The Costs of Wind's Variability: Is There a Threshold?

Managing wind's intermittency entails costs even when wind power supplies a small fraction of load. If electric power systems evolve efficiently as wind capacity grows, the costs of managing intermittency will grow smoothly with increasing penetration, allowing wind power to provide deep reductions in CO_2 emissions at costs that are competitive with other mitigation options.

Joseph F. DeCarolis and David W. Keith

Joseph F. DeCarolis, *who recently obtained a Ph.D. at the Department of Engineering and Public Policy at Carnegie Mellon University in Pittsburgh, recently accepted a post in Research Triangle Park, NC, in the Atmospheric Protection Branch of the Office of Research and Development at the Environmental Protection Agency (decarolis.joseph@epa.gov).*

David W. Keith *holds the Canada Research Chair in Energy and the Environment in the departments of Chemical and Petroleum Engineering and Economics at the University of Calgary, where he works on climate-related energy-technology, science, and related public policy. His recent work has focused on the capture and storage of carbon dioxide and the economics and climatic impacts of large-scale wind power. He can be contacted at keith@ucalgary.ca; his Web site is www.ucalgary.ca/~keith.*
The work described in this article was not funded by the U.S. Environmental Protection Agency, nor has the article been reviewed by the Agency. The contents reflect the views of the authors alone, not those of the Agency, and no official endorsement should be inferred.

I. Introduction

Global wind power capacity is roughly 40 GW, with annual capacity additions approaching 8.2 GW and annual equipment sales exceeding \$9 billion.[1] Construction of wind farms has been driven by government regulation or subsidies in combination with steady declines in unit costs. At good sites, the average cost of wind power at the turbine is currently 4–6 ¢/kWh without credits or subsidies, and advances in turbine design may plausibly reduce the cost to 3 ¢/kWh within two decades.[2] Although wind energy currently serves about 0.1 percent of global electricity demand,[3] it has the fastest relative growth rate of any electric generating technology: capacity has increased by roughly 30 percent annually for the five years ending in 2002.[4]

Two factors—the spatial distribution and intermittency of wind resources—raise the effective cost of wind above the average cost of electricity from a single turbine. In this article, we focus on understanding how the cost imposed by wind's intermittency scales with the amount of wind power in an electric power system. Many authors assert, either implicitly or explicitly, that a threshold exists (expressed as

 1040-6190/\$–see front matter doi:/10.1016/j.tej.2004.12.006

the fraction of demand served by wind energy), below which wind imposes negligible costs on grid operation and above which wind imposes substantial costs. Perhaps the most important role for wind power is in supplying electricity without CO_2 emissions. Long-range energy system models used in climate policy analysis often limit the penetration of wind power in response to carbon constraint using such thresholds. We contend that no such threshold exists. Wind's intermittency imposes non-negligible costs even when wind serves only a tiny fraction of demand, but if the electric power system evolves as wind capacity is added, these costs grow monotonically from zero and need not be prohibitive even when wind serves more than half of demand.

II. Background: Managing Variability in Electric Power Systems

Wind must be converted to electricity where wind resources are located. While not addressed here, the spatial distribution of wind resources will often require long-distance transmission lines that increase the cost of electricity from wind.[5,6] Unlike conventional capacity, wind-generated electricity cannot be reliably dispatched or perfectly forecasted, and exhibits significant temporal variability. The uncontrollable nature of wind makes it less valuable to system operators than dispatchable power. In restructured electricity markets, for example, wind operators choosing to participate in markets for scheduled energy may have to settle schedule deviations at the real-time price, which decreases revenue.[7,8] Such penalties are not simply arbitrary financial mechanisms, but reflect, however imperfectly, the cost of managing variations in wind output.

