Memorandum

To: Keith Hofseth, U.S. Army Corps of Engineers

From: Kevin Henrickson and Wesley Wilson

Date: January 25, 2004

Subject: Parametric and Non-Parametric Elasticity Estimates of Barge Demand

Due to the useful comments made by you at the Transportation Research Board's annual meetings this January, we have investigated the assumption of constant elasticity made in the Henrickson and Wilson work and I would like to share the preliminary results of this effort with you. Our plan is to refine some of this work and include in the final report to be sent to you for inclusion in the NETS paper series.

The most obvious way of checking how elasticity might vary with river location is to use a varying coefficient model wherein the elasticity variable (log(barge rate)) is interacted with river mile. When this is done we obtain reasonable results that suggests the elasticity becomes "less elastic" as we travel upriver. The following graph reflects the estimated elasticity and river mile.

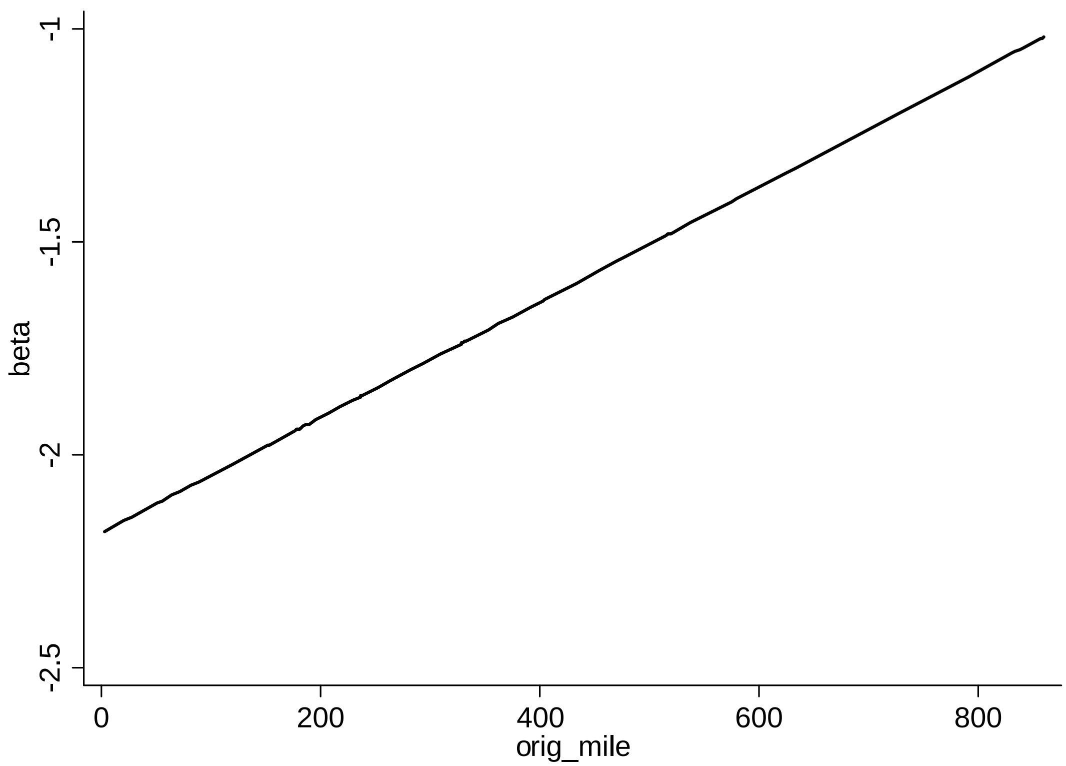

The figure illustrates that the magnitude of the estimated elasticity decreases with river mile. This suggests that shippers located further north on the river have relatively inelastic demands for barge than their southern counterparts.

