Supplement to

The Competitive Effects of Transmission Capacity in a Deregulated Electricity Industry

by Borenstein, Bushnell and Stoft
October 1998

Thin Lines and Mixed-Strategy Equilibria

We have shown that there is a threshold level of capacity, above which two otherwise isolated markets are effectively merged. It remains to examine what happens to prices and output levels for line capacities that are below the threshold level that induces the Cournot duopoly result. In this section we examine competitive behavior for markets connected by these relatively "thin" lines.

By Nash's theorem, the fact that a mixed strategy is a firm's best response implies that all of its pure components are equally profitable. So we have found a Cournot-Nash equilibrium if all of each player's q's are equally profitable and are more profitable than any other q, holding constant the opponent's strategy.

Even without finding the actual mixed-strategy Cournot-Nash equilibria, it seems likely that the expected price will decline as k increases from zero. With a very thin line, for instance, the expected price must be very close to the monopoly level. If it were not, then either firm could improve its expected profits by simply admitting imports of k and producing its optimal passive output as a pure strategy. With k near $k*$, the lower bound on price provided by the alternative of producing the optimal passive output is much weaker and the mixed strategy is more likely to result in a lower expected price.

The difficulty with finding the Cournot-Nash equilibrium in this particular model is that each player can choose from an infinite number of different pure strategies (q's). All that we can do numerically is approximate this infinite set of q's with a few hundred possible values, say $q = \{0, 0.1, 0.2, \ldots 10\}$. But even with 100 q's the vector that gives the probabilities of playing each particular q will be a 100 dimensional vector. Searching over such a space is not feasible since even with a few possible probability values the set of possible strategies runs into the billions.

Instead of solving for mixed strategies analytically, we create an efficient search algorithm based on a learning procedure that might be applied by two players actually playing this game repeatedly, an approach known as "fictitious play."[1]

[1] See Fudenberg and Levine, 1998, chapter 2, for a full description of fictitious play.

The procedure starts by assigning each player a reasonable pure strategy. For example, we could assign each player its monopoly quantity in a single market or its Cournot quantity when the markets are fully integrated. In the simulations that we carried out, the choice of starting strategy made no difference in the answer to which the algorithm converged, and even picking a very bad strategy delayed convergence by only a few iterations. So we assume each player starts by playing a starting strategy, which we call S_0. Strategy S_0 is simply to produce q_0.

Step 1. Player 1 finds her optimal strategy under the assumption that player 2 will always produce q_0. Player 1 plays this best response strategy, which we call S_1. S_1 could be a mixed strategy in theory, but in practice, as we *approach* a mixed-strategy equilibrium (and due to the numerical analysis that the computer does out to 32 digits), this will always yield a single most profitable quantity. S_1 is then a pure strategy of producing a certain q_1.

Step 2. Player 2 finds his optimal strategy assuming player 1 will always play q_0 with 50% probability and q_1 with 50% probability. Call this optimal strategy S_2. Again, theoretically S_2 could be a pure or mixed strategy, but a pure strategy best response always emerges. This results in player 2 producing some q_2.

Step 2 can be interpreted as player 2 assuming that q_0 and q_1 are both part of a mixed strategy being pursued by player 1. In this case, player 2's best estimate of the probabilities being used by player 1 is 0.5 for each quantity. Player 2 then optimizes with respect to this best estimate of player 1's mixed strategy. This philosophy is now simply pursued throughout the iterative process, so the general step in the iteration can be described as follows:

Step N. Player i, assumes that all of her opponent's past plays are part of a mixed strategy and assigns each play an equal probability of being repeated in the future. Based on this assumption, player i find her optimum strategy and plays that.

Thus if the opponent had played $q = 4$ twice and $q = 5$ twice and $q = 6$ once, player i would assign a probability of 40% to both $q = 4$ and $q = 5$ and 20% to $q = 6$. Player i would then optimize with respect to that estimated mixed strategy.

