LBNL-1007269

Using Probability of Exceedance to Compare the Resource Risk of Renewable and Gas-Fired Generation

Authors:

Mark Bolinger

Energy Analysis and Environmental Impacts Division
Lawrence Berkeley National Laboratory

March 2017

This work was supported by the U.S. Department of Energy's Wind Energy Technologies Office and Solar Energy Technologies Office within the Office of Energy Efficiency and Renewable Energy under Lawrence Berkeley National Laboratory Contract No. DE-AC02-05CH11231.

Disclaimer

Acknowledgements

The work described in this study was funded by the U.S. Department of Energy's Wind Energy Technologies Office and Solar Energy Technologies Office within the Office of Energy Efficiency and Renewable Energy under Lawrence Berkeley National Laboratory Contract No. DE-AC02-05CH11231.

The author would like to thank James Preciado and Thomas Lee (U.S. Energy Information Administration), Morgan Bazilian (The World Bank), Stephen R. Jones and Alex Byrne (DNV GL), Galen Barbose and Ryan Wiser (Lawrence Berkeley National Laboratory), Jeffrey Logan and Thomas Jenkin (National Renewable Energy Laboratory), David Loomis (Illinois State University), and Matt McCabe (Clear Wind) for their review of and comments on this report. Of course, any remaining errors or omissions are my own.

Table of Contents

Table of Figures

List of Tables

Executive Summary

Of the myriad risks surrounding long-term investments in power plants, resource risk is one of the most difficult to mitigate, and is also perhaps the risk that most-clearly distinguishes renewable generation from natural gas-fired generation. For renewable generators like wind and solar projects, resource risk manifests as a *quantity* risk—i.e., the risk that the quantity of wind and insolation will be less than expected.[i] For gas-fired generators (i.e., a combined-cycle gas turbine or "CCGT"), resource risk manifests primarily as a *price* risk—i.e., the risk that natural gas will cost more than expected. Most often, resource risk—and natural gas price risk in particular—falls disproportionately on utility ratepayers, who are typically not well-equipped to manage this risk. As such, it is incumbent upon utilities, regulators, and policymakers to ensure that resource risk is taken into consideration when making or approving resource decisions, or enacting policies that influence the development of the electricity sector more broadly.

This paper presents a new framework, grounded in statistical concepts related to probability of exceedance (and confidence intervals more broadly), to incorporate resource risk into decision-making processes. This framework recognizes that the same probability of exceedance concepts that are regularly used to characterize the uncertainty around annual energy production for wind and solar projects can also be applied to natural gas price projections, allowing one to develop a probabilistic range of projections for not only wind and solar capacity factors, but also natural gas prices.

Importantly, these probability distributions have markedly divergent characteristics. Renewable resource risk is symmetrical about the mean or "P50" projection and declines when considered over longer time horizons (due to mean reversion in the inter-annual variability of the resource). In contrast, natural gas price risk is asymmetrical (skewed towards higher prices) and increases when considered over longer time horizons (reflecting the fact that it is easier to project where natural gas prices will be three months from now than three years from now). Converting these distinctly different probability distributions into directly comparable levelized cost of energy ("LCOE") terms reveals that even when gas-fired generation is competitive with, or cheaper than, wind and solar power on an expected or P50 basis—the basis on which these resources are most often compared—comparisons that are instead based on worse-than-expected outcomes (e.g., P25 or P1) often reach the opposite conclusion: that wind and solar are cheaper than gas-fired generation.

Figure ES-1 illustrates this concept by comparing the 25-year LCOE of a new wind project in the United States (without the benefit of the production tax credit or "PTC") to that of a new CCGT across P-levels ranging from P50-P1 and over time horizons ranging from one to 25 years. The range of time horizons along the x-axis warrants additional explanation to avoid confusion. Every data point shown on Figure ES-1—regardless of where it falls along the x-axis—represents an LCOE that is calculated over a 25-year period (in nominal dollars). These 25-year LCOEs are based on modeling inputs that are held constant

[i] Over shorter time periods there is also a temporal aspect to wind and solar resource risk—e.g., whether the wind will be blowing (or the sun shining) at times of high system demand and prices. This report, however, focuses on longer time frames, measured in years rather than in minutes, hours, days, or months.

in all cases, with two exceptions—the wind project's capacity factor and the CCGT's levelized fuel costs vary by P-level *and* by time horizon. The x-axis simply represents the time horizon (in number of years) over which these two important, but uncertain, inputs into the 25-year LCOE calculation are considered.[ii] For example, at year 12 on the x-axis, wind's 25-year LCOE range reflects 12-year P50 and 12-year P1 capacity factors used as inputs to the 25-year LCOE calculation; similarly, the range of gas-fired LCOE reflects 12-year P50 and 12-year P1 gas price projections that are levelized over 12 years and then used as the fuel price inputs in the 25-year LCOE calculation.

In Figure ES-1, wind (without the PTC[iii]) is more expensive than gas-fired generation on a P50 basis over all time horizons of less than 24 years (the two P50 curves converge at 24 years). But on a P25 basis, the cost of wind falls below the cost of gas-fired generation for all time horizons longer than 15 years. This "break-even" point—where the wind and gas-fired LCOE curves for each P-level cross—drops to 10, 8, and 2 years for P10, P5, and P1 levels, respectively.

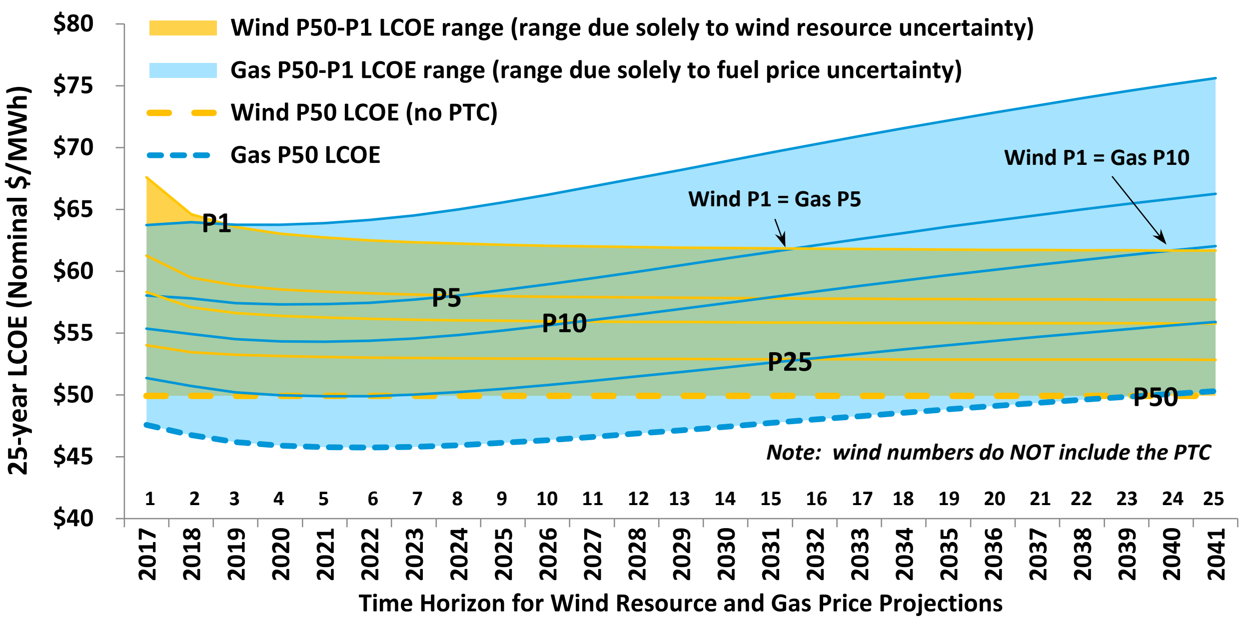

Figure ES-1. P50-P1 wind (no PTC) and gas-fired LCOE ranges over different time horizons

---

[ii] Although investment decisions about new generation sources are typically made with long-term time horizons (e.g., 20-25 years) in mind, the shorter time horizons that are also covered by this new framework may nevertheless be relevant in certain situations. For example, a utility might have a temporary or short-term (e.g., 5-year) need for energy, in which case the price stability offered by wind or solar likely won't be as important of a consideration as it is over longer time horizons. Different types of investors in renewable energy projects also tend to focus on different time horizons (as well as different P-levels); for example, project sponsors may be most concerned about 25-year time horizons (presuming they intend to be long-term owners), while tax equity investors in wind projects may be most interested in a 10-year time horizon, given the 10-year PTC window (after which tax equity investors will often exit a project), and lenders will be most interested in even shorter time horizons when assessing the risk of default. Finally, ratepayers may also have a relatively short-term focus if, for example, they are not certain how long they will reside (or remain in business) within the utility's service territory. Varying both the P-levels *and* time horizons for the two uncertain inputs allows a user who is concerned about resource risk to make an informed decision in all cases, regardless of the time horizon.

[iii] The full report provides additional comparisons that include the PTC, and that compare wind's LCOE (with the PTC) to just the operating costs (as opposed to LCOE) of gas-fired generation.

In other words, Figure ES-1 presents an illustrative example where wind, without the PTC, is not cost-competitive with new gas-fired generation (except over a 24-year or longer time horizon) when evaluated on a P50 basis as is typically done. But when considering the possibility of worse-than-P50 outcomes (i.e., higher than-expected natural gas prices and/or a lower-than-expected wind resource), wind looks more competitive—particularly the lower the P-level and the longer the time horizon—and in many cases is cheaper than gas-fired generation.[iv] The "wedges" that begin where the respective wind and gas-fired LCOE curves at each P-level cross and then widen over longer time horizons illustrate wind's "hedge value," which increases with both the level of risk aversion (assumed to be negatively correlated with the P-level—i.e., a lower P-level suggests greater risk aversion) and the time horizon.

Figure ES-2 shows much the same story for a utility-scale solar photovoltaic project. In this example, solar (with the 30% investment tax credit or "ITC") is always more expensive than gas-fired generation on a P50 basis, regardless of time horizon shown. But, as with wind, worse-than-P50 comparisons reveal solar to be more competitive: the solar and gas-fired P25 LCOE curves converge at a 25-year time horizon, while the P10, P5, and P1 curves show solar's hedge value starting to accrue at progressively shorter time horizons.

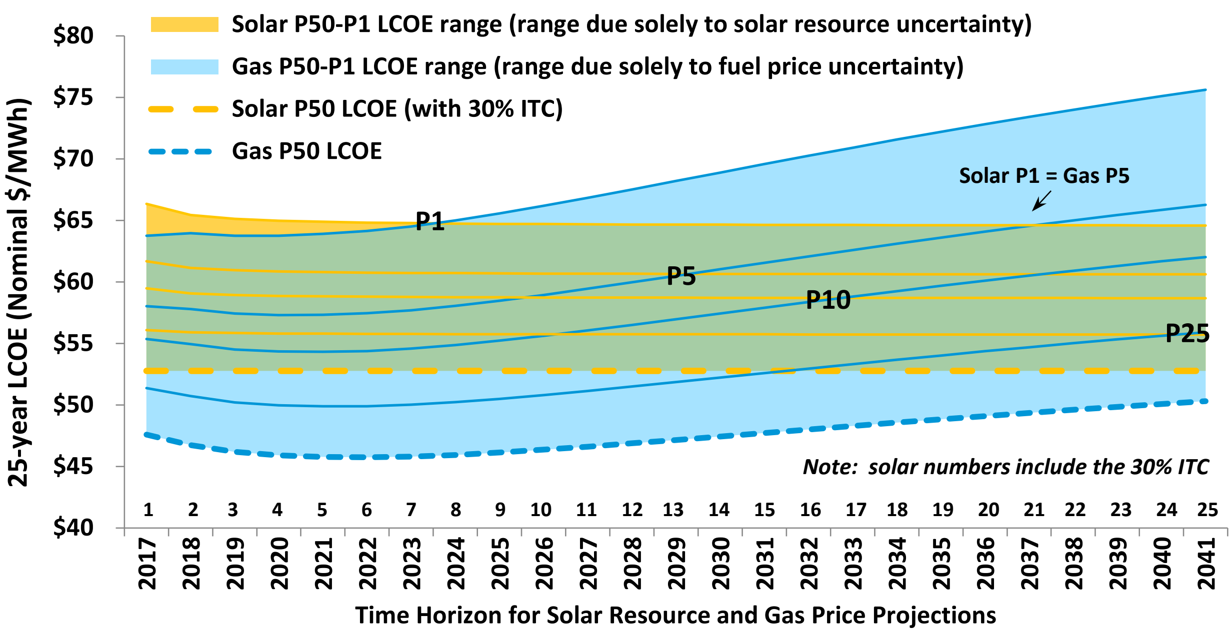

Figure ES-2. P50-P1 solar (with 30% ITC) and gas-fired LCOE ranges over different time horizons

Another related way to interpret Figures ES-1 and ES-2 is that higher-than-expected gas prices are riskier than lower-than-expected wind or solar output. This suggests that from a ratepayer perspective, we should perhaps be more concerned about gas price risk than about wind or solar resource risk. In other words, in a case where two scenarios—one focusing on higher-than-expected gas prices and another focusing on lower-than-expected wind or solar resources—may be considered to have the

[iv] Also notable in Figure ES-1 is that lower-probability wind resource projections still, in some cases, yield lower LCOEs than higher-probability gas price projections. For example, over time horizons of 24 years or longer, P1 wind is competitive with P10 gas, and P10 wind is competitive with P25 gas.

same probability (i.e., the same P-level), the resulting impact of the high gas price scenario may be more harmful to ratepayers than the impact of the low wind/solar resource scenario.

Although the discussion surrounding Figures ES-1 and ES-2 has so far focused on LCOE comparisons at distinct P-levels, by definition, each P-value has an associated probability, thereby enabling a more formal probabilistic assessment. For example, although probability of *exceedance* does not necessarily imply probability of *occurrence*, the P50 outcome can nevertheless be thought of as carrying a 50% weight, while the P1 outcome can be given a 1% weight, with all other P-values that fall in between these two extremes (e.g., P49, P48, P47…P4, P3, P2) weighted accordingly (i.e., 49%, 48%, 47%...4%, 3%, 2%). Hence, within this framework, one can easily "probability-weight" the full range of outcomes across the full P50-P1 spectrum, or even some subset thereof—e.g., perhaps just the P50-P25 range for those who are less risk averse.[v]

The probabilistic nature of this new framework is one of its key advantages over previously proposed approaches to account for the price stability benefit of wind and solar power. Other advantages include its fairness (recognizing that wind and solar also face resource risk), familiarity (probability of exceedance is already widely used within the energy industry), simplicity (just a few key inputs are needed to set up these comparisons), and flexibility (this framework caters to any level of risk aversion over any time horizon).

Of course, cost is only one side of the equation (value being the other), and few if any resource decisions within the electricity sector are made based on LCOE alone. Instead, the cost of competing resources must be considered along with the value that each provides, which is most often determined by sophisticated models that endogenously assess energy and capacity value as well as integration and transmission costs—all in addition to the LCOE of the generator itself. In this sense, it should be recognized from the start that this report is focused on just one side of a two-sided coin.

[v] These "probability-weighted" outcomes, along with the unweighted comparisons shown in Figures ES-1 and ES-2, focus exclusively on P50-or-worse scenarios, and ignore the very real possibility of better-than-expected outcomes. This focus on the negative (at the exclusion of the positive) is driven by several considerations. First and foremost, risk is most commonly defined as something that leads to a negative (rather than a positive) outcome, and that is how it is most often thought of within the electricity sector as well. Second, for renewable generators, a better-than-expected wind or solar resource—particularly out at the very tail end of the positive side of the distribution—may not translate to correspondingly higher amounts of wind or solar generation, due to technological limitations (e.g., maxing out the wind turbine power curve or solar inverter capacity). This asymmetry at the very tails of the renewable generation distribution complicates an assessment of outcomes across the full P1-P99 range, but does not similarly affect assessments of only worse-than-expected (i.e., P50-P1) outcomes (i.e., a worse-than-expected wind or solar resource will always lead to lower-than-expected wind or solar generation—there is no leveling off like there is at the positive tail). Finally, with natural gas prices recently at historically low levels, and unlikely to fall below zero, both the likelihood and impact of positive outcomes (i.e., lower-than-expected gas prices) are perhaps less than the likelihood and impact of negative outcomes (i.e., higher-than-expected gas prices).

1. Introduction

The fixed-price nature of renewable generation sources like utility-scale wind and solar has long been contrasted with the volatile and unpredictable price of fossil-fueled, and in particular gas-fired, generation. Because wind and solar power have no fuel costs, they can be (and most often are) sold at fixed prices through long-term (e.g., 20-year) contracts. Gas-fired generation, in contrast, is often sold on the spot market or through short-term contracts that are linked to the underlying price of natural gas, or else through so-called "tolling contracts" whereby fuel procurement is the power purchaser's (rather than the generator's) responsibility. In all three of these cases, the risk of natural gas prices—and hence also the cost of gas-fired generation—exceeding expectations is largely or exclusively borne by the power purchaser. Offtakers can partially mitigate this "resource risk" by adding fixed-price wind and solar power to their resource portfolios.[1]

The value of long-term price certainty provided by renewable generation has been studied extensively over the past quarter century, using a variety of methods (which, for interested readers, are summarized in the text box on the next page). Historically, this literature has been motivated by a desire among the renewable energy industry to justify policy support for the deployment of what has, in the past, been an above-market source of generation. More recently, however, the cost of wind and solar power has fallen dramatically (Wiser and Bolinger 2016, Bolinger and Seel 2016), at the same time as advances in horizontal drilling and hydraulic fracturing have opened up vast new natural gas reserves in the United States and driven natural gas prices to all-time lows. This convergence of trends—i.e., low-cost wind and solar projects that can largely compete with gas-fired generation, even as natural gas prices mark historical lows—has given rise to a new audience interested in price stability: utility and corporate offtakers, who see this as an ideal time to lock in cost-competitive prices via long-term fixed-price renewable power purchase agreements ("PPAs"). The text box on page 4 illustrates, in their own words, that price certainty has become an important driver of demand for renewable generation among utilities and corporate offtakers alike.

Wind and solar power, however, are not immune to resource risk. Whereas resource risk is primarily a *price risk* for gas-fired generation, it manifests instead as a *quantity risk* for renewable generation—i.e., the risk that the wind or solar resource will not be as strong as expected over the life of the project (Matsui et. al 2016). Along with their other limitations (as briefly described in the text box on the next page), previous approaches to measuring the price stability benefit of renewable generation have largely overlooked this uncertainty over the renewable resource.

This paper develops a probabilistic framework for assessing the relative resource risk of renewable and gas-fired generation. Grounded in statistics, this new framework makes use of confidence intervals or, more specifically, "probability of exceedance" levels to compare the levelized cost of renewable and

---

[1] Other non-gas resources with relatively stable prices such as nuclear and coal-fired generation, as well as energy efficiency, can play a similar mitigating role. In addition, gas price risk can be hedged, at least in the short term, using conventional hedging instruments.

gas-fired generation across a range of statistical outcomes attributable to resource risk, with a particular focus on worse-than-expected outcomes. Modeling a representative wind, solar, and gas-fired combined-cycle plant, this framework demonstrates that while gas-fired generation is competitive with or cheaper than wind and solar generation under modeled conditions on an expected or "P50" basis, worse-than-expected scenarios often yield the opposite conclusion—i.e., that wind and solar power are cheaper than gas-fired generation. The probabilistic nature of this framework allows for not only useful visualizations, but also numerical quantification, of the benefits of price certainty.

