QUANTIFYING THE VARIABILITY OF SOLAR PV PRODUCTION FORECASTS

Steve R. Dean, ASA, P.E.
DAI Management Consultants, Inc.
1370 Washington Pike
Bridgeville, PA 15017
e-mail: sdean@daimc.com

ABSTRACT

The actual power produced by a solar photovoltaic power system varies according to the refraction, reflection, and absorption of radiation by the atmosphere. Standard production forecasts, however, do not address production uncertainty probabilistically. Using thirty years of historical data from the National Solar Radiation Database ("NSRDB"), we use stochastic simulation to evaluate the production uncertainty that is otherwise ignored by traditional production forecasts. In the case documented herein, we find significant differences between the forecasts made by conventional production forecasting models and those designed explicitly to reflect the uncertainty found in actual historical experience. Having more accurate information about production uncertainty should facilitate more project-appropriate financial structures, reducing risk and increasing investor returns.

1. INTRODUCTION

The growing popularity of photovoltaic ("PV") technology to power investors has led to a keen interest in understanding the risks associated with commercial operation of solar projects. Although the sun has a perception of stability that is attractive to investors, the electricity actually produced by solar power projects can vary widely. In the face of concerns over production variability, lenders often limit the amount of financing they are willing to provide to such projects, constraining returns and presenting challenges to the continuing development of PV projects.

Uncertainty, of course, is pervasive throughout the power complex. In response, risk measurement and management tools have been developed to inform project stakeholders about these uncertainties and assist in the creation of strategies for responding to them. In this paper, we seek to advance the state of risk measurement and management as it pertains to one particular aspect of PV project risk: production uncertainty.

1.1 Production Models and Historical Data

Typical PV energy production forecasting tools consist of models that take meteorological inputs and convert them first to radiation forecasts and then to electrical production. The National Renewable Energy Laboratory ("NREL") developed the Typical Meteorological Year ("TMY") database to serve as the underlying source of meteorological inputs for a given location. The *PVWatts* model developed by NREL is one such model, which uses the *TMY2* dataset, which was constructed using historical data from the NSRDB (1, 2). For purposes of this analysis, we relied on *PVWatts* and *TMY2*.

There is relatively little disagreement over the functionality of the solar production model itself. Although other such models exist (*e.g.*, PVsyst, Polysun), the mechanics underlying the measurement of global radiation are well-understood and common. As a result, the model itself is not a significant source of production uncertainty. The "uncertainty" that is often attributed to such models may be more accurately termed an issue of resolution. The nature of the calculations performed may result in conclusions that are only known to within a certain degree of numerical error or time partition.

In contrast, the data component of the forecasting tools contributes a vastly greater degree of uncertainty to the measurement of production. Although the *TMY2* dataset contains thirty years of historical data, its construction and aggregation for modeling purposes introduces uncertainty. In addition, the components of the dataset themselves exhibit uncertainty.

We are clear, then, to draw a distinction between model-derived uncertainty and data-derived uncertainty in our results here. We focus on data-derived uncertainty. This focus is deliberate and, we believe, appropriate in light of the empirical evidence.

1.2 Sources of Uncertainty

At the most general level, we characterize production uncertainty probabilistically. We assume that a random process determines production, and therefore our forecasts of future values take the form of a probability distribution. Production uncertainty, then, can be measured as the parameters of the probability distribution. Specifically: the location of the mean production level and the spread of potential values around that level.

The PV production model (*PVWatts*) provides an estimate of the mean. However, we seek to evaluate whether the estimated mean determined by *PVWatts* reflects any bias produced by the manner in which the *TMY2* data is aggregated and used. In addition, we seek to evaluate the spread of possible values around the mean production level determined by the inherent variability in the various meteorological inputs relied upon.

Because these objectives make significant use of the historical record contained within *TMY2*, let us briefly review exactly what *TMY2* is.

