the right to buy or sell without the obligation to buy or sell. For this right, the holder must pay an up-front premium. The amount of the premium should reflect the value of the option to the potential holders. The worth of an option may vary from trader to trader due to risk preferences, makeup of portfolios (collection of assets and contracts), etc. So the question is, how does one determine the value of an option? The Black-Scholes equation has been used in many markets to value options. It's usage assumes many things about the traded physical commodity that may not be true about electricity.

The approach taken in this research is to allow computerized agents to develop their own valuation formulas as they participate in a simulated option markets. The agents with proper options valuation should achieve higher profit than do other agents with poor valuations. The computerized agents evolve in a genetic algorithm. Those with poor valuations are replaced with new agents that are based on the successful ideas of the better agents.

The remainder of this paper is organized as follows. Section II discusses how options can be used in the deregulated electric marketplace. Section III discusses agent-based economics and how computerized agents trading in simulated options markets evolve in a genetic algorithm. Section IV describes an experiment where the agents optimize their valuation formula and presents the results of that experiment. Finally, Section V presents a summary and conclusion of this paper.

II. ELECTRICITY MARKETS AND OPTIONS

As mentioned in the introduction, it is quite likely that regional commodity exchanges in which buyers and sellers participate in a double auction will soon exist. Such exchanges are utilized in other markets and are essentially an extension of the electric market operating in California. A centralized exchange allows many and varied traders to easily trade a common commodity and derivatives based on that commodity for various periods.

In the cash (spot and forward) market, buyers and sellers interact through an independent contract administrator ICA, who matches the bids subject to all operational constraints. GENCOs and ESCOs cooperate with the ICA which is responsible for ensuring that the energy transactions resulting from the matched bids do not overload or render the electrical transmission system insecure. The ICA monitors and responds to the power system limits and transmission capacities.

The *spot market* is what we are most familiar with in the electrical industry. A seller and a buyer agree (either bilaterally or through an exchange) upon a price for a certain number of MWs to be delivered sometime in the near future (e.g., 10 MWs from 1:00 p.m. to 4:00 p.m. tomorrow).

An *options contract* is a form of insurance that gives the option purchaser the right, but not the obligation, to buy (sell) a contract at a given price. For each options contract, there is someone "writing" the contract who in return for a premium, is obligated to sell (buy) at the strike price. See Figure 2. Both the options and the futures contract are financial instruments designed to minimize risk. Although provisions for delivery exist, they are not convenient (e.g., the delivery point is not located where you want it to be located). The trader ultimately cancels his position in the futures market either with a gain or loss. The physicals are then purchased on the spot market to meet demand with the profit or loss having been locked-in via the futures contract.

"Long" denotes ownership; to go long means to purchase the item in question. In the figure, long indicates that the trader has purchased the option and now has the right to buy (call) or the right to sell (put) the future. A trader who writes the option is "short"; to go short is to sell the item in question. Let's assume that the item in question is a MWh of electricity In the long call diagram, the long trader has paid a premium (e.g., $1) to the option writer for the call option. This call option gives the trader the right to buy a MWh for the strike price (e.g., $7). At any price greater than the strike price, the option is exercised and the trader is said to be "in-the-money" regardless of whether he gains or loses. If the price goes above the strike price plus the premium (e.g., $8), the trader has made a profit. The long trader has reduced risk by limiting his losses to the premium. On the other side of the figure is shown what happens from the option writer's point-of-view. The option writer receives the premium for assuming the risk and is obligated to sell the MWh at the strike price even though the market price is higher. The bottom half of the figure shows how the "put" works. The long trader pays a premium to lock in a maximum price (exercise price) that he will have to pay for the MWh. The short trader takes that premium in return for promising to sell the MWh for that same exercise price.

