The Valuation of Power Options in a Pirrong-Jermakyan Model

Craig Pirrong
University of Houston
Houston, TX 77204
713-743-4466
cpirrong@uh.edu

December 8, 2005

Abstract. This article uses a Pirrong-Jermakyan framework to value options on electricity, including daily strike, monthly strike, and spark spread options. This framework posits that power prices depend on two state variables–load and fuel prices. Although variations in load explain a large fraction of variations in power spot prices, the model implies that power option prices do not vary strongly with load except very close to the expiry of daily strike and spark spread options due to the strong mean reversion in load. Load mean reversion also affects time decay, and the evolution of implied volatilities over time. I also discuss how to modify the model to take into account the impact of factors other than load and fuel prices (e.g., outages) that affect power prices; I show that a relatively simple modification can capture the effect of these factors on value if they are not very persistent and not priced in equilibrium.

1 Introduction

The valuation of contingent claims on electricity (such as power forwards and options) presents acute challenges. As noted in Pirrong-Jermakyan ("PJ") (2005), the traditional approach to valuing derivatives is highly problematic when applied to power markets. In this approach, the modeler posits a stochastic process for the price of the claim underlying the derivative. Power spot prices exhibit complex features that are very difficult to capture in an SDDE, however. Even complicated models (e.g., Geman-Roncoroni, 2006) exhibit some unrealistic features. Moreover, it is very difficult to estimate the numerous parameters in these models. Furthermore, even if one estimates the parameters for a reasonable power spot price model in the physical measure, electricity markets are inherently incomplete because power is not properly an asset; derivatives valuation therefore requires determination of an equivalent measure. In addition, many interesting power contingent claims (such as the value of a load serving contract or a power plant) are dependent on both price and output. Valuing such claims in the traditional framework requires the grafting of price and output processes. Given that price-output relations are inherently non-linear, this is a complex endeavor, which is complicated even further by the fact that output is not a traded asset and therefore requires determination of an associated market price of risk.

Given these obstacles inherent in the traditional approach, Pirrong and Jermakyan ("PJ") advance an alternative power derivative pricing model. The PJ model posits that power prices depend on two state variables, load (i.e., demand/output) and fuel prices. It captures the non-linearities (e.g., spikes) in power spot prices via the non-linear relation between load and prices. Moreover, the model reproduces the seasonality, mean reversion, time varying and random power price-fuel price correlations, and time varying load and power price correlations found in electricity prices. Since load is a non-tradeable, the model requires estimation of a market price of risk function, but PJ show how to use standard inverse techniques for solving ill-posed problems to estimate this function from observed forward prices.

Pirrong-Jermakyan (2005) focuses on the pricing of linear power contingent claims (e.g., forward and futures prices). This article investigates the pricing of options in this framework. I derive a valuation PDE that is applicable to any power derivative, and solve this PDE numerically for a variety of commonly traded power options, including monthly strike, daily strike, and spark spread options under the realistic assumption that the correlation between load and fuel forward prices is zero. This assumption allows the computationally efficient solution of the PDE using a combination of finite difference and quadrature methods.[1]

The most interesting finding of this analysis is that the strong mean reversion in load exerts a decisive influence on the pricing of power options. Even though variations in load are the single most important cause of variations in power spot prices, monthly strike options, and daily strike and spark spread options with more than a few days to expiration, exhibit almost no dependence on load. This occurs because strong load mean reversion implies that the conditional distribution of load at expiry (and hence the conditional

[1] When the correlation is non-zero, the option value is most effectively determined using Monte Carlo techniques.

distribution of payoffs) is almost unaffected by current load shocks unless (a) expiration is imminent, and (b) the claim underlying the option (e.g., a forward price) matures soon after expiry. This further implies that monthly strike options, and daily strike options with more than a few days to expiry are effectively options on the price of fuel. Any time decay associated with these options is attributable to the fuel factor, rather than load. Hence, options with values that are linear in the fuel price, such as a spark spread option, exhibit very little time decay until expiry looms.

Mean reversion also impacts the behavior of volatilities implied from option prices. In particular, it implies that: (a) implied volatilities for daily strike options tend to rise sharply as time to expiration falls; (b) implied volatilities for daily strike options depend on load only shortly before expiration; and (c) power options (especially daily strikes) exhibit volatility "smirks" that become more pronounced as time to expiration approaches.

The article also addresses the implications of the fact that although load and fuel prices explain a large fraction of the variation in power spot prices, they do not completely explain these variations. I show that if the differences between spot prices and the best-fitting (non-linear) function of load and fuel prices exhibit very little persistence (as is plausible), the model will produce accurate option valuations if option payoffs depend on a forward price with a day or more to expiry (e.g., a daily forward price). If the option payoff depends on a shorter term forward price (e.g., an hour ahead forward), I show how to modify the initial conditions used to solve the PDE to generate more accurate option values, and identify the circumstances under which this modification is justified.

The remainder of this article is organized as follows. Section 2 sets out the PJ valuation approach. Section 3 describes the types of power options that I will value in this framework. Section 4 presents the computational technique that I use to value these options, and section 5 describes the main findings. Section 6 discusses various extensions of the model to handle real world complications, such as the fact that other factors (e.g., generation outages) also impact power prices. Section 7 summarizes the work.

2 The Pirrong-Jermakyan Valuation Approach

Although the traditional approach to valuing options is to write down a stochastic process of the Price of the underlying claim, PJ (2005) point out numerous deficiencies in this approach as applied to the valuation of power derivatives. Instead, they propose a methodology that exploits the transparency of fundamentals in power markets. In this approach, the price of power is a function of two fundamental driving forces: the demand for power ("load"), and the price of fuel used to generate it. Both variables are observable. Moreover, both clearly are important determinants of power prices. It is more costly to produce electricity with relatively inefficient generating units, as is necessary when demand is high. Similarly, it is more costly to produce electricity when fuel prices are high. Since there should be a strong relation between cost and price in deregulated markets, there should be a strong relation between load and fuel prices on the one hand, and power prices on the other.

That said, PJ also recognize that other factors impact power prices. For instance, outages of generation and transmission assets, spatial variations in load, and path dependence in generation unit economics can cause prices to vary in the absence of variations in fuel prices or load. Although these factors are arguably of little importance when pricing linear power contingent claims (e.g., forward contracts), they may be material when pricing options. Fortunately, under certain assumptions the framework presented here can be modified to address this issue. This modification is deferred to section 6. Until then, I focus on the basic model in which the power price is a function of load and fuel prices alone.

I treat load as a controlled process. Defining load as q_t, note that $q_t \leq X$, where X is physical capacity of the generating and transmission system.[2] If load exceeds this system capacity, the system may fail, imposing substantial costs on power users. The operators of electric power systems (such as the independent system operator in the PJM region) monitor load and intervene to reduce power usage when load approaches levels that threaten the physical reliability of the system.[3] Under certain technical conditions (which are assumed to hold herein), the arguments of Harrison and Taksar (1983) imply

[2] This characterization implicitly assumes that physical capacity is constant. Investment in new capacity, planned maintenance, and random generation and transmission outages cause variations in capacity. This framework is readily adapted to address this issue by interpreting q_t as capacity utilization and setting $X = 1$. Capacity utilization can vary in response to changes in load and changes in capacity. This approach incorporates the effect of outages, demand changes, and secular capacity growth on prices. The only obstacle to implementation of this approach is that data on capacity availability is not readily accessible. In ongoing research I am investigating treating capacity as a latent process, and using Bayesian econometric techniques to extract information about the capacity process from observed real time prices and load. The analysis of price-load relations in section 3 implies that load variations explain most peak load price variations in PJM prices, which suggests that at least over the short run ignoring capacity variation in this market is not critical. This may not be true for all markets.

[3] See various PJM operating manuals available at www.pjm.com for information on emergency procedures in PJM.

that under these circumstances the controlled load process will be a reflected Brownian motion.[4] Formally, the load will solve the following SDE:

$$dq_t = \alpha_q(q_t, t)q_t dt + \sigma_q q_t du_t - dL_t^u \tag{1}$$

where L_t^u is the local time of the load on the capacity boundary.[5] The process L_t^u is increasing (i.e., $dL_t^u > 0$) if and only if $q_t = X$, with $dL_t = 0$ otherwise. That is, q_t is reflected at X.

The dependence of the drift term $\alpha_q(q_t, t)$ on calendar time t reflects the fact that output drift varies systematically both seasonally and within the day. Moreover, the dependence of the drift on q_t allows for mean reversion. One specification that captures these features is:

$$\alpha_q(q_t, t) = \mu(t) + k[\ln q_t - \theta_q(t)] \tag{2}$$

In this expression, $\ln q_t$ reverts to a time-varying mean $\theta_q(t)$. $\theta_q(t)$ can be specified as a sum of sine terms to reflect seasonal, predictable variations in electricity output. Alternatively, it can be represented as a function of calendar time fitted using non-parametric econometric techniques. The parameter $k \leq 0$ measures the speed of mean reversion; the larger $|k|$, the more rapid the reversal of load shocks. The function $\mu(t)$ represents the portion of load drift that depends only on time (particularly time of day). For instance, given $\ln q_t - \theta_q(t)$, load tends to rise from around 3AM to 5PM and then fall from 5PM to 3AM on summer days.

