PROBABILISTIC OPTIMALITY IN LONG-TERM ENERGY SALES

R. A. Campo

ABSTRACT

This article introduces a graphic methodology that makes it possible to compare over a wide range of probabilities the revenues that a generator obtains using different policies to decide how much energy to sell under long-term contracts and how much on the spot market. Examples are provided based on results for the Colombian Power System.

KEYWORDS

Electric Markets Policies Probabilistic Comparison.

I. INTRODUCTION

The power sector's new regulatory environment in many countries and, particularly, in Colombia, introduces two markets for energy commercialization: a short-term one (the energy spot market or energy exchange), and a long-term one. The exchange is marked by notable volatility which subjects generators to considerable risks. One way of handling these risks is through long-term energy sale contracts. As an example, in the Summer of 1998 the spot price of electricity hit the 5000 $/MWh mark in some places of the U.S. Midwest. Some distribution utilities without the hedging provided by long term purchase contracts lost hundreds of millions of dollars in a few days. Rumor has it that, as a result, many marketing people lost their jobs. The methodology developed in this paper would have helped these utilities in determining the right amount of long term risk coverage.

It is thus of vital importance for generators to decide on the most suitable amount of energy to be sold on each type of market, in keeping with their financial objectives. For this several approaches are possible. Some workers propose, in particular for mainly hydro power systems, a Monte Carlo simulation analysis, based on the assumption that short-term market volatilities are due to hydrological uncertainties that impact on short term marginal costs. In fact, as can be seen in figure 1, marginal costs of mainly hydro systems exhibit a large degree of volatility, being "almost always" low, with the exception of brief periods of time when unusually low run-offs (caused, for example, by meteorological phenomena like "El Niño" in some regions or "La Niña" in others), translate into very high marginal costs.

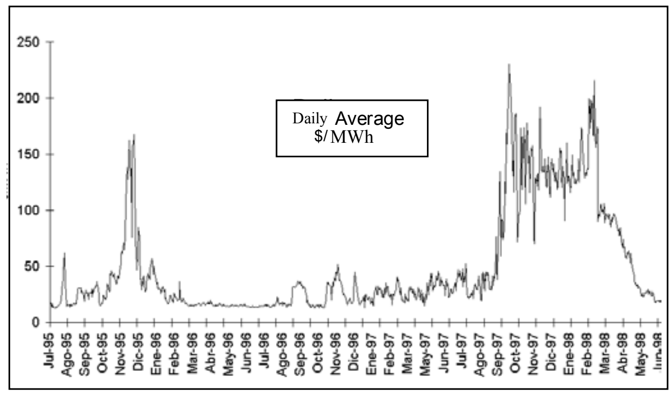

Fig. 1. Marginal prices in the Colombian spot market.

If the rules of the power exchange are such that no generator has market power, the dispatch resulting from the supply on the exchange should be least-cost and can be obtained for several hydrology scenarios with the aid of a suitable minimum cost dispatch model, for example the one included in program **SUPER,** described in detail in reference [1]. On the basis of these scenarios you can determine the most convenient amount of energy to be sold under long-term contracts as will be explained shortly. Whatever energy is not sold under long-term contracts is sold on the exchange.

II. DECISIONS FACING GENERATORS IN COMPETITIVE MARKETS

As an illustration, let´s analyze the decisions that a hydro generator that participates in the Colombian energy exchange (say GEN1) faces regarding how much energy to sell in the long term market for the month of November 1999. The Colombian power system is a mainly hydro one (about 68% of installed capacity) and includes several reservoirs with seasonal or multi annual regulation. Roughly 32% of its generation is thermal, mainly based on gas (25% of total installed capacity) but including some coal units as well. Installed capacity in december 1999 was about 11,600 MW. Energy consumption for 1999 was about 42,000 GWh. A dispatch simulation made by **SUPER** produces for GEN1 generations in GWH as well as system marginal costs for each one of 37 equally likely hydrologies that range from very dry to very wet. These hydrologies constitute the available historical record of monthly inflows to the reservoirs for the past 37 years[1]. Of course GEN1 does not know in advance which hydro scenario will materialize. According to SUPER, the generation of GEN1 for November 1999 ranges between 300 GWH and 1000 GWH, depending on the hydrology. GEN1 can therefore decide in advance to sell in the long term market any amount of energy between 0 and 1000 GWH, at a price that we assume to be the average (under all hydro scenarios) system marginal cost produced by SUPER. This price is (at least in principle) known in advance by all market participants, since all of them have access to the same data and can run the same least cost dispatch simulation program. Furthermore, this price being produced by a least cost dispatch program, corresponds to a Bertrand equilibrium, in the language of game theory ([2]).