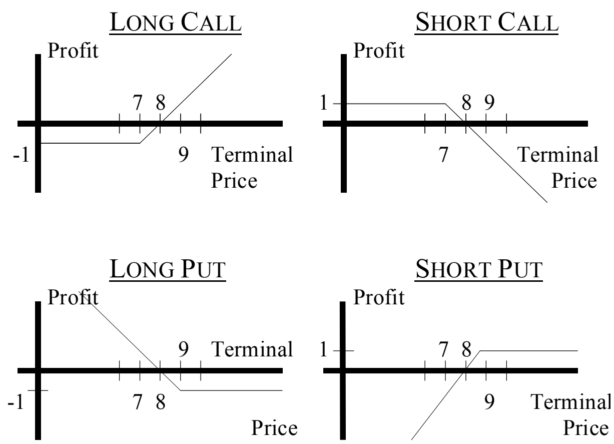

Figure 2 Using put and call options

Determining the value of the option has been the subject of some debate. A couple of decades ago, Black and Scholes put together their formula which has been widely used for valuing options in other commodity markets. Marshall [5] states that Black-Scholes requires:

- The short-term interest rate is known and constant
- The underlying asset pays no dividends
- The underlying asset is efficiently priced
- The option is of the European type
- There are no transaction costs (for buying and selling)
- Can borrow any fraction of underlying asset value
- No artificial restrictions on, or penalties for short selling

The Black-Scholes equation for valuing a put option is as follows [3]:

$$p = \left[X \cdot \exp\left(-r \cdot (T-t)\right) \cdot N(-d2) - S \cdot N(-d1)\right]$$

where:

$X = strike \cdot price$

$S = spot \cdot price$

$r = risk \cdot free \cdot rate$

$N(dn) = Cumulative \cdot Normal \cdot Distribution$

$T = \exp iration \cdot date$

$t = current \cdot time$

$$d1 = \frac{\ln\left(\frac{S}{X}\right) + \left(r + \frac{\sigma^2}{2}\right)\cdot(T-t)}{\sigma \cdot \sqrt{T-t}}$$

$$d2 = d1 - \sigma \cdot \sqrt{T-t}$$

III. AGENT-BASED COMPUTATIONAL ECONOMICS

Market participants (suppliers and consumers) and their responses to prices can be quite complex, changing with time and with market conditions. Realistic models allow agents to modify their behavior as time goes along, varying their strategies to optimize their position. Although some researchers model market agents with fixed rules, and model the market responses using control theory, it is quite likely that they make some unrealistic assumptions about the players or the market. In reality the solutions reached by a market are a function of the agents who are participating in the market at that point in time.

Agent-based computational economics (ACE) allows one to model each part of these complex systems without having to make possibly unrealistic assumptions in order to insure that a solution to the model can be found. As suggested by its name, ACE uses interacting intelligent/adaptive agents to do economic computations. Agents are well suited to economic problems for many reasons. First, an economy is essentially a collection of participants each acting in their own way, hence each player should be modeled independently. Secondly, the agents can be allowed to interact in myriad ways, without those interactions being constrained by the necessity of formulating them as a set of equations. Thirdly, the agents' behavior need not be modeled as a differentiable function; it can be discontinuous and non-differentiable if that is what best models the system of interest. The power and flexibility afforded by agent-based computational models make an ACE approach very worthwhile.

With ACE, one can simulate thousands of possible systems and classify their behavior. Because of advancements in computer technology and the computational power it gives researchers, it is possible to search for an optimal system of functions rather than trying to optimize a contrived system.

The number of different scenarios and possibilities is almost limitless. Starting with a simple model in which the producers and consumers interact via a simple pricing mechanism, one quickly begins to think of things to add and different things to test. The modular design that comes with using agents makes it possible to add to the model. One can model spot market as well as futures and/or options market. The agents then use these markets to maximize their utility. Trying to formulate an optimal futures market and complex interactions with the spot market and other derivative markets using linear algebra would be difficult if not impossible.

There are many different ways of programming adaptive agents (e.g., neural networks, genetic algorithms, etc.). Neural networks (NN) are nice for problems that can be reduced to finding an optimal function for some task. They are well suited for problems in which the search space is relatively smooth and there are no localized sub-optimal solutions. Genetic algorithms (GA) are another search mechanism in which the parts of the solution are coded as a "gene" and one searches for the best collection of genes. They are less susceptible to being trapped in localized sub-optima than are neural nets because one normally tests an entire population of genes all at the same time, which essentially examines discrete points in the search space and selects those grouping of genes that best solve the problem.