[4] The conditions are (1) there exists a "penalty function" $h(q)$ that is convex in some interval, but is infinite outside the interval, and (2) in the absence of any control, q would evolve as the solution to $dq = \mu dt + \sigma dW$. The penalty function can be interpreted as the cost associated with large loads. If $q > X$, the system may fail, resulting in huge costs. I thank Heber Farnsworth for making me aware of the Harrison-Taksar approach.

[5] This is an example of a Skorokhod Equation.

The load volatility σ_q in (1) is represented as a constant, but it can depend on q_t and t. There is some empirical evidence of slight seasonality in the variance of q_t.

The second state variable is a fuel price. For some regions of the country, natural gas is the marginal fuel. In other regions, coal is the marginal fuel. In some regions, natural gas is the marginal fuel sometimes and coal is the marginal fuel at others. I abstract from these complications and specify the process for the marginal fuel price. The process for the forward price of the marginal fuel is:

$$\frac{df_{t,T}}{f_{t,T}} = \alpha_f(f_{t,T}, t) + \sigma_f(f_{t,T}, t)dz_t \tag{3}$$

where $f_{t,T}$ is the price of fuel for delivery on date T as of t and dz is a standard Brownian motion. Note that $f_{T,T}$ is the spot price of fuel on date T.

The processes $\{q_t, f_{t,T}, t \geq 0\}$ solve (1) and (3) under the "true" probability measure $\mathcal{P}$. To price power contingent claims, we need to find an equivalent measure $\mathcal{Q}$ under which deflated prices for claims with payoffs that depend on q_t and $f_{t,T}$ are martingales. Since $\mathcal{P}$ and $\mathcal{Q}$ must share sets of measure 0, q_t must reflect at X under $\mathcal{Q}$ as it does under $\mathcal{P}$. Therefore, under $\mathcal{Q}$, q_t solves the SDE:

$$dq_t = [\alpha_q(q_t, t) - \sigma_q \lambda(q_t, t)]q_t dt + \sigma_q q_t du_t^* - dL_t^u$$

In this expression $\lambda(q_t, t)$ is the market price of risk function and du_t^* is a $\mathcal{Q}$ martingale. Since fuel is a traded asset, under the equivalent measure $df_{t,T}/f_{t,T} = \sigma_f dz_t^*$, where dz_t^* is a $\mathcal{Q}$ martingale. The change in the drift functions is due to the change in measure.

Define the discount factor $Y_t = \exp(-\int_0^t r_s ds)$ where r_s is the (assumed deterministic) interest rate at time s. (Later we assume that the interest rate is a constant r.) Under $\mathcal{Q}$, the evolution of a deflated power price contingent claim C is:

$$Y_t C_t = Y_0 C_0 + \int_0^t C_s dY_s + \int_0^t Y_s dC_s$$

In this expression, C_s indicates the value of the derivative at time s and Y_s denotes the value of one dollar received at time s as of time 0. Using Ito's lemma, this can be rewritten as:

$$Y_t C_t = C_0 + \int_0^t Y_s(\mathcal{A}C + \frac{\partial C}{\partial s} - r_s C_s)ds + \int_0^t [\frac{\partial C}{\partial q} du_s^* + \frac{\partial C}{\partial f} dz_s^*] - \int_0^t Y_s \frac{\partial C}{\partial q} dL_s^u$$

where $\mathcal{A}$ is an operator such that:

$$\begin{aligned} \mathcal{A}C &= \frac{\partial C}{\partial q_t}[\alpha_q(q_t,t) - \sigma_q \lambda(q_t,t)]q_t \\ &\quad + .5\frac{\partial^2 C}{\partial q_t^2}\sigma_q^2 q_t^2 + .5\frac{\partial^2 C}{\partial f_{t,T}^2}\sigma_f^2 f_{t,T}^2 + \frac{\partial^2 C}{\partial q_t \partial f_{t,T}}\sigma_f \sigma_q \rho_{qf} q_t f_{t,T}. \end{aligned} \tag{4}$$

For the deflated price of the power contingent claim to be a $\mathcal{Q}$ martingale, it must be the case that:

$$E[\int_0^t Y_s(\mathcal{A}C + \frac{\partial C}{\partial s} - r_s C_s)ds] = 0$$

and

$$E[\int_0^t Y_s \frac{\partial C}{\partial q} dL_s^u] = 0$$

for all t. Since (1) $Y_t > 0$, and (2) $dL_t^u > 0$ only when $q_t = X$, with a constant interest rate r, we can rewrite these conditions as:

$$\mathcal{A}C + \frac{\partial C}{\partial t} - rC = 0 \tag{5}$$

and

$$\frac{\partial C}{\partial q} = 0 \text{ when } q_t = X \tag{6}$$

It is obvious that (5) and (6) are sufficient to ensure that C is a martingale under $\mathcal{Q}$; it is possible to show that these conditions are necessary as well.

Expression (6) is a boundary condition of the Neumann type. This boundary condition is due to the reflecting barrier that is inherent in the physical capacity constraints in the power market.[6] The condition has an intuitive interpretation. If load is at the upper boundary, it will fall almost certainly. If the derivative of the contingent claim with respect to load is non-zero at the boundary, arbitrage is possible. For instance, if the partial derivative is positive, selling the contingent claim cannot generate a loss and almost certainly generates a profit.

Expression (5) can be rewritten as the fundamental valuation PDE:[7]

$$\begin{aligned} rC &= \frac{\partial C}{\partial t} + \frac{\partial C}{\partial q_t}[\alpha_q(q_t,t) - \sigma_q \lambda(q_t,t)]q_t \\ &+ .5\frac{\partial^2 C}{\partial q_t^2}\sigma_q^2 q_t^2 + .5\frac{\partial^2 C}{\partial f_{t,T}^2}\sigma_f^2 f_{t,T}^2 + \frac{\partial^2 C}{\partial q_t \partial f_{t,T}}\sigma_f \sigma_q \rho_{qf} q_t f_{t,T} \end{aligned} \tag{7}$$

This PDE must be solved subject to an initial condition that relates the payoff to the relevant contingent claim to the state variables at option expiry. There are many different types of electricity options (described in more detail in section 3 below), but for the options considered herein, the payoff is a function of some forward price (which could be an hourly, daily, or monthly forward price, for instance) or spot price. Under the assumption

[6] If there is a lower bound on load (a minimum load constraint) there exists another local time process and another Neumann-type boundary condition.

[7] Through a change of variables (to natural logarithms of the state variables) this equation can be transformed to one with constant coefficients on the second-order terms.

that spot prices depend only on load and the fuel price, this implies that the payoff to any option is also a function of the load and fuel price.

For claims with payoffs that depend on a spot price, the methodologies set out in PJ (2005) can be used to relate the spot price to the state variables.[8] Similarly, for a claim with a payoff that depends on forward prices, the PJ model can be used to solve for these forward prices as a function of load and fuel price. These forward prices in turn can be used to establish the relation between option payoff and the state variables.

In either case, before valuing electricity options, it is first necessary to implement the PJ inverse problem solution to determine the market price of risk function $\lambda(q_t)$. The solution to the PDE (7) depends on this market price of risk, which is not directly observable. Instead, it can be implied from the observed forward prices in the marketplace. See PJ (2005) for details of the calibration process.

3 Commonly Traded Power Options

There are a variety of electricity options traded (primarily on the OTC market.) Among the most common are daily strike options, monthly strike options, and spark spread options. I consider each in turn.

3.1 Daily Strike Options

A daily strike option has a payoff that depends on the price of power on a given day. Typically, these options have a payoff that depends on the price

[8] PJ set out two approaches, one using the bids of generators, the other using econometric techniques, to establish this relation.

of power for delivery during peak hours of a given day.

Daily strike options can by physically settled or cash settled. For a physically settled daily strike call option, upon exercise the owner effectively receives a long position in a daily forward contract that entitles him to receive delivery of a fixed amount of power during the peak hours on that day. Upon exercise, the owner of a put establishes a put position in a daily forward contract. The option owner must decide to exercise prior to the beginning of the delivery period (e.g., the day before delivery.)

A cash settled daily strike option can be constructed in many ways. For instance, one can have a cash settled daily strike call in which the owner is paid an amount equal to the maximum of zero or the difference between the relevant daily forward as of the some date prior to the delivery period and the strike price. As an example, the call owner's payoff (determined on Tuesday) may depend on Tuesday's forward price for delivery on Wednesday. Alternatively, a daily strike call can pay the difference between the average spot price observed on the pricing date and the strike. For instance, the daily strike call can pay the maximum of zero or the difference between the average spot price observed on Wednesday and the strike price. In a market with a centralized real time market (such as PJM) it is eminently feasible to construct options with such a payoff structure.

