Solving Transport Networks in which Supply and Demand are Elastic

Jack Devanney

Sisyphus Beach

Tavernier, Florida

2007

2007-12-29

Contents

Abstract

This paper shows that a wide range of non-linearities can be incorporated into transportation network problems and still be feasibly solved by linear programming techniques. These include source supply curves, sink demand curves, and non-linear fuel consumption curves.

Keywords

Transportation Network; Non-Linear supply and demand; Linear Programming; Tanker Market Model;

1 Introduction

In the classical transportation network problem, the supply at each source is assumed to be fixed as is the demand at each sink. The problem is to meet the fixed demand from the available supply at minimum overall transportation cost.

A somehwat more realistic situation is that in which the supply provided at each source is dependent on the fob price at that production point, and the demand in each consuming area is dependent on the delivered price in that area. In other words each producing area has its own little supply curve and each consuming market has its own little demand curve, neither of which need be totally inelastic. We will call such elastic transportation networks *elnets*.

The purpose of this note is to show that despite the fact that the individual supply and demand curves may be non-linear, under some not very restrictive assumptions, the elnet problem can be converted into a linear program, for which extremely efficient algorithms exist.

We also point out that in a similar manner highly non-linear fuel consumption curves can be incorporated into a linear program. This is crucial to modelling the tanker market.

2 A Familiar Problem

Let's start out with a network in which there is only one source node and only one demand node and zero transhipment cost. The source node has a non-decreasing supply curve. The consuming node has a non-increasing demand curve. In this case, our problem is nothing more than determining the intersection of a supply and demand curve. Figure 1 sketches this classical problem. The important point for our purposes, is that the market clearing price is the price which maximizes the difference between the total area under the demand curve (the gross amount consumers would be willing to pay if each were driven to their indifference price) less the total area under the supply curve (the gross resource cost). This is the gray area in Figure 1. If the price is somehow held above the market clearing price, some demand will be choked off, and the gray area cutoff at that point. If the price is somehow fixed below the market clearing price, some supply will be choked off, and the gray area cutoff at that point. This observation allows us to turn the problem into an optimization problem: find the price that maximizes this area, the sum of the producers and consumers surplus, which we will call simply the *surplus*.

Almost all supply and demand curves are manifestly non-linear, but we want to use linear methods. So we approximate both the supply and demand curves by a series of steps, as shown in Figure 2. Let D_i be the incremental demand which comes on at price P_i. Let S_j be the incremental supply which comes on at price C_j. In a competitive market C_j is also the marginal resource cost associated with this increment of supply.

Let d_i be the amount of demand step i which is actually consumed at equilibrium and let s_j be the amount of supply step j which is actually supplied at equilibrium, then the problem of maximizing the surplus can be written as:

$$max \sum_i P_i d_i - \sum_j C_j s_j$$

subject to

$$\sum_i d_i = \sum_j s_j$$

$$0 <= d_i <= D_i \quad for \quad all \quad i$$

$$0 <= s_j <= S_j \quad for \quad all \quad j$$

This is ridiculously round about way of solving a very simple problem; but it does have some interesting features:

- It is a linear program for which extremely efficient solution algorithms exists. In fact, it is a particularly simple form of linear program known as a transportation problem. The underlying supply (demand) curves can be highly non-linear – in fact they can be any non-decreasing (non-increasing) function for which the prices are non-negative – but as long as we are willing to approximate them with a series of steps, we have rid ourselves of these non-linearities from a computational point of view.[1]
- As long as the supply curves are non-decreasing, and the demand curves are non-increasing, the optimization routine does not have to be told the order of the steps. At optimality, it will never supply a lower priced demand while a higher price demand goes unsupplied, for by making the appropriate switch the objective function could be increased. Nor at optimality for the same reason will it ever produce from a higher cost supply when lower cost supply is available.

As our little linear program makes clear, once we have approximated the supply and demand curves by a series of steps, we have a simple network problem which has J sources where J is the number of supply steps and I markets where I is the number of demand steps. The amount of supply available at source j is S_j and its cost is C_j. The amount of demand available at sink i is D_i and this demand is willing to pay P_i. The problem is to determine the profit maximizing distribution of supply to demand. In so doing, it is obvious that the ordering of the demand and supply nodes is irrelevent.

[1] The curves must also be bounded, that is, there is a price so high that demand goes to zero, and a price so low that supply goes to zero. In the real world, this is not much of a restriction. There is no requirement that the slopes of the demand or supply curve be increasing. In fact the slopes can go to zero. (If both slopes are zero at the market clearing price, the quantity supplied is not completely determinant, but the price and surplus are.)