TMY2 is a dataset of hourly values of solar radiation and meteorological elements for a one-year period that defines a “typical” year for a particular location. The meteorological elements include such items as global horizontal radiation, direct normal radiation, dry bulb temperature, dew point temperature, and wind speed. The *TMY2* database contains this data for more than 230 locations across the country. Thirty years of data (1961 – 1990) were evaluated for each location.

For each month, an algorithm is used to select the “most typical” month of the thirty years in the database. The algorithm minimizes the difference between the year in question and the long-run average for each parameter. The parameters are then weighted according to their importance to the determination of location solar resource availability. Certain data is then excluded on the basis of persistence or unusual occurrence. For example, the volcanic eruptions of El Chicon in Mexico in 1982 and Mt. Pinatubo in the Philippines in 1991 materially altered the level of solar radiation observed. As a result, those years were considered atypical and excluded from the database (the *TMY3* version of the database spans 1976 – 2005).

Having identified the “most typical” month for each of the twelve months, they are concatenated to produce the typical year. The meteorological characteristics of this “typical year” are used to generate the forecasted production level for a given location.

2. UNDERSTANDING VARIABILITY

Our concern is that models such as *PVWatts* may have issues with both the estimated mean production level and the variance around that level. We state this question succinctly: is “typical” central?

The use of the “typical year” as the production forecast implies a degree of probabilistic meaning that does not actually exist. There is no statistical reason to believe that the “typical” production level is the mean, median, or modal production level, even though one may intuitively perceive the typical level to be the *expected* (in the statistical sense) production level.

We find this problematic.

PV project developers and investors tend to rely on these production forecasts in cash flow models to evaluate project economics. As with any cash flow model, users expect the model to reflect either the modal outcome (the most likely state of the world, as in a scenario analysis setting) or the mean outcome (the expected state of the world, as in a probabilistic setting). Instead, there is no formal reason to believe that a “typical year” forecast has any claim to central tendency at all.

Consider some basic concerns that reveal problems with the “typical year” model:

- The data selection process in *TMY2* is asymmetric with regard to outliers. Years in which solar output was adversely affected by unique circumstances (*e.g.*, volcanic activity) are excluded, but unusually “good” years are not excluded.
- Months are constructed by a weighting of parameters related to solar resource availability, rather than actual experience.
- Years are constructed in a manner that neglects the potential for sequential dependence. One especially “cold” month, for example, may more likely be followed by another “cold” month.
- The likelihood of an entire year of “typical” months may be overstated.

The overuse of typicality is the most problematic element of the *PVWatts* forecast. Because years are constructed as the concatenation of twelve typical months, no unusual periods can *ever* be incorporated into the forecast. A cash flow model based on such a "Panglossian" year must be grossly misleading to those considering investment. It is also patently unrealistic. Since it is a construction of months culled from different years, the "typical year" also reflects a year that has *never* actually occurred – a seemingly strange outcome for something deemed "typical."

Just as a year of "typical" months "one size fits all" approach may be inappropriate, the same principle applied to location and design may not be appropriate either. *PVWatts* proposes a blanket error level of "40% for individual months and up to 20% for individual years" with regard to production forecasts in light of "uncertainties associated with weather data and the model used to model the PV performance." (2) Whether or not 20% is the right annual error estimate, it is very unlikely that it remains constant across geographic location and across system design (*e.g.*, fixed or tracking). A more accurate (and useful) investment analysis would be based on results that reflected the unique circumstances of each particular installation. In addition, to generically bound the production forecast provides little insight into how annual production varies within the ±20% annual range. Is an error of 10% just as likely as one of 20%, or are most errors within ±5% and 20% errors represent improbable events?

It must be noted that the *TMY2* data is unquestionably useful, and our concerns here should not be interpreted as calling into doubt the importance of that data. We assert simply that the *TMY2*-driven *PVWatts* approach may be inappropriate from a cash flow modeling standpoint. We believe a probabilistic approach, as outlined in the remainder of this brief paper, would provide investors with a more accurate and informative perspective on the risk-reward tradeoff of a PV investment.