The option payoff may depend appreciably on how the contract is written. Specifically, as detailed in section 6, variations in realized spot prices driven by highly transitory factors (other than load and fuel prices) would tend to cause the expected payoff to the option that is based on realized spot prices to exceed that for the option that is based on the forward price measured some time prior to the delivery period, which is assumed to depend only on load and the fuel price. That section details some potential solutions to this difficulty, but until then I focus on daily strike options with payoffs that depend on a forward price. For such an option, the call payoff at exercise is $(F_{t',T}(q_{t'}, f_{t',T}) - K)^+$ and the put payoff is $(K - F_{t',T}(q_{t'}, f_{t',T}))^+$.

3.2 Monthly Strike Options

Upon exercise, the holder of a monthly strike call receives a long position in a monthly forward contract. For instance, upon exercise at the end of June, the holder of a July monthly strike call receives a forward contract for delivery of a fixed amount of power during the peak hours of the coming July. Denoting the forward price as of exercise date t' for delivery of peak power on day j in the option month as $F_{t',j}$, the payoff to the monthly strike call is:

$$(\frac{\sum_{j \in \mathbf{M}} F_{t',j}}{\sum_{j \in \mathbf{M}} \delta_j} - K)^+$$

where $\mathbf{M}$ is the set of delivery dates in the contract month and δ_j is an indicator variable taking a value of 1 when $j \in \mathbf{M}$ and zero otherwise.

3.3 Spark Spread Options

A spark spread call option has a payoff equal to the maximum of zero or the difference between a forward price and the price of fuel multiplied by a contractually specified heat rate. The heat rate is measured in terms of megawatts (MW) per million British Thermal Units (mmBTU). The heat rate measures the efficiency of a generating plant. The marginal cost of generating power from that plant equals its heat rate multiplied by its fuel price.

Therefore, a spark spread option can be viewed as an option to burn fuel to produce power because its payoff is based on the difference between the price of power and the cost of generating it at a given heat rate. For this reason, power plants are often viewed as bundles of spark spread options, although spark spread options are also traded as stand-alone financial products.

Spark spread options raise some of the same issues relating to the timing of exercise and physical settlement and cash settlement as daily strike options. Specifically, if the spark spread option must be exercised at some time t' prior to the power delivery date T, the call payoff is $(F_{t',T} - f_{t',T}H^*)^+$ where H^* is the contractually specified heat rate, which effectively determines the strike. If the payoff to a cash-settled option is based on realized spot prices over the delivery period, the valuation approach applied herein may underestimate its value because it ignores short-term price fluctuations driven by variables other than load and fuel prices. Again, this issue is discussed in more detail in section 6.

4 Valuation Methodology

4.1 Daily Strike and Monthly Strike Options

I value daily strike and monthly strike options by solving the PDE (7) using a combination of finite difference and quadrature methods. I do so under the assumption that the load-fuel correlation ρ equals zero. This assumption is a reasonably accurate characterization of circumstances in many markets. For example, during 2000-2005, the correlation between PJM load and the NYMEX front month natural gas futures price is .04. Moreover, this assumption eases computation. Most common approaches to solving 2D PDEs, such as the alternating direction implicit (ADI) method or the method described below, do not readily handle $\rho \neq 0$. Although it is possible to create orthogonalized state variables based on eigenvectors and eigenvalues, this transformation destroys the economic information contained in the boundary conditions inherent in the underlying problem. For instance, the von Neumann boundary condition in the problem in q-f space communicates valuable economic information that is lost under this transformation. Therefore, when the correlation is reliably different from zero, a Monte Carlo approach is preferable to the finite difference method employed here.[9] However, in the (reasonable) case where ρ is close to zero, the finite difference and quadrature approach is computationally efficient, and allows more accurate calculations of the relevant "Greeks." Moreover, it suffices to identify the effect of load dynamics on the pricing of power options.

In general, the value of any contingent claim is the expectation under the equivalent measure of the discounted present value of its payoffs. Thus, the value at t of a daily strike call exercisable at t', with $t \leq t' < T$, for instance, is:

$$C_t = e^{-r(t'-t)} \tilde{E}_t [F_{t',T} - K]^+$$

where $\tilde{E}_t$ indicates the time-t expectation under $\mathcal{Q}$. In the PJ model, moreover, the forward price is a multiplicatively separable function of the fuel price and some function of load, t, and T.[10] Specifically, $F_{t',T} = f_{t',T} V(q_t, t, T)$.

[9] In the case of $\rho \neq 0$, combination of Monte Carlo and finite difference methods may be useful. The value of a claim with $\rho = 0$ determined using finite differences can be used as a control variate in the Monte Carlo simulation which uses $\rho \neq 0$.

[10] This result obtains because PJ assume that the spot price of power is a multiplica-

Then:

$$C_t = e^{-r(t'-t)}\tilde{E}_t[f_{t',T}V(q_{t'}, t', T) - K]^+$$

Given zero correlation between $f_{t,T}$ and q_t, this becomes:

$$C_t = e^{-r(t'-t)}\tilde{E}_t^f\tilde{E}_t^q[f_{t',T}V(q_{t'}, t', T) - K]^+$$

where $\tilde{E}_t^f$ indicates $\mathcal{Q}$-expectations over $f_{t',T}$ (conditional on $f_{t,T}$) and $\tilde{E}_t^q$ indicates $\mathcal{Q}$-expectations over $q_{t'}$ (conditional on q_t). Furthermore, the Feynmann-Kac Theorem implies that there exists a function

$$u(q_t, t, T|f_{t',T}) = \tilde{E}_t^q e^{-r(t'-t)}[f_{t',T}V(q_{t'}, t', T) - K]^+$$

that satisfies:

$$ru = \frac{\partial u}{\partial t} + \frac{\partial u}{\partial q_t}[\alpha_q(q_t, t) - \sigma_q\lambda(q_t, t)]q_t + .5\frac{\partial^2 u}{\partial q_t^2}\sigma_q^2 q_t^2 \tag{8}$$

This is a one dimensional PDE that is very easy to solve using an implicit finite difference method. Given this solution, the value of the option is:

$$C_t = \int_0^\infty u(q_t, t, T|f_{t',T})\phi(f_{t',T}|f_{t,T})df_{t',T}$$

where $\phi(.)$ is the (lognormal) density of the date-T fuel forward price at exercise date t' conditional on the current fuel price $f_{t,T}$.

This derivation motivates the numerical methodology. First, I form a valuation grid in time $t = t_1, \ldots, t'$, fuel price f, and load q. The increments in t equal δt, the increments in f equal δf, and the increments in q are

tively separable function of fuel price and load. This specification is motivated by the conventional way of characterizing generation economics, in which the price of energy is equal to the product of a fuel price and a heat rate, where the heat rate is a function of load.

denoted δq.[11] For each value of f in the fuel price grid, I determine the time t' payoff to the option for all values of load in the q grid. Using this vector of payoffs as an initial condition, I use a standard implicit solver to determine $u(q_t, t, T|f_{t',T})$ for all $t = t_1, \ldots, t' - 1$. There is one solution vector for each fuel price and each t in the valuation grid.

For each t, and each $f_{t',T}$ in the fuel grid, I integrate these valuation vectors multiplied by $\phi(f_{t',T}|f_{t,T})$ using a Gaussian quadrature.[12] The result of this integration produces the option value at each $\{t, f_{t,T}, q_t\}$ in the grid.

This approach, which combines finite difference methods with quadrature techniques similar to those of Andricopolous *et al,* (2003) is computationally efficient and produces more stable results than an implementation of standard ADI; when using ADI, the magnitudes of the load drift and volatility parameters for the PJM market result in instability around the locus of q-f points where the option is just at the money, and around the upper boundary for load.

For daily strike options, the initial conditions are determined as follows. It is assumed that the option holder must decide to exercise the option the day prior to the power delivery date, i.e., $t' = T - \delta t$.[13] Upon exercise, for a

[11] δt is set equal to one day (1/365 years) to capture effectively the seasonality inherent in th valuation of a power derivative.

[12] When using Gaussian quadrature, it is necessary to interpolate the $u(.|f_{t',T})$. I use a cubic spline interpolation to mitigate the false volatility problem identified by Tavella and Randall (2000). Moreover, using this quadrature method sometimes requires determination of C for values of $f_{t',T}$ that are outside the valuation grid. Such values are estimated by exploiting boundary conditions motivated by the economic characteristics of the particular instrument. For instance, the value of the call converges to zero as the fuel price approaches zero, and converges to $e^{-r(t'-t)}(f_{t,T}V(q_t, t, T) - K)$ as $f_{t,T}$ approaches infinity.

[13] This assumption can be readily modified.

load q and fuel price f the holder of the call receives a payment equal to the maximum of zero, or the difference (a) between the day-ahead forward price $F_{t',T}(q, f)$ implied by the solution to the Pirrong-Jermakyan model calibrated to the observed curve, and (b) the strike price.[14]

For monthly strike options, the delivery days in the month are first determined at each f and q in the grid. For simplicity, I assume that delivery occurs during the peak hours of each business day of the month. The option is assumed to be exercisable on the business day prior to the first day of the delivery month. On this date, the Pirrong-Jermakyan model forward price for each day of the delivery month (calibrated to the forward curve as of the valuation date) is determined. For instance, the prices of forwards expiring on business days falling between 1 July and 31 July are determined as of the expiry date of 30 June. The proceeds to the exercise of the call equal the maximum of zero, or the difference between the average of these forward prices and the option strike price.