Even without wind, managing electric supply and demand requires sufficient flexibility to respond to time-varying demand, forecast inaccuracies, and contingencies. Three time scales concern system operators on a day-to-day basis: minute-to-minute, intra-hour (5–60 minute time scale), and inter-hour. System operators typically schedule energy each hour using economic dispatch to meet forecasted demand. The schedule is typically drawn up the day before scheduled dispatch. Sub-hourly differences between scheduled energy and forecasted demand during each hour are met by load-following units that can ramp output quickly to balance supply and demand. In restructured electricity systems, load-following units participate in a real-time (intrahour) market. For example, the New York, New England, and PJM independent system operators (ISOs) determine load imbalance on five-minute intervals and use supply curves to dispatch the load-following units participating in the real-time market.[9] Typically, any generating unit deviating from its schedule must pay the imbalance at the real-time price. Load-following units are also known as spinning reserve because they are synchronized to the grid and either idle or operate at less than full capacity.

Even without wind, managing electric supply and demand requires sufficient flexibility to respond to time-varying demand, forecast inaccuracies, and contingencies.

System operators employ an automatic generation control (AGC) system to manage minute-to-minute load imbalances—an ancillary service known as regulation. Units participating in AGC are equipped with governors that sense a change in frequency and automatically adjust output. Intra-hour dispatch every few minutes allows the units providing regulation to return to their nominal set points. There are three important distinctions between regulation and load-following: (1) regulation takes place over a shorter time scale (minute-to-minute versus every several minutes), (2) load centers have uncorrelated variability on the regulation timescale, but exhibit significant correlation on the load-following time scale, and (3) load-following changes often follow predictable diurnal cycles while regulation does not.[10] These time scales are illustrated in **Figure 1**.

 1040-6190/$–see front matter doi:/10.1016/j.tej.2004.12.006

In order to provide AGC and spinning reserve, some generating units must operate at lower power output than would be dictated by optimal economic dispatch without the requirement to follow changing loads; this adjustment forces the system operator to dispatch higher marginal cost units to make up the difference, which raises the average cost of electricity. Additional costs arise from the degraded efficiency that results when generators are operated at partial power or are forced to follow rapidly changing loads.

In addition to making minor corrections to load forecasts or small schedule deviations, system operators must also have enough generating capacity to meet system contingencies, such as a forced outage of a particular generating unit or transmission line. Operating reserve, which consists of spinning and non-spinning reserves, represents capacity that can be dispatched within minutes to meet demand in the event of a system contingency such as failure of a generating unit. Non-spinning reserves consist of quick-start units that are not operating, but can be brought online in a matter of minutes. The requirements for operating reserves are generally set by deterministic criteria, such as a fraction of the forecasted maximum peak demand, to ensure that they are large enough to compensate the most likely or largest contingencies.

III. Wind at Small Scale

Several analyses suggest that there is a threshold below which wind has a negligible effect on