The basic genetic algorithm as described by Goldberg [2] can be written as follows (see Figure 3 for a block diagram):

1. Randomly initialize a population and set the generation counter to zero.
2. Until done or out of time, do the following:
   - Calculate the fitness of each member of the population.
   - Select parents using some fitness bias.
   - Crossover the parents to create candidate offspring.
   - Mutate these new offspring.
   - Replace the less fit members with the offspring.
   - Increment the generation counter and go to step 2.

The authors have had some success in previous research in using GAs for constructing adaptive agents. Genetic algorithms are used for the results presented later.

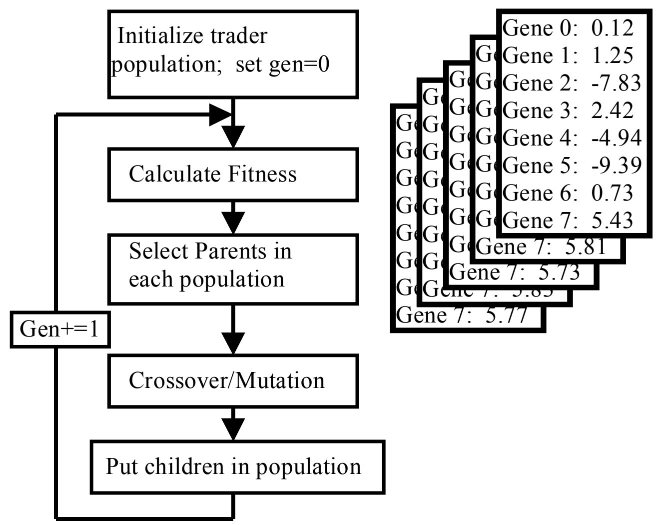

Figure 3. Genetic algorithm to evolve a population of trading agents

IV. VALUING OPTIONS WITH AGENTS

A. General

A simple electricity market with four generators that provide power to an inelastic demand is modeled. Generators are dispatched to meet demand and a market price is determined from the aggregate marginal cost curves of the dispatched generators. Four put options with strike prices of \$15, \$20, \$25, and \$30 are offered with valuations determined using Black-Scholes and the market price data. GA-based agents then buy and sell the options at these Black-Scholes prices. Implicit in the generation of buy and sell signals is a valuation of the put options by each of the agents.

B. Data Preparation

1) Demand data: Hourly demand data for an extended period was provided by a large Midwestern utility and was used as a source of realistic load data in this simulation. See Figure 4.

2) Electricity market price data: Before evolving strategies for buying and selling, price data was needed with which the put option prices were calculated using Black-Scholes.

The hourly demand data was used in conjunction with the generator parameters to determine the market price in an iterative procedure reminiscent of unit-commitment. Each of the suppliers has a unit that is modeled with a quadratic cost curve ($Cost = a + bP + cP^2$). See Table 1 for the values of the coefficients. The supplier produces power as long as the market price does not fall below their minimum marginal cost (which is determined by their minimum production level).

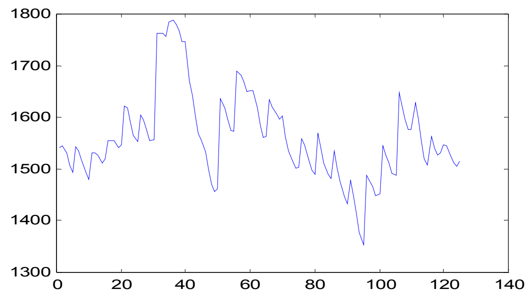

Figure 4 Demand on vertical axis (MW) versus time (hours)

Table 1 Generator parameters

Generator	a	b	c	P_{min}	P_{max}	λ_{min}	λ_{max}
1	100	6	0.005	100	600	7.0	12.0
2	150	7	0.004	120	700	8.0	12.6
3	200	9	0.006	150	750	10.8	18.0
4	250	8	0.007	200	800	10.8	19.2