A brief review of prior approaches to account for the price stability benefit of renewable energy

Awerbuch (1994, 1993) proposed the use of risk-adjusted discount rates to inflate less-predictable expenses—e.g., natural gas fuel costs—within discounted cash flow models. This practice requires use of the Capital Asset Pricing Model ("CAPM")—borrowed from the financial sector—to adjust the risk-free discount rate based on the "beta," or correlation, of the expense in question with other costs. CAPM is also integral to mean-variance portfolio theory, which was subsequently adopted and used to demonstrate that adding even high-cost (but fixed-price) renewable generation to a portfolio of riskier resources can move the overall portfolio closer to the "efficient frontier" (Awerbuch 2003, Awerbuch and Berger 2003; Bazilian and Roques 2008). Stirling (1994, 2010) took an altogether different approach, grounded in the natural sciences rather than financial analysis, and focused on measuring the level of diversity within a generation portfolio (under the assumption that diversity is inherently desirable). Other attempts to quantify the price stability benefit of renewable generation include measuring the "avoided cost" of hedging fuel price risk through more conventional means, like futures contracts (Bolinger et al. 2006), and comparing the fixed prices within renewable power purchase agreements to uncertain natural gas price forecasts (Bolinger 2013). Finally, more-familiar tools and concepts like Monte Carlo simulation, scenario analysis, sensitivity analysis, decision analysis, and certainty equivalence have all been used, to varying degrees, to study the fixed-price attribute of renewable generation (Jenkin et al. 2013, Bolinger and Wiser 2005, Brower et al. 1997).

Though often innovative and informative, none of these approaches have been entirely satisfactory on their own. For example, risk-adjusted discount rates and mean-variance portfolio theory both only make sense for this purpose if natural gas has a negative beta, which can be problematic to measure—e.g., is the correlate the stock market or the broader economy?—and can vary over time. Diversity indices can measure the diversity of a portfolio, but not the value of diversity. Some academics claim that hedging is costless, calling into question an avoided cost approach; conventional gas hedges are also typically much shorter in duration than a renewable power purchase agreement, leading to a temporal mismatch that complicates a fair comparison. Comparing renewable PPA prices with gas price projections reveals more about expected long-term savings than the value of price certainty itself. To be effective, scenario analysis requires modeling well-developed (almost prescient) scenarios that account for critical inter-dependencies among variables, that reflect the full distribution of possible outcomes, and that are not prematurely weeded out in a multi-step modeling process (Bolinger and Wiser 2005). Decision analysis and certainty equivalence require assumptions about utility functions and risk-aversion coefficients when, in fact, consumers' level of aversion to higher-than-expected electricity bills is not well-understood (Costello 2016). Finally, few (if any) of these approaches explicitly acknowledge or account for the fact that renewable generation also faces resource risk (in the form of a quantity, rather than price, risk).

In other words, despite the wide variety of approaches that have been employed over the past quarter century, no consensus has emerged on how best to measure or quantify the value of price certainty that renewable generation can provide. Far from remedying this situation, this paper contributes yet another approach for consideration, in the form of a new statistical framework to analyze the relative resource risk of renewable and gas-fired generation.

Of course, cost is only one side of the equation (value being the other), and few if any resource decisions within the electricity sector are made based on the levelized cost of energy (“LCOE”) alone. Instead, the cost of competing resources must be considered along with the value that each provides, which is most often determined by sophisticated models that endogenously assess energy and capacity value as well as integration and transmission costs—all in addition to the LCOE of the generator itself. In this sense, it should be recognized from the start that this report is focused on just one side of a two-sided coin.

This paper proceeds as follows. Chapter 2 begins by discussing why, out of the many different types of risk facing both renewable and gas-fired generation, resource risk is perhaps the most important and worthy of study. Having established the rationale for an acute focus on resource risk, Chapter 2 then moves on to the important question of who bears resource risk under a variety of procurement options. The answer—overwhelmingly ratepayers—establishes ratepayers and those who look after their interests, such as regulators and ratepayer advocates, as the primary audience for this work, along with utility resource planners and procurement teams. Chapter 3 introduces the new framework by first reviewing how probability of exceedance is already widely used within the wind and solar power industries to characterize the resource, and then developing a related extension of this methodology for natural gas prices and gas-fired generation. Chapter 4 then applies the new framework to compare the LCOE of wind and solar generation with that of gas-fired generation from a combined cycle gas turbine (“CCGT”) across a range of statistical outcomes attributable to resource risk. Chapter 5 summarizes and concludes.

Wind and solar's inherent price stability clearly motivates buyers

Two surveys conducted in 2016 found price stability to be one of the top three drivers (along with sustainability goals and attractive return on investment) of the recent surge in corporate purchases of wind and solar power (Mortenson Construction 2016, PwC 2016). The quotes below—which are attributed to representatives of investor-owned utilities, publicly owned utilities, and corporations who are buying wind and solar power in part as long-term insurance against the possibility of rising natural gas prices—provide further supporting evidence of the importance of this demand driver.

Utility offtakers:

- "This solar energy center adds diversity to WPPI Energy's power supply portfolio in a way that's more cost-effective than other opportunities currently available to us." – *WPPI Energy, 2017*
- "When we're buying wind at $25, it's a hedge against natural gas." – *Xcel Energy, 2015*
- "We like wind because it's a hedge against fossil prices…and wind, with no fuel costs associated, can keep those rates stable." – *MidAmerican Energy, 2015*
- "The latest addition of 150 megawatts of low-cost wind energy provides AECC with a hedge against fluctuating natural gas energy prices." – *Arkansas Electric Cooperative Corp, 2013*
- "We think of this wind contract as an alternative fuel, with known contract pricing over 25 years that will displace fuels where the pricing is not yet known. That is the essence of the fuel hedge" – *PSCo, 2012*
- "[Wind energy power purchase agreements] decrease our exposure to natural gas, provide a hedge against any future global warming legislation, and help us give our customers lower, more stable prices." – *Empire District Electric Company, 2008*
- "Wind generation provides value simply for the insurance it furnishes in insulating customers from some of the aspects of unexpectedly high and volatile fuel and wholesale energy prices." – *Westar Energy, 2007*

Corporate offtakers:

- "Investing in large-scale renewable power…helps Lockheed Martin hedge against the volatility of the electricity market and lower our energy costs…This is a nice addition to our current hedging strategy…This gives us the ability to hedge out in a different way, for a much longer term." – *Lockheed Martin, 2016*
- "Electricity costs are one of the largest components of our operating expenses at our data centers, and having a long-term stable cost of renewable power provides protection against price swings in energy." – *Google, 2016*
- "Cost savings are the main driver, but price stability is a close second." – *General Motors, 2013*
- "We see value in getting a long-term embedded hedge. We want to lock in the current electricity price for 20 years. We are making capital investment decisions on the order of 15 to 20 years. We would like to lock in our costs over the same period." – *Google, 2011*

2. Resource Risk: Why focus on it and who bears it?

Whether renewable or fossil-fueled, utility-scale power plants are long-term, capital-intensive investments. Not surprisingly, those who build, finance, own, and operate power plants face a myriad of risks along the way. At a high level, these risks—as described here, largely post-development risks[2]—can be grouped into seven over-arching, and in some cases inter-related, categories:

- **Financing risk:** The risk that financing (or refinancing) terms will not be as favorable as expected (or that financing may not be available at all).
- **Capital Expense ("CapEx") risk:** The risk that it will cost more than expected to build the plant.
- **Resource risk:** The risk that the underlying energy resource (e.g., the wind speed or insolation for a wind or PV project, or natural gas for a combined-cycle gas plant) will not be as plentiful as expected, or will cost more than expected.
- **Technology risk:** The risk that the prime mover technology (i.e., wind turbine, PV module and inverter, or gas turbine) will not convert the underlying energy resource into electricity as efficiently as expected (or at all, in the case of technology failure).
- **Operating Expense ("OpEx") risk:** The risk that it will cost more than expected to operate and maintain the plant (not including the cost of fuel, which is captured under resource risk).
- **Environmental or regulatory risk:** The risk that new or stricter environmental or regulatory requirements imposed after a project is built will increase costs or possibly even challenge the right to operate.
- **Offtake risk:** The risk that the plant will have to sell its electricity at a lower-than-expected price (or that it may not be able to sell its electricity at all).

Despite these many different risks that deserve careful consideration, this paper singles out just one of the seven categories listed above—resource risk—for further investigation.

2.1 Why focus on just resource risk?

This report's narrow focus on resource risk is driven by two main considerations. First, among the seven categories of risk listed above, resource risk is one of just two categories (environmental or regulatory risk being the other) that clearly impacts renewable and fossil (e.g., gas-fired) generation differently. The other five risk categories—CapEx, financing, technology, OpEx, and offtake risk—all affect renewable and gas-fired generation in similar ways (though not necessarily to similar degrees). But resource risk (along with environmental/regulatory risk) is inherently different for renewable and gas-fired generators. For renewable generators like wind and solar projects, resource risk manifests as a *quantity* risk—i.e., the risk that the quantity of wind and sunshine will be less than expected. For gas-fired generators, resource risk manifests primarily as a *price* risk—i.e., the risk that natural gas will cost

[2] Power plant development is an inherently risky endeavor that faces a different array of risks that fall outside of the scope of this report.

more than expected.[3]

The second reason to focus on resource risk is that it is perhaps the most-problematic—for *both* renewable and gas-fired generators—of the seven risk categories listed above, in the sense that it is difficult to mitigate and will persistently impact a project over its entire life (see the text box on page 8). In contrast, the other six risk categories can either be largely eliminated or resolved prior to or at commercial operations (e.g., financing and CapEx risk), can be mitigated contractually (e.g., technology, OpEx, and offtake risk), and/or are not really even all that applicable to renewable generation (e.g., environmental or regulatory risk). For example:

- **Financing risk:** The risk of initial financing either not being available or having worse-than-expected terms will be recognized and resolved (one way or another—even if that means killing the project) early in a project's life, most likely even before the start of construction. As such, financing risk does not create long-term uncertainty that will remain with a project for years to come. Refinancing does present a longer-term risk, though is likely more applicable to gas plants than to wind and solar projects, given the greater tendency of gas plants to utilize (and repeatedly roll over) shorter-term, partially-amortizing debt as a principal component of the capital stack (Feldman and Bolinger 2016).
- **CapEx risk:** For both renewable and gas-fired generators, the risk of cost overruns while building a project is typically handled contractually through an engineering, procurement, and construction ("EPC") contract that fixes or caps construction costs.[4] Moreover, as with financing risk, any CapEx-related problems will be recognized and resolved early in the project's life—most likely even prior to the start of commercial operation—and therefore do not represent long-term risks or unknowns that could plague a project for years to come.
- **Technology risk:** Technology risk is typically mitigated contractually through a combination of general equipment warranties and performance guarantees from the manufacturer (perhaps backed by insurance policies), as well as a long-term maintenance plan, perhaps involving a service contract with the equipment manufacturer or a third party. For example, wind turbine supply agreements often include a 3- to 5-year general equipment warranty, solar modules typically have a 10- to 20-year physical equipment warranty, and solar inverter warranties range from 10 to 25 years. With respect to conversion efficiency, wind turbines often come with availability guarantees and power curve warranties,[5] while solar module manufacturers will typically guarantee an initial level of performance that will not degrade by more than a certain

---

[3] Though quantity or supply risk has also been a concern (in addition to price risk) for natural gas at times in the past, advances in horizontal drilling and hydraulic fracturing techniques in recent years have largely mitigated domestic natural gas supply concerns for the foreseeable future. That said, pipeline transportation constraints may still restrict the supply of natural gas in some locations; in these cases, the quantity risk is most often reflected via higher local prices, and so can be considered a form of price risk.

[4] A fixed-cost EPC contract may, of course, cost more than one that does not provide a fixed-cost guarantee (as the EPC contractor prices the risk into the contract), but that is somewhat beside the point, which is simply that this risk can be contractually mitigated—even if at a cost.

[5] Power curve warranty claims, however, can be difficult to substantiate. In contrast, availability guarantees have become more sophisticated in recent years, evolving from simple time-based availability metrics (e.g., the wind turbine is guaranteed to be available to generate electricity a certain percentage of the time) to more-complicated production-based availability metrics (e.g., the wind turbine is guaranteed to be available to generate electricity during a certain percentage of the windiest times). Revenue-based availability metrics are perhaps the next logical step.

amount over a 25-year period.[6] For gas-fired generators, the corollary is the "heat rate"—i.e., how many BTUs of natural gas are required to generate one kWh of electricity—which can be contractually guaranteed.

- **OpEx risk:** Unlike financing and CapEx risk, the risk that a project might cost more to operate and maintain than was initially expected is a long-term risk that is of significant concern. Like technology risk, however, OpEx risk is typically mitigated (though not eliminated) contractually through a combination of manufacturers' warranties and medium-to-long-term service contracts. For example, many utility-scale project owners enter into service contracts with O&M providers (including the original equipment manufacturers) to operate and maintain their projects over periods of as long as 10 years (and in rare cases, up to 20 years). These service contracts typically do not cover major repair and replacement costs, however (hence the "technology risk" category above).
- **Environmental or regulatory risk:** Although wind and solar power are certainly not devoid of environmental impacts (e.g., both technologies have had problems with avian mortality), they are largely immune to what are likely the greatest regulatory risks facing power plants in the future: the risk of new or stricter regulations limiting emissions of criteria pollutants and greenhouse gases. Even though this risk (like resource risk) clearly differentiates renewable from gas-fired generation, the fact that it does so by being largely one-sided—i.e., falling disproportionately on gas-fired generators while not being particularly applicable to renewable energy—makes it of less interest to this study, which is focused on comparing the relative cost impacts of worse-than-expected outcomes for *both* types of resources. In addition, the relative environmental benefits of renewable generation have already been widely studied elsewhere (Millstein et al. 2017), and cannot readily be pulled into comparative cost analyses without making a series of assumptions about the value of externalities.
- **Offtake risk:** The vast majority of renewable generators operating in the United States have entered into long-term power purchase agreements or medium-term financial hedge agreements in order to mitigate offtake risk. Gas-fired generators, on the other hand, are often significantly exposed to spot market prices, though in most regions of the country with organized markets, gas-fired generators benefit from a built-in or natural hedge, in the sense that gas-fired generation often sets the market clearing price, meaning that wholesale electricity prices will increase or decrease commensurately with underlying natural gas prices, leaving the generator's "spark spread" (i.e., the spread between the price of wholesale power and the cost of burning natural gas to generate power) largely insulated.

This is not to say that these other six risk categories that are not included in this analysis are not important—they are. Instead, the point is that, in most cases, these other risks either (A) do not impose the same type of long-term uncertainty that resource risk imposes, (B) do not as clearly differentiate between renewable and gas-fired generation as does resource risk, or (C) do not really even impact renewable generation to any significant extent (e.g., in the case of environmental risk). As such, these other risks are less interesting for the purpose of this paper, which will focus solely on resource risk—i.e., wind and solar resource risk for wind and solar projects, and fuel price risk for gas-fired generators.

---

[6] For example, although details vary (e.g., some guarantees are linear while others step down in various increments over time), most solar modules sold today are warranted to generate at least 97% of rated capacity initially, declining over 25 years to at least 80% of initial rated capacity.

Resource risk is relatively difficult to mitigate

Of the various categories of risk discussed in this chapter, resource risk is perhaps the most difficult to mitigate or hedge—for both gas-fired and renewable generation. For gas-fired generation, resource risk primarily means natural gas price risk. It is, of course, possible to hedge natural gas price risk through the use of futures contracts (or, more broadly, forward contracts), but typically only over the short-term, as the natural gas futures market is not very liquid out beyond a few years, where uncertainty is greatest (Bolinger 2013). One could string together a series of short-term hedges in order to span a longer time period, but each time the short-term hedges roll over, the hedger will be exposed to any price movements that have occurred in the interim. In an attempt to hedge gas price risk over longer terms, some utilities have gone so far as to purchase ownership stakes in proven natural gas reserves, but this approach has its own challenges (Costello 2016).

For wind and solar generation, resource risk is a quantity risk rather than a price risk. Weather derivatives that are based on relatively simple metrics like "heating degree days" have been used to hedge the vagaries of the weather in other industries, but new instruments based on wind speed or insolation have been slow to catch on in the renewable power sector for a variety of reasons, including greater measurement challenges, basis risk (between the measured index location and the project site), and cost (particularly relative to the likely benefit under most PPA structures). That said, in 2016, at least two wind farms under construction in the United States entered into 10-year "proxy revenue swaps" that hedge not only the price of electricity that the project will receive for its generation (electricity price hedges are common among merchant wind projects in Texas and elsewhere), but—for the first time—also the amount of electricity generated (Allianz 2016a, Allianz 2016b). Specifically, these first-of-their-kind financial instruments swap out "the floating revenues of a wind farm—those driven by the hourly wind resource and power price—for a fixed annual payment" (Allianz 2016a). More recently, a solar performance data provider has argued for the provision of production or revenue "puts" that would similarly guarantee solar generators a minimum amount of production or revenue (Matsui 2017, Matsui et al. 2016).

While the financial and insurance sectors have made progress towards enabling wind and solar generators to more-cost-effectively hedge resource risk, the uncertainty surrounding resource risk can also be reduced (but not eliminated) in two simpler ways: by assessing resource risk across a diverse portfolio of projects rather than at just a single project, and by supplementing pre-construction energy assessments with operational energy assessments once projects have been operating for ideally at least a year. Both of these strategies are discussed later in the text box on page 36.

2.2 Who bears resource risk?

Having established an acute focus on resource risk, an important question is who bears the risk that gas prices might be higher than anticipated or that the wind or solar resource will be lower than anticipated? The answer will help to shape the analysis and determine the relevant audience for this work.

As shown in Table 1, the range of potentially impacted parties includes utility ratepayers, utility shareholders, and/or independent power producers ("IPPs") that own projects that sell electricity to utilities.[7] As also shown, the incidence of resource risk varies somewhat depending on whether the project is owned by a utility or by an IPP, and whether the risk in question is high gas prices or low solar/wind output.

Table 1. Qualitative heat map of resource risk impact on relevant stakeholders

	High Gas Prices		Low Solar/Wind Output	
Stakeholder	**Utility-Owned**	**IPP-Owned**	**Utility-Owned**	**IPP-Owned**
Independent Power Producers (IPPs)	None	None/ Low	None	Moderate/ High
Utility Ratepayers	High	High	Moderate	Moderate/ Low
Utility Shareholders	Low	Low	Low	Low

Higher-than-expected gas prices: As shown in Table 1, the risk of higher-than-expected gas prices is borne primarily by utility ratepayers. IPPs face very little of this risk because most gas-fired PPAs are structured either as tolling agreements (where the electricity purchaser is independently responsible for procuring and delivering gas to the generating plant) or as short-term contracts with electricity prices that are either linked to natural gas prices (in which case the buyer again bears the risk) or, in some cases, fixed (in which case the IPP bears the risk, but can cheaply hedge it in the gas futures market, which is very liquid over the short-term). Alternatively, if an IPP sells gas-fired generation on a merchant basis, the fluctuating spot price of wholesale power often provides a built-in or natural (though not entirely perfect) hedge against gas price risk, rising as gas prices increase and falling as gas prices decrease, leaving the "spark spread" (i.e., the spread between the price of wholesale power and the cost of burning natural gas to generate power) relatively stable regardless of gas price levels.