In order to use concrete numbers, suppose that the average marginal cost for the month of November 1999 is 20 $/MWh. Let´s assume that HD (Hydro Dry) is a dry scenario for which the marginal cost is 50 $/MWh, while HW (Hydro Wet) is a wet scenario for which the marginal cost is 10 $/MWh (the dryer the hydro scenario, the smaller the total generation available and, therefore, the higher the marginal costs).

Suppose GEN1 sells 300 GWH in the long term market and that HW materializes. For this scenario, assume that GEN1 produces 500 GWH. The difference (200 GWH) is sold in the exchange at the marginal cost that corresponds to HW. The total revenue of GEN1 would then be:

$$300 * 20 + 200 * 10 = 8\,000 \text{ K } \$$$

This is also the operating income of GEN1, under the assumption that its variable operating cost is zero, a valid assumption for hydro generators. On the other hand, suppose that GEN1´s generation under the dry hydrology HD is 100 GWH. If Hydro scenario HD materializes, GEN1 would have to buy 200 GWH in the spot market in order to fulfill its long term contract, at a cost of 50 $/MWh. Its income would then be:

$$300 * 20 - (200 * 50) = -\,4\,000 \text{ K } \$$$

(a loss of 4 million dollars).

Obviously, GEN1 would have sold more energy than 300 GWH in the long term market, had it known that scenario HW would occur and less than 300 GWH if it had expected scenario HD to materialize. Hydro uncertainty translates into income uncertainty. Uncertainties entail risks and there are several ways of dealing with risks, as will be indicated for the present context in the next section.

III. RISKS IN LONG-TERM CONTRACTS

Figure 2 indicates the November 1999 revenues (or operating incomes) of a given generator (let´s call it GEN1) participating in the Colombian electricity market for several long term sales policies ranging from 0 (i.e., all the energy is sold on the energy exchange) to 1000 GWH, which corresponds to the maximum value for which the income of GEN1 is positive for all 37 hydro scenarios. Policies are defined by how much energy the generator decides to sell in the long term market. For each policy, maximum, minimum and average operating incomes (called incomes from now on for simplicity) are displayed. For example, if GEN1 decides to sell all its output in the spot market (in which case there would not be any long term sales), its maximum income would be about 310 million dollars, corresponding to a very dry hydrology and its minimum would be about 10 million dollars, corresponding to a very wet hydro condition. If, on the other hand, GEN1 decides to sell 1000 GWH in the long term market, the maximum income it would obtain is close to 75 million dollars and the minimum income would be close to zero.

[1] Data needed to run program SUPER include technical parameters of all thermal and hydro plants, demand projections, fuel costs, hydrologies, etc., as is usual in programs that simulate minimum cost dispatch of power systems. For the Colombian power system, these data can be obtained from the local Independent System Operator, called ISA, with web address www2.isa.com.co

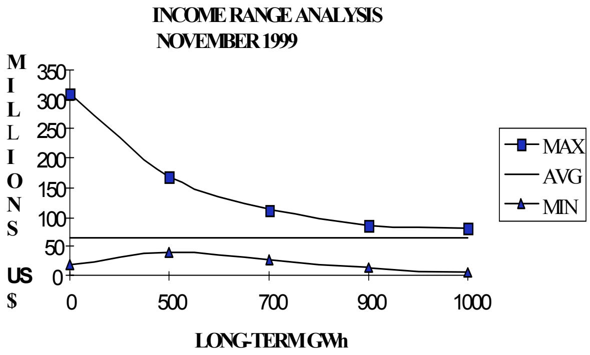

Fig. 2. Income range analysis for November 1999.