4.2 Spark Spread Options

In the PJ model the forward price is a multiplicatively separable function of the fuel forward price and a function of load. In this case, the payoff to the spark spread option can be re-expressed as:

$$(F_{t',T} - f_{t',T}H^*)^+ = (f_{t',T}V(q_{t'}, t', T) - f_{t',T}H^*)^+ = f_{t',T}(V(q_{t'}, t', T) - H^*)^+.$$

Therefore, the payoff to the spark spread option is multiplicatively separable in load and fuel. Consequently, it is possible to utilize the PJ (2005)

[14] The daily strike put payoff is defined analogously.

decomposition to write the value of the spark spread option as another multiplicatively separable function of the current fuel forward price and current load. Specifically, denoting the spark spread option value as $H(.)$:

$$H(q_t, f_{t,T}, t, T, H^*) = f_{t,T}\Phi(q_t, t, T, H^*).$$

The $\Phi(.)$ function can be determined using a standard implicit solver with $(V(q_{t'}, t', T) - H^*)^+$ as an initial condition.[15]

5 Results

The behavior of power options prices implied by this model is best understood through the use of various figures and focus on a few salient results. The behavior of the "Greeks" in particular sheds light on the economic factors driving the option values.

In this regard, it bears noting that due to the two-dimensional nature of the problem, there are a set of Greeks for each of the state variables. For instance, there is a "load Delta" ($\partial C/\partial q$) and a "load Gamma" ($\partial^2 C/\partial q^2$), and a "fuel Delta" ($\partial C/\partial f$) and a "fuel Gamma" ($\partial^2 C/\partial f^2$). The behavior of the Gammas is of particular interest.

All option values in the figures are based on a calibrated PJ model. The model is calibrated using estimates of load volatility σ_q, mean reversion parameter k, and average log load $\theta_q(t)$ estimated from PJM data for 1 January, 2000-31 May, 2005; see PJ (2005) for a description of the estimation methodology. The model is calibrated to PJM power forward prices (from

[15] Due to the multiplicative separability, using the transformation presented in Pirrong-Jermakyan (2005) it is possible to solve for $\Phi(.)$ even when $\rho \neq 0$.

the NYMEX ClearPort system) and natural gas forward prices for Texas Eastern Pipeline Zone M-3 observed on 7 June, 2005 using the method of Pirrong-Jermakyan. The fuel volatility is the implied volatility from the at-the-money NYMEX natural gas futures options with delivery months corresponding to the maturity of the option being analyzed, as observed on 7 June, 2005.

The valuation grid has 100 points in the load and fuel dimensions. The minimum fuel price is \$1.00, and the maximum is \$25.00. The minimum load is the smallest PJM load observed in 1999-2005, and the maximum load is the total amount of generation bid into PJM on 15 July, 2004 (the date used to determine the payoff function for July forwards in the model calibration–PJM bid data are available only with a six month lag.)

Figure 1 depicts the value of a daily strike call option expiring on 15 July, 2005, measured two days prior to expiration, as a function of fuel price and load. The strike of this option is \$85, which was the at-the-money strike on 7 June, 2005. The horizontal plane dimensions are the fuel price f and the load q (running into the chart from front to back). The option value is increasing in fuel price and load, as would be expected. Thus, load and fuel Deltas are both positive. Note too that there is noticeable convexity of the option value in both f and q. That is, both load and fuel gammas are positive. The load Gamma is noticeably large and positive for high levels of load, and for high fuel prices.

Figure 2 depicts the value of the same option on 7 June, 2005, or approximately 38 days prior to expiry. In the figure, the positive fuel Delta and Gamma are readily apparent; the convexity in fuel price is especially evident for intermediate fuel prices (where the option is near-the-money).

However, the option value exhibits little dependence on load. In fact, the load Delta and load Gamma are effectively zero. (When one plots the option value as a function of load for a given fuel price in Matlab, the change in the option value across the range of load values is smaller than the minimum increment that can be depicted by the Matlab plotting function.) Indeed, the zeroing out of the load Delta and Gamma occurs as time maturity falls to as little as 7 or 8 days. Thus, despite the strong dependence of spot power prices on load, daily strike options with maturities of more than a few days exhibit virtually no dependence on load.

This phenomenon reflects the strong mean reversion in load. Due to this strong mean reversion, the distribution of load for future dates conditional on current load converges quite quickly to the unconditional load distribution. Thus, for maturities beyond a few days, variations in current load convey very little information about the distribution of load at expiry, and thus such variations have little impact on the daily strike option value.

This analysis implies that for a week or more prior to daily strike option expiration, such options are effectively options on fuel. Until expiration nears, these options can be hedged using fuel forwards (to hedge fuel Delta) and fuel options (to hedge fuel Gamma). In the last few days before expiry, however, the option value exhibits progressively stronger dependence on load (especially when load is high), and hedging requires the use of load-sensitive claims (e.g., a forward to hedge load Delta, or another load-sensitive option to hedge load Gamma).

The effects of load mean reversion on power option value is especially evident when one examines monthly strike options. Figure 3 depicts the value of a July, 2005 monthly strike call option one day prior to expiry. Even given this short maturity, there is only a slight load Delta, and virtually no load Gamma. However, the non-zero fuel Delta and Gamma are evident. The lack of load dependence reflects the fact that the payoff to the monthly strike option depends on forward prices for delivery dates that are half-a-month on average after option expiry. For all but the forward contracts maturing a few days after the monthly strike option's expiry, load has little impact on the forward price. Hence, variations in load at expiry have little effect on most of the daily forwards included in the monthly bundle.

Mean reversion also impacts option time decay. This is most evident for a spark spread option. Note that due to their multiplicative separability in load and fuel (and the separability of the forward price in these variables in the PJ framework), conditional on q spark spread option values are linear in the fuel price and hence have a fuel Gamma of zero. Thus, in contrast to what is observed for monthly and daily strike options, this implies that there is no time decay attributable to the fuel factor for a spark spread option. Any time decay for this type of option is attributable to the impact of load.

With this in mind, consider figure 4, which depicts the value of $\Phi(q_t, t, T)$ for a spark spread call option with $H^* = 10$ as a function of time to expiration and load (with the load dimension running into the chart).[16] The maximum time to expiration on the chart is 60 days, and hence corresponds to a mid-

[16]The spark spread option value is extremely high when load is high close to expiration. Therefore, to highlight the lack of time decay and avoid the impact of option values for very high loads on the scaling of the figure, spark spread option values are presented only for loads that are less than 15 percent above the mean load.

August 2005 expiration date. Note that the option value is virtually constant until a few days short of expiration. Thus, there is very little time decay until very close to expiration. As the option nears expiry, however, for low loads the option value declines precipitously. Conversely, for high loads (especially very high loads) the value of the option increases dramatically.

These characteristics again reflect mean reversion in load. Well before expiry, due to mean reversion the conditional distribution of load (the only payoff relevant variable for the spark spread claim) changes virtually not at all as time passes. This contrasts with the value of an option with a payoff determined by a geometric Brownian motion (GBM), where the dispersion in the conditional distribution of the payoff-relevant variable declines monotonically as time passes. The stationarity of load translates into little time decay.

Similar influences affect time decay for daily and monthly strike options. These options exhibit time decay, but this reflects the dependence of payoffs on a GBM–the fuel price. The dispersion in payoffs declines as time passes for monthly and daily strikes due to the fall in the dispersion of fuel prices at expiry. Holding fuel price at expiry constant, the passage of time does not affect the variability in payoffs attributable to load. That is, $\partial u/\partial t$ is very close to zero when a daily strike option has more than a few days prior to expiry (regardless of the level of load), and is very close to zero immediately prior to expiry even when a monthly strike option is at-the-money.

The strong mean reversion in load also impacts the behavior of implied volatility for power options. Although the Black model is not well-suited for pricing power options, practitioners still employ it for that purpose, and option values are often quoted in terms of implied volatilities.

One impact of load mean reversion is to cause implied volatilities for daily strike options to rise systematically as expiration nears. This is depicted in figures 5 and 6. Figure 5 depicts implied volatility as a function of q and f when a daily strike option (struck at \$85) has a month to expiration. Figure 6 presents the implied volatility surface for the same option with only 2 days to expiration. The implied volatilities set the Black formula for an option value with strike \$85 and a forward price given by the PJ model for the appropriate q and f equal to the daily strike option value implied by the solution to (9) for that q and f.

Note that the implied volatility surface is markedly higher with shorter time to expiration, especially for large values of f and q. This again reflects mean reversion. Volatility measures the rate of information flow (Ross, 1989). The constant volatility in a GBM process (that underlies the Black model) means that the rate of information flow is constant over time. This is wildly misleading for electricity. Strong mean reversion in load means that a load shock today confers very little information about the distribution of load even a few days hence. That is, one learns little new about the distribution in load in a month based on an observation of current load. Virtually all of the load-related information flow occurs in the last few days prior to expiration (and variations in load explain upwards of 65 percent of PJM spot power price fluctuations). The Black implied volatility effectively calculates an average rate of information flow. For a power option, the average rate of information flow over a long time prior to expiry is small, whereas the average rate of information flow over a short time leading up to expiry of a daily strike option is large, because virtually all of the information flow occurs in these last few days.