Figure 1: A Familiar Problem

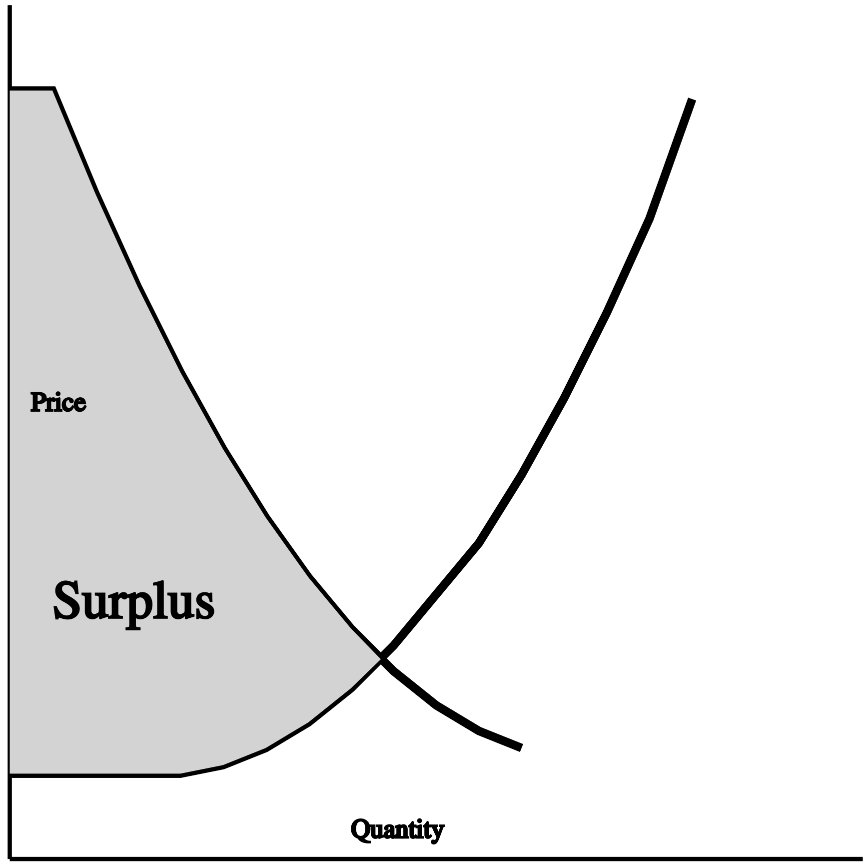

Figure 2: Supply/Demand Curve Approximated by Steps

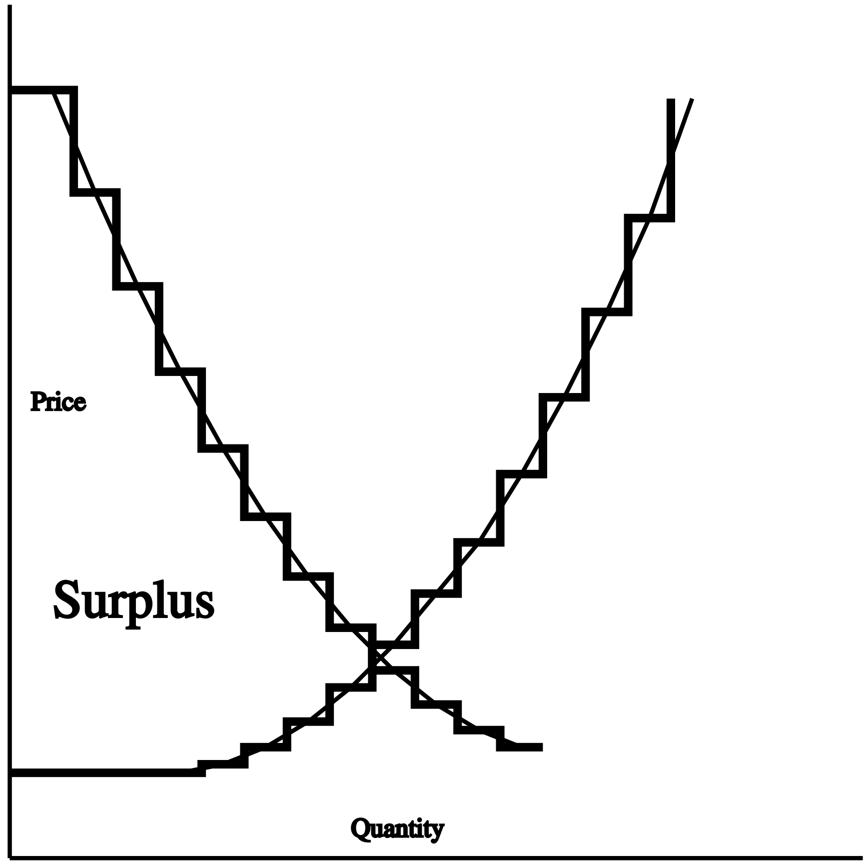

3 A Real Network

The value of these observations come into play when we move to a real network. Suppose we have a network of M supply nodes or sources indexed by m. Associated with each supply node is a non-decreasing supply function which we will approximate by a series of ascending steps. Let S_{mj} be the incremental supply from source m which comes on at fob price C_{mj}. And suppose we have N consuming nodes or markets indexed by n each of which has a non-increasing demand function which we will approximate by a series of descending steps. Let D_{ni} be the incremental demand at market n which comes on at cif price P_{ni}. Let T_{mn} be the unit cost of transportation from m to n.