The marginal cost is found by taking the derivative of the cost curve ($\lambda = b + 2cP$). The marginal cost curves for each generator are shown in Figure 5. Note that each generator has both a minimum and a maximum operating level. (Startup and shutdown costs, ramp rates, and minimum up and down-time constraints were not considered in this simulation.) If the market price is below the minimum marginal price for a generator, that generator is removed from consideration and the market price recalculated. This process is repeated until demand is balanced by a set of generators for whom it is profitable to produce at the discovered price. If price discovery does not occur after 20 iterations, the market price from the 20th iteration is taken as the market price. (Under this simple scheduling scheme, it is possible that a unit could be forced to produce below its minimum marginal cost but a check showed that this never happened.)

A brief clarification at this point may be in order to prevent confusion in the use of the term "spot". The market price is referred to here as the spot price. This is in keeping with the terminology used in finance (i.e. options prices are determined by spot prices); this is not to imply that the hourly market here is the same as the spot electricity market (i.e. the "spot" electricity market as the real-time electricity market). The spot price data for a typical week is shown in Figure 6.

3) Standard deviation of spot price: The standard deviation (sigma) of the spot price is used when calculating the Black-Scholes formula. For a given hour, sigma is calculated using a window of the last 25 peak-period hours prices. The standard deviation of the market price is shown in Figure 7.

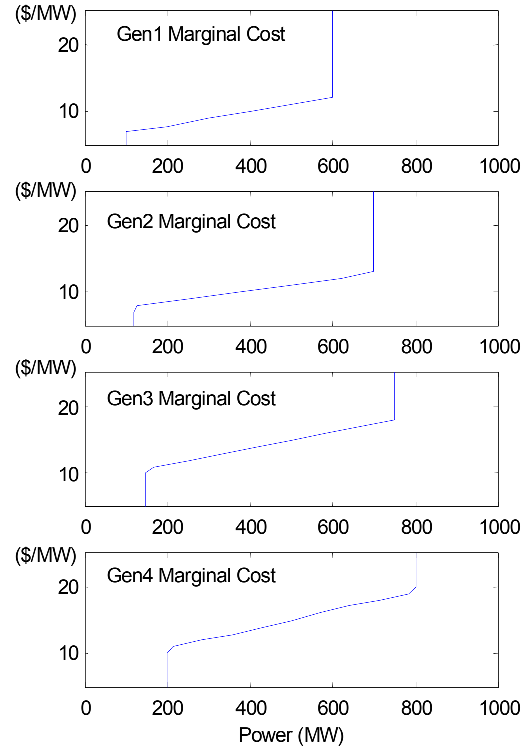

Figure 5 Marginal costs on vertical axes ($/MW) vs. MW for each generator

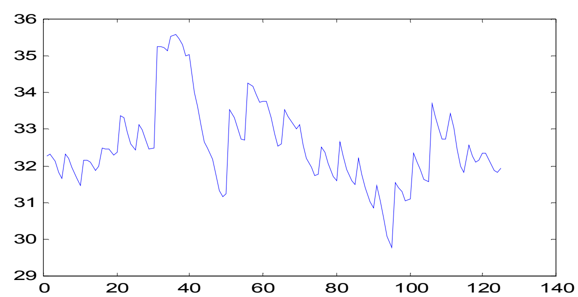

Figure 6 Spot price on vertical axis ($/MWh) versus time (hours)

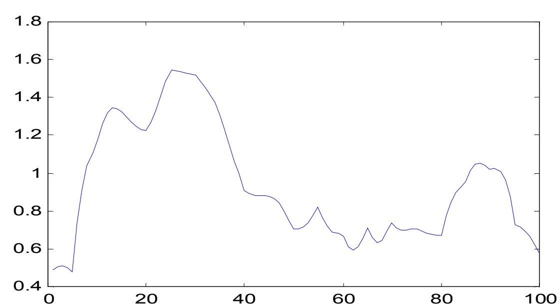

Figure 7 Standard deviation of spot price on vertical axis versus time

3) Put options price data: There are four put options which can be bought and sold having strike prices of $15, $20, $25, and $30. The market valuation (price) of each of these is calculated using the Black-Scholes formula for put options, as presented earlier.