In most cases, then, fuel price risk falls upon the utility buyer or owner of gas-fired generation, which in turn typically allocates it to ratepayers rather than shareholders via rate case proceedings, as well as "fuel adjustment clauses" that allow electric rates to track fuel prices in between regularly scheduled rate cases.[8] If allowed or instructed by regulators to do so, a utility may hedge a portion of its fuel price risk in an attempt to protect ratepayers from extreme price swings, but as noted in the text box on page 8, natural gas price hedging is typically short-term in nature and leaves ratepayers exposed to long-term

---

[7] Table 1 was designed with investor-owned utilities in mind, but can easily be adapted for publicly owned utilities like municipal utilities or cooperatives, as well as competitive retail electricity providers in states that allow retail choice. In the case of publicly owned utilities, there is no distinction between ratepayers and shareholders, so one can simply ignore the bottom row ("Utility Shareholders"). In a competitive retail environment, one can think of "Utility Ratepayers" more broadly as "electricity customers" and "Utility Shareholders" as "retail electricity provider shareholders." Swapping these labels does not alter the basic findings that "electricity customers" or "ratepayers" bear the brunt of the risk of higher-than-expected fuel costs whether the project is utility- or IPP-owned, while wind/solar resource risk tends to be more moderate all around and, in the case of an IPP-owned project, is shared by the IPP and the customer/ratepayer.

[8] In some instances, fuel price risk might be allocated among both ratepayers and shareholders, either on an ex ante basis (e.g., through regulatory agreement to cap the recovery of fuel cost increases) or on an ex post basis (e.g., through regulatory disallowance of some portion of fuel cost increases).

price increases. In addition, the cost of hedging and the negative impact of any unprofitable hedges are typically passed along to ratepayers as well (Costello 2016).

Lower-than-expected solar or wind resource: Utility shareholders also largely escape the risk of lower-than-expected solar or wind output, which is borne primarily by ratepayers and IPPs. In the case of a utility-owned project, shareholders likely earn a guaranteed rate of return on the investment regardless of how well the project performs (within bounds),[9] while ratepayers will likely shoulder the burden of any electricity or renewable energy credit ("REC") shortfall that might arise from poor solar or wind resources (though this risk is likely modest assuming that replacement energy and RECs may not actually be required, or if required can perhaps be purchased at roughly comparable costs).

In the case of an IPP-owned project, sub-par wind or solar resources will eat into the IPP's profit margin, while requiring ratepayers to pay the cost of any replacement energy or RECs that are required (just like with a utility-owned project). Ratepayers are largely protected from severe under-performance under a PPA, however, through minimum output guarantees that have become a regular feature of most wind and solar PPAs, and that impose financial penalties on the IPP if minimum generation levels are not achieved.[10] These output guarantees are the reason why the ratepayer risk is deemed slightly lower in Table 1 for an IPP-owned project than it is for a utility-owned project (and are also why the IPP risk can be considered moderately high as opposed to just moderate). From a shareholder perspective, meanwhile, a PPA with an IPP is a primarily a pass-through, and as such imposes little or no risk.

Overall: Viewing Table 1 in its entirety, two general observations arise. First, utility ratepayers bear considerably more resource risk than do utility shareholders (though in the case of an IPP-owned wind or solar project, ratepayers share some of the resource risk with the IPP project owner). As such, utility ratepayers and those who represent their interests, like ratepayer advocates and utility regulators, are the primary audience for this work. To the extent that regulators require utilities to also be mindful of resource risk when planning for and procuring generation resources, utility planners and procurement teams are a secondary audience.

Second, higher-than-expected gas prices seem riskier (to ratepayers) in general than lower-than-expected wind or solar output. This suggests that from a ratepayer perspective, we should perhaps be more concerned about gas price risk than about wind or solar resource risk. In other words, in a case where two scenarios—one focusing on higher-than-expected gas prices and another focusing on lower-than-expected wind or solar resources—may be considered to have the same probability of occurrence, the resulting impact of the high gas price scenario may be more harmful to ratepayers than the impact

---

[9] As an example, when evaluating MidAmerican Energy's request for approval of the utility's "Wind IX" project, Iowa regulators stated that "…once Wind IX is added to rate base, MidAmerican's customers will be under an obligation to pay increased rates to allow MidAmerican to recover the return on equity specified in the ratemaking principles if the project does not deliver the anticipated results even if it were for no fault of theirs, despite MidAmerican's stated goal of no harm to customers. Thus, MidAmerican's risks associated with this project are relatively low." (Iowa Utilities Board 2016)
[10] PPAs also often specify a maximum amount of generation that will be paid the contract price; any generation above this level receives a lower price, or in some cases is not paid for at all (Schnitzer and Thienpont 2014).

of the low wind/solar resource scenario. The remainder of this document will develop a methodology that enables one to test this hypothesis (in terms of the levelized cost of energy or “LCOE” of each type of generation), as well as to compare LCOE impacts across scenarios that have the same, as well as different, probabilities.

3. Probability of exceedance as a tool to compare resource risk

This chapter introduces a framework for analyzing, visualizing, and quantifying the risk that natural gas prices might be higher-than-expected and/or the wind or solar resource might be lower-than-expected. This framework is based on a statistical concept known as probability of exceedance, which, with just a few key inputs, enables one to estimate the probability that, for example, the wind resource at a given site will fall below a given level, or the price of natural gas will rise above a given level. As such, it allows one to compare these risks in similar, probabilistic terms.

This chapter begins by discussing the application of probability of exceedance concepts within the renewable energy sector, where it has been well-developed and widely used to help characterize the quality of the resource at a given site. It then goes on to describe and develop a less-well-known but corollary application in the natural gas market.

3.1 Use of probability of exceedance within the renewable energy sector

Resource assessment is a crucial part of the development process for utility-scale wind and solar projects. Though approaches vary based on site terrain, whether the project in question is wind or solar, and other factors, such campaigns typically include some amount of on-site measurement of the resource, perhaps augmented by satellite data and meteorological modeling, and likely combined with longer-term measurements from offsite (but ideally nearby) reference stations.

The resulting central or "P50" estimate of wind speed or insolation at the site comes with two primary sources of potential error or bias. One stems from the resource measurement and/or modeling techniques employed (and is considered to be a *systematic* bias) while the second relates to the inherent inter-annual variability of the resource over time (considered to be a *random* error that is normally distributed about the mean). In addition to these two sources of error that are specific to the strength of the renewable resource, a third source of potential error (also *systematic*) is introduced when a production model is used to convert the estimated resource strength (i.e., wind speed or insolation) into a projected amount of generation from the proposed project.

Statistical error propagation techniques dictate that in order to arrive at the *total* uncertainty surrounding the central estimate of wind or solar generation, one must combine these three sources of error in quadrature (i.e., by taking the square root of the sum of squares), as follows:

$$\text{Total uncertainty}_{\text{generation}} = \sigma_T = \sqrt{\sigma_a^2 + \left(\sigma_b/\sqrt{\#\,years}\right)^2 + \sigma_c^2} \qquad (1)$$

where σ_T=total uncertainty surrounding energy generation, σ_a=measurement uncertainty, σ_b=inter-annual variability, and σ_c=production modeling uncertainty (all standard deviations are expressed as a percentage of the means). Because inter-annual variability in the wind or solar resource (σ_b) is considered to be random and normally distributed about the mean, it tends to cancel out somewhat over longer time periods, decaying at the rate of $1/\sqrt{\#\,years}$. As a result, the total uncertainty around the central generation estimate also decreases over longer time horizons, even though the other two error terms (σ_a and σ_c) are considered to be systematic, and do not decay over time.

For example, Table 2 provides a breakdown from two different renewable energy consultancies—Black & Veatch (Black & Veatch 2011) and AWS Truepower (Schnitzer et al. 2012)—of the typical range of energy uncertainty for utility-scale PV projects. The Black & Veatch ("B&V") breakdown includes just the three terms shown in Equation 1, while the AWS Truepower ("AWS") breakdown provides additional granularity on two of those three terms (which are both subtotaled to enable easier comparison with the B&V breakdown). Although there are differences, the overall illustrative ranges of uncertainty (and in particular the midpoints of those ranges) are broadly consistent with one another, with uncertainty declining over longer time horizons as inter-annual variability decays. Later sections of this report that model a utility-scale PV project use energy uncertainty estimates that fall within the middle of these ranges: specifically, 7% insolation error, 4% inter-annual variability, and 3.5% production model error combine (via Equation 1) for a total energy uncertainty of 8.79% in any single-year period, 7.93% in any 10-year period, and 7.87% in any 25-year period.

Table 2. Two estimates of energy uncertainty* for utility-scale PV projects

B&V Terminology & Breakdown	B&V Uncertainty*		AWS Terminology & Breakdown	AWS Uncertainty*	
	Min	Max		Min	Max
			Spatial Variability	0%	1%
			Representativeness of Monitoring Period	0.5%	2%
			Measurement Uncertainty	2%	15%
Insolation Error	5%	10%	*Subtotal:*	*2.1%*	*15.2%*
Inter-Annual Variability	2%	6%	Inter-Annual Variability	2%	5%
			Annual Degradation	0.5%	1%
			Transposition to Plane of Array	0.5%	1%
			Energy Simulation & Plant Losses	3%	5%
Production Model Error	3.5%*	3.5%*	*Subtotal:*	*3.1%*	*5.5%*
Total Energy Uncertainty*			**Total Energy Uncertainty***		
1-Year:	**6.4%**	**12.2%**	**1-Year:**	**4.2%**	**16.9%**
10-Year:	**6.1%**	**10.8%**	**10-Year:**	**3.8%**	**16.2%**
25-Year:	**6.1%**	**10.7%**	**25-Year:**	**3.7%**	**16.2%**

*Uncertainty is expressed as the coefficient of variation—i.e., the standard deviation divided by the mean.

**Black & Veatch (2011) refers to a single production model and does not provide a range of uncertainty.

Uncertainty levels for a typical wind project are similar to those shown for PV in Table 2, but are generally a little larger.[11] For example, Bailey and Kunkel (2015) provide a typical range of total wind energy uncertainty of 5.1% to 14.6% (with a mean of 7.5%) over a 10-year period. Anderson (2010) also

[11] Solar uncertainty is less because the solar resource is generally more predictable than the wind resource over every time scale (i.e., diurnal, seasonal, and—most relevant in this instance—inter-annual).

provides a 10-year uncertainty estimate, which totals 8.1%. Similarly, a confidential pre-construction energy assessment for a wind project that is now operating in Oklahoma reflects 1-year total uncertainty of 11.2%, which declines to 8.4% over any 10-year period and 8.2% over any 25-year period.[12] Because these uncertainty estimates from Oklahoma are derived from a real (rather than indicative) project, they will be used when modeling a wind project in later sections of this paper.

The central or median estimate of wind or solar generation from the project (which is also equal to the mean, given the assumption of a normal distribution) is called the "P50" estimate, meaning that there is a 50% chance that actual generation from the project will exceed the P50 estimate (and hence, in this case, also a 50% chance that actual generation will fall short of the P50 estimate). Armed with this central P50 projection and its total propagated uncertainty, in conjunction with the well-known statistical properties of a normal distribution, the resource analyst can also calculate other "probability of exceedance" levels ("P-levels") ranging, for example, from P1 (i.e., there is only a 1% chance that actual generation will exceed the P1 estimate) to P99 (i.e., there is a 99% chance that actual generation will exceed the P99 estimate). For any probability of exceedance level P_α, the equation is as follows:

$$P_\alpha = P_{50} * [1 - (z_{\alpha,\infty} * \ \sigma_T)] \qquad (2)$$

Where
P_{50} = P50 energy production estimate
$z_{\alpha,\infty}$ = Standard normal distribution value for $(1-\alpha)$ confidence level with infinite degrees of freedom
σ_T = total uncertainty surrounding the central estimate of wind or solar generation (from Equation 1)

Although Equation 2 implies that comparable P-levels above and below the P50 generation estimate (e.g., P5 and P95) will be symmetrical around the P50 estimate, this is not entirely the case—particularly out at the tails of the distribution. While P-levels that pertain to resource strength alone should be symmetrical around the P50, once those resource estimates are translated into generation estimates, technology constraints come into play and limit the upside potential from a stronger-than-expected resource. For example, once a wind turbine is operating at full capacity (i.e., operating in the flat portion of its power curve), any subsequent increase in wind speed will not result in incremental power generation (in fact, if wind speed continues to increase unabated, the turbine will eventually shut down for safety reasons and generate no power at all). Similarly, for a solar project, higher-than-expected insolation may simply lead to greater-than-expected "power clipping" (i.e., periods when the inverters are operating at full capacity and not converting excess DC generation from the array) rather than additional solar generation. In this way, a better-than-expected wind or solar resource may not necessarily lead to greater-than-expected wind or solar generation (particularly out at the tail of the distribution—e.g., P5-P1—where technology constraints kick in), while a worse-than-expected wind or solar resource will *always* lead to less-than-expected wind or solar generation, regardless of how bad it

[12] These energy uncertainty estimates are derived from a net energy production table that shows different probability of exceedance values over different time horizons. This derivation enables one to break out total uncertainty into its overall random (inter-annual variability of 7.9%) and systematic (8.1%) components, but does not allow the combined systematic error of 8.1% to be further disaggregated into measurement error and production model error.

gets (i.e., even in P99 situations). This asymmetry at the positive tail of the distribution complicates a full probabilistic assessment across *all* statistical outcomes (i.e., P1-P99), but can be easily sidestepped by focusing only on worse-than-expected outcomes – which is the approach we take here.[13]

This is also the approach taken by most project developers and investors, who will look at a fairly standard subset of worse-than-P50 outcomes—including P75, P90, P95, and P99—to gauge the resource risk. In addition, they will want to see these probability of exceedance values ("P-values") presented over different time frames that match up with their respective interests in the project—e.g., at least over one and 10 years, and perhaps also over 20-25 years.[14] While the P50 generation estimate does not vary across time horizons, all other P-values are time-dependent, gradually drifting towards the P50 estimate from above (if <P50) or below (if >P50) over longer time horizons due to the decay in inter-annual variability noted above.

Figure 1—expressed in terms of projected capacity factor rather than generation[15]—illustrates this narrowing for the hypothetical wind and solar projects modeled in this paper. The range of P-values below the expected (i.e., P50) capacity factors of 47% for wind and 32% for solar reflect the total energy uncertainty values discussed above—i.e., 11.2% in any single year declining to 8.2% over any 25-year period for wind and 8.8% in any single year declining to 7.9% over any 25-year period for solar. As a result of this reduction in uncertainty over longer time horizons, the various capacity factors drift upwards toward the P50 estimates over time. For example, the P99 wind capacity factor is 34.7% in any single year but increases to 38.1% over any 25-year period. Similarly, the P99 solar capacity factor is 25.5% in any single year but increases to 26.1% over any 25-year period. As expected, the P50-P99 capacity factor range is considerably narrower for solar than for wind—both in absolute and percentage terms. The capacity factors shown in Figure 1 are critical inputs into the LCOE analysis in Chapter 4.

---

13 As will be seen later, natural gas prices do not similarly suffer from this type of technology-limited asymmetry—although they are lognormally distributed and therefore skewed towards higher prices.

14 Different parties involved with a project will be most interested in different P-levels, and over different time horizons. For example, developers will use the P50 production estimate (which is the only P-value that does not vary over different time horizons) to help determine the minimum required price per MWh at which the project needs to sell its generation in order to reach a target return on investment. Tax equity investors in a wind project, meanwhile, may be most interested in seeing the suite of standard P-levels calculated over a 10-year period, given the 10-year production tax credit window (after which tax equity investors will often exit a project). Lenders to the project are exposed to the risk of lower-than-expected production but do not share in the benefits of higher-than-expected production, and so will be most interested in the 1-year P99 production estimate, which they will use to conservatively size the amount of debt that the project can likely support without defaulting in any single year. In addition to these other time horizons that are important to its partners or investors, project sponsors may also be interested in the 20- or 25-year P-values, presuming they intend to be long-term owners. Finally, ratepayers—i.e., those who ultimately pay for the electricity generated by the project—may have a relatively short-term focus if, for example, they are not certain how long they will reside (or remain in business) within the utility's service territory.

15 Projected capacity factor is defined as the projected amount of generation over a given time period divided by the maximum possible amount of generation over that same time period if the plant were running at full capacity over the entire period. The use of projected capacity factor rather than projected generation for this purpose enables wind and solar projects of different capacities to be compared on a normalized basis.

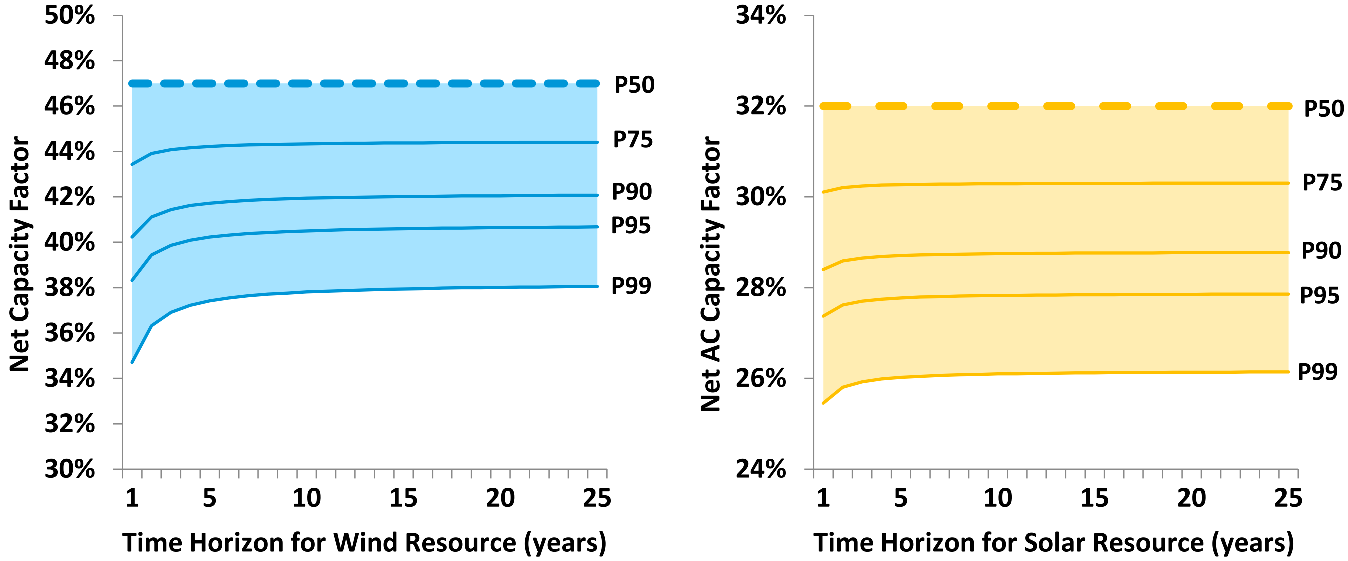

Figure 1. Illustrative probabilistic wind and solar capacity factors (P50-P99)

3.2 Applying probability of exceedance to natural gas price projections

The same probability of exceedance concepts that are used in the wind and solar industries to characterize the quality of the renewable resource are also used within the natural gas futures and options markets, where traders employ the concepts in conjunction with empirical market data to derive the market's view of gas prices moving above or below certain price levels by a given future date.