Note that the shape of the implied volatility surface also changes dramatically as one nears expiry. A month prior to expiration, the implied volatility depends on the level of the fuel price (with high fuel prices associated with higher implieds in an S-shaped form), but does not vary with load. With the short-dated option, however, the volatility surface is U-shaped in load for low-to-moderate fuel prices, but sharply increasing in load for high fuel prices. (The concavity at the upper load boundary reflects the von Neumann boundary condition at the boundary.) This reflects the non-linear relation between power prices and load. This non-linear relation causes (a) the volatility of the spot power price to depend on the level of load (i.e., heteroscedasticity), and (b) right skewness in the distribution of the spot price, with more pronounced right skewness for high loads. The heteroscedasticity induces smiles in volatilities (discussed in more detail below), and a greater the right skew causes higher call option values and hence higher implied volatilities.

Not surprisingly, the shift in the volatility surface over time is much less pronounced for monthly strike options. As noted earlier, much of the payoff for a monthly strike option is determined by forward prices for forward contracts with more than a few days to maturity. Thus, load shocks that occur even in the days immediately prior to maturity of the monthly strike option confer very little payoff-relevant information. The rate of information flow days before the monthly strike's expiry is therefore not markedly different than the rate weeks before maturity. Indeed, the information flow is almost entirely related to the price of fuel. Under the assumption that the fuel futures price is a GBM, this implies that the implied volatility is effectively the same regardless of time to expiry of the monthly strike option.

Mean reversion also affects the nature of volatility "smiles" and "smirks" in power options. Long-maturity daily strike options exhibit no smile or smirk–the implied volatility does not vary with strike.[17] However, figure 7 demonstrates that (a) implied volatilities smirk for short dated daily strike options, and (b) the smirk depends on load when time to expiry is small. The figure depicts 3 smiles for a daily strike option expiring on 15 July with 2 days to expiry. The highest curve is for a load that is 5 percent below the mean value (given by $\theta_q(t)$) on this date. The curve with the lowest values of implied volatility at the low strike is for a load that is at the mean value on this date. The curve that cuts across the other two, and which is somewhat U-shaped is for a load that is 5 percent above the mean. Each smirk is centered on the at-the-money strike; since the relevant forward price is different for different load levels this close to expiry, the at-the-money strike differs across options. The figure is centered at the at-the-money strike, with \$1 increments between strikes. The smile is calculated assuming a fuel price of \$7, and a time to expiration of 2 days. Note that the smirk is towards the call wing with loads at or below the mean (i.e., higher volatilities are associated with higher strikes) but that it smiles more symmetrically the

[17] These options should exhibit smiles if fuel options do, as would be the case when fuel prices exhibit stochastic volatility or jumps. In this case, the power option smile will be related to the smile in fuel options. To see this, rewrite the option value as $C = \int_0^\infty v(f_{t',T}, t, T, K|q_{t'})g(q_{t'}|q_t)dq_{t'}$ where $g(.)$ is the distribution of $q_{t'}$ conditional on q_t and $v(.)$ is the value of a contingent claim with initial condition given by its payoff. For instance, for a call this payoff is $(f_{t',T}V(q_{t'}, t', T) - K)^+$ which is the value of a call on $V(.)$ units of fuel and strike K. In the presence of stochastic volatility or jumps, the $v(.)$ function will exhibit a volatility skew, which will impact the skew of the power claim C.

load is well above the mean. It should be noted, however, that the behavior of the smile is also dependent on fuel prices. For some values of fuel price and load, implied volatility can smirk towards the put wing, for instance.

Due to the general lack of load dependence for monthly strike options, even when time to expiration is low, there is no pronounced smile or smirk for these options.

The model implies power options exhibit other features that deserve comment, but are which quite intuitive. These include:

- Daily and monthly strike option values are increasing in the volatility of the fuel price σ_f. Since spark spread option prices are linear functions of fuel forward prices, they do not vary with fuel price volatility.

- Daily and monthly strike and spark spread call option values are increasing in the volatility of load σ_q. The increase is due to two factors. First, an increase in load volatility increases the power forward price due to the effect of Jensen's inequality because the forward payoff is a convex function of load. Second, holding the moneyness of the option constant (by increasing the call strike to off-set the impact of the higher volatility on the forward price), the payoff to the option is a convex function of load, so again a Jensen's inequality effect implies that the higher volatility is associated with a higher option value. For puts, the effect of higher load volatility is ambiguous *a priori* because these two effects work in opposite directions. However, a strike-compensated increase in load volatility increases the put value.

6 Complications

As noted earlier, although fuel prices and load are crucial determinants of power prices, electricity spot prices depend on other factors as well. For instance, outages of transmission or generation assets can influence power prices. Thus, at any instant t, the power spot price reasonably has the form:

$$P_t = f_{t,t}\phi(q_t) + \varepsilon_t$$

where ε_t is orthogonal to the fuel price and load and has an unconditional mean of 0, and $\phi(.)$ is a non-linear function (given by the bid stack, for instance.)

Moreover, it is also plausible that ε_t mean reverts very rapidly. Forced outages, for instance, are typically of short duration.

Note that the date-T forward price at $t < T$ is:

$$F_{t,T} = \tilde{E}_t f_{T,T} V(q_T) + \tilde{E}_t \varepsilon_T$$

If ε_t mean reverts very rapidly (with a half-life measured in hours, for instance), and this risk is not priced, then if $T - t$ is as little as a day then $\tilde{E}_t \varepsilon_T$ (which is conditional on ε_t) is very nearly zero. Hence, $F_{t,T}$ can reasonably be considered a function of q_t and $f_{t,T}$ alone. In this case, the model implemented above, which assumes that exercise proceeds depend only on load and the fuel price, will give accurate option values.

However, in some cases this is problematic. For instance, consider a call option to purchase power during a particular hour that can be exercised shortly before that hour. In that circumstance, the hour ahead forward price for delivery in hour T as of t will depend on ε_t even if this shock mean reverts rapidly. Similarly, in the case of (say) a cash settled daily strike option where the payoff is calculated using the realized spot prices during some hours, the ε_t for these hours will affect the option payoff.

This problem can be addressed in a straightforward fashion based on knowledge of the distribution of ε_t conditional on information available at the time the option holder must exercise. Specifically, consider a call option with strike K that has a payoff that depends on the spot price during a given hour. For each q_t and $f_{t,t}$, calculate:

$$\hat{C}(q_t, f_{t,t}) = E_\varepsilon(f_{t,t}\phi(q_t) + \varepsilon_t - K)^+$$

where the subscript on the expectations operator denotes that the expectation is over ε. If the ε risk is unpriced, the expectation is taken under the physical measure. This measure can be calculated based on parameters derived from a statistical analysis of the errors in a (non-linear) model that relates observed power prices to observed load and fuel prices.[18] The resulting $\hat{C}$ function can be used as an initial condition in the finite difference-quadrature solver described in section 4 above.

Other complications are not so readily handled. For instance, the model valuations depend on the market price of risk function $\lambda(q_t)$. Given a $\lambda(q_t)$ function calibrated to observable derivative price information (e.g., visible forward prices), the solution to the PDE (7) solved subject to the appropriate initial condition will give an option value that is consistent with contemporaneous forward prices used for calibration. However, as Joshi (2004) notes, the market chooses $\lambda(q_t)$, and the market can change its mind. For instance,

[18] This relation can be estimated econometrically, or based on bid curves. Both methods are discussed in Pirrong-Jermakyan (2005).

changes in hedging pressure, driven perhaps by financial shocks to market participants, can affect risk premia in the forward market. That is, such shocks may affect $\lambda(q_t)$. As an example, the collapse of Enron and the subsequent deterioration in the financial condition of merchant energy firms plausibly affected the market price of risk. Similarly, Bessembinder-Lemon and Pirrong-Jermakyan (2005) note that changes in available generating capacity and the changes supply of risk bearing capacity by financial intermediaries can also affect the market price of risk function.

Variations in the market price of risk imply changes in the value of power contingent claims. Although it is not difficult to calculate the sensitivity in power claim values to changes in $\lambda(.)$, this is not sufficient to quantify fully the risk of a power option or forward position, as this risk depends on both this sensitivity and the dynamics of $\lambda(.)$. These dynamics are quite difficult to model and estimate because (a) the process for estimating this function is computationally expensive, (b) the function is typically non-linear, and (c) the function is estimated statistically, and is hence subject to sampling error. Thus, although the methodology set out here and in Pirrong-Jermakyan (2005) can give consistent valuations of many power contingent claims at a point in time, it cannot readily quantify all of the risks of power forwards and options.