Note that the risk free rate is taken to be constant throughout the simulation and that T-t is a constant 90 days.

This was done to prevent having to "roll over" the options position because the expiration date was reached.

Valuations for the put options are shown in Figures 8. One can see that they go up and down with swings in the underlying spot price of electricity and that the put options with higher strike prices have higher market valuations, as would be expected.

C. Evolving Trading Strategies with a Genetic Algorithm

Each agent in the population buys and/or sells the four put options. These agents act according to internally generated buy and sell signals. These signals are generated using a genetic algorithm to vary the coefficients in a modified Black-Scholes calculation. Options could be traded only for peak periods on weekdays, i.e. Monday-Friday, 11am-4pm.

1) GA Valuation of Options and Buy/Sell Signals: The genetic algorithm is coded as a string of real number genes. The number of genes is determined by the calculation being performed by the GA (described next). For these simulations each GA has 8 genes, each of which is a real number.

The equation currently used by the GA to generate a buy or sell signal is a modified Black-Scholes valuation. A signal to buy or sell an option will be generated if the GA valuation minus the market valuation is greater than some threshold. The terms d1 and d2 in the Black-Scholes formula are recalculated using a modified sigma called σ', where $\sigma' = (Gene2)\cdot\sigma$ and where σ is the "standard" calculation of the standard deviation of the spot price. A buy signal is generated if $\left[Gene\,0 \cdot X \cdot e^{(-r\cdot(T-t))} \cdot N(-d1) - (Gene\,1)\cdot S \cdot N(-d2)\right] + (Gene\ 2)$ > Market Price. Similarly, if a new d1 and d2 are calculated with $\sigma' = (Gene6)\cdot\sigma$ and the Market Price is > $[(Gene4)\cdot X \cdot \exp(-r\cdot(T-t))\cdot N(-d1) - (Gene5)\cdot S \cdot N(-d2)] + (Gene7)$ then a sell signal is generated.

2) Fitness measure: Fitness is determined by the amount of profit made from playing options market many times. The agents are speculators who attempt to buy when options are undervalued by the market and sell when they are overvalued.

A generation consists of 100 hours of buying and selling. The agents have 100 chances or time-steps to buy and sell options. A buy signal causes the agent to buy one option of the specified strike price, increasing their holdings of that option type by 1 and decrementing their bank account by the current market price of that option. A sell signal causes the agent to sell all their holdings of the type of option that generated the sell signal, causing their bank account to increase by (price * number of options held); their holdings for that type of option are then set to zero (since all options were sold). To repeat, only one option is bought in response to a buy signal, while all options are sold in response to a sell signal. This may seem like an oversimplification since real-world agents decide how many options to buy or sell. However, determining the quantity of the transaction adds another dimension of complexity to the problem with which the GAs are presented. Note that agents earn (pay) interest at the risk free rate when they have a positive (negative) bank account.

Fitness is calculated as money earned (or lost) per hour. It is the agent's bank account at the end of the generation divided by the generation length.

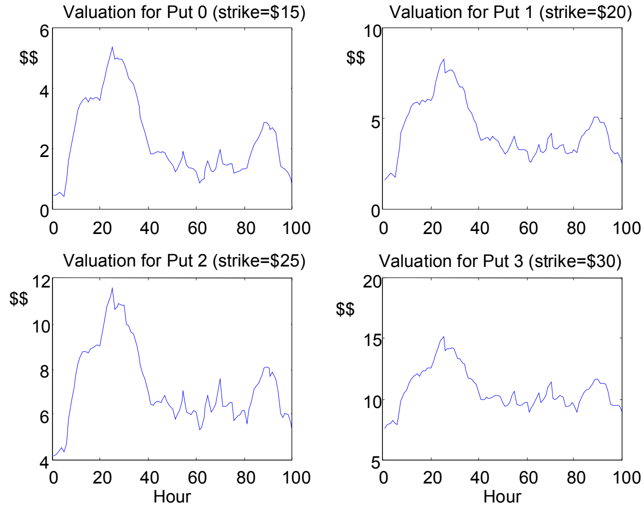

Figure 8 Market valuation for the put options ($ vs. hours)

3) Reproduction: After fitness is calculated, the agents are sorted according to their fitness. Reproduction is performed using single-point crossover of two parents selected from the best half of the population using rank selection. One child is created and replaces an agent in the worst half of the population.