The Energy Information Administration ("EIA") describes this practice in an October 2009 supplement to its monthly *Short-Term Energy Outlook* ("STEO") publication (Ryan and Lidderdale 2009), and subsequently began using this methodology within each monthly STEO release to calculate confidence intervals around the latest natural gas futures curve or "strip." Distilled down to its simplest form, the methodology described by Ryan and Lidderdale (2009) is based on two main pieces of information—the market's central projection of future natural gas prices and the market-implied uncertainty surrounding that central price projection.

Central projections of future natural gas prices come from the natural gas futures market, which currently consists of contracts for delivery in every month over the coming 12-13 years.[16] The full series of contracts over this 144- to 156-month period is often referred to as the "futures curve" or "futures strip." Settlement prices along the futures strip can be thought of as the market's central expectations for spot natural gas prices in any given month within the strip.

Uncertainty surrounding these central price projections can be derived from the observed price of

---

16 Each December the market tacks an additional 12 months onto the end of the futures strip, for 156 months in total. That total then gradually declines back to 144 months over the course of the ensuing year, as the first twelve nearby contracts expire in turn.

options on the underlying natural gas futures contracts. Using the Black-Scholes options pricing model, one can back into the amount of volatility implied by an empirical options price—this is known as "implied volatility." Armed with the implied volatility, the amount of time until the option expires, the price of the underlying futures contract on which the option is based, and the assumption that natural gas prices are lognormally distributed (supported by the text box on page 19), the EIA then calculates two-tailed confidence intervals according to the following formulas (reproduced directly from Ryan and Lidderdale (2009)):

$$\mathrm{E}(f_{\tau,k}) > f_{t,k} * \exp[-(z_{\alpha/2} * \sigma_k\sqrt{\tau_k})] \tag{3}$$
$$\mathrm{E}(f_{\tau,k}) < f_{t,k} * \exp[+(z_{\alpha/2} * \sigma_k\sqrt{\tau_k})] \tag{4}$$

Where
$\mathrm{E}(f_{\tau,k})$ = Expected month k futures contract price at expiration date
$f_{t,k}$ = Month k futures contract price at day t
σ_k = Implied volatility for option on month k futures contract
τ_k = Time to expiration for month k futures contract (in years)
$z_{\alpha/2}$ = Standardized normal distribution value for $(1-\alpha)$ confidence level

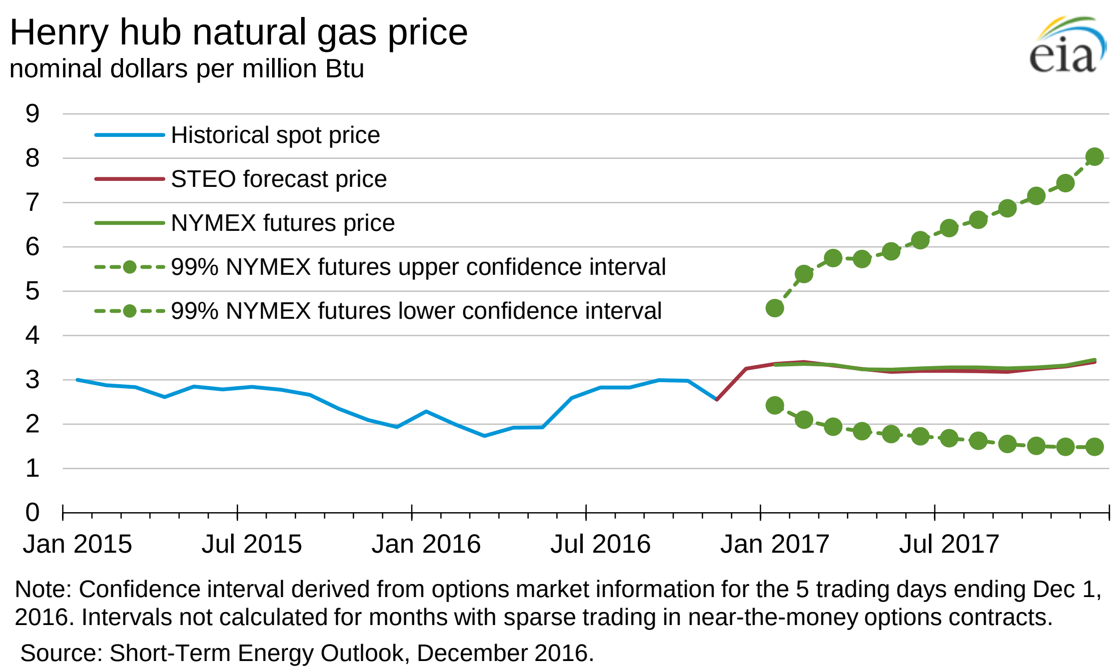

Figure 2. Illustration of 99% confidence intervals around natural gas futures prices, from the EIA's December 2016 Short-Term Energy Outlook

This process results in confidence intervals that generally look like those in Figure 2 above, which is pulled directly from the December 2016 edition of the EIA's *Short-Term Energy Outlook* and is generated by the method described in Ryan and Lidderdale (2009). As shown, the 99% confidence intervals around the futures strip start out close to the strip but progressively widen the further out in time one looks, reflecting the fact that it is easier to predict where natural gas prices will be one or two months from now than it is one or two years from now. These widening confidence intervals stand in stark contrast to the narrowing range of P-values for wind and solar capacity factors shown earlier in Figure 1. In other words, while uncertainty surrounding the quality of the renewable resource *decreases* over longer time horizons, uncertainty over natural gas prices *increases* with time. This key distinction is fundamental to the new framework for evaluating resource risk that is developed in this paper.

Astute readers will have noticed that we have previously characterized renewable resource risk using *probability of exceedance* values, while natural gas price risk is expressed here in terms of *confidence intervals*. Though the terminology is different, the two concepts are closely related, and in fact differ only in terms of whether the level of confidence is expressed in terms of a one-tailed or two-tailed interval. Probability of exceedance—i.e., in this case the likelihood that the amount of wind or solar generation will *exceed* a certain level—is grounded in one-tailed confidence (i.e., just one "tail" of the normal distribution or probability density function), whereas the natural gas confidence intervals derived from Equations 3 and 4 are two-tailed confidence intervals that express the likelihood that natural gas prices will fall *in between* two prices. In order to set up a fair comparison between renewable resource risk and natural gas price risk in later sections of this paper, we convert the two-tailed natural gas confidence intervals in Equations 3 and 4 to one-tailed probability of exceedance values by replacing $z_{\alpha/2}$ with z_{α}.

Natural gas prices are lognormally distributed; changes in gas prices are normally distributed

The EIA's approach to constructing confidence intervals, as described in Ryan and Lidderdale (2009), rests on the assumption that natural gas prices are lognormally distributed (given a minimum price of zero and a probability density function that is skewed to the right with a long tail). As shown in in the figure below, this assumption of lognormality seems to be borne out by the empirical distribution of daily first-nearby natural gas futures prices from 1990 through 2016.

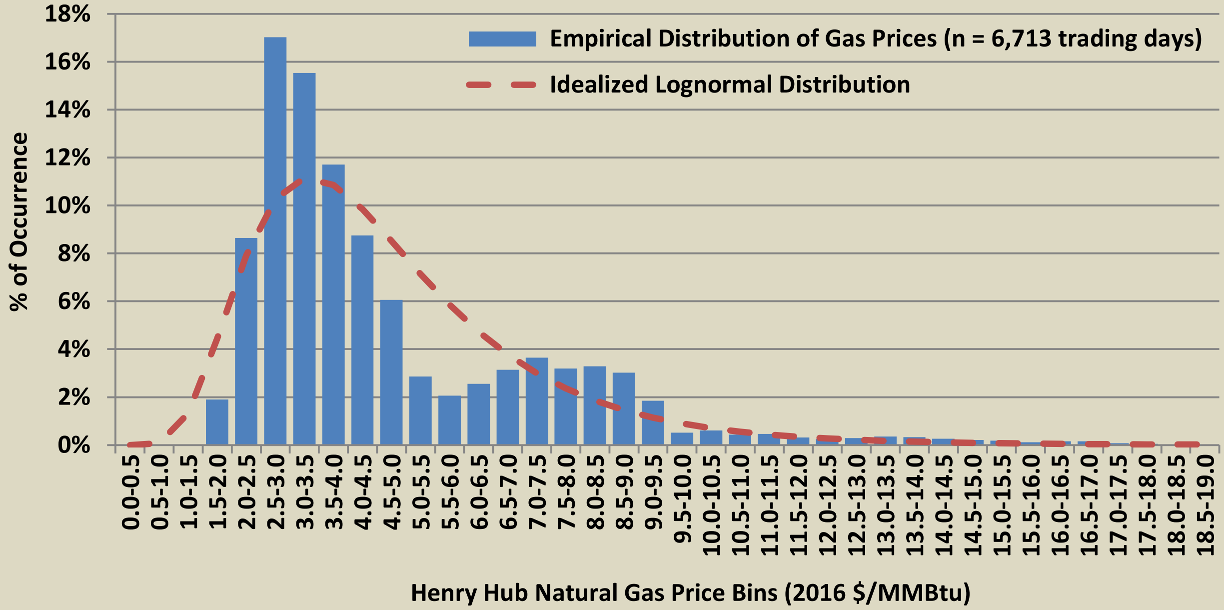

While natural gas prices are lognormally distributed, the year-on-year percentage changes in natural gas prices—from which historical volatility is calculated—are normally distributed (see the figure below), which enables the use of well-understood statistical properties and methods associated with normality, as described elsewhere in this paper.

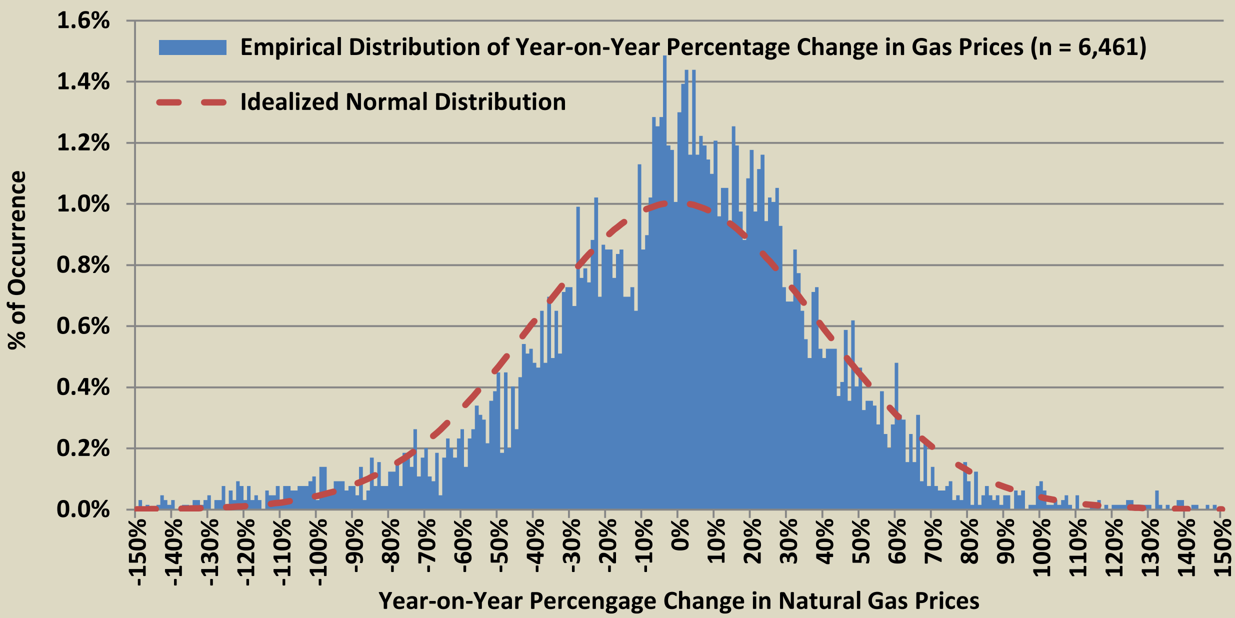

3.3 Extending the projected gas price ranges over longer terms

Before we can adopt the EIA's confidence interval approach for this purpose, however, we need to extend the projections further out in time—ideally to the 20- to 25-year timeframe over which most utility and IPP investment decisions are made (and for which we already have renewable resource P-values—see, for example, Figure 1). Within the STEO, the EIA has never published confidence intervals out beyond a maximum of three years. This is because the options market on natural gas futures contracts—from which implied volatility is derived—is not liquid out beyond a few years, which complicates price discovery, and hence also the calculation of implied volatility. If, however, one is willing to forego the direct link to market expectations that the use of implied volatility provides, then one can easily extend the confidence intervals generated by the EIA's approach by substituting *historical volatility* for *implied volatility*.

Historically volatility, which is calculated from past price movements, is clearly not the same thing as implied volatility, which is based on market expectations of future price movements. But in most instances, the two should be similar—i.e., absent some reason to believe otherwise, natural gas options traders will likely assume that future volatility will look very much like recent historical volatility. Numerous academic studies have tested whether implied or historical volatility is a better predictor of realized volatility; Ryan and Lidderdale (2009) provide a summary of this literature. Though the balance of these studies have found implied volatility to be the better predictor, most of this academic work has, understandably, focused on short time horizons—e.g., the period over which options markets are liquid. But over the longer time horizons of interest to this study (e.g., 10-20 years), one might reasonably question whether the market's predictive ability (which is not directly observable over such lengthy time frames) is really any better than simply taking a cue from past price movements. In short, although our use of historical volatility to extend the confidence intervals is admittedly a second-best solution, in the sense that it prevents us from tying the resulting confidence intervals to market expectations, it is nevertheless seemingly the sole option available over the longer time periods of interest to this study and, regardless, likely provides a good approximation of implied volatility over these longer periods (were it available).[17]

A relevant question, however, is over what historical period to calculate historical volatility? Natural gas prices have varied significantly ever since the market was deregulated back in the mid-to-late 1980s, providing nearly 30 years over which to observe historical price movements and calculate historical volatility. That said, as shown by the price curve (solid blue line) in Figure 3, the market has gone through three distinct pricing environments over this period: in the immediate wake of deregulation, the 1990s was a decade of relatively low and stable prices; the 2000s were, in contrast, a period of extreme volatility characterized by several severe price spikes and subsequent reversals; finally, the post-2008 period has largely seen a return to less-volatile pricing, in large part due to advances in horizontal drilling and hydraulic fracturing, which together have greatly increased domestic supply. As these same market conditions are expected to persist for the foreseeable future, this post-

[17] In fact, as shown later in Figure 5, historical volatility aligns well with the implied volatility estimates for 2017 contained in the December 2016 *Short-Term Energy Outlook*.

2008 period is arguably the most likely of the three different periods to be representative of price volatility going forward.

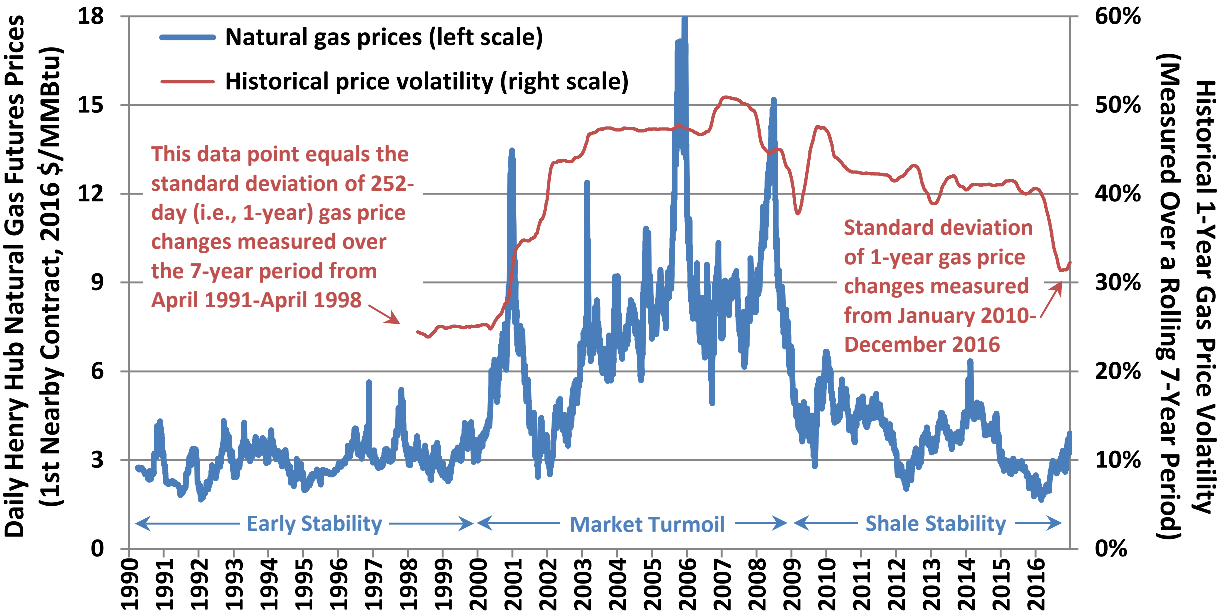

Figure 3. Daily first-nearby natural gas futures price history

That said, it should be recognized that focusing only on this post-2008 period is—at least within an historical context—a conservative choice for the purpose of assessing natural gas price risk and comparing it to wind and solar resource risk, in the sense that lower volatility assumptions result in narrower confidence intervals around the central gas price projection (thereby favoring gas-fired relative to wind and solar generation). Specifically, the end point of the rolling historical volatility curve (dashed red line) in Figure 3 represents the standard deviation of 252-day (i.e., 1-year) gas price changes measured over the 7-year period from January 2010 through December 2016, and as such reflects historical volatility—of 32.3%—over this post-2008 period.[18] Following this same curve back in time reveals that 32.3% volatility is relatively low compared to the 40-50% volatility that has persisted for much of this century (on the other hand, 32.3% is higher than the 24.4% historical volatility from the 1990s). Although focusing on post-2008 historical volatility is the most logical choice, it should nevertheless be recognized that this choice reflects an implicit assumption that bountiful shale gas will continue to keep prices and price volatility in check over the coming 25 years, and to the extent that this implicit assumption does not pan out, we will have most likely underestimated actual gas price volatility.

[18] While there are 365 calendar days in a normal year, there are only 252 trading days due to weekends and holidays. Even though the standard deviation of 252-day price changes is measured from 2010-2016, the resulting volatility of 32.3% represents the full 2009-2016 pricing period given that 2009 prices are part of the 252-day price change calculations that start in January 2010.