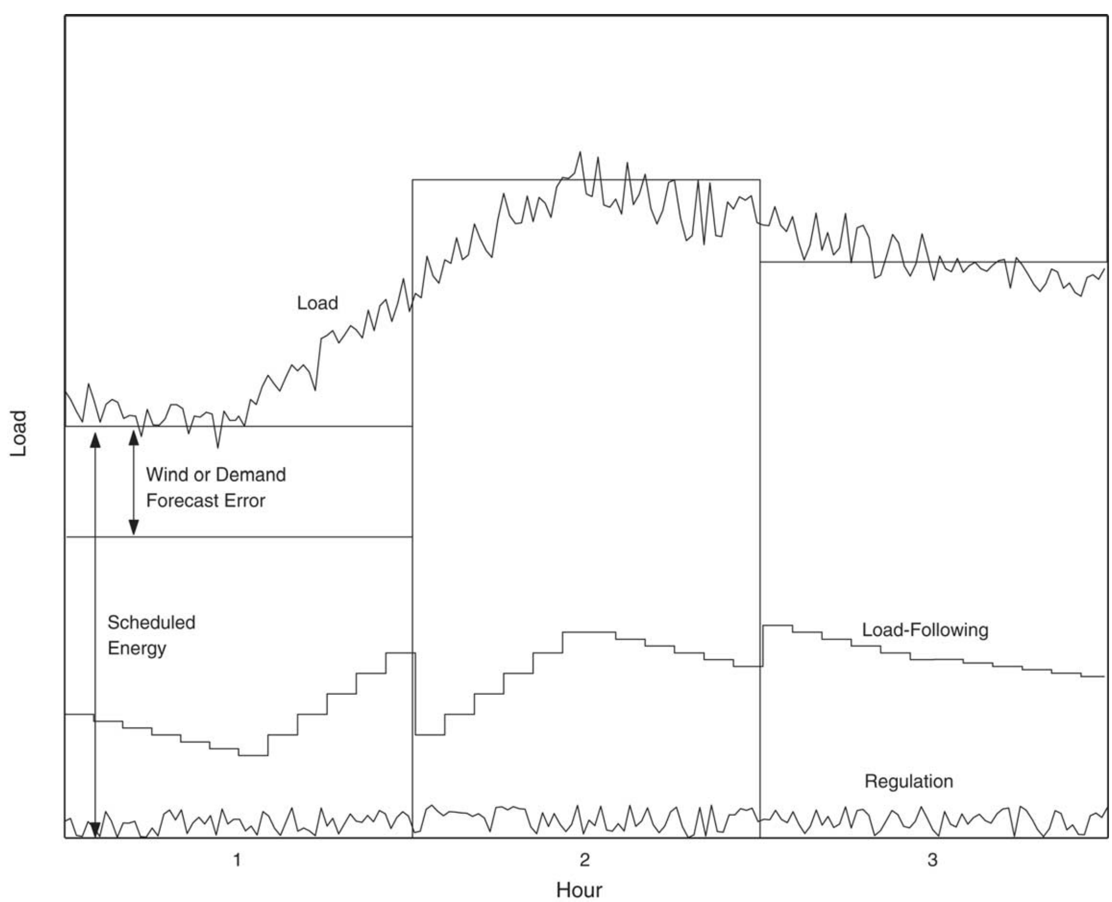

Figure 1: Stylized picture of supply and demand. In most control areas, energy is scheduled ahead of time on an hourly basis according to forecasted demand and unit availability, represented by the three bars. The noisy line represents actual demand and can be separated into the intra-hour load-following and regulation components. Load-following units (spinning reserve) are employed to correct the hourly energy schedule so that supply meets demand on a sub-hourly timescale (every 5–15 minutes), and units equipped with AGC perform regulation to meet the minute-to-minute variability. Regulation and load-following are displayed separately near the bottom of the plot. Inaccuracies in forecasted demand and/or wind can increase the need for load-following capability

 1040-6190/$–see front matter doi:/10.1016/j.tej.2004.12.006

grid reliability, and therefore imposes negligible costs.[11,12,13,14] Richardson and McNerney assert that ''if the generation displacement provided by the wind turbines is within the power-handling capabilities of the load-following units, then wind turbines should not affect system stability.'' Grubb and Meyer claim that ''with no significant measures taken either to make thermal units more flexible, or to predict wind energy better, then serious operational penalties could arise for wind contributions much above 10–15 percent of system energy,'' and also indicate that variability from wind at low levels of penetration are ''drowned out by errors in predicting demand, so there is no operational penalty at low wind penetrations.'' The European Wind Energy Association (EWEA) claims that ''numerous assessments involving modern European grids have shown that no technical problems will occur by running wind capacity together with the grid system up to a penetration level of 20 percent.'' In a final example, van Kuik and Slootweg claim that wind can serve 15–20 percent of electricity demand ''without special precautions to secure grid stability.''

These studies implicitly assume that small-scale wind does not affect reserve capacity and does not produce a measurable effect on grid operations. By this logic, wind's variability imposes no costs until it approaches the limit of the existing system's operating reserve capability. This assumption is unrealistic, however, because as we discussed above, anything that adds variability to load or supply—even if uncorrelated with existing load—will impose additional costs if the same level of reliability is to be achieved. If wind is a very small fraction of load then these costs will be small in absolute terms, but they may still be significant when compared to the cost of wind power itself.

It may be difficult, or impossible, to unambiguously partition the cost of wind's variability between various markets and market participants; it is nevertheless possible, at least in principle, to assess the overall cost of wind's intermittency.