7 Summary and Conclusions

The Pirrong-Jermakyan model, which posits that power prices are a function of load and fuel prices, can be used to price a variety of options on electricity. This article demonstrates that the behavior of one of these state variables–notably load–exerts a decisive impact on the pricing of these options. Specifically, load is strongly mean reverting. As a consequence, the conditional distribution of load at option expiration does not vary substantially with contemporaneous load with more than a few days to expiration even though variations in load are the single most important cause of variations in power spot prices. This causes the prices of daily strike options (i.e., options on the delivery of power on a single day that are exercised shortly before the delivery date) to vary little with load more than a few days to expiration. Monthly strike options (i.e., options on the delivery of power during a month that are exercised shortly before the delivery month) exhibit almost no load dependence even as expiry nears. Mean reversion also impacts option time decay; an option with a payoff that is proportional to the fuel price (e.g., a spark spread option) exhibits virtually no time decay until right before expiry.

The Pirrong-Jermakyan model assumes that variations in load and fuel prices explain all variations in power prices. In reality, although these factors are the most important determinants of power price movements, other variables impact power prices as well. Fluctuations in these variables are likely to be highly transitory, so they can be ignored when determining forward prices a few days before contract maturity, or when valuing options with payoffs that depend on the prices of forwards maturing more than a day or two after option expiry. This is not reasonable when valuing options with payoffs that depend on very short term forward prices (e.g., a hour ahead forward), or on spot prices. Under certain simplifying assumptions, however, it is possible to modify the initial conditions to the Pirrong-Jermakyan valuation PDE to take into account transitory fluctuations in power prices attributable to factors other than load and fuel prices, such as short-lived outages or out-of-merit dispatch driven by transmission constraints and fluctuations in the spatial pattern of load.