Shapley (1964) showed that fictitious play algorithms will not always converge to a Cournot-Nash mixed strategy equilibrium, but it appears to work well in this case. In all cases analyzed, the algorithm converged rapidly to an equilibrium. Once a possible Cournot-Nash equilibrium is found with this algorithm, it is straightforward to check numerically that it is in fact such an equilibrium.[2] In every case, the algorithm converged to an actual Cournot-Nash equilibrium.

One modification of the algorithm has been used, which appears to speed up convergence and which certainly simplifies the appearance of approximate equilibria. This is to eliminate pure strategies that are calculated at very early stages of the iterative process. At these early stages the optimized response quantity is based on estimates of the opponents mixed strategy that are quite crude. Consequently these early "optimal" responses are inaccurate and often involve q's that are not part of the true equilibrium distribution of q's. For instance, when the actual mixed strategy is simply a mixture of $q = 5$ and $q = 6$, early optimal responses based on inaccurate estimates of the true strategy may involve $q = 5.2$, $q = 5.5$, or even $q = 8$. By eliminating these early responses, we can often find a final estimate that involves only the q's that are actually part of the mixed equilibrium. Of course if we left the early ones in, the iteration process would still converge to the correct answer, because these values would never be repeated, and so as iteration continued their probabilities would diminish towards zero.

Using this method, we examined market outcomes for various line capacities using two different demand models. The first model uses a linear demand function of $q = 10 - p$ and has suppliers with a constant marginal cost of zero. The second model assumes demand with a constant elasticity[3] of -2 and suppliers with constant marginal costs equal to one.

The results from each of the models further support the analytical conclusions of the previous sections. Line capacity has a clear and pronounced effect on the output of both competitors. Figures 4 and 5 show the changes in expected prices in each market for the demand functions we modeled. In each of the cases, increases in line capacity cause a monotonic decline in the expected market price and a monotonic increase in expected quantities produced up to the point that the unconstrained Cournot duopoly equilibrium is reached. Beyond the line capacity that yields the unconstrained outcome, further increases have no effect.

In both of these models, we find that small increases in line capacities can yield output increases much larger than the added line capacity. This is consistent with the analytic conclusion that the line capacity necessary to completely merge the two markets

[2] We do this by establishing that the profits are equal for the strategies over which a firm is randomizing and greater than all other quantity choices.

[3] $q = ap^{-2}$ where a is set so that $q = 5$ when $p = 5$

is relatively small compared to the added output that such a merging produces. In both models, even small lines produced big benefits. In fact, the marginal effect of increased line size appears generally to be greatest when the line is very thin, though the slope does not appear to change monotonically.

With linear demand and a transmission line that isn't very thin, the mixed-strategy equilibrium that obtains from the search algorithm consists of mixing over two strategies.[4] As the line capacity declines, the difference between the two quantities over which the firm mixes in equilibrium also declines. At some line size (in the linear demand case, at about one-fifth of the line size that yields the unconstrained outcome), the number of distinct quantities over which the firm mixes in its equilibrium strategy rises to three. With still thinner lines, the number rises further. With constant elasticity demand, the search algorithm yielded similar results except that instead of two distinct quantity choices there appear to be two distinct quantity regions; a bimodal distribution with positive probability in only two narrow regions, but not so narrow that they can clearly be attributed to computer approximation.

References

Fudenberg, D. and D.K. Levine (1998). *Theory of Learning in Games*, (Cambridge, MA.: MIT Press).

Shapley, L. (1964). "Some Topics in Two-Person Games" in *Advances in Game Thoery*, Drescher, Shapley, and Tucker, eds., (Princeton: Princeton University Press).

[4] In some cases, the program indicated equilibrium strategies including two prices that were adjacent or nearly adjacent in the grid over which the program searched. We considered these to be a single point in the mixing distribution.