In short, to construct figure 2, the incomes of GEN1 are calculated for all long term sale policies in the range 0 to 1000 GWH, for each of the 37 hydro scenarios considered, in the way indicated in section 2. Then maximum, average and minimum incomes (under all hydro scenarios for all sales policies) are plotted. Notice that average incomes are equal for all policies and are not, therefore, suitable for comparing policies with respect to income risk.

One way of dealing with hydro (and income) uncertainties is to choose the policy that maximizes minimum income. In figure 2 this policy corresponds to selling 500 GWH in the long term market. In this way, GEN1 is guaranteed to obtain at least about 45 million dollars, which might be consistent with its financial objectives. An alternative (less conservative) criterion would be to maximize "firm" as opposed to minimum income. By definition, incomes are larger than "firm" incomes 95% of the time. It can be determined that under the "firm" criterion, GEN1 would sell 700 GWH in the long term market, as will be detailed in Section 5.

The approach proposed herein differs from others such as the one described above in that it enables a graphical comparison of the income derived from the use of different policies, over a wide range of probabilities. In the examples provided in section 3 (using actual information from the Colombian power system), situations in which a given policy seems preferable to another are identified, even if its minimum or firm income may be slightly lower, since it dominates for almost all probabilities.[2] The theoretical basis for the methodology developed here can be found in certain properties of the Total Time on Test Transform first described and demonstrated in [3]. A key theorem is established in the Appendix.

IV. TOTAL TIME ON TEST PLOTS

Total time on test plots provide a visual characterization of some of the statistical properties of a sample and allow stochastic comparison of several samples. As we will see, if the Total Time on Test plot of a sample is above the one corresponding to another sample, the first sample probabilistically dominates the second in a way that will be made explicit later on. These plots are based on the Total Time on Test Transform. In what follows, occasional reference is made to some basic statistical concepts that are described in detail in texts like [4].

The Total Time on Test Transform originates in system reliability studies. If at time 0 **n** identical units (for example, light bulbs) are put on test, their failure times, ordered from smallest to largest, yield an order statistics:

$$X_{n:1} \le X_{n:2} \le X_{n:3} \le \ldots \le X_{n:n} \qquad (1)$$

In (1) $X_{n:1}$ is the failure time of the first unit in failing, $X_{n:2}$ is the failure time of the second unit in failing and so on; $X_{n:n}$ is the failure time of the last unit in failing; n, of course, refers to the total number of units that are put on test.

The Total Time on Test (TTT) until unit i fails is defined as follows:

$$T(n:i) = n\,(X_{n:1}) + (n-1)\,(X_{n:2} - X_{n:1}) + \ldots + (n-i+1)\,(X_{n:i} - X_{n:i-1}) \qquad (2)$$

It can be demonstrated (see [3]) that {T(n:i) / i} is an unbiased estimator of minimal variance of the mean life of the units, calculated on the basis on the first i observations.

One generalization of the TTT is the Total Test Time Transform (let´s call it T4) that, for a non-negative random variable with an F distribution function is defined as:

$$T4\,(t) = \int_o^{F^{-1}(t)} [1 - F(u)]\,du \qquad (3)$$

The lower and upper limits of the integral in (3) are 0 and $F^{-1}(t)$, respectively; this last term being the inverse distribution function evaluated at t. If t = 1, then it can be proven that :

$$T4\,(1) = E[X] \qquad (4)$$

[2] The income streams for all 37 hydrologies analyzed can be obtained from the author at rcml1@attglobal.net

i.e, the expected value of X.