On this basis, we have several questions to answer: (*i*) is "typical" central? Does the typical year actually represent a year with statistical meaning? If not, can the historical data contained in the NSRDB be used to provide a more probabilistically-appropriate forecast? (*ii*) How does the shape of the probability distribution change based on location and technology? Is forecast confidence (*i.e.*, the "tightness" of the resulting probability distribution) higher (or lower) in some regions and for some technologies than others?

In the interest of space, we address only the first question in this short paper, although we note that the answer to the second question is 'yes, it does vary'.

3. METHODOLOGY

To be sure, other researchers have also recognized the uncertainty in PV production. *TMY2* users have been cautioned that the results may vary by as much as 20% on a year-to-year basis (2). Lohmann *et al*., provide a similar blanket estimate of 15% error over a one-year period (3). Others have noted empirically that *PVWatts* appears to overpredict production by approximately 10% (4). We also differentiate our work from studies investigating long-term structural shifts in PV production, such as those addressing the impact of global dimming and climate change-related factors (5, 6).

We focus exclusively on an examination of the underlying data, using a quantitative risk analysis ("QRA") approach, to explore the impact of parametric uncertainty on this dynamic system. We use a neural network to create a reduced form model of PV production as a function of certain key input parameters. Then, we use stochastic simulation to estimate the probability distribution of PV production based on distributions of input parameters estimated from actual historical data.

3.1 Quantitative Risk Analysis

QRA is a process for identifying, measuring, and evaluating the impact of uncertainty on dynamic physical or economic systems. Over the past century, it has been applied broadly, with applications ranging from the Manhattan Project to financial risk management to engineering design (see Bedford and Cooke (7) for a review of applications and the methodology). QRA was borne out of a recognition that large-scale, complex projects are subject to a variety of risks that may cut across disciplinary boundaries and combine elements of "natural" uncertainty (*i.e.*, the consequences of random processes) and as-yet-unknown future choices (*i.e.*, how a decision-maker in the future might act).

Taking such a holistic view was, in the past, often prohibitively expensive. Computing resources were limited and data were often unavailable in a format amenable to analysis. Thanks to advances in computational and data-retrieval tools over the past several decades, QRA is now used widely at the enterprise level for evaluating project risks.

QRA is typically conducted in four basic steps: (*i*) identify the sources of uncertainty, (*ii*) evaluate the probabilistic specification of the sources of uncertainty, (*iii*) apply simulation to the underlying model, and (*iv*) produce the output probability distribution. The objective is to enhance a basic model (a cash flow model, for example) by replacing single-point estimates of inputs with probability distributions reflecting the underlying uncertainties in those inputs. In addition, the uncertainty must be incorporated in a fashion mindful of relationships between the inputs (*e.g.*, if crude oil prices increase, gasoline prices are also likely to increase). In other words, the analysis must faithfully reproduce the covariance structure of the input variables.

The origins of QRA extend back more than a century, with interest in probabilistic analysis predating computers (8). Applications were limited, however, by the absence of computational support. It wasn't until computers made significant inroads into science and engineering that QRA began to be widely applied. Indeed, many of these early applications were found in the nuclear arena, with the very concept of Monte Carlo simulation advanced by Stanislaw Ulam (among many others) as part of the Manhattan Project. As computing power increased and costs fell, QRA began to be applied in other areas, spreading to the business world in the 1970s with the rise of management science as a discipline.

In fact, risk analysis has been used to evaluate investment and financing decisions for decades in such diverse areas as corporate planning (9, 10), leasing (11), petroleum investment (12), plant expansion (13), evaluating debt service risks in infrastructure investment (14), and electric power system planning (15, 16).

3.2 Reduced Form Models and Neural Networks

The underlying principle behind reduced-form modeling is that, for certain applications, extensive detail is unnecessary. It is this same principle, in fact, which results in such simplifications in traditional analysis as dropping higher-order terms in a Taylor-series expansion, quadratic approximations of highly-nonlinear functions, and local linearization of nonlinear systems. For most reduced-form modeling, the quest for simplification emerges as a result of the computational complexity of the underlying full-form model.