It may be difficult, or impossible, to unambiguously partition the cost of wind's variability between various markets (day ahead, real-time, and regulation) and market participants (producers, consumers, and transmission operators); it is nevertheless possible, at least in principle, to assess the overall cost of wind's intermittency.

Suppose an electric power system without wind supplies electricity at an average cost C_0 while wind power can be supplied at average cost C_W.[15] If wind power had the same temporal characteristics (e.g., dispatchability) as the conventional supply then the average cost of power for the combined system would be a simple linear combination of C_W and C_0 as the fraction of total power supplied by wind was increased. In practice, the average cost of electricity in an optimally dispatched system that combines wind and conventional capacity will rise above the simple linear combination of average costs. The system-level cost of wind's intermittency is the difference between actual costs and the linear average cost line that would apply if intermittency were neglected (**Figure 2**). The effective cost of wind power at the margin—including the cost of intermittency—is the derivative of the total cost curve evaluated at zero wind penetration (line A in Figure 2).

Supporting our assertion, Hirst and Hild find that the revenue received by the wind generators declines smoothly and steadily as the percent of wind serving demand increases and attribute the declining payments to several factors: the addition of supply to a small control area, forecast errors, interhour variability, intrahour energy imbalance, and regulation.[16] The authors estimate the marginal system costs imposed by wind, but do not address the issue of whether existing reserves are sufficient to maintain the pre-wind level of grid reliability. We argue that the portion of aggregate variability attributable to wind ties up a fraction of the