The derivative of T4 (t) evaluated at t = F(x) is the inverse of the failure rate corresponding to the F distribution function with probability density function f(t): 1/r(t). The failure rate (r(t) = f(t) / (1 – F(t)) can be interpreted as the tendency to fail at time **t** for a unit whose lifetime is a random variable with the F distribution function, given that it has "survived" until **t**. This is a central concept in system reliability theory, originating in actuarial sciences, in which the failure rate is known as the "force of mortality."

The scaled T4 (ST4) is defined as T4 divided by the mean. Notice that it varies between 0 and 1 when **t** goes from 0 to 1. Furthermore, ST4 is dimensionless. The ST4 characterizes a given probability distribution function in a one to one way. If the function is exponential (with constant failure rate), its ST4 is a 45^o line between 0 and 1 (remember that both axes vary between 0 and 1 and that ST4 is dimensionless). Notice that the plot is dimensionless. If F has an increasing failure rate (IFR), the ST4 is a concave function between 0 and 1. If F has a decreasing failure rate (DFR), its ST4 is a convex function, also between 0 and 1. Reference [3] offers graphs of the ST4 for various classical distribution functions (truncated normal, log-normal, Weibull, gamma, etc.), and shows that the empirical figure obtained on the basis of the data generated from a given distribution resembles the ST4 of the corresponding theoretical distribution. For n observations, ST4 can be estimated as equal to:

ST4 = T(n:i) / T(n:n) (5)

The (scaled) Total Time on Test Plot of several observations is obtained by ordering them from smallest to largest, applying (2) to each ordered value and dividing the result by T(n:n). The vertical axis thus represents scaled total time on tests and the abscissas the relative order of the observations.

To illustrate the calculations, we provide a simple example for 10 incomes, assumed to have been obtained for a given policy P under 10 equally likely hydrologies. By a "hydrology" or "hydro condition" we mean hydro scenario, as defined in Section II above. These hydro scenarios can be built on the basis of hystorical inflow data using well known statistical simulation models. Data and results can be seen in Tables 1 and 2.

HYDRO CONDITION (i)	**INCOME (MM$)**	**SORTED INCOMES (MM$)**	**T(n:i)**
1	290	10	100
2	150	50	460
3	80	60	540
4	300	80	680
5	250	150	1,100
6	10	170	1,200
7	50	200	1,320
8	170	250	1,470
9	60	290	1,550
10	200	300	1,560

Table 1: Calculation of T(n:i)

It can be verified that the average income (156 MM$) is equal to T(10:10) / 10.

On the basis of T(n:i) we can obtain the scaled values dividing by T(10:10) as indicated in Table 2.

i	**Scaled T(n:i)**
1	.0641
2	.295
3	.346
4	.436
5	.705
6	.769
7	.846
8	.942
9	.994
10	1.000

Table 2: Scaled TTT

The scaled TTT are, of course, the values that are plotted.

Plots corresponding to policies to be compared can be drawn on the same figure. As proved in the Appendix, when the plot corresponding to a particular policy P lies above the plot corresponding to another policy R, the first plot stochastically dominates the second. This means that for any income I_0:

$$\text{Prob}\ \{I_P > I_0\} > \text{Prob}\ \{I_R > I_0\} \qquad (6)$$

In other words, incomes under policy P are probabilistically larger than incomes under policy R. Policy P, therefore, is probabilistically better than policy Q. Policies can then be ranked with the aid of the TTT plots and probabilistically optimal policies can thus be identified.

Furthermore, it is proved in the Appendix that a policy can dominate another one only for a certain range of values. For example, the TTT plot of a given policy (S) could be above the one of another policy (T) only for the smallest values of the abscissa i, corresponding to low incomes. Choosing S over T therefore diminishes the risk of obtaining low incomes and, therefore, lowers the "Value at Risk" [3] of the income.