Figure 4 further explores the implications of focusing on this 2009-2016 subset of price history by plotting annualized historical volatility curves that have been calculated over four different price histories—i.e., over the full price history from 1990-2016, as well as over the three distinct shorter periods mentioned above (the 1990s, 2000-2008, and 2009-2016). Three aspects of Figure 4 are worth discussing:

1) In all four cases, gas price volatility clearly has a declining term structure—i.e., annualized volatility is high when measured over shorter time horizons (e.g., 1 year) but declines over longer time horizons. This decay in volatility (which, as explained below, is believed to be exponential in nature) reflects an element of mean reversion in the natural gas market—i.e., although gas prices may seem to move randomly over time, in some cases diverging significantly from "normal" price levels, sooner or later they eventually reverse course and revert back towards where they had been. Hence, as the time horizon lengthens, one is less likely to experience gas prices at their extremes and is more likely to experience gas prices closer to their means, resulting in lower volatility.
2) The relative ranking of historical volatility calculated over the three shorter price histories largely conforms to expectations (as well as to the rolling volatility curve in Figure 3): historical volatility is lowest when calculated using prices from the 1990s, highest when calculated using prices from 2000-2008, and is somewhere in between when calculated using prices since 2009 (which is the period chosen for this study).
3) The full-period volatility curve (blue line with circle markers) appears to be somewhat inflated (relative to an expected exponential decay pattern) over periods ranging from roughly four to nineteen years. This is presumably due to the presence of volatility calculations that span more than one of the three distinct shorter price history periods, thereby more often pitting low prices (e.g., from the 1990s, or post-2008) against high prices (from 2000-2008), leading to higher volatility when calculated over those mid- to-long-term time horizons from ~4-19 years.

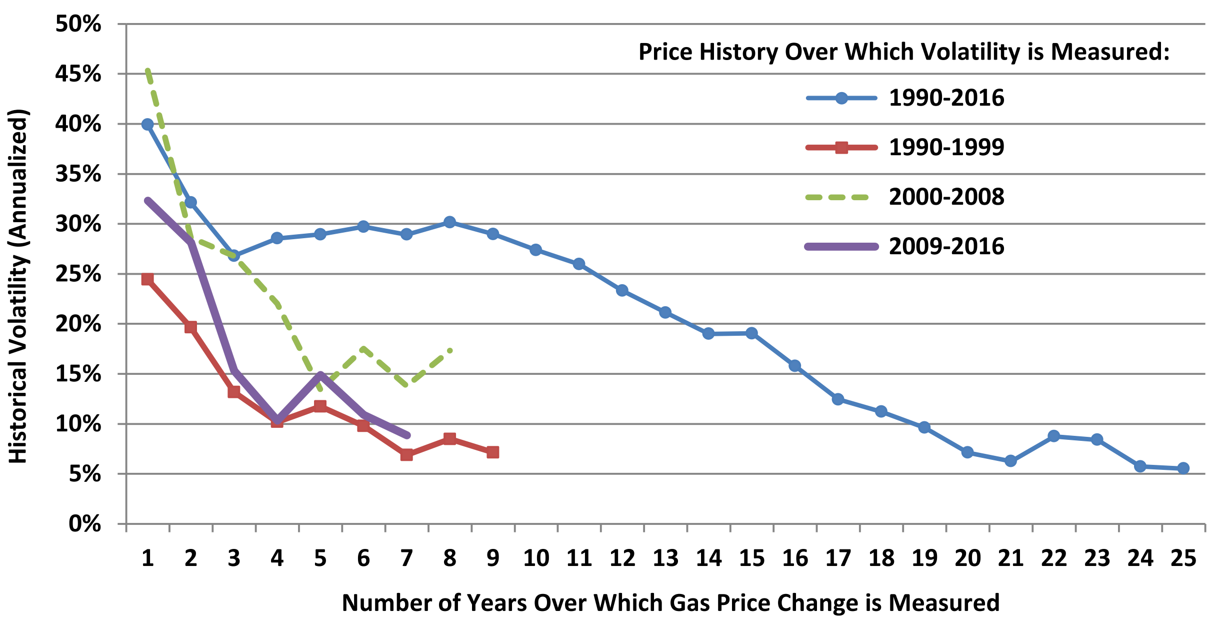

Figure 4. Historical natural gas price volatility measured over different time horizons and using different price histories

One downside of relying on just the 2009-2016 price history when calculating historical volatility is that it limits empirical volatility estimates to periods of up to seven years maximum (e.g., see the purple curve in Figure 4)—far shorter than the desired 25 years. Although historical volatility estimates over periods longer than seven years (and up to 25 years in Figure 4) are available when working with the full price history from 1990-2016, as noted above, these empirical volatility estimates are seemingly inflated by the highly-volatile 2000-2008 period, and are therefore likely not representative in this new era of shale gas. Fortunately, when working with random price data whose returns are normally distributed (i.e., natural gas prices), one can extrapolate shorter-term empirical volatility estimates to longer time horizons using a common convention: to convert an n-year volatility estimate to a q-year volatility estimate, multiply the n-year estimate by the square root of n and then divide by the square root of q.[19] This convention results in a volatility curve that decays exponentially over longer time horizons, similar in nature to the three short-term empirical curves shown in Figure 4.

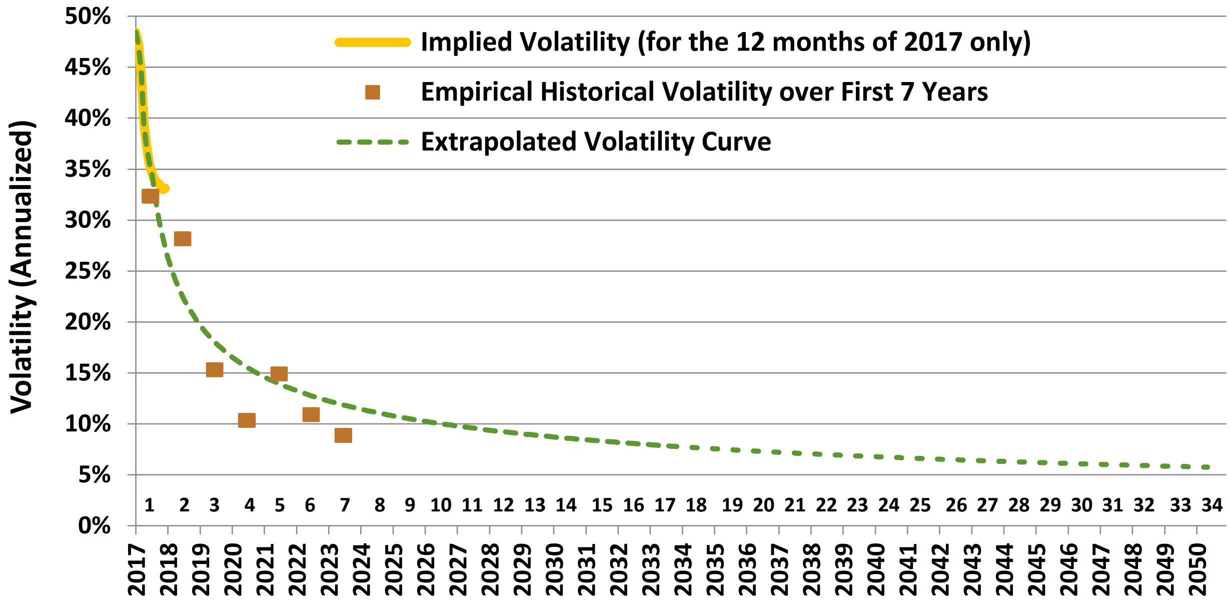

Figure 5. Implied, historical, and extrapolated natural gas price volatility over different time horizons

The dashed green line in Figure 5 shows such an exponential decay curve, in this case derived entirely from the one-year empirical volatility estimate of 32.3% (calculated from the 2009-2016 price history). As shown, this extrapolated volatility curve matches up reasonably well with the seven empirical data points (brown rectangles) calculated from the 2009-2016 price data (these seven data points are the

19 When starting with a one-year volatility estimate (i.e., n=1), this formula simplifies to dividing the one-year volatility estimate by the square root of q.

same as the purple line in Figure 4).[20] The first 12 months of the extrapolated curve, meanwhile, largely overlap with the implied volatility curve from the December 2016 STEO for the 12 months of 2017 (yellow line), suggesting a close linkage between implied and historical volatility. Finally, the long-dated part of the derived curve, with annualized volatility approaching 5%, matches up reasonably well with the long-dated portion of the 1990-2016 empirical curve from Figure 4. Given that the extrapolated curve provides a reasonable fit with both the implied and historical volatility estimates, and extends out to 25 years (and beyond), it is the curve used in our analysis.

Again, however, it is perhaps worth stressing that this extrapolated decay curve reflects an implicit assumption that current market conditions, characterized by low natural gas prices and moderate volatility, will remain largely unchanged over the next 25 years. If instead there is an unexpected regime change in natural gas prices and/or volatility—as has happened in the past—then this extrapolated volatility curve would likely underestimate actual volatility. As such, using this volatility curve to construct confidence intervals around natural gas price projections is perhaps conservative (from the standpoint of comparing the resource risk of gas-fired to wind and solar generation).

The decaying term structure of natural gas price volatility—clearly evident in Figures 4 and 5—raises an important question: if gas price volatility declines over longer time horizons, why do the confidence intervals shown earlier in Figure 2 progressively *widen* over increasing time horizons? The culprit is the time variable in Equations 3 and 4: at each successive point in time, the confidence interval formulas multiply a progressively lower volatility by the square root of a progressively longer time horizon. At least to date, time has emerged victorious in this struggle, causing the confidence intervals to widen over time, even as volatility declines.

In addition to the extended volatility estimate (i.e., the dashed green derived curve shown in Figure 5), adapting the EIA's approach described in Ryan and Lidderdale (2009) for our purpose requires an extended central price projection over the same period. As noted earlier, the underlying futures market extends out much further than shown above in Figure 2—it currently runs through 2029. And although it, too, is not liquid out beyond the first few years, the daily settlement prices nevertheless provide some information value over the full length of the futures strip (and would be subject to

[20] Specifically, the empirical historical volatility calculations are based on a data set consisting of daily first-nearby natural gas futures prices from January 2, 2009 through December 30, 2016. From that data set, seven time series of rolling daily logarithmic price returns (i.e., subtracting the natural log of the earlier price from that of the current price) over periods ranging from one calendar year (i.e., 252 trading days) up to seven years (i.e., 1,764 trading days) were constructed. Next, the annualized volatility of each of the seven time series was calculated by dividing the standard deviation of each time series by the square root of the number of calendar years in question (again, ranging from 1 to 7). The results are depicted by the solid purple line (with no markers) in Figure 4 as well as the brown rectangles in Figure 5. Imperfect alignment between the derived or extrapolated curve and the discrete historical volatility estimates—particularly over the 4-year period—likely reflects sampling issues, as well as the simple fact that empirical data do not always behave exactly as expected or according to theory.

arbitrage if they were way out of line with market expectations).[21] Beyond 2029, for lack of a more-defensible alternative, we simply extend the futures strip through 2050 at the same slope that existed in 2028-2029.

Figure 6 shows the end-result of this exercise in extending both the central price projection and confidence intervals (expressed now as P-values, having been converted from two-tailed to one-tailed terms). The solid green saw-toothed line that runs through 2029 represents the average natural gas futures strip for the five trading days leading up to December 1, 2016 (the sawtooth pattern reflects seasonal variation in gas prices—higher in winter, lower in summer).[22] The dashed extension of that line through 2050 is the aforementioned extension of the futures strip at the same slope as existed from 2028-2029. Collectively, the futures strip and its extension can be considered the central or P50 price projection.[23] The saw-toothed solid red lines, meanwhile, represent the P1, P25, P75, and P99 projections, as calculated using Equations 3 and 4 (though again, calculated as one-tailed lower and upper bounds, rather than as a two-tailed confidence interval). These should *not* be considered "market-based" P-values, given that they are derived from historical rather than implied volatility (except for perhaps during 2017, when historical and implied volatility overlap, as shown in Figure 5). Finally, included simply for the sake of comparison, the blue-shaded area represents the full range of gas price projections from the EIA's *Annual Energy Outlook 2017* ("AEO17"), with the dark-blue line running through the middle showing the AEO17 Reference Case projection (EIA 2017).[24]

---

21 Although the exchange provides daily settlement prices for each natural gas futures contract along the strip out through December 2029, it is important to recognize that many of the long-dated contracts rarely if ever trade. In these cases, the exchange calculates the settlement price according to an established protocol that relies on bids and offers (if available) and spread relationships. As such, long-dated futures prices should be considered merely indicative rather than firm and executable. For more information on how the exchange sets settlement prices for long-dated contracts, see http://www.cmegroup.com/confluence/display/EPICSANDBOX/Natural+Gas#NaturalGas-Contractsbeyondsixmonths

22 We take this 5-day average to be consistent with the EIA's approach in the STEO (EIA 2016). Taking a 5-day average (rather than a single day) helps to guard against the possibility of choosing an anomalous trading day.

23 P50 is not a strictly accurate label for the central gas price projection laid out by the futures strip, because the futures strip more likely represents a mean (i.e., expected value) than a median. The mean and median are the same under an assumption of normality, but recall from the text box on page 19 that natural gas prices are lognormally distributed, which means that the mean lies above the median. This issue of semantics does not, however, disqualify the approach to constructing confidence intervals around natural gas price projections as laid out by Ryan and Lidderdale (2009); it is instead primarily a labeling problem in the present application. Though not technically correct, I will nevertheless refer to the central gas price projection as a P50 projection in order to facilitate LCOE comparisons involving the central gas price projection and the central wind and solar resource projections (which do truly represent P50 projections).

24 Although there is no reason that one would expect the EIA projections to line up with any particular P-values, it is nevertheless interesting to see that the upper bound of the AEO17 gas price range is generally higher than the P1 projection, while the lower bound more or less follows the P50 projection over time. The Reference Case projection, meanwhile, is close to a P10 level (note: the P10 projection is not shown on the graph). In other words, the AEO17 gas price range appears to be skewed high, at least in a probabilistic sense based on where natural gas futures prices were trading in late November 2016 (AEO17 was published on January 5, 2017, but the modeling presumably occurred in late 2016). The fact that the P50 extrapolation (post 2029) has roughly the same slope as the AEO17 Reference Case projection (while remaining several $/MMBtu below that Reference Case projection) suggests that this extrapolation is appropriately conservative.

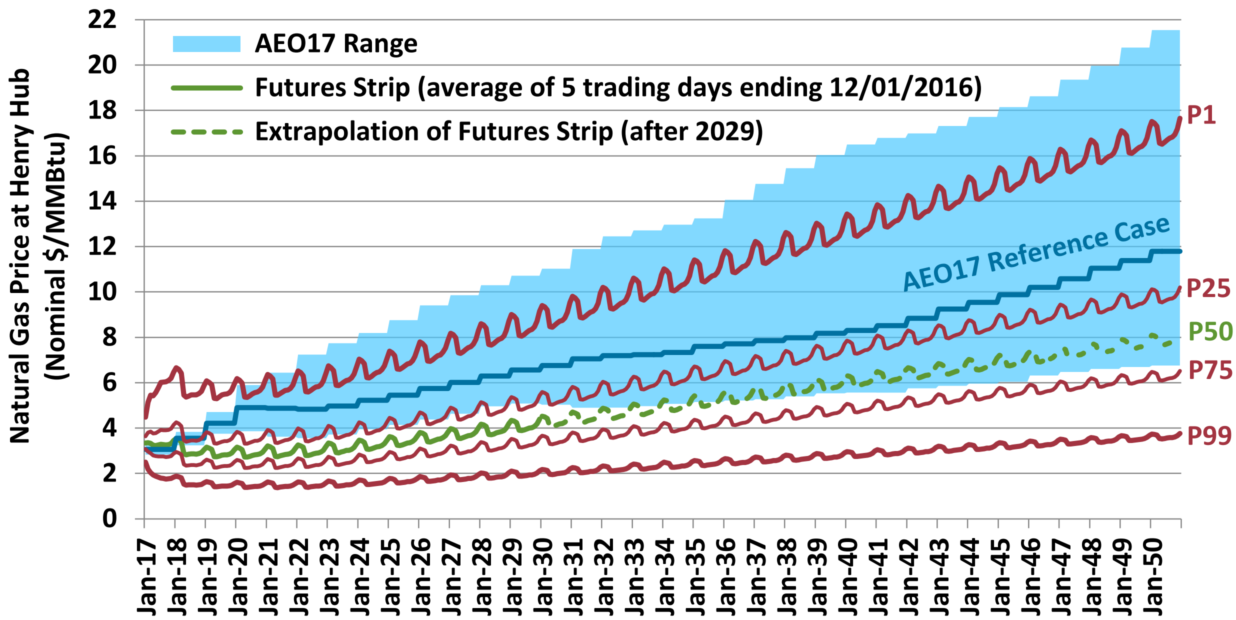

Figure 6. Projected natural gas prices over P1-P99 range and from AEO17

Focusing only on the P1, P50, and P99 projections—which form the core of the new framework developed in this paper—it is clear that gas price risk is skewed to the upside. The gap between the P1 and P50 projections is 1.8 times as wide as the gap between the P50 and P99 projections in 2017 (averaged across all 12 months), and increases over time to 2.2 times as wide in 2050. This upward skew,[25] which widens over time, stands in contrast to the symmetrical risk profile exhibited by wind and solar resource risk.[26]

Although they are a necessary starting point, the gas price projections shown in Figure 6 do not feed directly into the LCOE calculations in the next chapter. Instead, they are first levelized (at an 8% nominal discount rate[27]) over each time horizon from 1 to 25 years, as shown in Figure 7. For example, the 2025 values of the five curves (P1, P25, P50, P75, P99) in Figure 7 reflect the levelization of just the first nine years (i.e., 2017-2025) of gas prices from the same five curves shown in Figure 6. Similarly, Figure 7's values in 2041 reflect the levelization of the full 25-year price streams shown in Figure 6.

[25] It is perhaps worth noting that this upward skew is not simply a result of the present low gas price environment (i.e., where gas prices have much more room to rise than to fall). The EIA's historical log of confidence intervals around natural gas price projections pulled from past STEOs demonstrates that this upward skew existed even back when gas prices were much higher (see https://www.eia.gov/outlooks/steo/marketreview/pdf/uncertainty_hh_2007_2008.pdf).

[26] As described earlier in Section 3.1, although the distribution of wind and solar resource strength can be thought of as being symmetrical around the P50 projection, the resulting distribution of wind and solar *generation* is not fully symmetrical, due to technological limitations that prevent full conversion of the resource into electricity at the very tail end of the positive side of the resource distribution.

[27] This same nominal discount rate (or weighted average cost of capital) is used later to levelize the cost of electricity from wind and solar projects. Using a consistent discount rate across all technologies ensures that the vagaries of finance do not influence the cross-technology comparisons being made.

