Memo

From B.F. Hobbs, bhobbs@jhu.edu

Re: Comments on 9/24/2013 CSTF Presentation of James Wilson[1]

24 September 2013 (Revised)

Mr. Wilson has critiqued the assumptions of the Hobbs model, and applied it to a number of scenarios.

My initial comments on the critique and analyses are below.

1. Impact of Supply Curve Slope

Slide 4 of Mr. Wilson's presentation indicates that it focuses on Assumption 3, that all supply is offered at zero, and that the fixed quantity offered implies a vertical supply curve at that quantity.

That is indeed one possible assumption in the Hobbs model. However, I also simulated a stepped curve in which new capacity was offered at a positive price (http://www.pjm.com/~/media/committees-groups/task-forces/cstf/20130909/20130909-item-02b-capacity-market-simulation-results-update.ashx), as I did in the original 2005 affidavit. This curve has similar behavior as a sloped curve, in that it stabilizes prices for both vertical and sloped curves, the reason being that it puts a floor under the range of possible price responses over the quantity range of interest (between the previous capacity committed, and the maximum total capacity offered). A stepped supply function can be viewed as an approximation of a continuously sloped function. As in the original affidavit, Slide 10 (yellow highlighted cases) of the Sept. 9 2013 presentation of Hobbs model sensitivities summarizes those results, which indicate that the higher the bid by new capacity, the better the performance by the vertical curve, although the sloped curve results are essentially unaffected (at least when holding the other assumptions constant).

A key point (made in the 2005 affidavit as well) is that the performance of the vertical curve (in terms of probability of meeting target, average reserve, and consumer costs) remains less desirable than the sloped curve, although the differences between the two curves narrows.

The reason for the narrowing in the difference in performance is, in essence, explained in Slides 11 and 12 of Mr. Wilson's presentation. Mr. Wilson is correct that the range and variation of prices is the result of the interaction of the slopes (and, I add, the floors/ceilings) of both the supply and demand curves. As his Slide 6 illustrates, assuming that new capacity offers at zero results in the most variability. But what that slide omits is the effect of increasing the bid of new capacity, which has essentially the same effect as a sloped supply curve by stabilizing prices and improving the performance of the vertical curve.

[1] Any opinions or errors in this memo are the responsibility of the author alone, and are not the responsibility of PJM LLC or the Johns Hopkins University.

Mr. Wilson's point about the importance of supply curve slopes has been made previously, in the 2011 Brattle report, pp. 106-109.[2] In that report, Brattle examines the effect of alternative curve slopes given actual supply curves offered in eight BRAs (2007/08 through 2914.15). The authors note that the slope of the offers by suppliers of capacity results in more stability for the vertical demand curve results than in my original analysis, as Mr. Wilson does. However, they also conclude, as I did with my sensitivity analyses, that the sloped curve nevertheless results in higher stability of prices, compared to the vertical curve:

> p. 107: "*The left chart of Figure 21 shows actual BRA auction prices while the right chart shows prices that would have been realized with a vertical VRR curve. The comparison of these two charts shows that the volatility with the actual VRR curve is somewhat lower than under a vertical curve. For example, actual MAAC prices between 2008/09 and 2009/10 increased by $79/MW-day, while the price increase for the vertical VRR curve simulation was $204/MW day, or more than 2.5 times larger. However, overall, the reduction in price volatility due to the VRR curve slope is somewhat less than we would have expected. The more moderate impact is not as surprising, however, when considering the causes of price changes we have identified and discussed in Sections II. It appears that the slope of the VRR curve, while beneficial in reducing price volatility, has not been sufficiently flat to fundamentally reduce the impacts of other uncertainty factors, such as changes in CETL values or whether individual LDAs are modeled.* "

Of course, my analyses do not consider perturbations in the market due to the latter factors they mention. Brattle also showed additional analyses in which the demand curve slope is further flattened (halved slope), and with the actual sloped supply curves of offers, the advantage of the sloped curve, in terms of dampening price volatility, increases, as would be expected from my analysis as well.[3]

Thus, Mr. Wilson's point about the importance of supply curve slope does not change the robust conclusion of my analysis that a sloped curve stabilizes market prices relative to a vertical curve, and the precise assumptions about the supply curve affect the magnitude but not existence of this advantage.