V. TOTAL TIME ON TEST PLOTS IN RISK ANALYSIS OF LONG TERM CONTRACTS

Total time on test plots similar to the one in figure 3 are obtained for the incomes associated with different policies of long term contracts (generating company incomes are now the random variables and play the same role as light bulb failure times in the previous section). Figure 3 was calculated using the November 1999 income figures of a hypothetical generator participating in the Colombian energy exchange. The relative order of incomes appears in ascending order on the abscissas. Thus, for example, I = 1 corresponds to the lowest income among the 37 hydro scenarios used, and I = 37 corresponds to the highest. For the sake of comparison, figure 3 includes the graphs corresponding to the following policies: "sell everything to the exchange," "sell 500 GWh on a long-term basis" ("500 LT"), and "sell 700 GWH on a long term basis" ("700 LT"). As was discussed in the previous section, the Appendix demonstrates that when the TTT graph for a given policy T is above the one for another policy R, the first dominates the second in probabilistic terms; in other words, for any income I :

$$\text{Prob } \{I_T > I\} > \text{Prob } \{I_R > I\} \qquad (7)$$

where I_T is the income obtained under policy T and I_R is the income that results from applying policy R. In this case,
of course, policy T is preferable to policy R. It was also discussed previously that it is possible for a policy to dominate another for a certain range of probabilities but not for a different one, as will be explained in the Appendix.

[3] The Value at Risk (VaR) summarizes in one number the total risk in the income of a given portfolio of financial assets. It makes statements similar to: "We are at least X percent sure that the income will be above I dollars in the next N days"

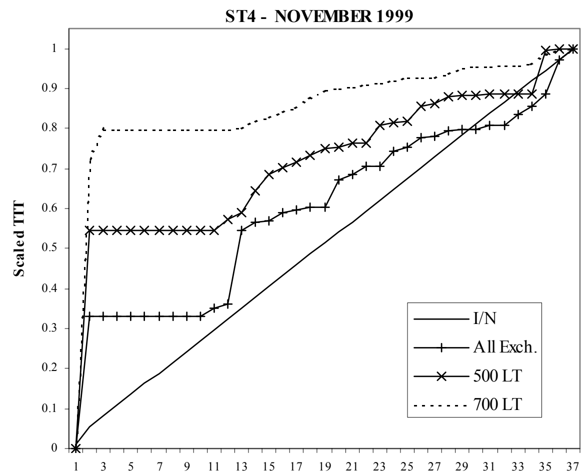

Fig. 3. ST4 for several November 1999 policies.

Thus, in the case of figure 3 it can be seen that the "500 LT" policy dominates the "everything to the exchange" policy. In addition, for all incomes except for the lowest, the "700 LT" policy dominates the "500 LT" policy and is therefore preferable.

The 37 values of incomes for any given policy can be ranked from smallest to largest. Notice that the lowest value has a probability of 36 / 37 (= 97%) of being exceeded, while the next-to-the-lowest value has a probability of 35 / 37 (= 95%) of being exceeded and can then be taken as the "firm" income of policy P, with a 95% probability. Firm incomes are thus obtained for all policies ranging from long term sales of 0 to 1,000 GWH and plotted in figure 4 in millions of U.S. dollars. The best policy in terms of maximizing firm income is "700 LT". However, figure 5 makes it possible to conclude that the "900 LT" policy dominates over "700 LT" in all cases except for the lowest incomes and might therefore be preferable, depending on the generating company's financial objectives.

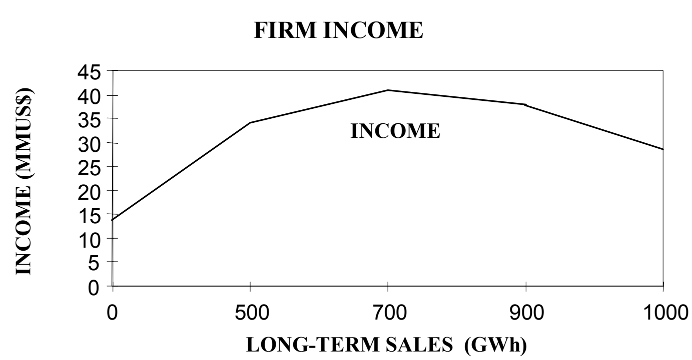

Fig. 4. Firm income in November 1999.