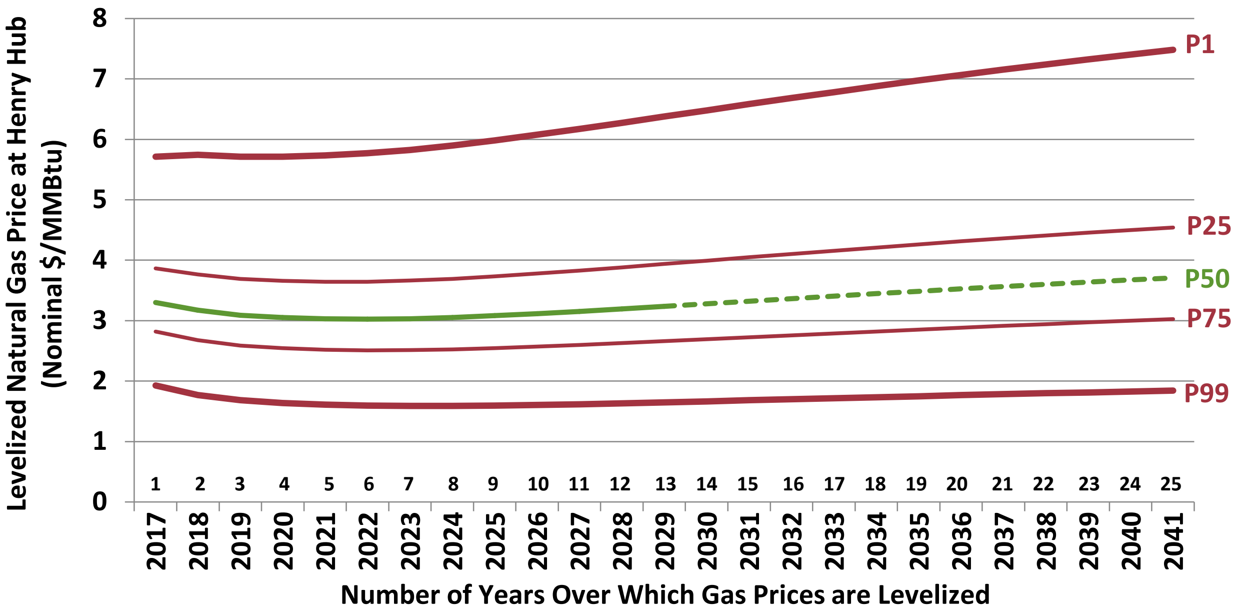

Figure 7. Levelized gas price inputs to the LCOE calculation

Because they are levelized, the gas prices in Figure 7 are much less extreme (and variable) than those in Figure 6. For example, while the P1 gas price curve in Figure 6 rises to almost $14/MMBtu in 2041 (i.e., in year 25), the levelized P1 curve increases to a little more than half as much over that same 25-year time horizon, because the high prices in later years are more-heavily discounted than the earlier, lower prices. It is these levelized price curves in Figure 7, rather than the raw price projections from Figure 6, that feed into the LCOE modeling in Chapter 4.[28]

For some readers, the range of natural gas prices shown in Figure 6 (or possibly Figure 7)—as well as the methodology used to derive the range—may be the key takeaway from this work. For example, utility resource planners, energy forecasters, and energy analysts may be more interested in the probabilistic range of natural gas prices developed in this chapter than in the LCOE comparisons constructed in the next chapter, given that gas prices are typically a key input into modeling processes that consider many factors other than LCOE. Integrated resource plans typically attempt to incorporate gas price uncertainty in some way; for example, by considering a range of published gas price projections (like the AEO17 projections shown in Figure 6) or by using Monte Carlo simulation to generate a range of natural gas price outcomes. The methodology developed in this report is yet another way to establish a gas price range—one that is primarily statistical in nature, which has both positive and negative implications.

On the positive side, this approach is probabilistic, thereby allowing the resource planner to look at as wide or as narrow of a range of gas price probabilities as is deemed appropriate. This framework can also be used to benchmark internal or third-party gas price projections, in terms of aggressiveness (as is

[28] Specifically, these levelized gas price curves (not just the five curves shown in Figure 7, but all P-levels in between P1 and P99) are converted to levelized $/MWh terms at the assumed heat rate, and then added to 25-year levelized estimates of CapEx and non-fuel OpEx to arrive at total 25-year gas-fired LCOE.

done with the AEO17 forecast in Figure 6). For example, if a "high price" scenario only equates to a P35 projection, this context is worth knowing.

On the negative side, this approach largely ignores the myriad fundamental drivers of natural gas prices, including supply, demand, weather, storage, the economy, the price of competing fuels, and many others. Although the central or P50 price projection is based on the natural gas futures strip, and so does presumably reflect the market's mean assessment of these fundamental price drivers, the P1-P99 range around the P50 projection is purely a statistical construct based on an assumption that future volatility will match historical levels. Although dispensing with the fundamentals in this manner might cause unease among some resource planners and analysts, others may find it liberating, enabling them to focus on the forest rather than just the trees.

In contrast to utility resource planners, other target audiences—e.g., ratepayers, ratepayer advocates, regulators, corporate offtakers, policymakers, investors—may be more interested in seeing the comparison of renewable resource and gas price risk developed further, and translated into consistent $/MWh LCOE terms so that they can be more easily compared. The next chapter takes this step.

4. Modeling and results

In Chapter 3 we learned that the same probability of exceedance concepts that are regularly used to express the uncertainty around annual energy production (and, hence, capacity factors) for wind and solar projects can also be applied to natural gas price projections, allowing one to develop a probabilistic range of projections for both annual energy production (and capacity factor) and natural gas prices. In addition, these distributions have interesting (and largely opposing) characteristics: renewable resource risk is symmetrical and declines over longer time horizons, while natural gas price risk is asymmetrical (skewed towards higher prices) and increases over longer time horizons. In this chapter, we plug these divergent capacity factor and natural gas price distributions into an LCOE model to enable a direct comparison of the impact of resource risk in consistent LCOE terms.

4.1 Modeling assumptions

To accomplish this, we use a financial pro forma model capable of modeling wind and solar projects as well as a natural gas combined cycle plant. The critical modeling assumptions for each type of generator are shown in Table 3, and are drawn from a combination of empirical data found in Wiser and Bolinger (2016) and Bolinger and Seel (2016), as well as specifications from Lazard (2016) and NREL (2016). The wind and solar project capacity factors and the levelized natural gas fuel prices shown in Table 3 represent P50 projections, and will differ at other P-levels (and over other time horizons), according to the curves shown earlier in Figure 1 (for wind and solar capacity factors) and Figure 7 (for levelized gas prices).[29] These are the only three assumptions that vary as we analyze a range of possible outcomes to reflect resource uncertainty; all other inputs are held constant (including the 75% capacity factor assumed for the gas-fired generator, in order to isolate the impact of fuel prices[30]).[31] Finally, although in reality there are subtle differences in how these three generation projects would likely be financed (Feldman and Bolinger 2016), this analysis assumes a common weighted average cost of capital ("WACC") of 8% across all three technologies, so that the comparative results are not driven by financing assumptions.[32]

[29] As in Figure 7, the levelized natural gas prices used in the LCOE calculations represent prices at the Henry Hub in Louisiana. Although the price of gas delivered to electricity generators will no doubt differ from the Henry Hub price (e.g., due to transportation costs, pipeline constraints, local supply, etc.), it is difficult to project and generalize whether the delivered price will be at a premium or a discount to Henry Hub prices, as this locational "basis" has historically varied both over time and by location. To keep things simple, this analysis simply uses unadjusted Henry Hub prices.
[30] One might expect that as fuel prices rise (i.e., for scenarios < P50) the capacity factor of gas-fired generation might decline somewhat as a result (as gas becomes less competitive with other sources of generation), thereby compounding the increase in gas-fired LCOE (i.e., fewer MWh over which to amortize higher fuel costs). Hence, holding the gas-fired capacity factor constant at 75% across all P-levels is likely a conservative assumption in this comparison.
[31] This is obviously a simplifying assumption that ignores linkages between resource strength and other factors. For example, if a wind project encounters a stronger-than-expected wind resource (e.g., P10) on a regular basis, its turbines would be subjected to increased loads and stresses, perhaps leading to higher-than-expected O&M costs.
[32] This common 8% WACC is equal to the 8% nominal discount rate used to levelize natural gas prices. In addition, we assume that the wind and solar project sponsors have sufficient tax appetite to benefit from the tax benefits created by these projects in the years in which they are generated.

Table 3. Modeling assumptions for generators starting operations on January 1, 2017

	Wind	Solar	Gas (CCGT)
CapEx	$1.50/W-AC	$1.80/W-AC	$1.10/W-AC
Fixed Non-Fuel O&M	$40.0/kW-year	$15.0/kW-year	$6.0/kW-year
Variable Non-Fuel O&M	None		$3.0/MWh
Fuel Price (25-Yr Nominal Levelized)	N/A		$3.71/MMBtu (25-yr P50)
Heat Rate	N/A		6,700 Btu/kWh
Net Capacity Factor	47.0% (P50)	32.0% (P50)	75.0%
Degradation	None	0.4%/year	None
Tax Depreciation	5-year MACRS		20-year MACRS
Nominal Discount Rate (WACC)	8.0%		
Inflation	2.0%/year		
Tax Rate	35% federal, 7.7% state (40% combined)		
Modeled Plant Life	25 years		

4.2 Visual results

Figure 8 shows the first LCOE comparison of wind (with the PTC) to a gas-fired combined cycle gas turbine. The blue-shaded area around the dashed blue line represents the CCGT's P1-P99 LCOE range around its P50 LCOE curve. Meanwhile, the gold-shaded area above the dashed gold P50 line represents wind's P50-P1 LCOE range (wind's P99-P50 range, which includes all lower-than-expected LCOEs resulting from better-than-expected wind resources, is omitted due to the difficulties described earlier in translating better-than-expected wind resources to appropriate LCOEs when way out at the tail end of the distribution where power curves and capacity factors level off). Astute readers will notice that in Figure 8, wind's P1 LCOE curve is derived from the P99 capacity factors shown earlier in Figure 1. This makes sense: all else being equal, a worse-than-expected capacity factor results in a higher-than-expected LCOE. Yet the use and lexicon of *probability of exceedance* requires that the higher-than-expected LCOE be re-labeled as a P1, rather than a P99, LCOE.

The x-axis requires additional explanation to prevent confusion over the treatment of time. Every data point shown on Figure 8 and the graphs that follow—regardless of where it falls along the x-axis—represents an LCOE that is calculated over a 25-year period (in nominal dollars). These 25-year LCOEs are based on the modeling inputs shown in Table 3, with two exceptions—the wind project's capacity factor and the gas-fired generator's levelized fuel costs vary across different P-levels *and* different time horizons (as per Figures 1 and 7). The x-axis simply represents the time horizon over which these two important, but uncertain, inputs into the 25-year LCOE calculation are considered.[33] For example, at

[33] Although investment decisions about new generation sources are typically made with long-term time horizons (e.g., 20-25 years) in mind, the shorter-term time horizons that are also covered by this new framework may nevertheless be relevant in certain situations. For example, a utility might have a temporary or short-term (e.g., 5-year) need for energy, in which case the price stability offered by wind or solar likely won't be as important of a consideration as it is over longer time horizons. Similarly, ratepayers—i.e., those who ultimately pay for the electricity generated by the project—may have a relatively short-term focus if, for example, they are not certain how long they will reside (or remain in business) within the utility's service territory. This framework reflects these simple truths. Varying both the P-levels *and* time horizons for the two uncertain inputs allows a user who is concerned about resource risk to make an informed decision in all cases, regardless of the time horizon.

year 12 on the x-axis, wind's 25-year LCOE range reflects 12-year P50 and P99 capacity factors used as inputs to the 25-year LCOE calculation; similarly, the range of gas-fired LCOE reflects 12-year P1 and P99 gas price projections that are levelized over 12 years and then used as the fuel price inputs in the 25-year LCOE calculation.

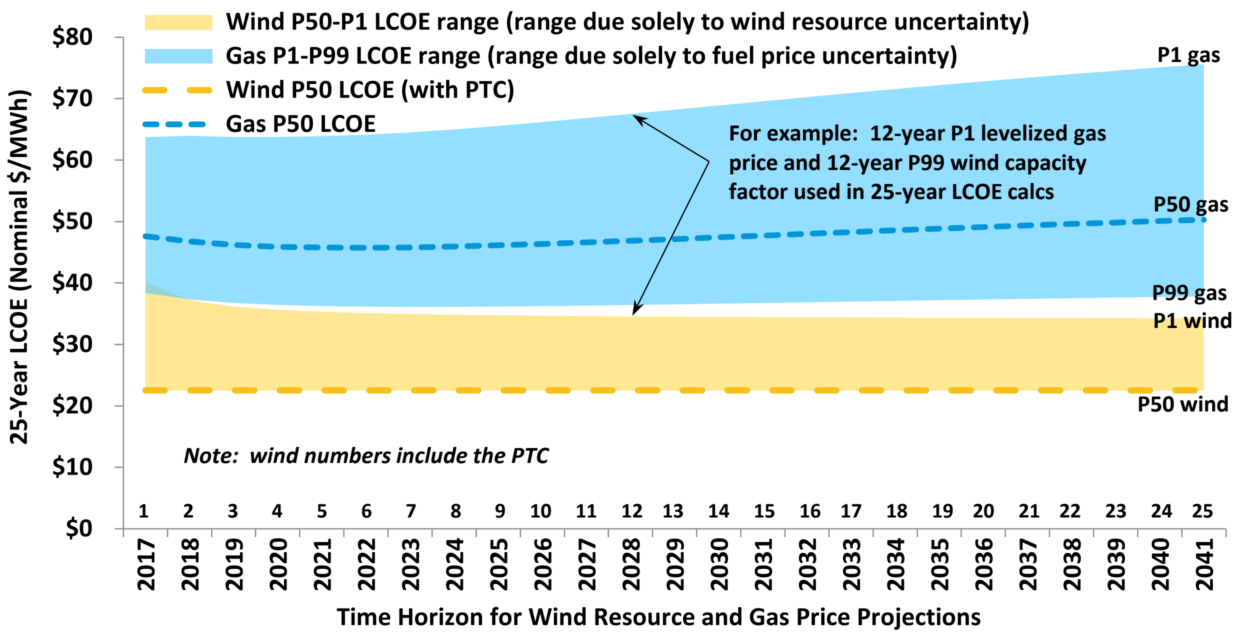

Figure 8. P50-P1 wind (with PTC) and P1-P99 gas-fired LCOE ranges over different time horizons
Note: Wind numbers include the full PTC as it existed on 12/31/2016

Figure 8 shows that with the PTC, wind is currently *very* competitive with new gas-fired generation. In fact, in this example, wind's P1 (i.e., worst-case) LCOE is below the gas-fired generator's P99 (i.e., best-case) LCOE for all time horizons longer than two years. In other words, at least in this example, there is essentially zero probability that gas-fired generation will have a lower LCOE than wind (with the PTC). This low probability helps to explain the recent surge in both utility and corporate purchases of wind power among buyers looking to lock in low prices (before the PTC phases out) and hedge future price risk (as described in the text box on page 4).

Without the PTC—i.e., a more likely scenario in a few years once the PTC phase-out is complete—Figure 9 shows a more even comparison between wind and gas-fired generation, in this case focusing only on worse-than-expected outcomes, and with particular attention paid to the P25, P10, P5, and P1 projections (depicted by the thin solid lines). Not surprisingly, the P50-P1 range of wind LCOE starts out wide but narrows over longer time horizons (similar to the range of wind capacity factors shown earlier in Figure 1), while the P50-P1 range of gas-fired LCOE starts out narrow but widens over longer time horizons (similar to the range of fuel prices shown earlier in Figures 6 and 7). Also evident from Figure 9 is that, over longer time horizons, the uncertainty surrounding the LCOE of gas-fired generation swamps the uncertainty surrounding the LCOE of wind. For example, the P50-P1 range of gas-fired LCOE is more than $25/MWh wide when considering the full 25-year gas price projection. In contrast, the same P50-

P1 range of wind LCOE (when considering the 25-year wind resource projection) is less than half as wide, at less than $12/MWh. Said another way, worse-than-expected outcomes with equal probability—e.g., P5 scenarios—have a much larger negative impact on the LCOE of gas-fired generation than they do on the LCOE of wind generation.

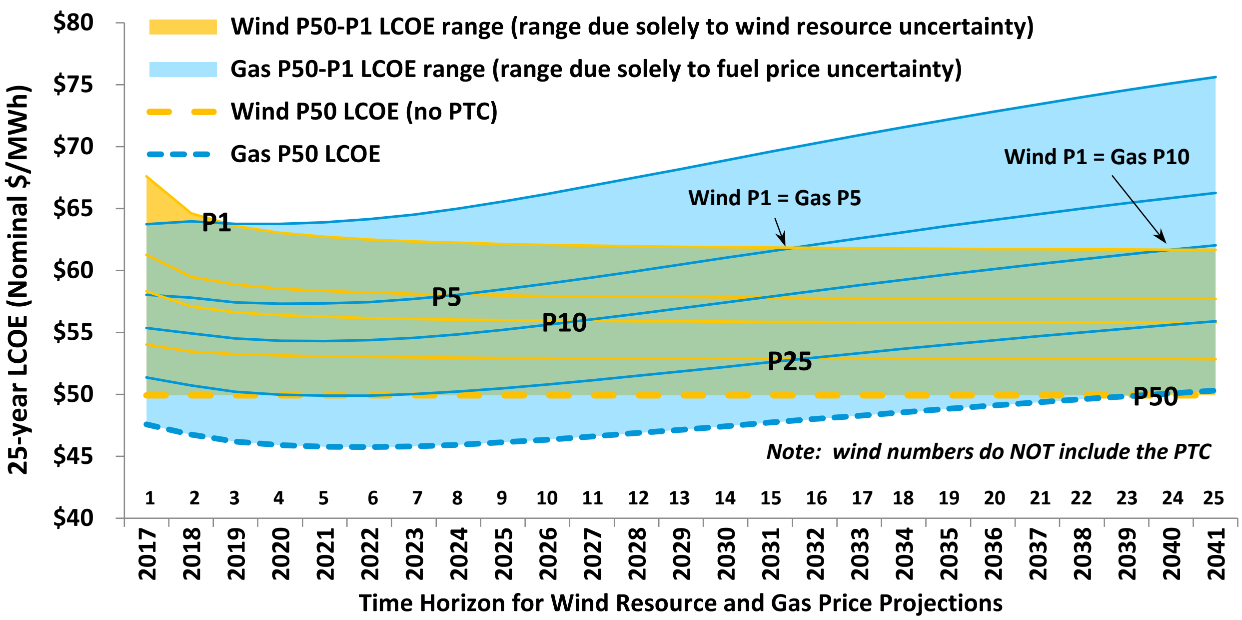

Figure 9. P50-P1 wind (no PTC) and gas-fired LCOE ranges over different time horizons

In Figure 9, wind (without the PTC) is more expensive than gas-fired generation on a P50 basis over all time horizons of less than 24 years (the two P50 curves converge at 24 years). But on a P25 basis, the cost of wind falls below the cost of gas-fired generation for all time horizons longer than 16 years. This “break-even” point—where the wind and gas-fired LCOE curves for each P-level cross—drops to 10, 8, and 2 years for P10, P5, and P1 values, respectively. In other words, Figure 9 presents an illustrative example where wind, without the PTC, is not competitive with new gas-fired generation (except over time horizons longer than 24 years) when evaluated on a P50 basis as is typically done. But when considering the possibility of worse-than-P50 outcomes, wind looks more competitive (particularly the lower the P-level and the longer the time horizon) and in many cases is cheaper than gas-fired generation.[34] The “wedges” that begin where the respective wind and gas-fired LCOE curves at each P-level cross and then widen over longer time horizons illustrate wind’s “hedge value,” which increases with both the level of risk aversion (represented by, and negatively correlated with, the P-level) and the time horizon.