Each child's genes can be mutated in four different ways (bearing in mind that the genes are real-valued):

- 2% of the time the gene is replaced randomly
- 5% of the time the gene is multiplied or divided by 1.5
- 10% of the time the gene is multiplied or divided by 1.05
- 1% of the time the sign of the gene is changed

Random genes are generated according to the relation: $NewGene = GeneMin + Random[0..1]\cdot(GeneMax - GeneMin)$ where *GeneMin* and *GeneMax* are the max and min values of that gene over the entire population. (This was tried as a reasonable way of generating new genes without discarding what the population has collectively learned about the "reasonable" range for a coefficient. This was devised because the space for real numbers is infinite.)

This process is repeated until every agent in the worst half of the population has been replaced. (A variation on this theme is to replace the worst quarter of the population with randomly generated agents, in an effort to introduce new genes into the gene pool and prevent stagnation.)

4) Results: The genetic algorithm was able to evolve a strategy that consistently made a profit buying and selling put options in this market. As shown in Figure 9, the fitness of the best agent is positive and improves over the generations, ultimately reaching a value of $1.5 per trade (with one trade allowed each hour).

Figure 9 also shows the fitness of the worst agent and the average fitness for the whole population. One can see that at the start of the run most agents actually lose money (make a negative profit) but by the end of the simulation the average fitness has risen to nearly zero.

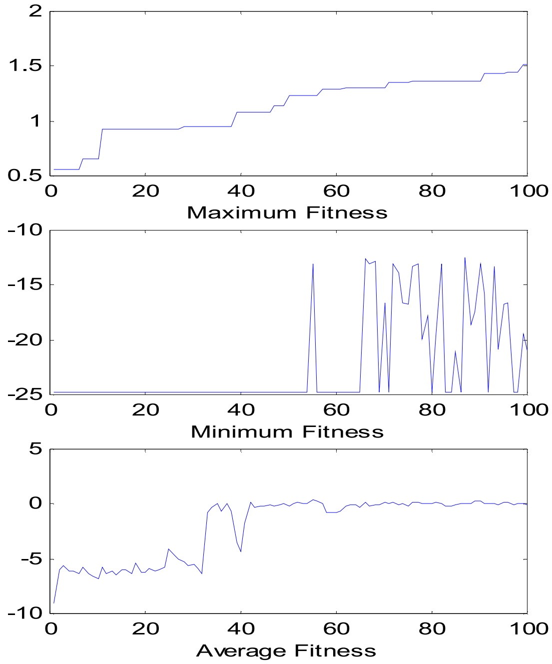

Figure 9 Maximum, minimum, and average fitness over a typical run. The vertical axis measures profit per generation; the horizontal axis counts generations.

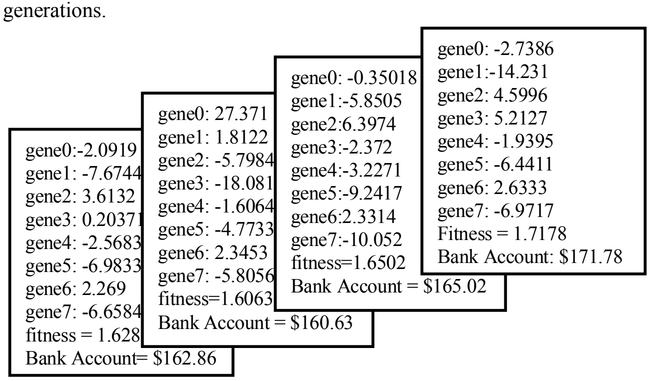

Figure 10 The best genes after 100 generations from 4 different runs

V. SUMMARY AND CONCLUSIONS

The use of derivative financial instruments such as futures and options will be an important and useful tool for managing risk in the deregulated electricity industry. The Black-Scholes valuation for options may not be appropriate for electricity due to the since some of its assumptions may not hold for electricity, most notably a log-normally distributed price of the underlying security. Electricity prices may not follow a jump-diffusion process, especially in the newly deregulated markets. Therefore, another method of valuing electricity options needs to be found. Genetic algorithms are a powerful method for finding successful strategies even when the details of the problem are somewhat obscure.