 1040-6190/$–see front matter doi:/10.1016/j.tej.2004.12.006

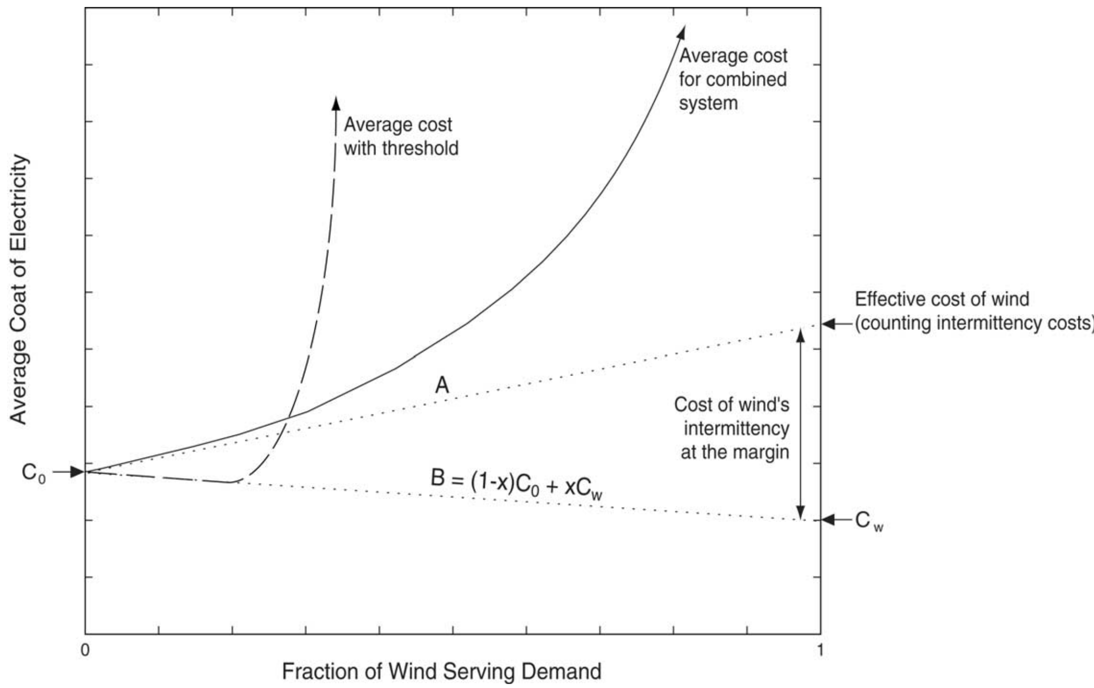

Figure 2: Schematic illustration of the economics of intermittent wind. The vertical axis is the average cost of meeting demand, including both capital and operating costs. The horizontal axis is the total energy supplied by wind divided by the total supplied energy from all generating sources. If wind were dispatchable, then the average cost of power for the combined system would be a simple linear combination of C_W and C_0 as the fraction of total power supplied by wind (x) was increased, as illustrated by line 'B'. Line 'A' includes both the generation cost of wind and the cost of reserve capacity for wind. The curve shows the minimum cost of supplying demand as a function of the amount of wind energy, where we assume that the demand and system reliability are held constant. several studies on the cost of wind power suggest that the cost of intermittency is negligible below some threshold beyond which it rises steeply, as illustrated in the heavy dashed curve

existing regulation and load-following capacity, which reduces the amount of reserve available for system contingencies. If reliability is held constant as wind power is added to the system, this requirement for additional reserve capacity necessarily adds to overall costs.

When wind is a small fraction of demand, operators (sensibly) manage its variability by treating it as negative load, but this does not mean that the cost of variable wind is negligible. Moreover, wind is in several respects more variable than typical loads. At the minute-to-minute or regulation time scale, the AGC requirement can be treated as a random variable with a Gaussian distribution and mean of zero.[17] For a sense of perspective, the regulation component is roughly 0.1 percent of total load in PJM.[18] For comparison, the regulation component for wind in isolation is much larger; one study demonstrates its decline from 10 to 6 percent of rated wind capacity (assuming a 3σ risk level) as the wind capacity grows from 10 to 100 MW.[19] Another study performed in Germany finds that the regulation burden from wind declines from 4.5 to 1 percent of rated wind capacity (or 14.5 to 3 percent assuming a 3σ risk level), for wind capacities of 2.8 and 44.6 MW, respectively.[20] The regulation required for wind grows more slowly than wind capacity because fluctuations on the minute time scale are weakly correlated. In the case of a single wind farm, the minute-to-minute change in each turbine's output is neither perfectly independent nor perfectly correlated with the other turbines. If wind farms are scattered over a large control area, then the regulation requirement for each wind farm is roughly independent of the others, and the total regulation requirement would scale as the square root of the sum of squares from each of the wind farms. For small-scale wind serving less than a few percent of demand, the growth in the regulation requirement for wind can be approximated as linear. But as the level of wind on

 1040-6190/\$–see front matter doi:/10.1016/j.tej.2004.12.006

the system increases, the regulation requirement grows slower than wind capacity and the regulation requirement per unit of wind energy decreases. As such, the cost of regulation—while important—is unlikely to place a strong constraint on the future growth of wind.

Wind is also more variable than typical loads at the inter-hourly load-following time scale, and this can lead to underestimates of the cost of wind's variability. Milligan, for example, employs the 3σ rule as a simple proxy to estimate the hourly load-following requirement for wind.[21] (N.B., the actual amount of AGC and load-following capacity must be sufficient to meet NERC's CPS1 and CPS2 reliability standards, respectively, which translates into a different capacity requirement for each system operator depending on the particular characteristics of the control area.) Analysis of PJM aggregate hourly load data suggests that load-following requirements have a sub-Gaussian distribution in which the actual number of hours which exceed the 3σ-rule is much less than the 0.3 percent that would occur if the variability of load were normally distributed, making the 3σ-rule conservative for loads. Inter-hour changes in wind power, on the other hand, have a super-Gaussian distribution.[22] This result suggests that Milligan's analysis may substantially underestimate the amount of load following capacity necessary to maintain system reliability because wind increases system variance and fattens the tail of the load-following distribution. More generally, it cannot be assumed that wind power time series have the same statistical characteristics as load time series. While Hirst and Hild find that the imbalance charge for intrahour load-following is very modest, even with wind serving ~25 percent of

demand, they acknowledge that reliability will be degraded but do not estimate the cost to upgrade reserves.[23] The cost of adding system reserve to cover the higher variance with wind is real and should be accounted for by system planners.