[34] Also notable in Figure 9 is that lower-probability wind resource projections still, in some cases, yield lower LCOEs than higher-probability gas price projections. For example, over time horizons exceeding 15 years, P1 wind is cheaper than P5 gas, while beyond 24 years P1 wind is cheaper than P10 gas.

Figure 10 shows a similar P50-P1 comparison between wind and gas-fired generation, but with two key differences. First, the wind LCOE now reflects the benefit of the PTC, and so is markedly lower. Second, the natural gas numbers reflect only operating costs ("OpEx"), which includes variable O&M as well as fuel costs (but not fixed O&M or CapEx), and so are also markedly lower. Whereas the previous comparison shown in Figure 9 is intended to reflect a future environment in which new wind generation will not have access to the PTC and will be competing (at least in some cases) against new gas-fired generation, the comparison shown in Figure 10 is instead meant to reflect current conditions, whereby new wind generation still likely has access to the PTC (by meeting "start construction" deadlines by the end of 2016) but is competing primarily against *existing* gas-fired generators at their marginal operating costs (under the assumption that existing gas-fired generators have already recouped their capital costs).

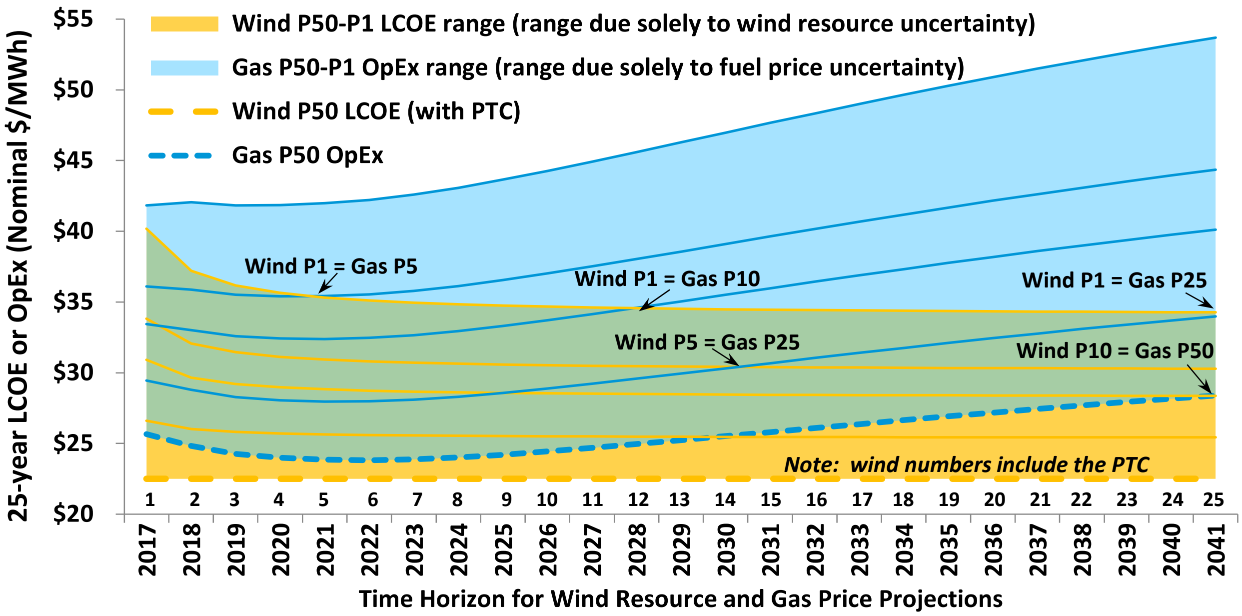

Figure 10. P50-P1 wind LCOE (with PTC) and gas-fired OpEx over different time horizons
Note: Wind numbers include the full PTC as it existed on 12/31/2016

In this case, wind is projected to be cheaper than gas-fired generation at all P-levels over all time horizons. In addition, as shown by the various notations on the graph, wind at lower P-levels outcompetes gas-fired generation at higher P-levels over various time horizons—e.g., the P1 wind LCOE is cheaper than the P10 gas OpEx when evaluated over time horizons longer than 12 years, and the P10 wind LCOE is similar to the P50 gas OpEx when evaluated over a 25-year time horizon. In other words, wind scenarios that are more conservative (probabilistically) than gas scenarios still yield cheaper power in this example.

Figure 11 shows yet another comparison, this time between new solar (with the 30% ITC) and new gas-fired generation (CCGT). In this example, the renewable resource (utility-scale PV) is more expensive than new gas on a P50 basis over all time horizons shown, but more risk-averse comparisons conducted on a worse-than-P50 basis reveal that solar (like wind) can provide significant hedge value. For example, on a P25 basis, the solar and gas LCOEs are equivalent when evaluated over a 25-year horizon, while P10 solar is cheaper than P10 gas over all time horizons longer than 16 years. More conservative solar P-levels also out-compete less-conservative gas P-levels—e.g., P1 solar is cheaper than P5 gas when evaluated over time horizons longer than 21 years.

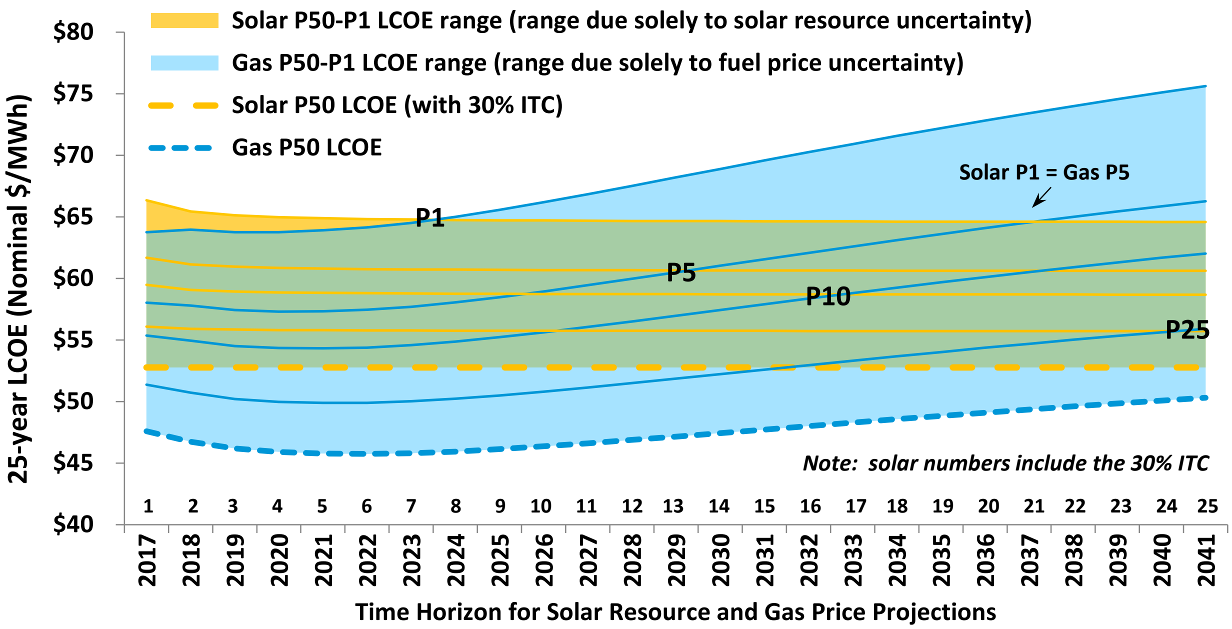

Figure 11. P50-P1 solar (with 30% ITC) and gas-fired LCOE ranges over different time horizons
Note: Solar numbers include the 30% ITC

In Figures 9-11, the "wedges" that start at the point where the wind or solar LCOE initially crosses the gas-fired LCOE for any given P-level and then widen out over longer time horizons are a visual representation of wind and solar's hedge value. Although these "hedge wedges" are discernible from the discrete P-level curves plotted in Figures 9-11, Figure 12 on the next page nevertheless depicts them more clearly—in this case, for the same wind scenario shown earlier in Figure 9—by removing other clutter that potentially obscures the wedges. Each of these "hedge wedges" shows how much cheaper wind is projected to be than gas-fired generation over a range of time horizons and based on a particular P-level comparison (as labeled). There are also a few overlapping portions of wedges where wind at a lower P-level is projected to be cheaper than gas at a higher P-level. Not surprisingly, the lower the P-level, the earlier the time horizon at which hedge value begins to accrue, and the greater the hedge value that exists over the full 25-year time horizon.

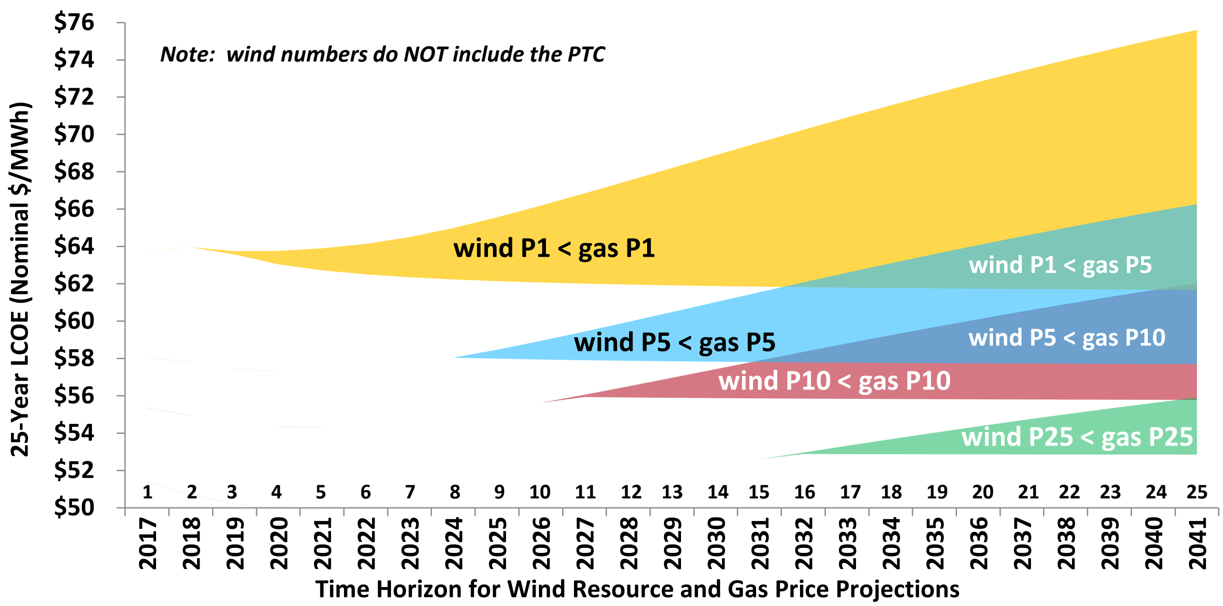

Figure 12. Wind's "hedge value" over different time horizons and at different P-levels (no PTC)

As explained earlier in footnotes 14 and 33, there is good reason to vary the time horizon along the x-axis, as is done in Figures 8-12. Yet when considering investments in *new* power plants in particular, the most relevant time horizons will likely be over longer terms. With this in mind, and in an attempt to consolidate results from Figures 8-11 onto a single graph, Figure 13 shows the cost difference between gas-fired and renewable generation across five common P-levels and considering only 25-year time horizons. Regardless of whether or not renewables are cheaper than gas (indicated by a positive value) on a P50 basis, all four comparisons show renewable generation to be at least at parity with gas-fired generation on a P25 basis, and progressively more competitive at even lower P-levels.

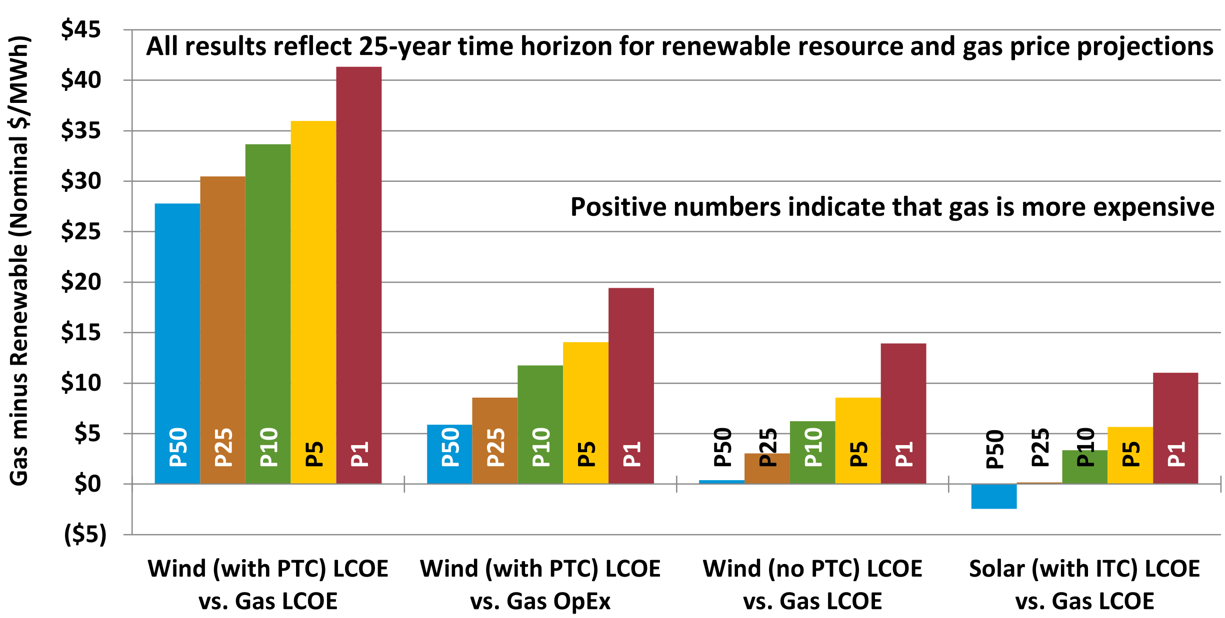

Figure 13. Consolidated results for five common P-levels over 25-year time horizon

Finally, it is worth highlighting that Figures 8-13 compare the LCOE of *individual* wind or solar projects against the LCOE of a *single* combined cycle gas plant. As explained in the text box below, however, one could also assess the impact of resource risk across a portfolio of wind or solar projects (or even across a mixed portfolio of wind and solar projects), which—due to the benefits of diversification—would result in lower overall uncertainty in annual energy production ("AEP") than if assessing each project in the portfolio individually. There is no comparable reduction in gas price risk across a portfolio of gas-fired generators.

Reducing renewable resource risk through portfolio effects and operational energy assessments

The uncertainties in annual energy production ("AEP") that are modeled in this report—i.e., 8.8% for solar and 11.2% for wind in any single year, dropping to 7.9% and 8.4%, respectively, over any 10-year period—are based on typical pre-construction projections for individual projects. For sponsors that own, or for utilities that purchase power from, a diverse portfolio of wind and/or solar projects, however, AEP uncertainty can be reduced by assessing it across the entire portfolio rather than for each individual project. And if some of the projects in that portfolio have been operating for at least twelve months, AEP uncertainty can be reduced even further by supplementing the pre-construction energy assessment with an operational energy assessment.

Rose and Apt (2015) found that inter-annual variability of the wind resource in the Great Plains states fell from 5-12% at individual sites to 3% when assessed across all sites, due to the portfolio effect created by geographic diversity (recall from Equation 1 that inter-annual variability is just one component of total AEP uncertainty). Likewise, in its SEC filing for an initial public offering, Pattern Energy noted a similar reduction in uncertainty—though in this case expressed as an *increase* in P75 AEP—within its portfolio of wind projects: "the sum of our individual projects' P75 output levels is approximately 92% of the aggregate P50 output level (which is unaffected by the portfolio effect), while the P75 output level, when taking into account the portfolio effect, is approximately 95% of our aggregate P50 output level" (Pattern Energy Group Inc. 2013). AWS Truepower (Optis 2016) finds a similar benefit on a P99 basis, with P99 AEP estimates typically increasing by 3-5% (but ranging from as low as 0.1% to as much as 9.3%, depending on the diversity of the portfolio) when evaluated on a portfolio basis rather than on an individual basis. The credit rating agency Fitch also notes that "the portfolio effect may result in an increase in the aggregate P90 estimate by some 2% to 5%, compared to the sum of the P90s of single projects" (FitchRatings 2016).

Operational energy assessments that are based on actual energy production over a period of (ideally) twelve months or longer can yield further reductions in AEP uncertainty. DNV GL (Smith and Byrne 2016) notes that "Unlike pre-construction energy assessments, energy assessments based on operational data can more accurately account for long-term effects, environmental losses, electrical collection system losses, wake effects and terrain effects of windflow." Although the benefits will vary from site to site and the range is wide, DNV GL estimates that an increase in the one-year P99 energy value of 5-6% on average is reasonable (Byrne 2017). On the other hand, an operational energy assessment may also find that the pre-construction P50 AEP estimate has overstated actual production; in fact, there has been a highly-publicized history of overly optimistic P50 projections within the U.S. wind industry (FitchRatings 2014, Mason 2013, Bernadett et al. 2012, Allevato 2011, Jones 2008, Johnson et al. 2008, Randall 2008). Even in these cases, however, the reduced uncertainty surrounding the downwardly revised P50 AEP estimate may still result in an improvement in the P99 AEP estimate relative to the pre-construction estimate (Lightfoote 2016).

As such, the probabilistic LCOE values calculated for wind and solar in this chapter (and in particular those at higher P-levels) could potentially be reduced somewhat if considering a portfolio of projects rather than individual projects, and/or if operational energy assessments were considered. For example, in the case of a wind project without the PTC (e.g., as shown by Figure 9), a 5% increase in the 1-year P99 AEP estimate—which is seemingly achievable through either the portfolio effect or operational energy assessments—reduces the corresponding 25-year LCOE by roughly $3/MWh. In contrast, the LCOE values for gas-fired generation are unlikely to benefit from either portfolio effects or operational energy assessments—neither of which is likely to impact the market price of natural gas.

4.3 Numerical results

While the visual manifestations of hedge value provided in Figures 8-13 are intuitively useful, of potentially greater value would be a way in which to convert these visualizations into numerical or quantitative assessments that facilitate decision-making. The framework developed herein provides quite a bit of flexibility in how one might do this.

First, one could simply compare the LCOE of two sources of generation at whatever P-level is of interest. This approach mirrors what already happens by default in most instances, where LCOE is typically compared solely on a P50 basis. But the framework developed here also enables more risk-averse comparisons of worse-than-P50 outcomes, such as comparing the two resources on a P25, P10, P5, or P1 basis, depending on the desired level of risk aversion. For example, Table 4 shows the difference in 25-year LCOE between new wind (without the PTC) and gas-fired generation, as depicted earlier in Figure 9 (the numbers are calculated as gas minus wind, so a positive number in Table 4 means wind is cheaper than gas-fired generation, and vice versa). The top five rows of Table 4 compare these two resources at the same P-level. For example, when considering the 20-year wind resource and gas price projection, wind is \$0.8/MWh more expensive than gas-fired generation on a P50 basis, but is \$1.5/MWh cheaper than gas-fired generation on a P25 basis and \$11.1/MWh cheaper on a P1 basis.