While this is an intuitively appealing result, applying a restriction that elasticity is linear in origin mile is an assumption just as much as placing a constant elasticity assumption on the demand for barge. Because of this we use two different estimation techniques aimed at letting the data uncover the relationship between elasticity and origin mile without specifying a functional form.[1]

The first of these techniques is known in the time series literature as rolling regressions. Using this technique, we order the data according to river mile. We run the barge demand equation on a "window" of data. The window is arbitrary (and we tried different windows). Essentially, we run the barge demand equation is run on the first x observations and the demand elasticity is recorded (the first x observations correspond to the shippers located furthest south).[2] The barge demand equation is then run on observations 2 through x+1 and the demand elasticities of this equation are again recorded. In essence, we are taking a window of size x and moving it along the river one position at a time estimating the demand elasticity in each window location. Using a window size of 30 the demand elasticity is estimated as:

[1] It should be noted that neither of these techniques are aimed at estimating elasticities that will be used. Instead, they are aimed at simply indicating the patterns of elasticities over the river, so that we can account for them in our model.

[2] x is arbitrarily chosen, and the only restriction on it is that it must be large enough to estimate the equation. However, the smaller x is the more we can determine local elasticity patterns across the river.

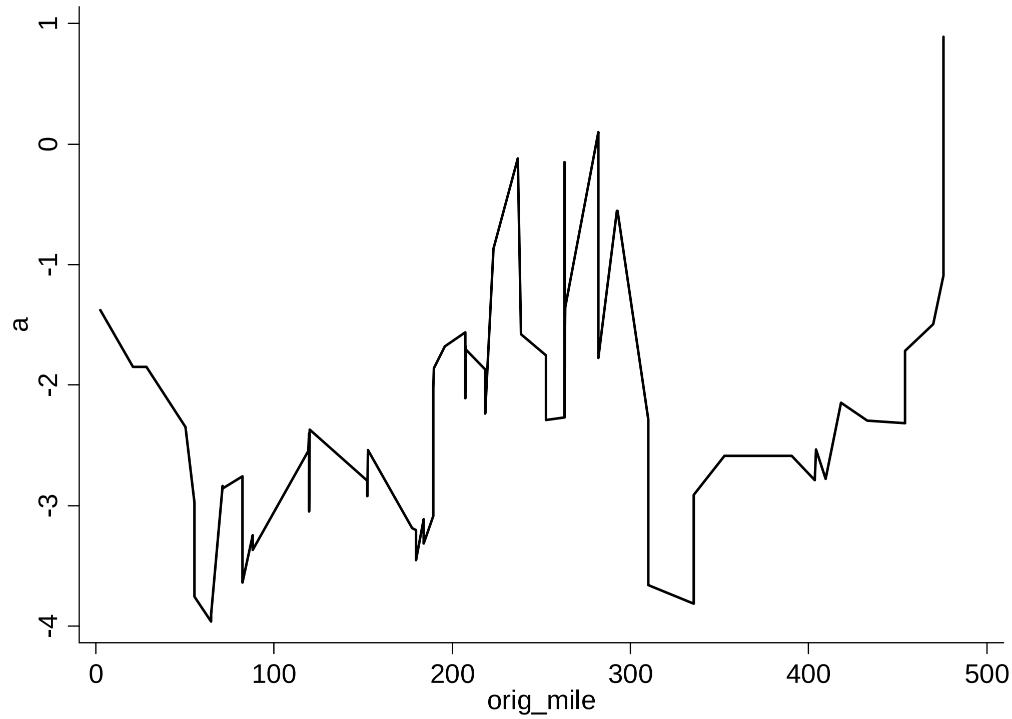

a
1
0
-1
-2
-3
-4
0
100
200
300
400
500
orig_mile

The results with a window size of 40 are presented graphically with:

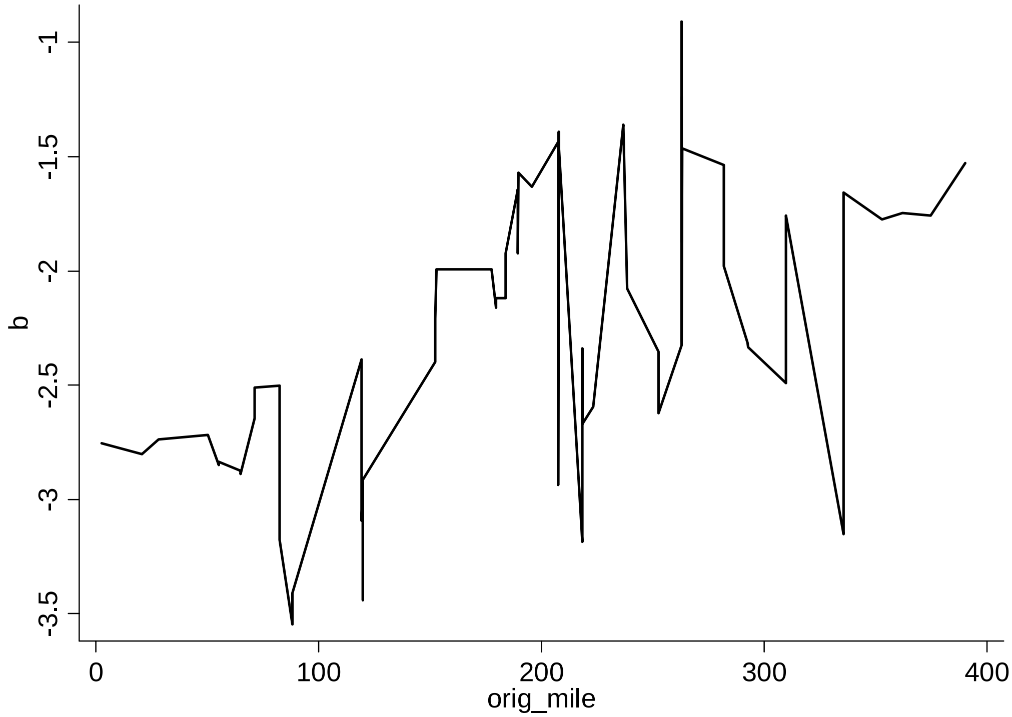

The results with a window size of 50 are presented graphically with:

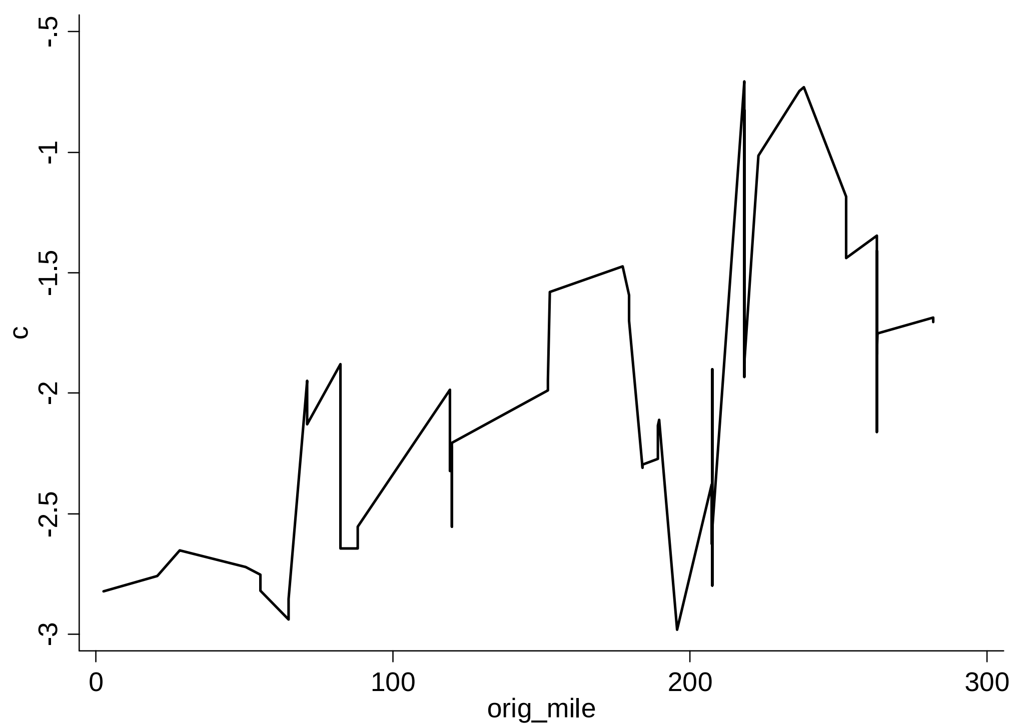

With window sizes of 40 and 50, elasticity appears to be trending down for locations further up river, although the trend is bumpy. With a window size of 30, it almost appears as if elasticity is decreasing in the middle and then increasing again, indicating that perhaps a parabolic term in the estimation equation would be appropriate.

A second technique used to examine elasticity over space is a locally weighted regression. This technique is similar to the rolling window technique with one notable difference. Again, we must specify a window size in which the demand equation will be run and again we move the window up the river one position at a time. The key difference is that we weight the observations in the window such that the middle position gets the highest weight and each position away from this middle gets subsequently lower weights. For example, if we had a window size of 5, the middle position would be the 3rd observation in the window and it would receive a weight of 1, indicating that it is fully weighted. Positions 2 and 4 would receive a weight of .89 each, positions 1 and 5 would receive a weight of .35 each, and positions 0 and 6 would receive a weight of 0 meaning that they are not included in the regression. Weighted least squares is then run to estimate the demand elasticity for the given middle location and window size. The estimated elasticity is then recorded and the window is moved up river one location. When this is done for a window of size 40:

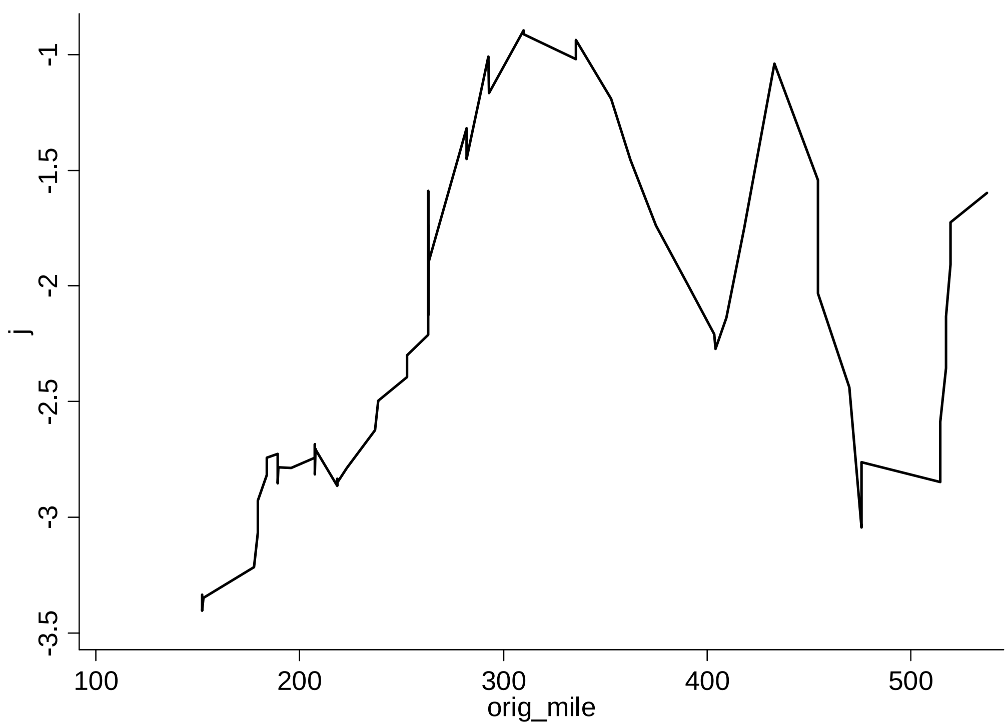

For a window size of 60:

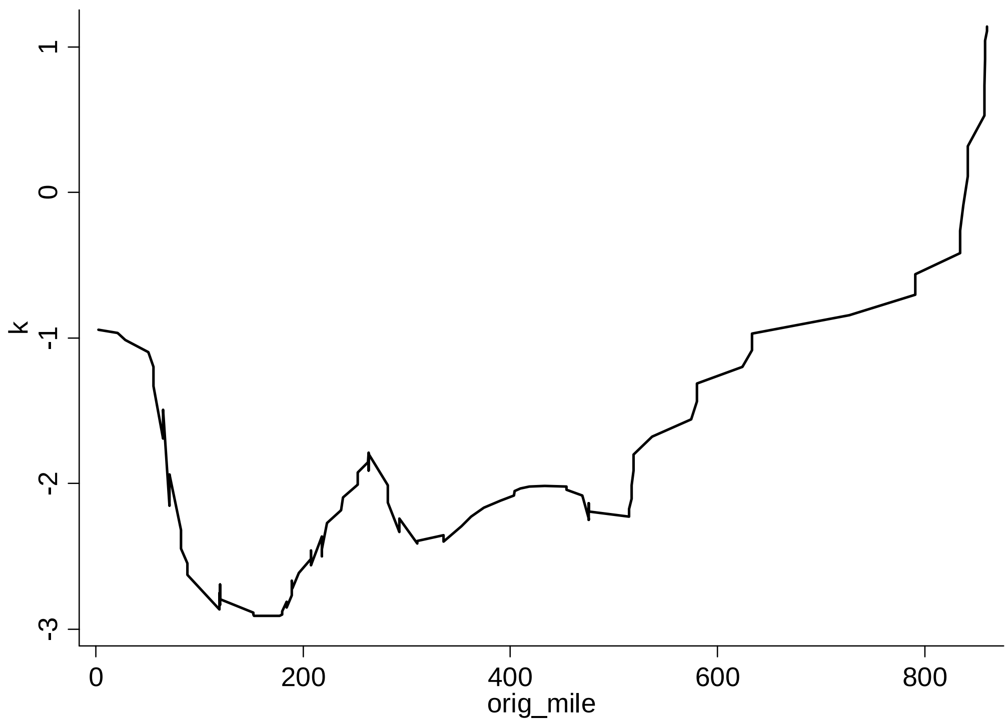

And for 80:

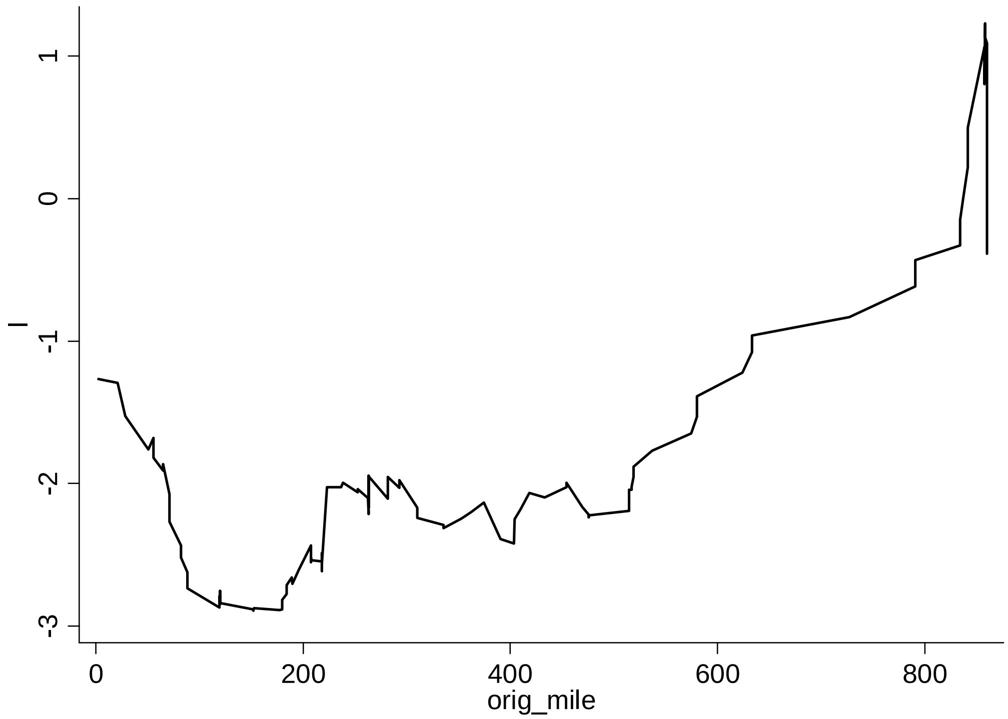

These three graphs again indicate bumpy downward trends in elasticity as one moves up river.

The final set of estimation is conducted using these results. Specifically, these are nonparametric approaches wherein a functional form for the pattern of elasticity across distance is not specified. In both cases, it appears that a reasonable parametric form may involve higher order interactions between river mile and elasticity. We now estimate a variety of different forms. First we estimate a second order interaction term and obtain the following elasticity estimates:

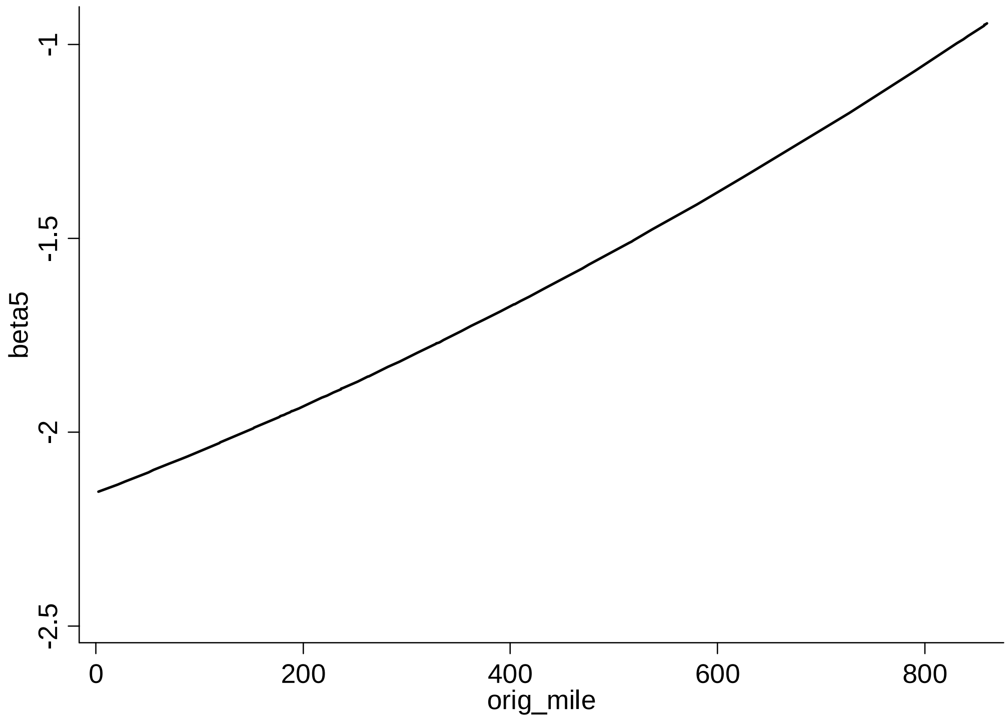

This graph shows that the second order term offers little difference from the linear interaction term.