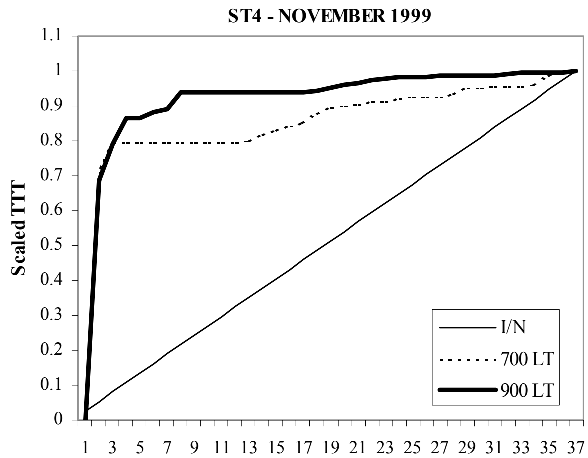

Fig. 5. ST4 for two November 1999 policies.

It can be appreciated that the methodology introduced here is more powerful than alternative methodologies, for example the ones that maximize minimum (or firm) income, since it enables complete probabilistic comparisons of different long-term sales policies and, consequently, provides better criteria for making decisions more in keeping with the financial objectives of the generating company.

VI. EXTENSIONS AND APPLICATIONS

A comparison of TTT graphs of marginal costs for February through November 1998 with the theoretical plots corresponding to lognormal distributions, suggests that a log-normal adjustment would not be unsuitable (figure 6). In fact, it is observed in [3] that theoretical TTT plots of lognormal distributions have at most two curvatures: concave-convex, or are totally concave, as the empirical plots of figure 7 appear to be. If this hypothesis were proven, results obtained for financial derivatives (call and put options, for example) could immediately be applied for the electric power market, on the basis of the popular Black-Scholes analysis [5]. Perhaps the supposed log-normality of marginal costs is due to the log-normality of hydrological inflows, which are the random variables underlying their calculation. It is worthwhile to emphasize that the hypothesis of the log-normality of marginal costs is only preliminary and must therefore be demonstrated with more exhaustive analyses than those herein, for example, using the Kolmogorov-Smirnov test.

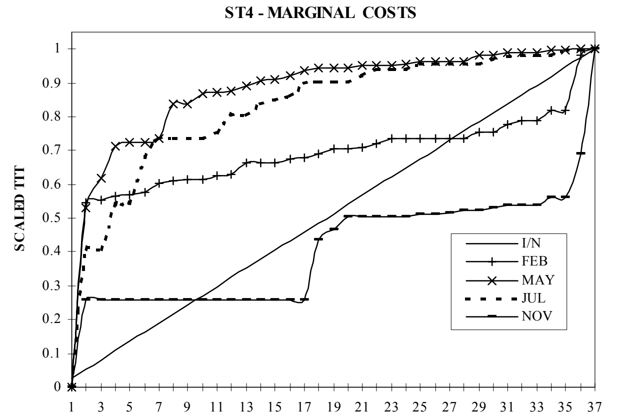

Fig 6. ST4 of marginal costs.

Furthermore, if the log-normality hypothesis of the incomes is accepted, better (higher) estimates of firm incomes can be calculated, as can be seen in figure 7, where higher firm incomes are obtained under the lognormality assumption (LN) than just using the raw data (ED, for Empirical Distribution). Notice that higher lower bounds are better bounds.

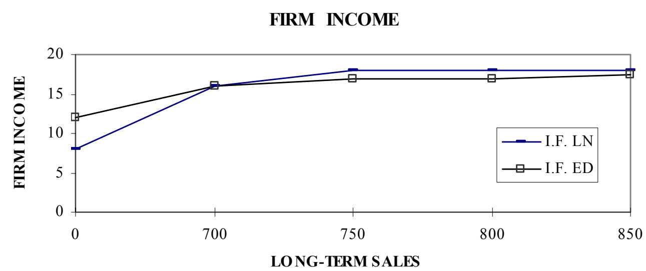

Fig. 7. Firm incomes under a lognormal hypothesis and based on the empirical distribution function.