Areas such as climate modeling (17), population dynamics (18), environmental planning (19), and even credit risk modeling (20) have all made extensive use of reduced-form modeling techniques – in most cases to facilitate the use of simulation analysis.

In an abstract sense, we may view a full-form model as an arbitrarily complex function *f* that takes a set of inputs $\mathbf{X} = \{x_1, x_2, \cdots, x_n\}$ and returns an "exact" solution $f(\mathbf{X}) = S$. For any one of a number of reasons, *f* is difficult, time-consuming, or computationally costly to evaluate. The objective, then, is to identify a function $\tilde{f}$ that takes $\tilde{\mathbf{X}} \subseteq \mathbf{X}$ and returns $\tilde{f}(\tilde{\mathbf{X}}) = \tilde{S}$, such that $\tilde{S} \cong S$. The function *f* need not be an analytic function; it may instead be an algorithm, as it is with most PV production models. Nevertheless, it may be treated *as if* it were an arbitrarily high-dimensional analytic function.

This characterization is enormously useful, because it sets out a structure to our problem: minimize the computational cost of $\tilde{f}$ such that $\tilde{S} \cong S$ (or, as a constraint, $\left|\tilde{S} - S\right| < \varepsilon$). In words, find a function with performance equivalent to the full-form model, but that is easier to compute. Although this is a clear problem specification, the actual task of searching for such a function is complicated by the fact that we don't necessarily know in advance what form such a function might take. Theoretically, the space of functions is unbounded and these functions may be high-dimensional, nonlinear, and discontinuous.

Fortunately, a result of Kolmogorov (21), subsequently developed by Cybenko (22), by Funahashi (23), and by Hornik, Stinchcombe, and White (24), demonstrated the ability of (multilayer feed forward) artificial neural networks to serve as *universal function approximators*. Neural networks are connectionist models of neural activity, initially designed to replicate the manner in which the human brain processes information and engages in problem-solving activity. Subsequently, they emerged as a computational tool for solving complex problems, particularly in pattern recognition (a traditional strength of the human brain) (25).

3.3 Simulating Production with a Reduced Form Model

The standard PV production modeling process takes the *TMY2* data and uses it as an input to calculate power production. We create a neural network using six inputs (direct radiation, diffuse radiation, wind speed, dry bulb temperature, hour of the day, and month) and one hidden layer with seven nodes to estimate power production. The network is then trained on data from *TMY2*. The fully-trained network is then able to accurately replicate *PVWatts* within a single, albeit complex, function.

The trained neural network constitutes our reduced from model of *PVWatts*. The reduced form model is able to capture more than 99% of the variability in *PVWatts* while only using six inputs. Figure 1 illustrates the ability with which the reduced form model is able to replicate *PVWatts*.

With the reduced form model, we are then able to explore PV production with datasets other than *TMY2*. In addition, we are able to replace static input parameters with probability distributions to facilitate simulation of PV production.