Table 4. Different ways to compare the 25-Year LCOE of wind and gas within this framework

25-Year LCOE Difference (gas-wind) Nominal \$/MWh	**5 Years**	**10 Years**	**15 Years**	**20 Years**	**25 Years**
Same P-Level Comparisons					
P50	-4.2	-3.6	-2.2	-0.8	0.4
P25	-3.2	-2.1	-0.3	1.5	3.1
P10	-1.9	-0.3	2.0	4.3	6.2
P5	-1.0	1.0	3.7	6.3	8.6
P1	1.2	4.1	7.7	11.1	13.9
Different P-Level Comparisons					
P10 gas – P1 wind	-8.4	-6.4	-4.0	-1.6	0.4
P5 gas – P1 wind	-5.4	-3.1	-0.3	2.3	4.6
P10 gas – P5 wind	-4.0	-2.3	0.1	2.4	4.3
P10 gas – P25 wind	1.2	2.7	5.0	7.2	9.2
Probability-Weighted Comparisons					
Some Negative Outcomes (P50-P25)	-3.8	-3.0	-1.5	0.1	1.4
All Negative Outcomes (P50-P1)	-3.5	-2.6	-0.9	0.8	2.2

Note: All comparisons in Table 4 pit new wind (without the PTC) against new gas-fired generation on a 25-year LCOE basis, considering the wind resource and gas price projections over the five different time periods shown. The results in Table 4 follow directly from the values depicted in Figures 9 and 12.

Alternatively, as noted earlier, one could also compare these two resources across different P-levels if so desired—e.g., P1 wind versus P5 gas, or even P25 wind versus P10 gas (i.e., these comparisons can be made with wind at either a higher or lower P-level than gas-fired generation, depending on the user's goals and motivation).[35] The middle rows of Table 4 show four such cross-P-level comparisons.

Comparisons across different time horizons are also possible *and* important given that the relative competitiveness of resources changes depending on the time horizon over which the wind resource and gas price projections are considered. For example, focusing on the "P10 gas– P1 wind" comparison in Table 4, wind is more expensive by $4.0/MWh over a 15-year time horizon, but is cheaper by $0.4/MWh over a 25-year time horizon.

Though certainly workable, the simple (and somewhat subjective) comparisons at the same or different P-levels described in the previous paragraphs fail to take advantage of the probabilistic nature of this framework. By definition, each P-value has an associated probability, thereby enabling a more formal probabilistic assessment. For example, although probability of exceedance does not necessarily imply probability of occurrence, the P50 outcome can nevertheless be thought of as carrying a 50% weight (which is a compelling reason not to completely ignore it by focusing solely on some other lower-probability P-value). Similarly, the P1 outcome can be given a 1% weight, with all other P-values that fall in between these two extremes (e.g., P49, P48, P47…P4, P3, P2) weighted accordingly (i.e., 49%, 48%, 47%...4%, 3%, 2%). Hence, within this framework, one can easily "probability-weight" the full range of outcomes across the full P50-P1 spectrum, or even some subset thereof – e.g., perhaps just the P50-P25 range for those who are less risk averse.

The bottom two rows of Table 4 show these two "probability-weighted" comparisons. When considering the wind resource and gas price projection over a 25-year period, wind is $0.4/MWh cheaper than gas-fired generation on a straight P50 basis (top row of Table 4), but is $1.4/MWh cheaper than gas when all outcomes ranging from P50-P25 (inclusive) are "probability-weighted," and is $2.2/MWh cheaper when all outcomes from P50-P1 (inclusive) are so weighted. In other words, moving from a straight P50 comparison to a "probability-weighted" P50-P1 comparison makes wind $1.8/MWh more attractive in this example.

Finally, though not shown in Table 4, one can also use this framework as a tool to answer a variety of customized "what if" questions. For example, harkening back to the utility-scale PV comparison shown in Figure 11, one can use this framework to answer questions like how much (in terms of P-level) would natural gas prices have to rise before gas-fired generation is no longer cheaper than P50 solar over a 20-year time horizon?[36] Or, as already noted in the discussion surrounding Figure 8, one can use this framework to answer questions such as what is the probability (considering only resource risk) of new gas-fired generation having a lower LCOE than wind (with the PTC) over a 25-year time horizon?[37]

---

[35] Comparisons across different P-levels acknowledge that wind (or solar) and gas-fired generation will seldom reach P25 or P10 or P5 levels (for example) simultaneously, and even if or when they do, the relative *impact* of each occurrence could be significantly different. For example, gas-fired generation reaching the P10 level may be significantly more harmful to ratepayers than wind or solar reaching the P10 level, in which case a comparison at the same P-level may not be appropriate. This framework allows the user to pick and choose different P-levels for each resource based on individual preferences, level of risk aversion, and perceived impact of various probabilistic outcomes.

[36] Answer: Gas prices would need to rise to P31 levels.

[37] Answer: 0% under the modeling parameters used for Figure 8.

5. Summary and conclusions

Of the myriad risks surrounding long-term investments in power plants, resource risk is one of the most difficult to mitigate, and is also perhaps the one that most-clearly distinguishes renewable generation from gas-fired generation. For renewable generation, resource risk is a quantity risk whose distribution narrows over longer time horizons; for baseload gas-fired generation, resource risk is primarily a price risk whose distribution broadens over time. Most often, resource risk—and natural gas price risk in particular—falls disproportionately on utility ratepayers, who in general are not well-equipped to manage this risk. As such, it is incumbent upon utilities, regulators, and policymakers to ensure that resource risk is taken into consideration when making or approving resource decisions, or enacting policies that influence the development of the electricity market more broadly. This paper presents a new framework, grounded in statistical concepts related to probability of exceedance (and confidence intervals more broadly), to incorporate resource risk into decision-making processes.

This framework recognizes that the same probability of exceedance concepts that are regularly used to express the uncertainty around annual energy production for wind and solar projects can also be applied to natural gas price projections, allowing one to develop a probabilistic range of projections for not only wind and solar capacity factors, but also natural gas prices. Importantly, these probability distributions have markedly divergent characteristics. Renewable resource risk is symmetrical about the mean or "P50" projection and declines when considered over longer time horizons (due to mean reversion in the inter-annual variability of the resource). In contrast, natural gas price risk is asymmetrical (skewed towards higher prices) and increases when considered over longer time horizons (reflecting the fact that it is easier to project where natural gas prices will be three months from now than three years from now). Converting these distinctly different probability distributions into directly comparable LCOE terms reveals that even when gas-fired generation is competitive with or cheaper than wind and solar power on an expected (P50) basis—the basis on which these resources are most often compared—comparisons that are instead based on worse-than-expected (e.g., P25 or P1) outcomes often reach the opposite conclusion: that wind and solar are cheaper than gas-fired generation.

Though not without shortcomings (e.g., the gas price ranges are purely statistical, and are derived from historical rather than implied volatility), this new framework nevertheless has a number of potential advantages over many of the previous approaches that have been proposed (and that are described earlier in the text box on page 2) to account for the benefit of price certainty that wind and solar power can provide. For example, this new framework is:

- **Fair:** Unlike most other approaches, which consider wind and solar to be "riskless" (de facto assuming only P50 outcomes for wind and solar), this new framework recognizes that wind and solar also face resource risk (though in terms of quantity, not price). As such, it places wind and solar on more equitable terms as natural gas with respect to resource risk, leading to a fairer comparison.

- **Familiar:** This new framework is grounded in probability of exceedance (and confidence intervals more broadly)—i.e., concepts that are well-known to energy analysts and policymakers alike, and that are already widely used in the renewable energy and natural gas industries to capture uncertainty.
- **Simple:** The calculations require just a few key and accessible parameters. For wind and solar, required inputs include only estimates of energy uncertainty or, alternatively, a table of P-values for annual energy production—either of which should be readily available from each project's wind or solar resource study. For natural gas, a central price projection (e.g., the futures strip) along with estimates of price volatility over different time horizons (whether estimated from historical price data or implied by options prices) are all that are needed.
- **Flexible:** This framework caters to any level of risk aversion (e.g., from P50 to P1) over any time horizon (e.g., from 1 year to 25 years or longer), allowing the user to pick and choose the combination that most closely aligns with present needs. While long-term time horizons will often be the most appropriate in utility settings, short-term comparisons may also be useful in some case—e.g., when a utility needs to decide how to fill a temporary energy deficit.
- **Intuitive:** The visual representations of hedge value (e.g., the "hedge wedges" highlighted in Figure 12) are intuitively easy to grasp.
- **Probabilistic:** Perhaps the greatest advantage over other methods is that this framework is grounded in statistics, thereby enabling the user to probability-weight the outcomes rather than relying on more arbitrary decision rules like comparing resources across just a few specific P-levels.

Regardless of its relative advantages or shortcomings, however, this framework reaches similar conclusions to all that have come before it. Namely, that gas-fired generation is riskier than renewable generation when it comes to resource risk, and that accounting for this risk in cost comparisons—whether through the probabilistic framework presented here or using some other method—will favor renewable generation. This is particularly the case the longer the time horizon and the greater the level of risk aversion.

Of course, cost is only one side of the equation (value being the other), and few if any resource decisions within the electricity sector are made solely on the basis of LCOE. Instead, the cost of competing resources must be considered along with the value that each provides, which is most often determined by sophisticated models that endogenously assess energy and capacity value as well as integration and transmission costs—all in addition to the LCOE of the generator itself. In this sense, it should be recognized that this report has focused on just one side of a two-sided coin.

References

Allevato, F. 2011. *Wind Power Project Underperformance White Paper*. DNV Report No.: IEWP0101-B. May 20, 2011.

Allianz. 2016a. "Allianz Risk Transfer and partners develop innovative swap solution to hedge volatile revenues of wind farms." Press release, May 5, 2016. http://www.agcs.allianz.com/about-us/news/art-innovative-swap-solution-wind-farms/ (accessed January 18, 2017)

Allianz. 2016b. "Allianz Risk Transfer and Apex Clean Energy Close Proxy Revenue Swap for Old Settler Wind Farm." Press release, July 25, 2016. http://www.agcs.allianz.com/global-offices/united-states/news-press-releases/allianz-risk-transfer-and-apex-clean-energy-close-proxy-revenue-swap-for-old-settler-wind-farm/ (accessed January 18, 2017)

Anderson, M. 2010. "Minimising Uncertainty in Production Estimates." July 15, 2010.

Awerbuch, S. 1993. "The Surprising Role of Risk in Utility Integrated Resource Planning." *The Electricity Journal* 6 (3): 20-33.

________. 1994. "Risk-Adjusted IRP: It's Easy!!!" In *Proceedings of the NARUC-DOE Fifth National Conference on Integrated Resource Planning*, Kalispell, Mont., 228-269.

________. 2003. "Determining the Real Cost: Why Renewable Power is More Cost-Competitive Than Previously Believed." *Renewable Energy World*, 6(2), March-April 2003.

Awerbuch, S. and M. Berger. 2003. *Energy Security and Diversity in the EU: A Mean Variance Portfolio Approach.* EET/2003/03. Paris: International Energy Agency (IEA).

Bailey, B. and J. Kunkel. 2015. "The Financial Implications of Resource Assessment Uncertainty." *North American Windpower*, Vol. 12, No. 3, April 2015.

Bazilian, M. and F. Roques. 2008. *Analytical Methods for Energy Diversity and Security*. Amsterdam: Elsevier Science.

Bernadett, D., M. Brower, S. Van Kempen, W. Wilson, B. Kramak. 2012. *2012 Backcast Study: Verifying AWS Truepower's Energy and Uncertainty Estimates*. AWS Truepower LLC, May 2012.

Black & Veatch. 2011. "Black & Veatch's Methodology for Calculating Confidence Intervals for Solar System Performance Estimates." Unpublished memo, July 20, 2011.

Bolinger, M. 2013. *Revisiting the Long-Term Hedge Value of Wind Power in an Era of Low Natural Gas Prices*. Berkeley, CA: Lawrence Berkeley National Laboratory. LBNL-6103E.

Bolinger, M. and J. Seel. 2016. *Utility-Scale Solar 2015: An Empirical Analysis of Project Cost, Performance, and Pricing Trends in the United States*. LBNL-1006037. Berkeley, Calif.: Lawrence Berkeley National Laboratory.

Bolinger, M.; Wiser, R.; Golove, W. 2006. "Accounting for Fuel Price Risk When Comparing Renewable to Gas-Fired Generation: The Role of Forward Natural Gas Prices." *Energy Policy* (34:6); pp. 706–720.

Bolinger, M. and R. Wiser. 2005. *Balancing Cost and Risk: The Treatment of Renewable Energy in Western Utility Resource Plans*. LBNL-58450. Berkeley, Calif.: Lawrence Berkeley National Laboratory.

Brower, M.C., K. Bell, P. Spinney, S. Bernow, M. Duckworth. 1997. *Evaluating the Risk-Reduction Benefits of Wind Energy*. A Report to the U.S. Department of Energy Wind Technology Division.

Byrne, A. 2017. E-mail communication with Alex Byrne of DNV GL, February 28, 2017.

Comstock, K., G. Randall, B. Day. 2011. “Tips for Making the Uncertain Less Costly.” *North American Windpower*, Vol 8, Number 10, November 2011. pp 1, 31-36.

Costello, K. 2016. *Vertical Arrangements for Natural Gas Procurement by Utilities: Rationales and Regulatory Considerations*. National Regulatory Research Institute. Report No. 16-04, February 2016.

Energy Information Administration (EIA). 2016. *Short-Term Energy Outlook December 2016*. Washington, DC: Energy Information Administration.

_________. 2017. *Annual Energy Outlook 2017*. Washington, DC: Energy Information Administration.

Feldman, D., and M. Bolinger. 2016. *On the Path to SunShot: Emerging Opportunities and Challenges in Financing Solar*. Golden, CO: National Renewable Energy Laboratory. NREL/TP-6A20-65638. http://www.nrel.gov/docs/fy16osti/65638.pdf

FitchRatings. 2016. “Rating Criteria for Wind Projects.” *FitchRatings Global Infrastructure & Project Finance*, March 31, 2016.

FitchRatings. 2014. “Wind Projects: High Risk of Production Shortfalls.” *FitchRatings Global Infrastructure & Project Finance Special Report*, June 12, 2014.

Iowa Utilities Board. 2016. “Order on Rehearing, In Re: MidAmerican Energy Company,” issued February 6, 2015, Docket No. RPU-2014-0002. https://efs.iowa.gov/cs/groups/external/documents/docket/mdaw/mjky/%7Eedisp/292948.pdf (accessed January 18, 2017)

Jenkin, T., V. Diakov, E. Drury, B. Bush, P. Denholm, J. Milford, D. Arent, R. Margolis, R. Byrne. 2013. *The Use of Solar and Wind as a Physical Hedge against Price Variability within a Generation Portfolio*. NREL/TP-6A20-59065, National Renewable Energy Laboratory, Golden, CO.

Johnson, C., A. Tindal, K. Harman, A.M Graves. 2008. “Validations of Energy Predictions by Comparison to Actual Performance. Presentation at WINDPOWER 2008, Houston, TX. June 4, 2008.

Jones. S. 2008. “Project Underperformance: 2008 Update.” Presentation at WINDPOWER 2008, Houston, TX. June 4, 2008.

Lazard. 2016. *Lazard’s Levelized Cost of Energy Analysis—Version 10.0*. December 2016. https://www.lazard.com/media/438038/levelized-cost-of-energy-v100.pdf (accessed January 23, 2017)

Lightfoote, S. 2016. “Operational Energy Assessments: Updated Methods to Better Predict Long-Term Energy Production.” AWS Truepower webinar, October 2016.

Mason, J. 2013. *2013 Actual versus Predicted Wind Power Project Performance White Paper*. DNV KEMA Report No.: IEWP0101-E. May 20, 2013

Matsui, R. 2017. “The Danger of Solar’s Hot-Hand Fallacy.” Greentech Media. January 13, 2017. https://www.greentechmedia.com/articles/read/the-danger-of-solars-hot-hand-fallacy (accessed January 18, 2017)

Matsui, R., J. Kaminsy, and J. Blanton. 2016. "New Product: Solar Revenue Puts." Chadbourne & Parke LLP's *Project Finance Newswire*, October 2016, pp. 18-22.

Millstein, D., R. Wiser, M. Bolinger, G. Barbose. 2017. "The Health and Environmental Benefits of Wind and Solar Energy in the United States, 2007-2015." Berkeley, Calif.: Lawrence Berkeley National Laboratory. https://emp.lbl.gov/publications/health-and-environmental-benefits (accessed March 23, 2017)

Mortenson Construction. 2016. "Wind Energy: Challenge Accepted (Thriving in Post-PTC Era)."

NREL (National Renewable Energy Laboratory). 2016. *2016 Annual Technology Baseline*. Golden, CO: National Renewable Energy Laboratory. http://www.nrel.gov/analysis/data_tech_baseline.html.

Optis, M. 2016 "Portfolio Benefit Analysis: Updated Methods to Incorporate Wind, Solar, and Hydro Projects." AWS Truepower webinar, November 2016

Pattern Energy Group, Inc. 2013. *Form S-1 Registration Statement Under the Securities Act of 1933*. As filed with the Securities and Exchange Commission on August 9, 2013. https://www.sec.gov/Archives/edgar/data/1561660/000119312513329269/d564947ds1.htm (accessed January 21, 2017)

PwC. 2016. *Corporate renewable energy procurement survey insights*. http://www.pwc.com/us/en/sustainability-services/publications/assets/pwc-corporate-renewable-energy-procurement-survey-insights.pdf (accessed January 18, 2017)

Randall, G. 2008. "Approaches to Understanding Project Performance." Presentation at WINDPOWER 2008, Houston, TX. June 4, 2008.

Rose, S. and J. Apt. 2015. "What can reanalysis data tell us about wind power?" *Renewable Energy*, Volume 83, November 2015, pp. 963-969.

Ryan, B. and T. Lidderdale. 2009. "Energy Price Volatility and Forecast Uncertainty." October 2009 *Short-Term Energy Outlook* Supplement, Washington, DC: Energy Information Administration.

Schnitzer, M. and P. Thienpont. 2014. "Is Overproduction Costing You? A review of the impacts of solar resource data on financing and revenues for PV plants" AWS Truepower webinar, November 18, 2014.

Schnitzer, M., C. Thuman, P. Johnson. 2012. "Reducing Uncertainty in Solar Energy Estimates: Mitigating Energy Risk through On-Site Monitoring" AWS Truepower white paper, July 3, 2012. http://www.awstruepower.com/assets/Reducing-Uncertainty-in-Solar-Energy-Estimates1.pdf (accessed January 23, 2017).

Smith, K. and A. Byrne. 2017. "Renew. Reuse. Repower." *North American Windpower*, Volume 13, Number 12, January 2017.

Stirling, A. 2010. "Multicriteria Diversity Analysis: A Novel Heuristic Framework for Appraising Energy Portfolios." *Energy Policy* (38:4); pp. 1622–1634.

_________. 1994. "Diversity and Ignorance in Electricity Supply Investment: Addressing the Solution Rather than the Problem." *Energy Policy* (22:3); pp. 195–216.

Wiser, R. and M. Bolinger. 2016. *2015 Wind Technologies Market Report*. LBNL-1005951. Berkeley, Calif.: Lawrence Berkeley National Laboratory.