However, a third order interaction term yields the following:

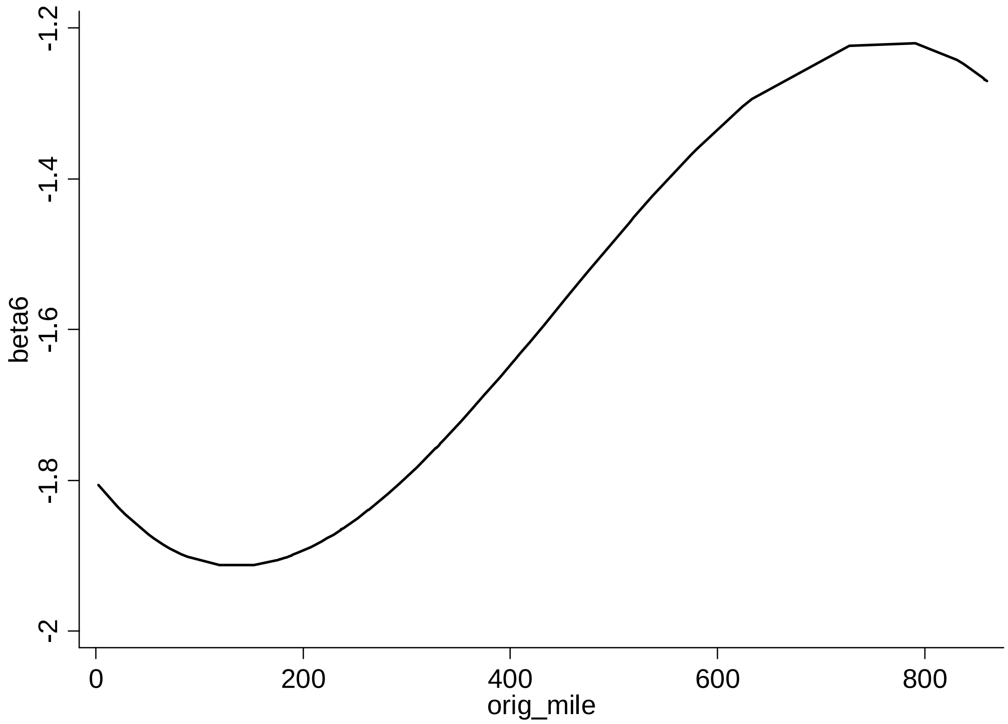

This latter seems to be the best representation of our observations from both the rolling regressions and the locally weighted regressions. The final two estimations relate to the use of a linear and quadratic in log functional forms for the elasticities. That is, elasticity=a0+a1*log(river mile) and elasticity=a0+a1*log(river mile)+1/2*log(river mile)*log(river mile). The graphs of each elasticity against river mile are:

Linear in logs

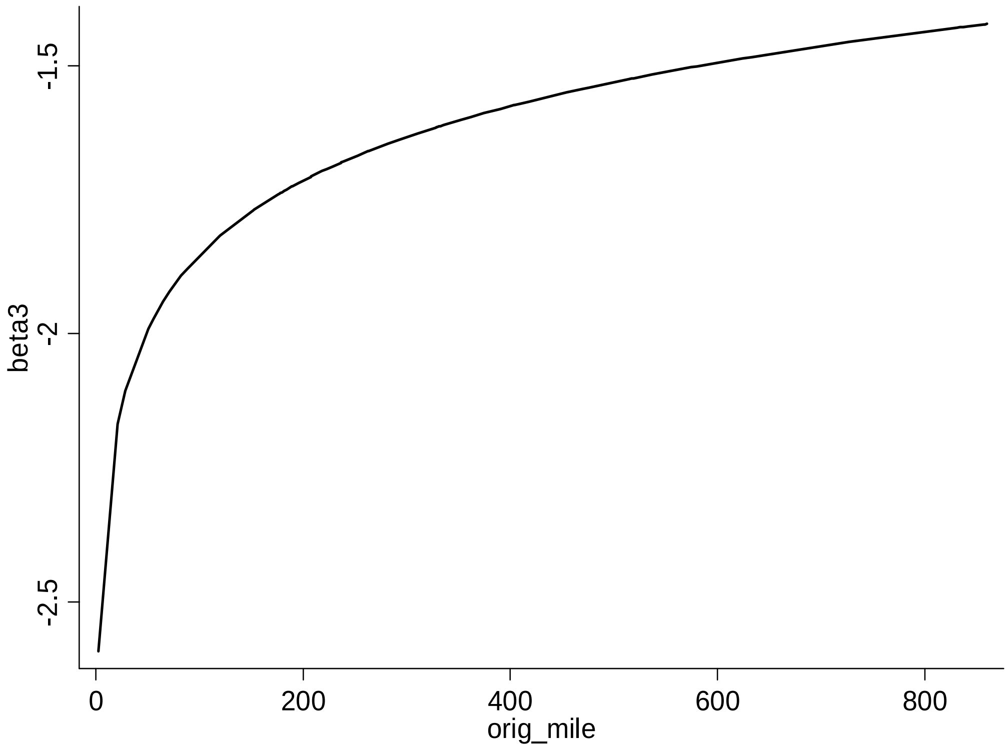

Quadratic in logs

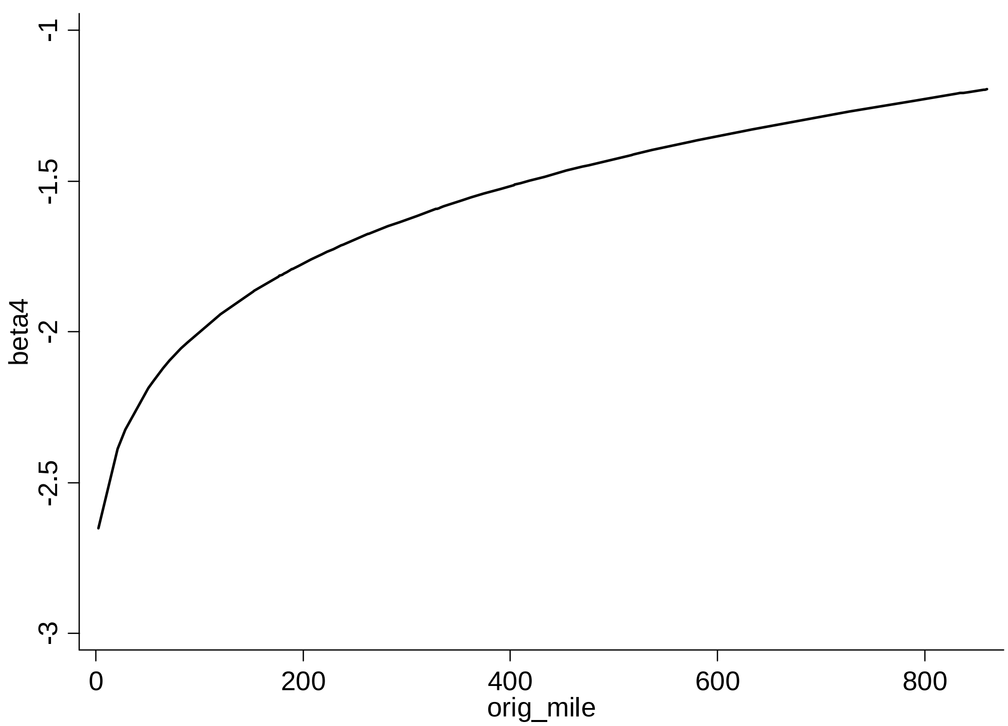

We also have plans for further refinements in the model. Specifically, the model from which the estimation is derived emanates from a model of spatial competition. From this model, we obtain an estimating equation that explains output of a firm in terms of spatial properties of the firm *and its competitors over space*. The competitors enter into the model as observed variables – these variables are *observed* spatial competition variables. However, in the econometrics literature, it is possible that the error terms of firms located near each other are correlated. In the plans for the final work, we plan to model the "spatial" correlation among firms. Theoretically, we have

$$Q_i = Q(r, characterics\ of\ the\ firm, characteristics\ of\ spatial\ competition) + \varepsilon_i$$

In a spatial autocorrelation model, the errors among competitors are correlated over space. Generally, this type of model has not been used much in economics, but we plan as a final piece of estimation to incorporate this correlation. Theoretically, this should help the efficiency of the estimates.