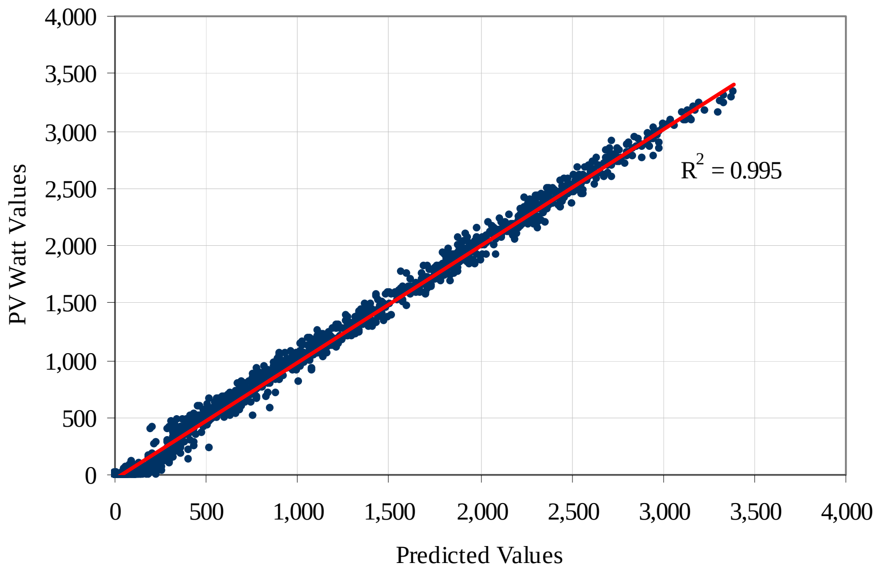

Fig. 1: Comparing the reduced form model to *PVWatts*.

Using these historically-derived input distributions, historically-derived covariance estimates, and our reduced form model, we then use Monte Carlo simulation to estimate the distribution of power production. Our analysis is based on a 4 kW reference system located in Newark, New Jersey, with a fixed tilt of 40.7° (latitude).

Our criticism of the conventional *PVWatts* approach is that the use of the *TMY2* data as the input source provides an unrealistic indication of actual experience. Instead, we rely on the historical data underlying *TMY2* (including any data previously deemed to be "atypical" or an outlier), but convert those input parameters to probability distributions. In addition, we estimate the covariance structure of these inputs. We make no effort to assess or seek "typicality." Rather, our modeling of the input data is designed to produce as historically accurate a representation as possible.

4. RESULTS AND CONCLUSIONS

Having developed a reduced form model of *PVWatts* and performed a simulation analysis using the actual historical data contained within the *TMY2* database, we return to our primary question: *is "typical" central*?

Not only does the *PVWatts* result not appear to be central, but it also appears to be a relatively extreme outlier. Figure 2 and Figure 3 illustrate the probability distribution of annual and January power production, respectively, along with the *PVWatts* forecasts for those periods. The month of January was selected since it experienced the largest divergence from the PVWatts estimate of all the months. In each case, the *PVWatts* forecast corresponds to the 99th percentile of the estimated distributions.

In each case, *PVWatts* overestimates production by approximately 6%. More importantly, however, developers

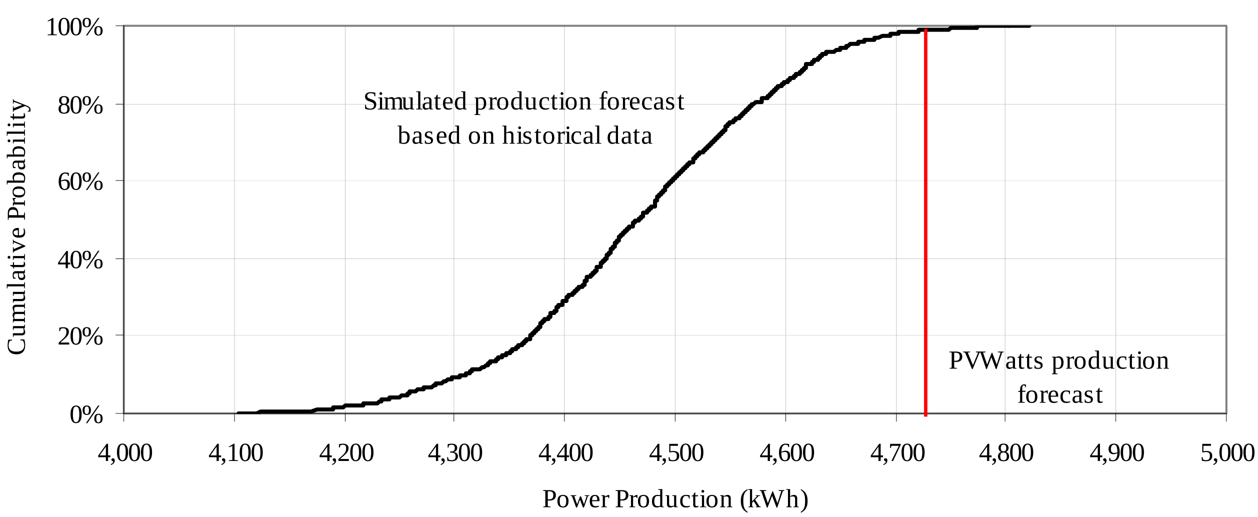

Fig. 2: Simulated annual power production compared to the *PVWatts* forecast.