---

[2] "Second Performance Assessment of PJM's Reliability Pricing Model," by Johannes P. Pfeifenberger, Samuel A. Newell, Kathleen Spees, and Attila Hajos, The Brattle Group, Inc., August 26, 2011 (available from http://www.brattle.com/Experts/ExpertDetail.asp?ExpertID=67).

[3] The Brattle analysis cannot assess the long run entry and consumer cost implications of sloped supply curves because by usually actual offers in each year, it implicitly assumes that offers by suppliers would not be changed by previous years' market settlements or by changes in expectations concerning the volatility of prices and profits. They recognize this limitation explicitly in the report. This however does not diminish the basic point, which is that a sloped demand curve does lower variation of prices, although not as much as in my original affidavit under the base case assumptions I made.

2. Risk Aversion

Turning to the simulations Mr. Wilson performs, I did not have time to duplicate his runs, but I do have the following comments.

First, regarding risk aversion, Mr. Wilson simulates a (nearly) risk neutral case (.52 risk aversion parameter rather than 0.7). As I state in my original affidavit, risk neutrality largely erases the differences in supplier response to the vertical vs. sloped curves, which is because the beneficial effect of the sloped curve is a result of combining the impact of stabilizing prices with the assumption that suppliers are risk averse (i.e., will offer more supply if the variance of profits decreases, for a given level of expected profit). Complete risk neutrality (i.e., all that matters is expected profit) would mean that suppliers would not care about profit variability, which is widely viewed as an unrealistic characterization of supplier behavior in real markets.

My base case choice of risk parameter is 0.7, but whether this value is itself realistic is difficult to ascertain, so I considered other values in both my affidavit and in Slide 10 of the Sept. 9, 2013 presentation (violet rows). As predicted, as risk aversion decreases, the difference between the performance of the two curves also diminishes. The value that Mr. Wilson chooses (0.52) is almost indistinguishable from pure risk neutrality (which is a parameter of 0.5, equivalent to a linear utility function), and so there is no surprise that the two curves perform very similarly under that value. (They won't necessarily perform the same because performance also depends on where along the quantity axis the vertical demand curve is located.) The below figure shows this by indicating that Mr. Wilson's implied utility curve (@ 0.52) on slide 14 cannot be meaningfully distinguished from full risk neutrality.