Figure 1

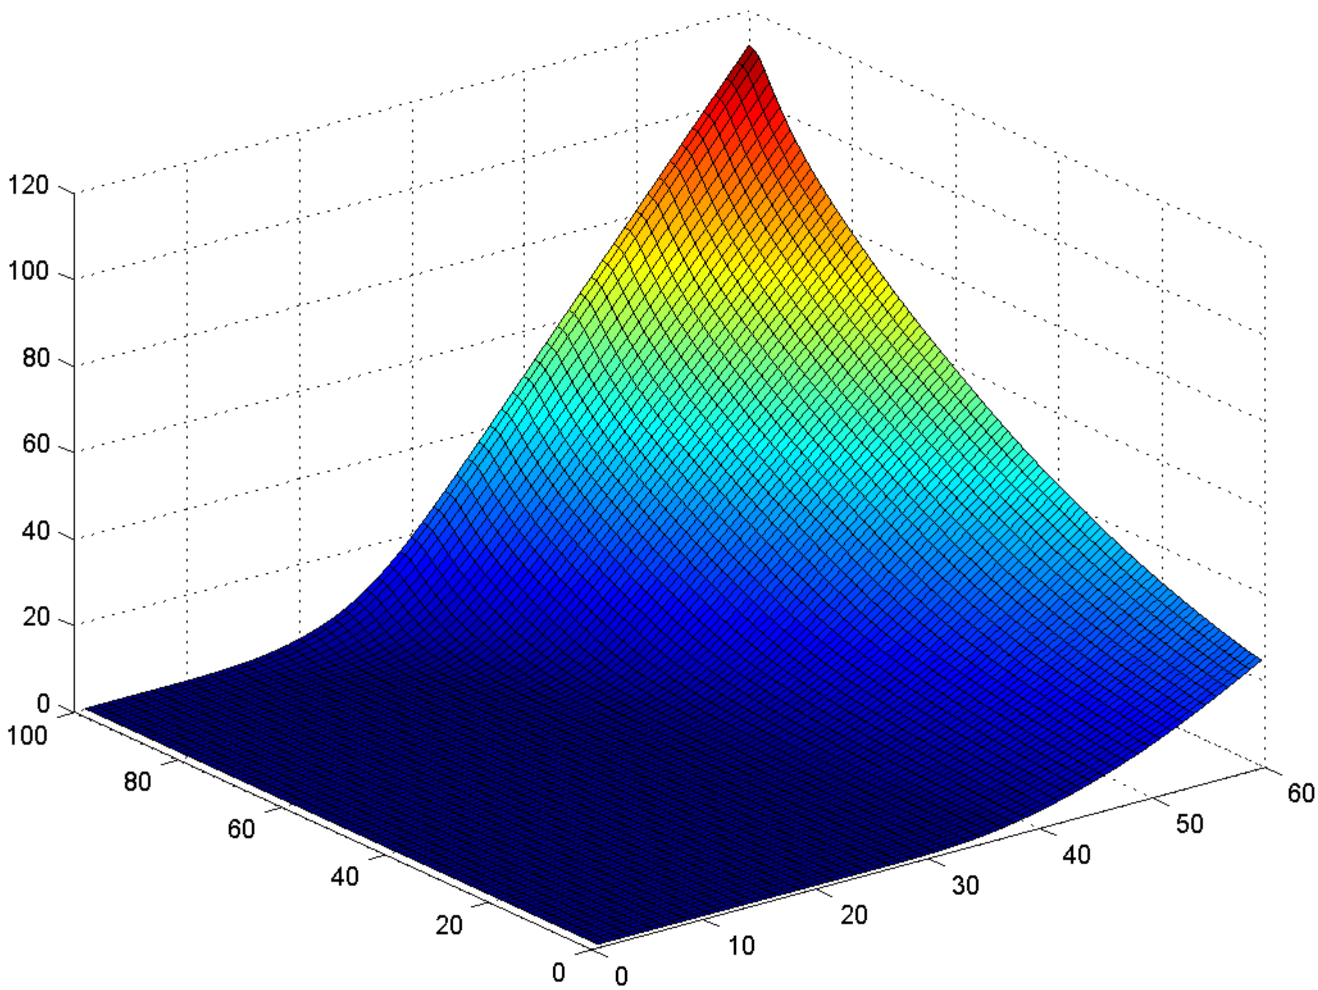

Figure 2

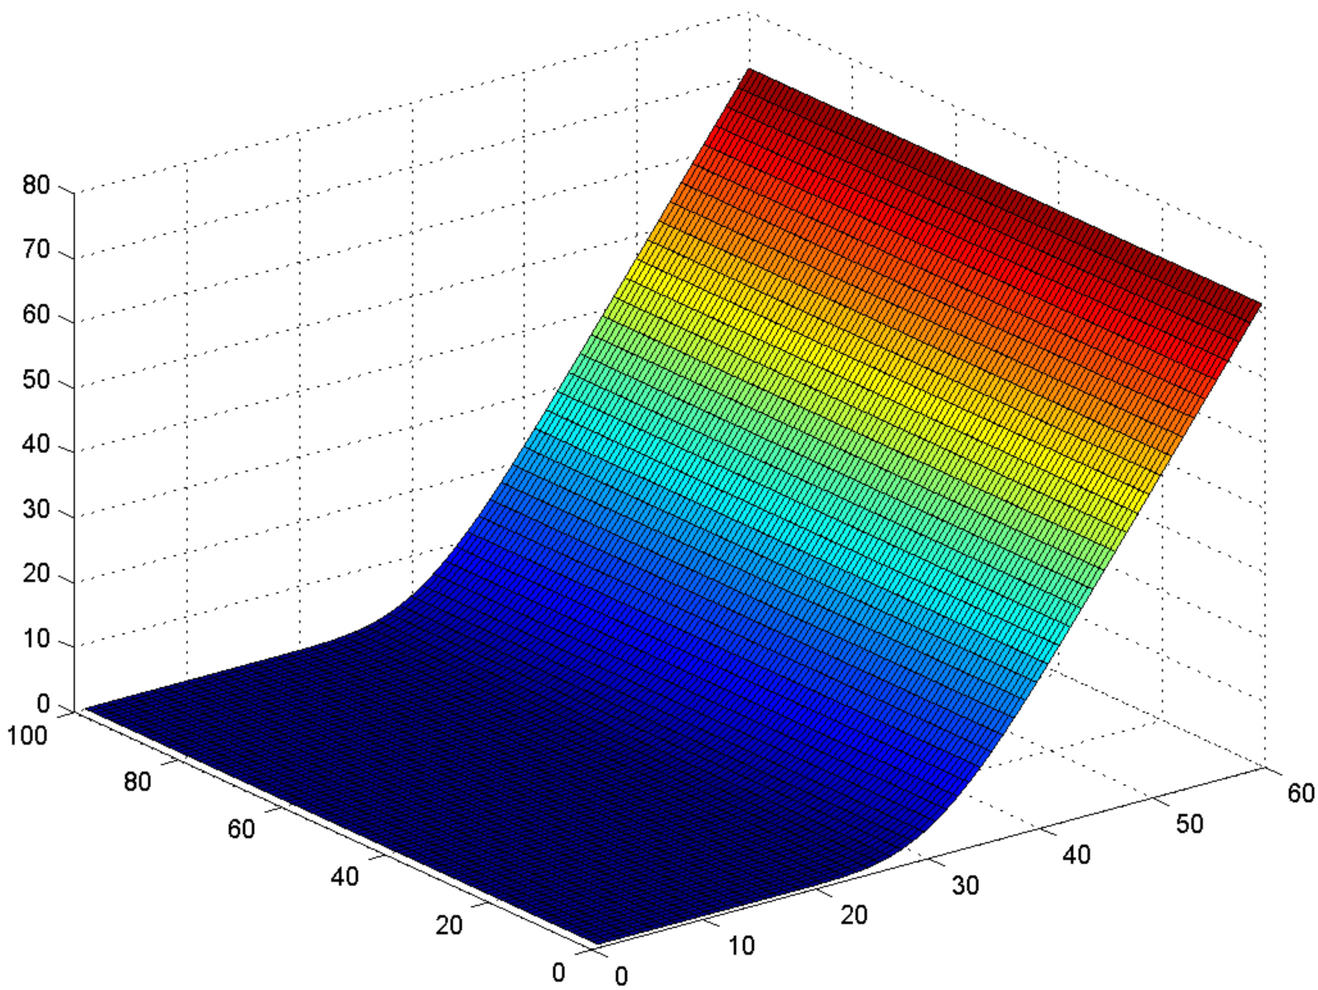

Figure 3

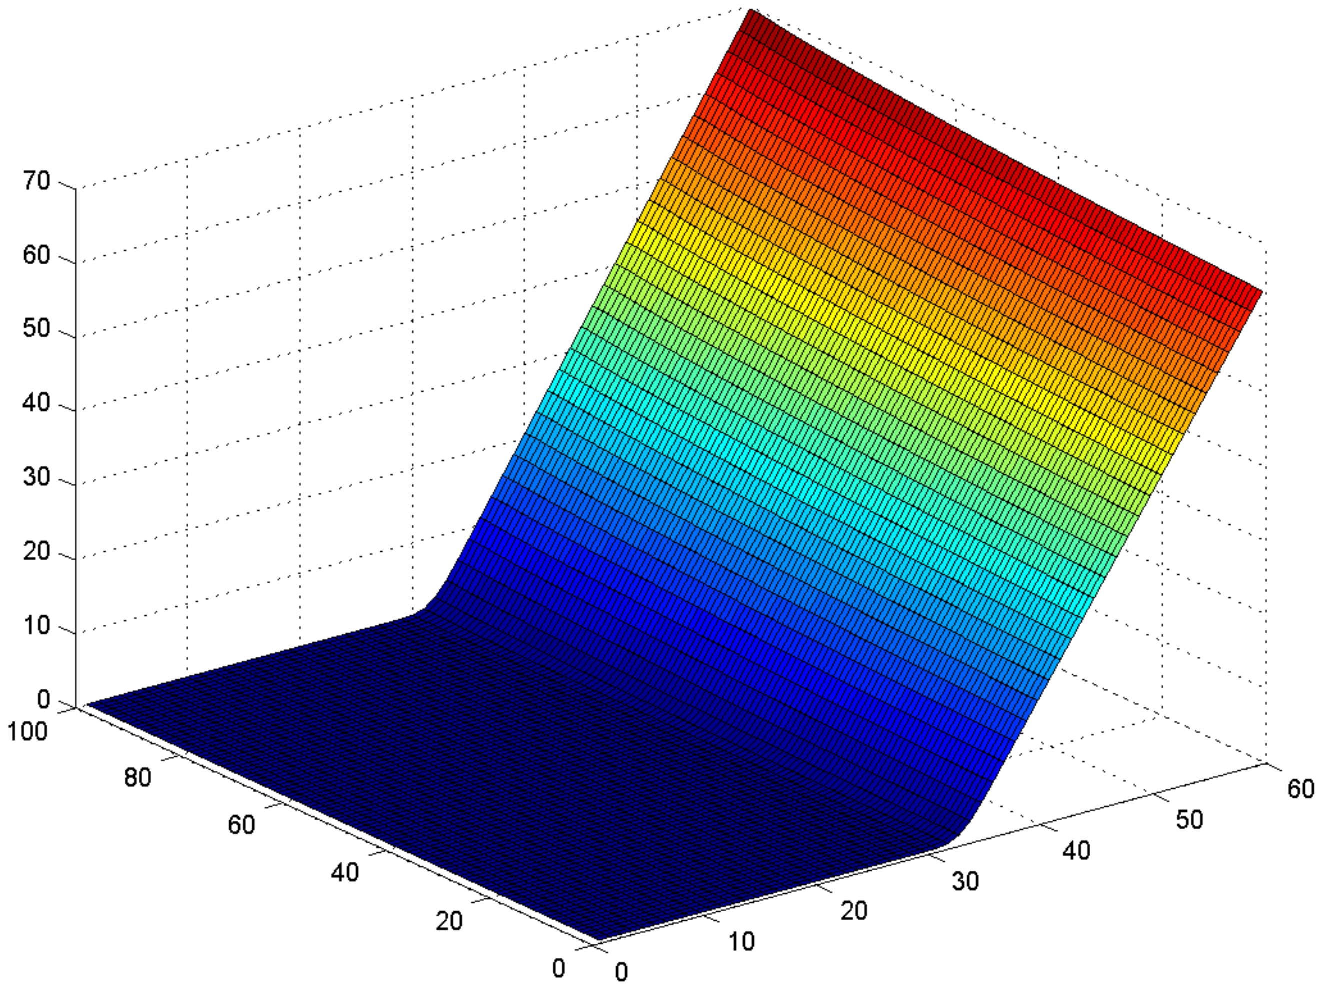