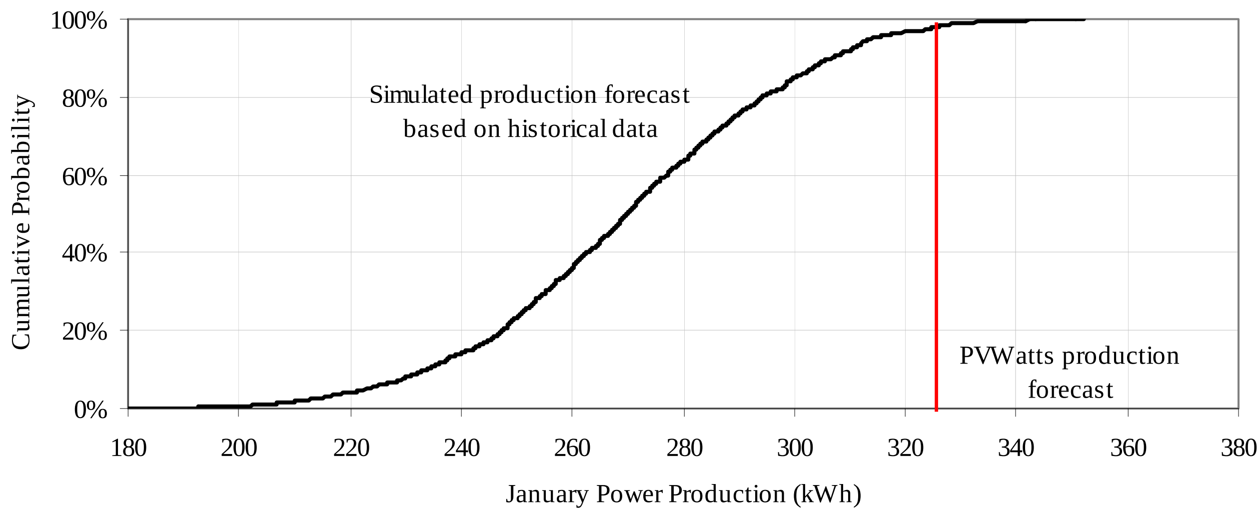

Fig. 3: Simulated January power production compared to the *PVWatts* forecast.

and investors can now obtain specific probabilistic production estimates (*e.g.*, the P95, P99 type estimates common among wind projects) and can therefore incorporate more empirically appropriate production forecasts into their cash flow models.

We consider one final step, however, to investigate the origins of this divergence between *PVWatts* and a probabilistically-inclined alternative. It may be asserted, for example, that our creation of the reduced form model or estimation of the input parameter distributions contributed to the discrepancy.

To examine such a hypothesis, we turn to the underlying *TMY2* data. *TMY2* does not contain power production. Rather, it contains measures of direct and diffuse radiation. We can examine, for example, thirty years of actual direct radiation from the NSRDB in comparison to the direct radiation of a "typical" year. One might *a priori* expect the typical year to fall in the middle of the collection of actual historical years (since we are assuming that a thirty-year period is sufficient to provide a representative example (26)).

Once again, however, we are starkly reminded that "typical" has no relevance to determinations of central tendency. We compare the historical and "typical" direct radiation data as cumulative distribution functions ("CDFs"). In other words, the percent of the year that direct radiation is below any given level of radiation (in Wh/m^2). Figure 4 compares the CDFs of the actual thirty years, the "typical" year, and the true average year.