IV. Wind at Large Scale

The discussion above assumed that, except for marginal additions to capital stock to cover AGC and load following, the electric power system remains static as wind is added. This assumption is reasonable for small amounts of wind, but as the fraction of wind serving demand increases, it becomes less plausible. Because wind serving a substantial fraction (e.g., more than a third) of demand will take (at least) several decades to achieve, the mix of generating units is likely to change significantly during this long period of wind development. Studies that assume wind will grow to serve 20 percent of demand or more while the existing infrastructure remains static may falsely produce a threshold. The dashed curve in Figure 2 represents such a scenario: wind added to a static system does not affect cost until a certain threshold, at which intermittency exceeds the system's operational flexibility, and the cost of electricity rises sharply. Any economic limit on the amount of large-scale wind in a given system will depend on how wind coevolves with the rest of the electric power system. All else equal, the cost of intermittency will be less if the generation mix is dominated by gas turbines (low capital costs and fast ramp rates) or hydro (fast ramp rates) than if the mix is dominated by nuclear or coal (high capital costs and slow ramp rates). In many parts of the world, the rapid growth in gas turbine capacity is likely to continue, thereby supplanting older coal capacity and making the economics increasingly attractive for wind. In a non-static system, low cost reserve can also be added to the wider grid to account for the increased variance from wind.

Three factors lower the economic value of wind as the wind penetration level increases, assuming a static system: (1) the reduced cost of marginal fuels (increasing wind generally saves fuel from progressively lower fuel cost thermal plant), (2) operational losses (repeated plant starts or partial plant loading), and (3) discarded wind energy (primarily due to operational constraints).[24] Discarded wind energy, even without operational constraints, lowers wind's marginal contribution to serving load as the supply of wind energy exceeds demand and is wasted[25]. The effect of discarded wind energy can be seen in Figure 2, where the average cost of wind diverges upward from the line A.

Grubb defines two (somewhat arbitrary) penetration limits: (1) the marginal fuel savings have dropped by one-quarter and (2) the marginal fuel savings have been halved. Grubb considers (1) to be an economic target and (2) to be a ''maximum credible penetration level.'' In terms of the percent of wind energy serving demand, Grubb finds that (1) is 17 percent and (2) is 26 percent for the British system. However, Grubb assumes a static system, and the results would change significantly—increasing or eliminating the threshold—if the rest of the electric power system was free to change as well.

More recently, we investigated the cost of large-scale wind in a non-static system. We used a time-resolved simulation model in which distributed wind farms interconnected via long-distance transmission lines, storage, and gas turbines meet a time-varying load. The installed capacity of various system components was then adjusted to minimize the average cost of electricity under a carbon tax.[26] In this system, cost of intermittency, as defined above, is only 1–2 ¢/kWh when wind serves 50 percent of demand. Our analysis does not, of course, resolve the issue. In addition to using a (relatively)

simplistic electric system model, our analysis assesses greenfield costs, examining an optimal endpoint while ignoring the temporal evolution of the electric power system from a current to future state.

V. Conclusions

Undispatchable wind energy imposes real costs on grid operations, even at the scale of a single wind farm. We posit that these costs increase smoothly and monotonically as the fraction of wind serving demand increases. Studies that assume reserve capacity is free up to a certain threshold are not taking into account the degraded reliability stemming from increased system variance. Even at small scale, wind adds to variable load, which reduces reserve margins by forcing fast-ramping capacity to correct wind-induced imbalances. Threshold arguments for wind are likely to be overly optimistic at low wind penetration levels (by ignoring the degraded reliability stemming from wind intermittency) and overly pessimistic at high wind penetration levels (by assuming that serious operational penalties will suddenly arise in a static system). While it is imperative to consider the system reliability implications of wind at all scales, we do not believe that the addition of operating reserve to the wider grid to counter variable wind will result in prohibitive costs. We stress that the costs imposed by large-scale wind serving more than a quarter of demand cannot be estimated by taking a static system view, but rather will depend on how the underlying system architecture changes over time as the amount of installed wind gradually increases.