To test a genetic algorithm's applicability to valuing electricity options, we simulated an electricity market and used Black-Scholes to set the market valuation of put options on electricity. The genetic algorithm was then used to evolve agents whose fitness was measured by their ability to make a profit trading put options. With even a simple representation of the solution, the genetic algorithm was able to earn a profit over the course of a trading run.

While this result was found using admittedly simple electricity and options markets, the ability of a genetic algorithm to find a profitable solution suggests that it may also be successful with more realistic and complex market simulations or even real market data.

VI. ACKNOWLEDGEMENTS

The authors wish to acknowledge the National Science Foundation for their support in researching this topic.

VII. REFERENCES

[1] D. Ashlock and C. Richter, "The Effects of Splitting Populations on Bidding Strategies," Proceedings of the 1997 Conference on Genetic Programming, Denver, CO, 1997.
[2] D. Goldberg, Genetic Algorithms in Search, Optimization & Machine Learning, Reading, MA: Addison-Wesley Publishing Company, Inc., 1989.
[3] J. C. Hull, Options, Futures, and Other Derivatives, USA: Prentice Hall, 1997.
[4] J. Kumar and G. Sheblé, "Framework for energy brokerage system with reserve margin and transmission losses," 1996 IEEE/PES Winter Meeting, 96 WM 190-9 PWRS, NY: IEEE.
[5] J. Marshall, Futures and Option Contracting: Theory and Practice, USA: South Western Publishing, 1989.
[6] C. Richter, D. Ashlock, and G. Sheblé, "Effects of tree size and state number on GP-Automata bidding strategies," Proceedings of the 1998 Conference on GP, Denver, CO, Morgan Kaufmann, 1998.
[7] C. Richter and G. Sheblé, "Genetic algorithm evolution of utility bidding strategies for the competitive marketplace," 1997 IEEE/PES Summer Meeting, PE-752-PWRS-1-05-1997. New York: IEEE.
[8] C. Richter and G. Sheblé, "Building fuzzy bidding strategies for the competitive generator," Proceedings of the 1997 North American Power Symposium, 1997.
[9] G. Sheblé, "Simulation of discrete auction systems for power system risk management," Frontiers of Power, OK, 1994.
[10] G. Sheblé, "Electric energy in a fully evolved marketplace," Presented at the 1994 North American Power Symposium, Kansas State University, KS, 1994.
[11] G. Sheblé, "Priced based operation in an auction market structure." Paper presented at the 1996 IEEE/PES Winter Meeting. Baltimore, MD, 1996.
[12] A. Wood and B. Wollenberg, Power Generation, Operation, and Control. New York: John Wiley & Sons, 1984.

VIII. BIOGRAPHIES

Derek W. Lane received his BS in Engineering Physics from Abilene Christian University and his PhD in Experimental High Energy Physics from ISU. He has been a researcher at Los Alamos National Lab, the Fermi National Accelerator Lab, and the European Laboratory for High Energy Physics. He is presently doing research in systems engineering with an emphasis on finance and economics.

Charles W. Richter, Jr. received his BS-EE from South Dakota State University in 1992. After working at a nuclear power plant and a consulting firm, he studied power systems at Iowa State University where he earned an MS (1996) and PhD (1998) in power systems. Presently he is a temporary assistant professor at ISU. Professional interests include economics of power system operations, market modeling and complex adaptive systems.

Gerald B. Sheblé is a Professor of Electrical Engineering at Iowa State University, Ames, Iowa. Dr. Sheblé received his BS and MS degrees in Electrical Engineering from Purdue University, and his Ph.D. in Electrical Engineering from Virginia Polytechnic. His more than fifteen years of industrial experience include projects with public utilities, research and development, computer vendors and consulting firms. His research interests include power system optimization, scheduling and control.