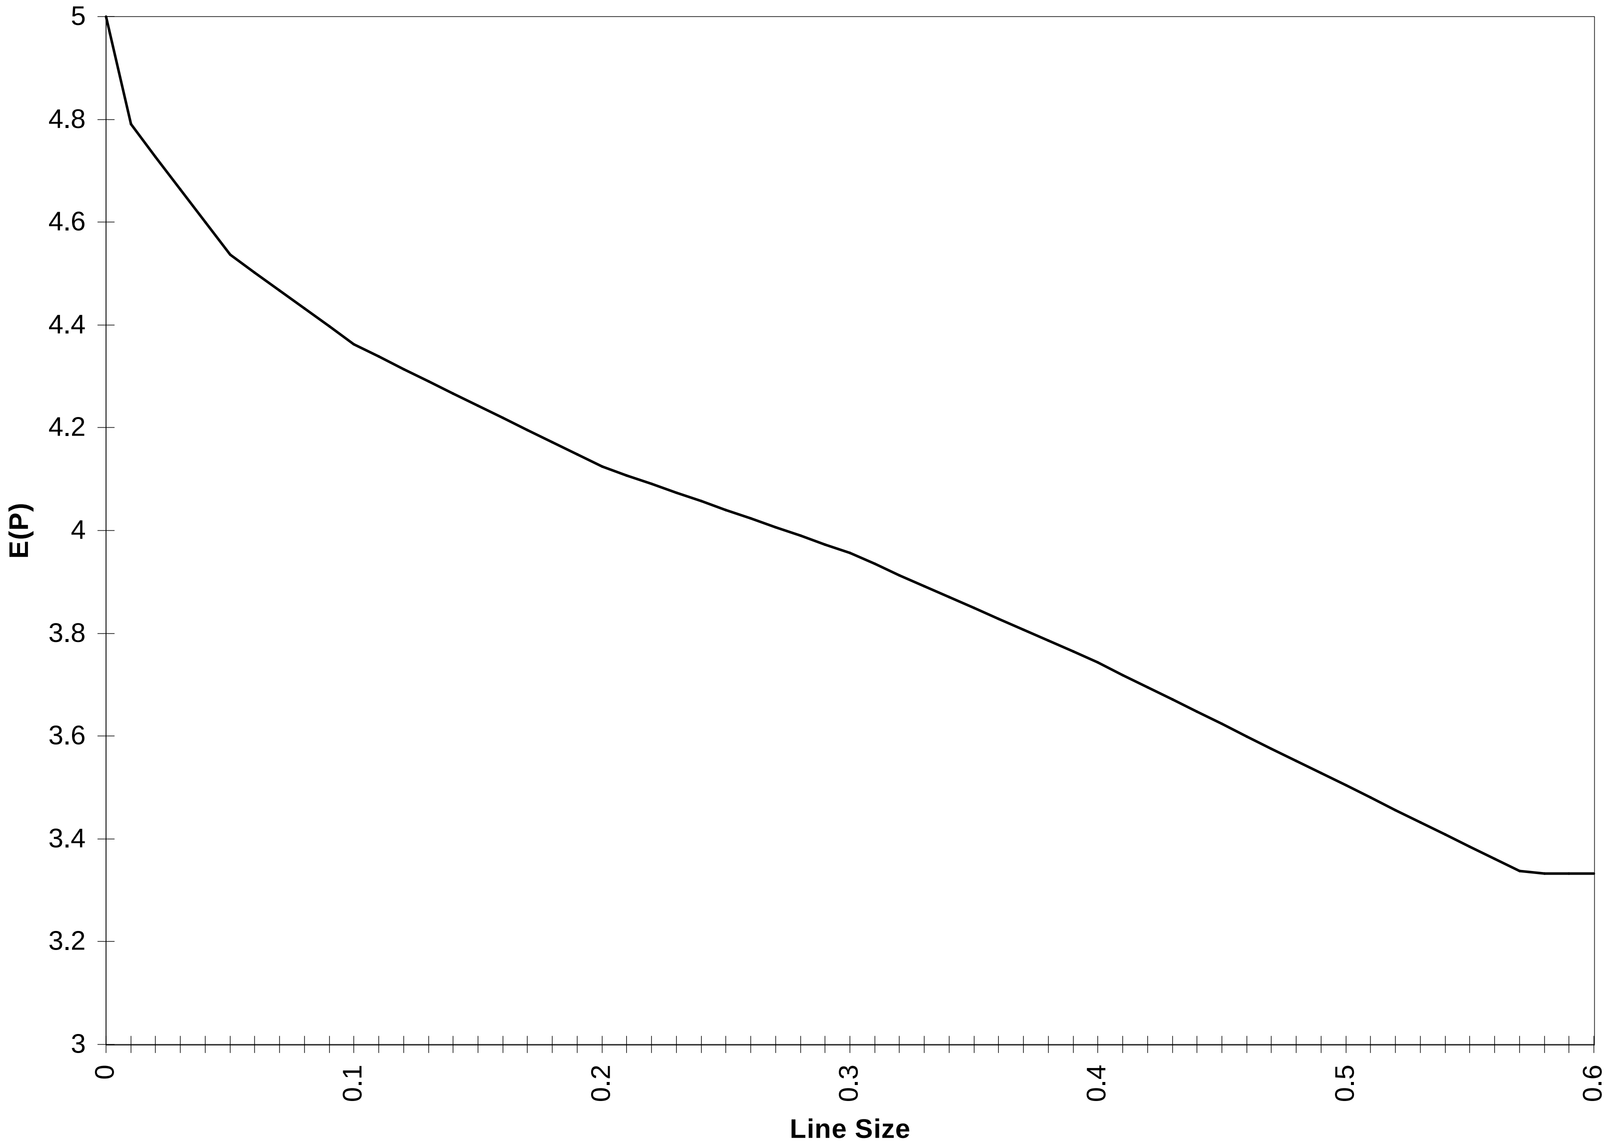

FIGURE 4: Expected Price as a Function of Line Size
Linear Demand (Q=10-P) and Constant Marginal Cost (MC=0)