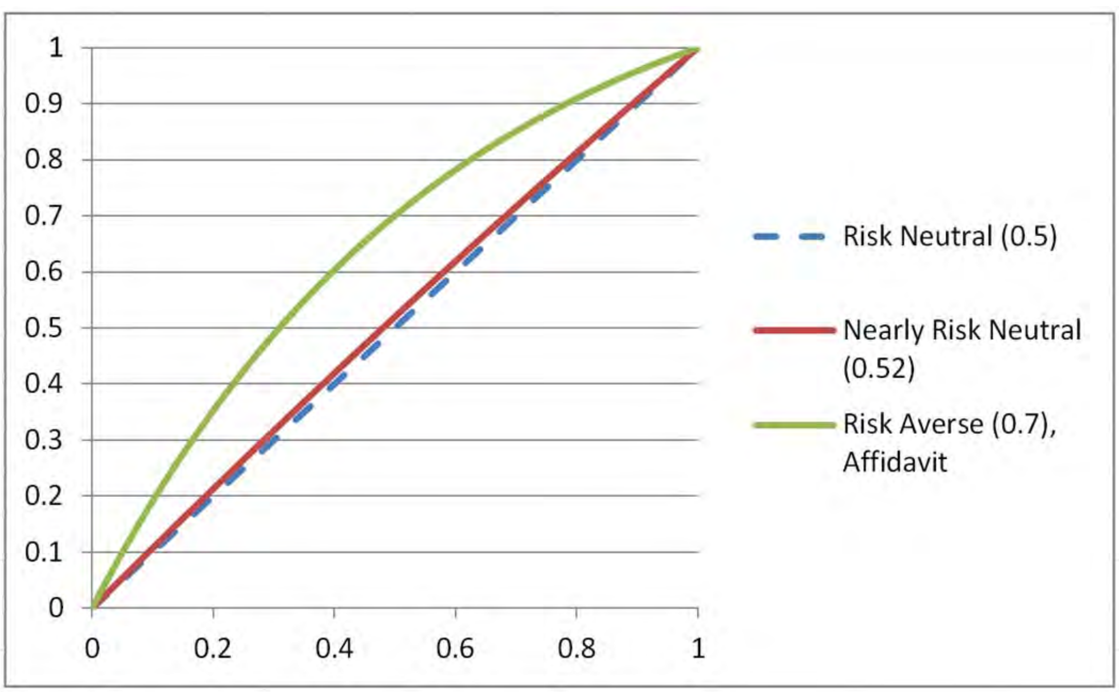

3. Impact of Sloped Curve

Mr. Wilson has a long appendix (starting on slide 15) in which he simulates the effect of increasing the slope of one-half of the curve (first the left half, and then the right), and shows that a combination of a steep (but not vertical) slope on one side combined with the original slope on the other side results in reliability performance close to the full sloped curve. These analyses illustrate the nonlinear response of the system to the slope. It would be easy to conclude from a superficial reading of this analysis that only a slight slope is needed to obtain the benefits of a sloped demand curve; indeed Mr. Wilson says "Hobbs model says: Vertical is Bad, but Only a Little Slope is Needed."

This conclusion (p. 5) is unjustified by the analysis, because Mr. Wilson's analyses did not consider curves where there was little slope throughout; each of the curves considered in the appendix had a significant slope in at least half of the curve. No curves with "only a little slope" were considered.

Therefore, I did some runs which varied the slope for the entire curve. The results that I summarize below shows that Mr. Wilson's conclusion is incorrect; a significant slope is needed to obtain those benefits, although there are diminishing returns to increasing the amount of slope. Consumer benefits from the slope are roughly proportional to the slope. Most of the reliability benefits are obtained by a curve with half the slope of the original sloped curve, but those are quickly lost after that as the curve approaches vertical.

The reason my analysis differs from Mr. Wilson's is that I consider the slope of the entire curve, but Mr. Wilson preserves the slope of half of the curve. When half of the curve is made vertical, most of the reliability benefits (in terms of the percentage of years that the target is met) are obtained. If you have a near vertical slope in both halves (which Mr. Wilson does not show, but I do), all the cost and reliability benefits are lost. Thus, the statement "Hobbs model says: Vertical is Bad, but Only a Little Slope is Needed" is highly misleading; a more correct characterization of his analysis is that significant slope is needed for at least half of the curve, which is a very different statement than "only a little slope is needed".

I now summarize the analysis. To be as straight-forward as possible, I considered a range of slopes for a simplified version of the entire curve, ranging from (a simplified) single-segment sloped curve whose performance is similar to the more complex sloped curve to the vertical curve. The most sloped curve is shown below, with a price of 180,963 $/MW/yr for a quantity less than 0.97, a price of zero for a quantity more than 1.03, and a linear slope in between. To focus on the effect of slope, I keep the center point of all curves at 1.0, and steepen the slope by squeezing the range to [0.98, 1.02], [0.99,1.01], [0.995,1.005], and finally to [1.0,1.0] (vertical curve, which is to the right of the original vertical curve, which is at [0.982,0.982]). The vertical version is shown below after the full sloped version.