VII. CONCLUSIONS

The methodology introduced in this paper was tested using actual data on the Colombian power sector, a generator, and several months in the 1998-2000 period. Through its use, it was possible to prove the following:

(a) Besides being pessimistic, the criterion of maximizing minimum income (or firm income) is sometimes very flat in long-term energy sales policies and, therefore, does not enable choosing among them.

(b) Sometimes policies that do not maximize minimum (or firm) income are preferable because, over the full range of probabilities, they are stochastically dominant. Dominant policies can be identified using TTT plots.

(c) By permitting a comparison for the full range of probabilities, the methodology introduced in this article provides more complete criteria for decision-making than alternative methodologies for risk handling, for example the ones that maximize minimum or firm income.

In addition to proving (or rejecting) the log-normality hypotheses introduced in Section 6, it is suggested that a study be done on the way that additional TTT properties may be used to decide on optimal long-term energy sales policies, accounting for the generating company's specific financial objectives.

Notice that even if the methodology described in this paper was developed from the point of view of generators (GENCOS), i.e., energy sellers, it is straightforward to extend it and apply it from the point of view of buyers like distribution companies (DISTCOS).

VIII. REFERENCES

[1] Millán, Campo and Sánchez. "A Modular System for Decision-Making Support in Generation Expansion Planning (SUPER)," IEEE Trans. Power Systems, vol. 13, no. 2, pp. 667–671.

[2] Gibbons, R. "Game Theory for Applied Economists." Princeton University Press. (1992).

[3] Barlow and Campo. "Total Time on Test Processes and Applications to Failure Data Analysis." Society of Industrial and Applied Mathematics (SIAM), Sept. 1975, pp. 451-81.

[4] Barlow and Proschan. "Statistical Theory of Reliability and Life Testing – Probability Models" To Begin With (1981).

[5] Black, F. and M. Scholes, "The Pricing of Options and Corporate Liabilities" Journal of Political Economy, 81 (May-June 1973), pp. 637-59.

IX. APPENDIX

Theorem A1, which is an important theorem for this paper, is demonstrated in [3]. It requires the following definition:

Suppose F and G are continuous distributions in $[0,\infty)$. F has a higher "star order" than G ($F *< G$) when $\{G^{-1} F(x) / x\}$ is increasing (non decreasing) in $x \geq 0$. In particular, if G is exponential, then F is IFRA. If F is exponential, G is DFRA. The concepts of IFRA (increasing failure rate average) and DFRA (decreasing failure rate average) are generalizations of the IFR and DFR concepts, respectively, and are defined in [3].

Let us assume that ST4 (F) and ST4 (G) are the Scaled Total Time on Test Transforms for data with probabilistic distributions F and G, respectively.

Theorem A1

If $F *< G$, $F(0) = G(0) = 0$, and F and G are continuous, then:

$$ST4\,(F) \overset{st}{\geq} ST4\,(G)$$

In other words, ST4 (F) stochastically dominates ST4 (G). Remember that a random variable Y stochastically dominates another random variable X when:

$$\Pr(Y > a) > \Pr(X > a) \text{ for every a.}$$

Note that if $F *< G$, F takes an S shape with respect to G; that is, F(x) crosses G(x) maximum once and from below, in the case of an intersection.

Theorem A1 makes it possible to compare statistical samples such as those provided in this paper.

One consequence of Theorem A1 is that, if the graph of ST4 for given observations is found above the graph of ST4 for other observations, the first set of observations stochastically dominates the second. In the demonstration of this theorem, it can be seen that one set of observations can have partial stochastic domination over another one; in other words, for a certain range of observations.

X. BIOGRAPHY

Rafael Campo is a consultant in issues related to energy markets, privatization, asset valuations, project evaluation and generation expansion studies. Dr. Campo is Electrical Engineer from the Universidad Nacional de Colombia and has an M.Sc. and a Ph.D. in Industrial Engineering and Operations Research from the University of California at Berkeley.