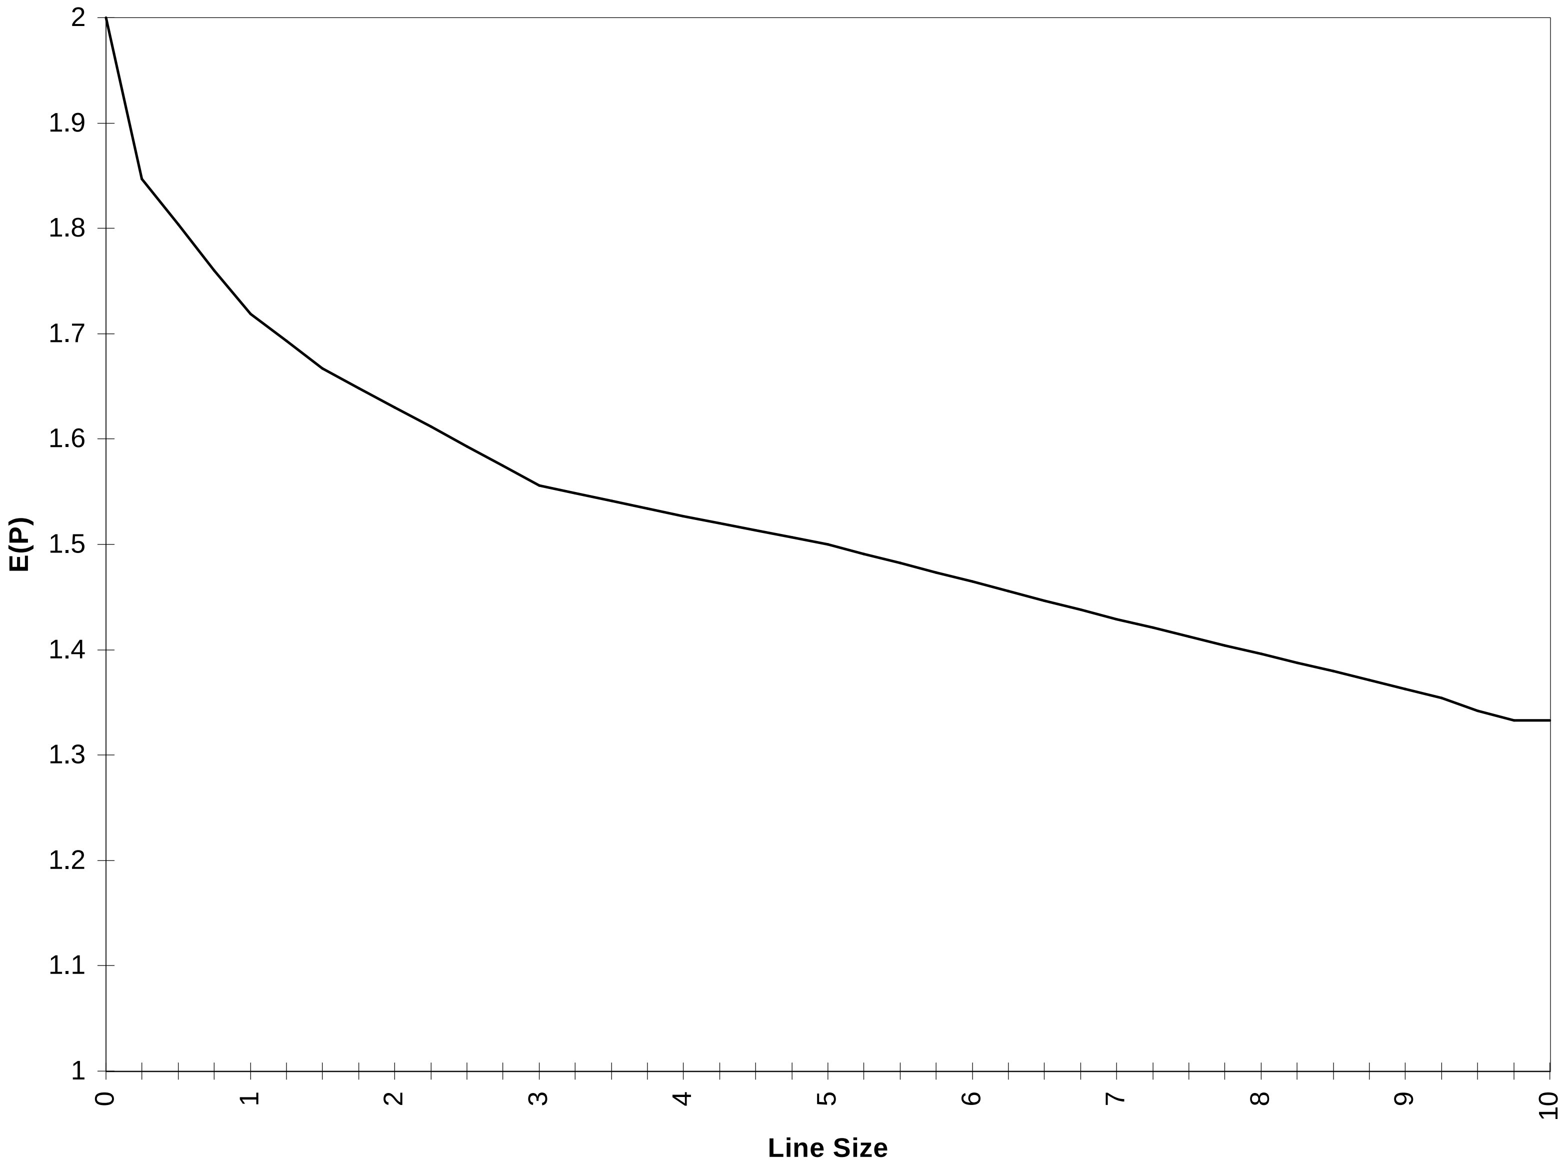

FIGURE 5: Expected Price as a Function of Line Size
Constant Elasticity Demand (Q=125/P**2) and Constant Marginal Cost (MC=1)