Sloped curve approximation:

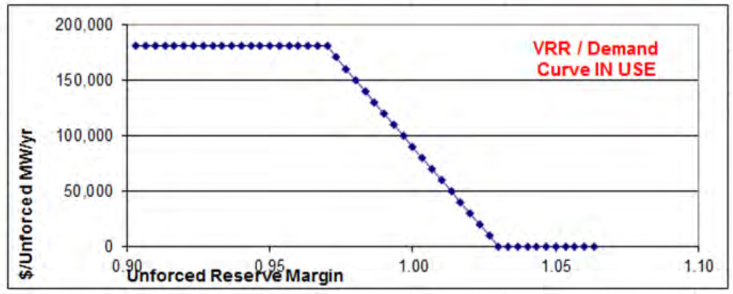

Vertical curve approximation:

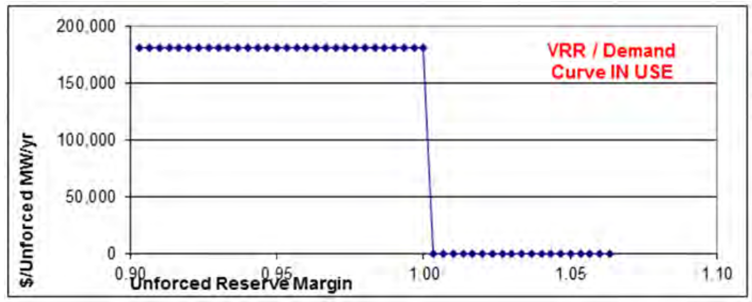

The numerical results are shown in the table on the next page. This shows that most of the benefits of the vertical curve are obtained at [.98,1.02] (which can be viewed as a weighted combination of the two curves, 2/3 weight on the full sloped curve and 1/3 on the vertical curve). However, the benefits deteriorate rapidly after that. At a 1/3 weight [0.99,1.01], the consumer costs are 147.9K $/peak MW/yr (compared to 141.2K for the full sloped curve and 150.6K for the vertical curve), the % of years that IRM is met is 85.8% (cf., 98.5% for sloped and 66.5% for vertical), although the average reserve margin is comparable to the sloped curve (but its standard deviation is over three times as large). At [0.995,1.005] (a weight of 1/6 on the sloped curve), the performance is similar to the vertical curve.

Thus, it would be mistaken to conclude, as Mr. Wilson does, that "Hobbs model says 'Only a Little Slope is Needed'" (slide 13).

Case	Reserve Indices		Generation Profit, $/Installed kW/yr	Components of Generation Revenue			Consumer Payments for Scarcity + ICAP $/Peak kW/yr	Standard Deviation				
	% Years Meet or Exceed IRM	Average % Reserve over IRM		Scarcity Revenue $/Installed kW/yr	E&AS Fixed Revenue $/Installed kW/yr	ICAP Pay-ment $/Installed kW/yr		Average % Reserve over IRM (s.d.)	Generation Profit, $/Installed kW/yr (s.d.)	Scarcity Revenue $/Installed kW/yr (s.d.)	ICAP Pay-ment $/Installed kW/yr (s.d.)	Consumer Payments for Scarcity + ICAP $/Peak kW/yr
1.5 CONE at 0.97, 0 CONE at 1.03	98.52	1.50	5.35	9.80	10.00	124.94	141.20	0.49	22	15	14	20
1.5 CONE at 0.98, 0 CONE at 1.02	98.08	1.87	6.27	8.79	10.00	126.87	142.59	0.65	28	13	23	27
1.5 CONE at 0.99, 0 CONE at 1.01	85.80	1.67	12.03	10.50	10.00	130.92	147.87	1.66	60	20	53	61
1.5 CONE at 0.995, 0 CONE at 1.005	70.80	1.13	14.94	14.01	10.00	130.32	149.65	2.32	79	29	67	81
Vertical at 1.0	66.52	0.98	16.20	15.27	10.00	130.32	150.59	2.51	86	32	73	87
Original vertical (vertical at 0.982)	42.16	(0.62)	23.32	24.91	10.00	127.80	155.99	2.43	98	50	74	98
Original sloped	95.68	1.24	5.20	10.74	10.00	123.85	140.71	0.65	22	17	11	20