Figure 4

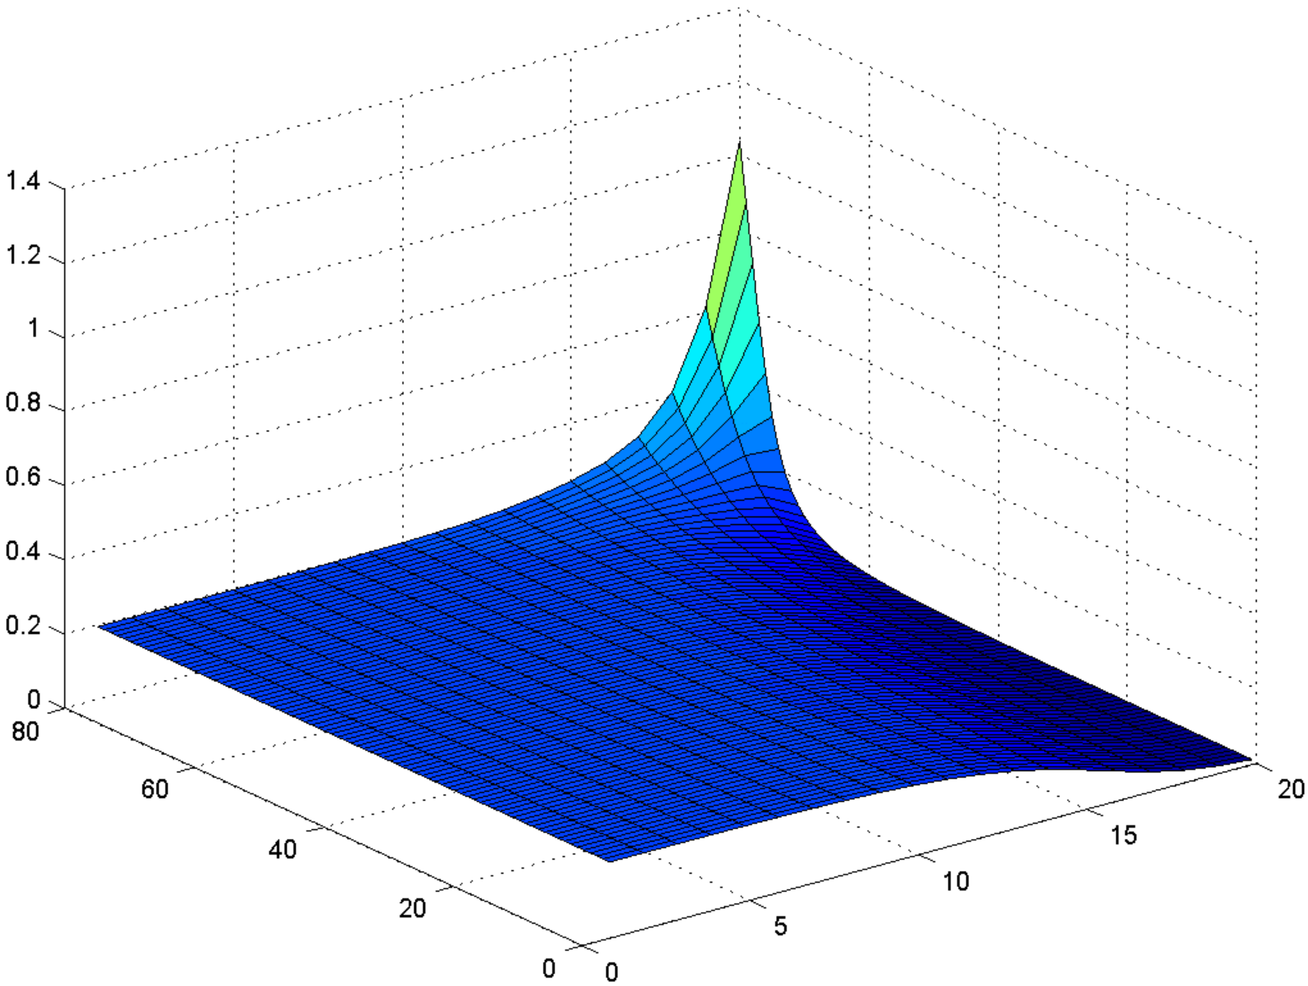

Figure 5

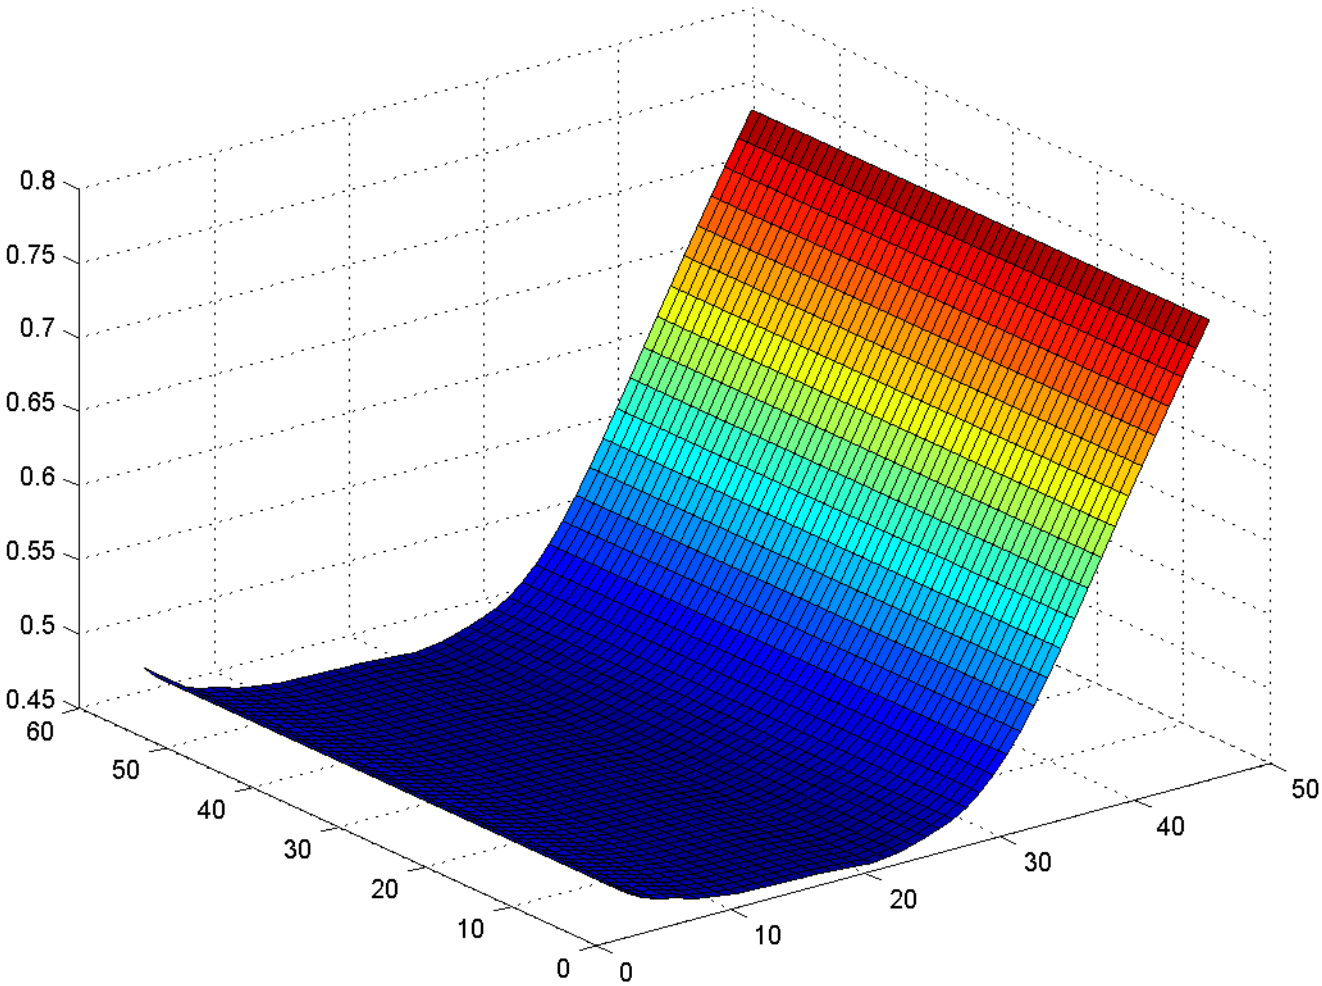

Figure 6

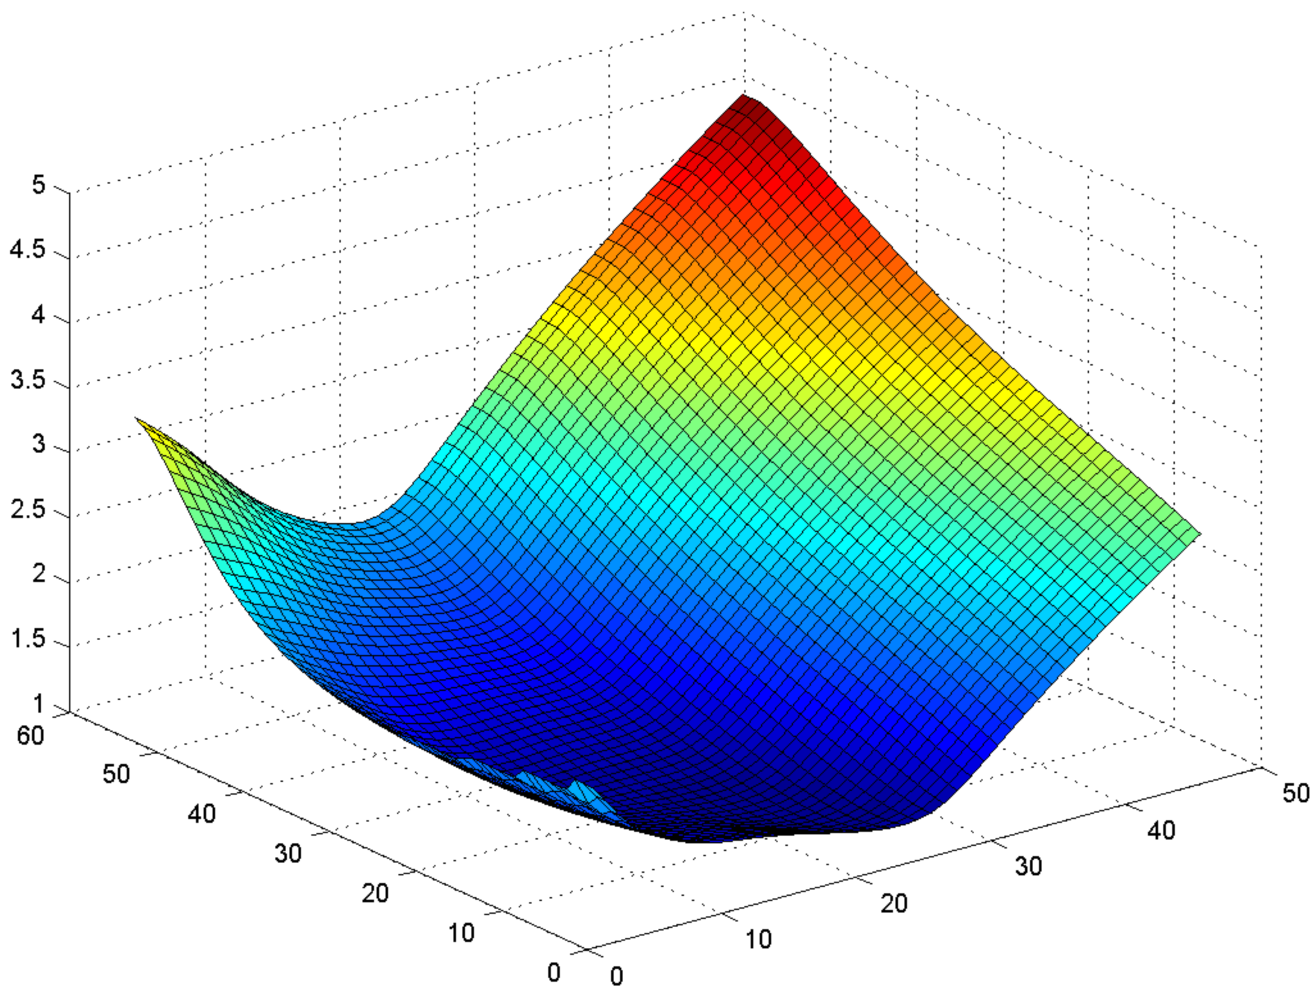

Figure 7