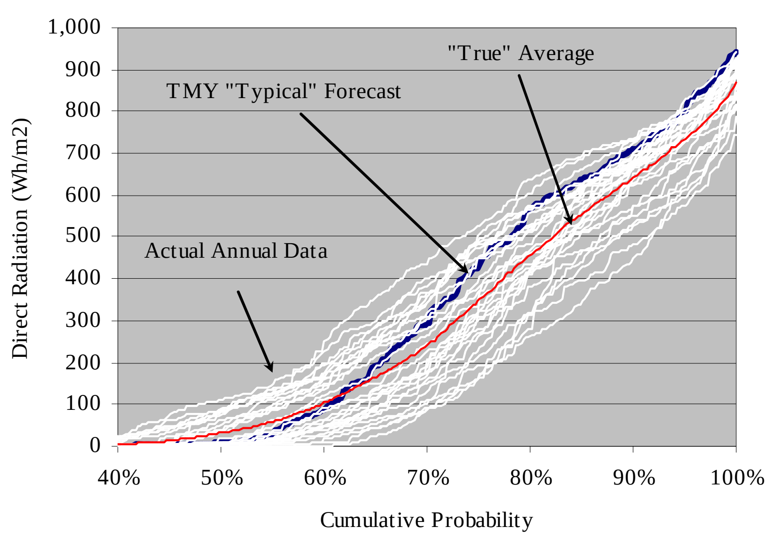

Fig. 4: The "typical" year compared to the actual and true average year.

The "typical" year is clearly *not* central to the actual historical data. Indeed, it significantly overstates direct radiation for nearly half of the year and the vast majority of the time during which solar radiation is present. Figure 5 illustrates this more directly by graphing the percent of the actual data exceeded by the typical year's direct radiation.

It appears, therefore, that the tendency of *PVWatts* to overpredict (in this instance) is a direct consequence of *TMY2*'s construction of the typical year. In contrast, our incorporation of empirically-derived input parameters provides a better estimate of what may be more properly described as an *average* year.

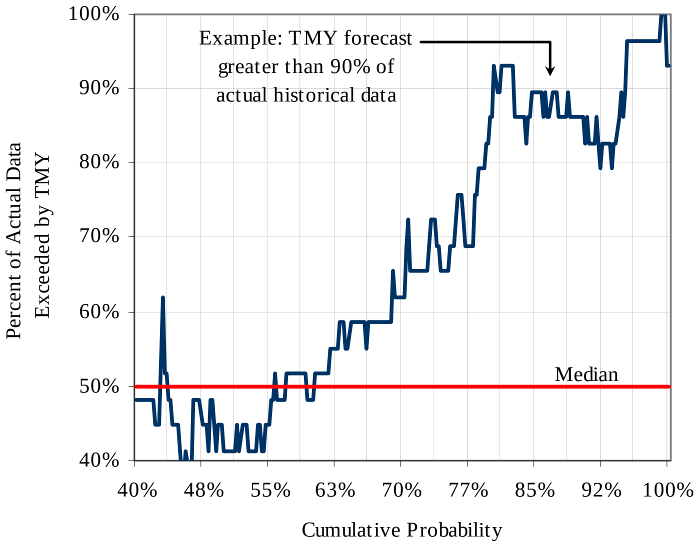

Fig. 5: Percent of actual historical data exceeded by "typical" year.

We assert that, for purposes of evaluating investment performance and for purposes of managing risk, use of an average year is preferable to use of a typical year. The "typical" year may still represent the level of production that might be expected in the event that everything "goes according to plan," but we do not believe that such optimistic assessments should be the objective of diligent investment evaluation (at least not for lenders).

Apart from concerns about accuracy, however, we note that the ability to have probabilistic information about expected project performance is in itself desirable. We write this at a time when attention to risk (and scrutiny by investors) has never been greater. The reduced form model QRA approach outlined above provides the ability for investors, lenders, and developers, to evaluate project performance and risk and to have that discussion on a common platform supported by direct reference to highly-relevant historical data.