We assert that credible estimates of the costs of wind's intermittency must assume that electricity is supplied with the same level of grid reliability with wind as without. While accepting a lower level of reliability could reduce the average cost of supplying electricity with wind power, lower reliability standards would enable roughly equivalent cost savings in the absence of wind. For the same reason, while

 1040-6190/$–see front matter doi:/10.1016/j.tej.2004.12.006

increasing the responsiveness of demand could reduce the overall costs of electric power, such measures entail roughly equal benefits with or without wind. Increasing the responsiveness of demand may make sense, but it is misleading to argue that the costs of wind's intermittency can be reduced simply because lower electricity costs can be achieved by increasing demand-responsiveness or reducing reliability.

The most credible driver for future wind development is a constraint on carbon emissions. Centralized ownership and management, significant experience with regulation, and large, manageable point sources of CO_2 make the electric power sector a prime target for deep cuts in CO_2 emissions. Even with the added cost to deal with intermittency, wind is roughly competitive with other generation technologies under a strong carbon constraint. While air pollution and energy-security are often cited as drivers for wind power, it is less plausible that wind power can provide a cost-competitive means of addressing these challenges.[27]

The role of wind in reducing CO_2 emissions over the long run (decades to a century or more) is addressed by energy-system models that attempt to compute the long-run costs of reducing CO_2 emissions across all economic sectors and energy technologies. Such models are integral to so-called integrated assessment models (IAMs) of climate change that play a central role in debates over long-term climate policy. Such models must necessarily use highly simplified representations of electric power systems and ignore the dynamics of generating system dispatch. These models often assume that there is a strong threshold beyond which wind power becomes uneconomic. In one of the most prominent of such models, for example, the fraction of electricity supplied by wind power is effectively limited to 10 percent.[28]

We suspect that by imposing arbitrary (and generally small) caps on wind power's penetration, such integrated assessment models may greatly understate the potential contribution of wind power to mitigating CO_2 emissions. The outputs of these models, which show comparatively small contributions from wind power, play important roles in debates about appropriate energy policies to manage climate change. It is important to objectively reassess wind's role through critical research on the implications of wind power's variability for large-scale electric power systems; research that connects the typically disparate communities of those who study near-term integration of wind power in existing markets with the community that does long-range energy modeling.

Future research on the intermittency cost of wind should include analysis of high-resolu-

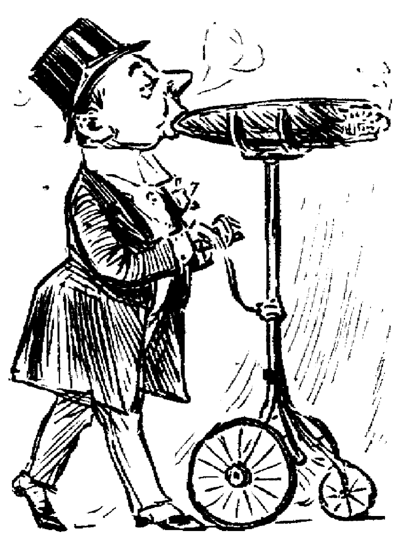

Even with the added costs to deal with intermittency, wind is roughly competitive with other generation technologies under a strong carbon constraint.

  doi:/10.1016/j.tej.2004.12.006

tion demand, supply, and wind power time series, consider plant retirement and the temporal development of the electric power system, and ensure that reliability is held constant as wind is added to the system. An important outcome of such work could be supply curves that provide cost estimates of mitigating carbon emissions with wind that do not impose an exogenous limit on wind development. Such supply curves could serve as input into integrated assessment models to achieve a fairer treatment of wind under a carbon constraint.

Endnotes:

1. American Wind Energy Association, *Global Wind Energy Market Report*. Washington, DC, 2004.

2. Stanley R. Bull, *Renewable Energy Today and Tomorrow*, Proceedings of the IEEE, 89(8), 2001, at 1214–1226.

3. Ralph E.H. Sims, Hans-Holger Rogner, et al., Carbon Emission and Mitigation Cost Comparisons between Fossil Fuel, Nuclear and Renewable Energy Resources for Electricity Generation, Energy Policy, 31, 2003, at 1315–1326.