5. REFERENCES

(1) *PVWatts*. National Renewable Energy Laboratory, www.pvwatts.org

(2) Marion, W., and K. Urban, User's Manual for TMY2's Typical Meteorological Years, Derived form the 1961-1990 National Solar Radiation Database, National Renewable Energy Laboratory, June 1995

(3) Lohmann, S., *et al*., Long-term Variability of Solar Direct and Global Radiation Derived from ISCCP Data and Comparison with Reanalysis Data, Solar Energy

(4) Pepper, J., Comments of Clean Power Markets, Inc. on CPUC Energy Division Staff Draft Proposal Phase I for the California Solar Initiative Design and Administration 2007-2016, Filed under Rulemaking 06-03-004, CPUC, May 16, 2006

(5) Pinker, R., B. Zhang, and E. Dutton, Do Satellites Detect Trends in Surface Solar Radiation? Science, May 6, 2005

(6) Schiermeier, Q., Cleaner Skies Leave Global Warming Forecasts Uncertain, Nature, May 12, 2005

(7) Bedford, T., and R. Cooke. Probabilistic Risk Analysis: Foundations and Methods, Cambridge University Press, 2001

(8) Covello, V., and J. Mumpower, Risk Analysis and Risk Management: An Historical Perspective, Risk Analysis, 1985

(9) Hertz, D., Risk Analysis in Capital Investment, Harvard Business Review, September – October 1979

(10) Coats, P., and D. Chesser, Coping with Business Risk through Probabilistic Financial Statements, Simulation, 1982

(11) Economos, A., A Financial Simulation for Risk Analysis of a Proposed Subsidiary, Management Science, 1968

(12) Newendorp, P., and P. Root, Risk Analysis in Drilling Investment Decisions, Journal of Petroleum Technology, July 1968

(13) Kryzanowski, L., P. Lustig, and B. Schwab, Monte Carlo Simulation and Capital Expenditure Decisions – A Case Study, Engineering Economist, 1972

(14) Jones, C., Financial Risk Analysis of Infrastructure Debt, Quorum, 1991

(15) Baughman, M., and D. Kamat, Assessment of the Effect of Uncertainty on the Adequacy of the Electric Utility Industry's Expansion Plans, 1983 – 1990, Interim Report EA-1446, EPRI, July 1980

(16) Jones, I., The Application of Risk Analysis to the Appraisal of Optional Investment in the Electricity Supply Industry, Applied Economics, 1986

(17) Mendelsohn, R., W. Nordhaus, and D. Shaw, The Impact of Global Warming on Agriculture: A Ricardian Analysis, American Economic Review, September 1994

(18) Lutz, W., S. Scherbov, A. Prskawetz, M. Dworak, and G. Feichtinger, Population, Natural Resources, and Food Security: Lessons from Comparing Full and Reduced-Form Models, Population and Development Review, Supplement, March 2002

(19) Webster, M., M. Babiker, M. Mayer, J. Reilly, J. Harnisch, R. Hyman, M. Sarofim, and C. Wang, Uncertainty in Emissions Projections for Climate Models, Atmospheric Environment, August 2002

(20) Duffie, D., and K. Singleton, Credit Risk: Pricing, Measurement, and Management, Princeton University Press, 2003

(21) Kolmogorov, A., On the Representation of Continuous Functions of Several Variables by Superposition of Continuous Functions of One Variable and Addition, Doklady Akademiia Nauk SSSR, 1957

(22) Cybenko, G., Approximation by Superposition of Sigmoidal Function, Mathematical Control, Systems, Signals, 1989

(23) Funahashi, K., On the Approximate Realisation of Continuous Mappings by Neural Networks, Neural Networks, 1989

(24) Hornik, K., M. Stinchcombe, and H. White, Multilayer Feed Forward Networks are Universal Approximators, Neural Networks, 1989

(25) Hopfield, J., Neural Networks and Physical Systems with Emergent Collective Computational Abilities, Proceedings of the National Academy of Sciences, 1982

(26) Vignola, F., and D. McDaniels, Value of Long-term Solar Radiation Data, Physics Department, University of Oregon