4. See AWEA, *supra* note 1.

5. Alfred Cavallo, *High-Capacity Factor Wind Energy Systems*, Journal of Solar Energy Engineering, 117, 1995, at 137–143.

6. Tom Factor and Tom Wind, Delivering 2,000 MW of Wind Energy to the Metropolitan Centers of the Midwest, Iowa Dept of Natural Resources.

7. Eric Hirst, *Interactions of Wind Farms with Bulk-Power Operations and Markets*, available at http://www.ehirst.com/publications.html.

8. Eric Hirst and Jeffrey Hild, The Value of Wind Energy as a Function of Wind Capacity, Elec. J., July 2004.

9. See Hirst, *supra* note 7.

10. Brendan Kirby and Eric Hirst, Customer-Specific Metrics for the Regulation and Load-Following Ancillary Services, Oak Ridge National Laboratory, ORNL/CON-474, 2000.

11. R. David Richardson and Gerald M. McNerney, *Wind Energy Systems*, Proceedings of the IEEE, 81(3), 1993, at 378–390.

12. Michael J. Grubb and Niels I. Meyer, *Wind Energy: Resources,*

Systems, and Regional Strategies, in Renewable Energy: Sources for Fuels and Electricity, T.B. Johansson and L. Burnham, Eds. (Washington, DC: Island Press, 1993).

13. European Wind Energy Association, *Wind Force 12*, available at http://www.ewea.org.

14. G.A.M. van Kuik and J.G. Slootweg, *Wind Energy Harvest and Storage Strategies*, in VGB Conference Proceedings, Brussels, Belgium, 2001.

15. Suppose that a perfect lossless electricity storage technology was available. Such a storage system would make wind dispatchable but would not alter the average capacity factor of the wind power system. In this case, the cost of the wind plus free storage, C_W, would be equivalent to the generation cost of intermittent wind, which could approach 3 ¢/kWh in the near future with continued design improvements and cost reductions.

16. See Hirst and Hild, *supra* note 8.

17. Randy Hudson and Brendan Kirby, *The Impact of Wind Generation on System Regulation Requirements*, in Proceedings of Windpower 2001.

18. See Hirst, *supra* note 7.

19. See Hudson and Kirby, *supra* note 17.

20. Bernard Ernst, Analysis of Wind Power Ancillary Services Characteristics with German 250 MW Wind Data (Golden, Colorado: NREL, 1999).

21. Michael Milligan, *Wind Power Plants and System Operation in the Hourly Time Domain*, in Proceedings of the Windpower 2003 conference, AWEA, Austin, TX.

22. Statistical analysis of PJM and simulated wind power time series was based on data described in Joseph F. DeCarolis and David W. Keith, *The Economics of Large-Scale Wind Power in a Carbon Constrained World*, Energy Policy, forthcoming. The hourly ramping requirements for PJM between 1997 and 2002 have a sub-Gaussian distribution: only 0.09 percent of the hours fall outside the 3σ risk level, substantially less than the 0.3 percent predicted by a Gaussian distribution. For comparison, the hourly differences in simulated wind power from five different wind sites exhibit a super-Gaussian distribution, i.e. roughly 2 percent of the hourly load-following requirements for each wind site fall outside the 3σ limit.

23. See Hirst and Hild, *supra* note 8.

24. Michael J. Grubb, The Economic Value of Wind Energy at High Power System Penetrations: An Analysis of Models, Sensitivities and Assumptions, Wind Eng'g 12(1), 1988, at 1–26.

25. DeCarolis and Keith, *supra* note 22.

26. *Id.*

27. For a detailed argument, see Joseph F. DeCarolis, *The Economics and Environmental Impacts of Large-Scale Wind Power in a Carbon Constrained World*, Ph.D. thesis (Pittsburgh: Carnegie Mellon University, 2004).

28. Steven Smith, personal communication, Apr. 20, 2004.

 1040-6190/$–see front matter , doi:/10.1016/j.tej.2004.12.006