The following linear program maximizes total consumers and producers surplus throughout the network.

$$max \sum_n \sum_i P_{ni} d_{ni} - \sum_m \sum_j C_{mj} s_{mj} - \sum_m \sum_n T_{mn} t_{mn}$$

subject to

$$\sum_j s_{mj} >= \sum_n t_{mn} \qquad for \ \ all \ \ sources \ \ m$$

.

$$\sum_i d_{ni} <= \sum_m t_{mn} \qquad for \ \ all \ \ markets \ \ n$$

$$0 <= s_{mj} <= S_{mj} \qquad for\ all\ j\ and\ for\ all\ m$$

$$0 <= d_{ni} <= D_{ni} \qquad for\ all\ i\ and\ for\ all\ n$$

where t_{mn} is the amount shipped from m to n.

The objective function is merely the sum of consumers and producers surplus throughout the network taking into account the transhipment costs. The first set of constraints insures that the total shipments from each source are no more than the amount actually produced at that node. The second set of constraints insures that the total shipments to each market are at least as much as the total amount consumed in that market. The third and fourth set of constraints enforce the step limits.

In the traditional transportation problem, supply and demand are totally inelastic. The supply curve at node m is a single step of size S_m and zero cost. In this situation the first and third set of constraints can be combined and s_m eliminated. The demand curve at node n is a single step of size D_n and some very high but unstated price. In this situation the second and fourth set of constraints can be combined and d_n eliminated. Since all demand must be satisfied, the first term in the objective function is a constant and can be thrown away and the second term is zero. In other words, the traditional inelastic network is a special case of this formulation.

This formulation can involve a lot of variables. If every supply and demand curve has ten steps, we will have ten times as many variables as the traditional counterpart. But the additional variables are all simple bounded variables. The number of real constaints stays the same.[2] With modern computational power and algorithms, the difference will rarely be critical.

The the fact that the above formulation is a perfectly good linear program and hence solvable can be regarded as another proof — as if we needed one — of the existence of a competitve market equilbrium. But unlike many classical proofs it explicitly incorporates transportation costs and does not depend on assuming convex supply and demand curves.

[2] If the LP is solved via a Simplex algorithm, computation is roughly linear in the number of variables and cubic in the number of constaints. A simple bound does not count as a constraint.

4 A More Real Network

In the real world, unit transport costs between m and n are almost never a single fixed number. More typically, we must allocate limited transport resources to the various routes. Modelling this allocation will depend on the particular problem; but our goal here is to show that we can handle all sorts of non-linearities within the transport sector, while staying within the confines of a linear program.

Take for example the world crude oil transportation network. In the short run, the deadweight tonnage available in each tanker size category k is fixed at say K_k. To move a ton of oil on route mn with ship category k requires W_{kmn} deadweight tons. In addition there will be a marginal cost (mostly fuel) F_{kmn} associated with employing a dwt ton of ship category k on route mn

The new problem is:

$$max \sum_n \sum_i P_{ni} d_{ni} - \sum_m \sum_j C_{mj} s_{mj} - \sum_k \sum_m \sum_n F_{kmn} t_{kmn}$$

subject to

$$\sum_m \sum_n W_{kmn} t_{kmn} <= K_k \quad for\ all\ tanker\ categories\ k$$

$$\sum_j s_{mj} >= \sum_k \sum_n t_{kmn} \quad for\ all\ sources\ m$$

$$\sum_i d_{ni} <= \sum_k \sum_m t_{kmn} \quad for\ all\ markets\ n$$

$$0 <= s_{mj} <= S_{mj} \quad for\ all\ j\ and\ for\ all\ m$$

$$0 <= d_{ni} <= D_{ni} \quad for\ all\ i\ and\ for\ all\ n$$

where t_{kmn} is the amount shipped from m to n on tanker category k.

The new set of constraints ensure that the total amount of deadweight employed in each tanker category is no more than the deadweight available. The dual of each such constraint is the momentary scarcity value of that size tanker. If the tanker market is competitive, this will be the category's spot rate.

The interesting numbers here are the W_{kmn}'s. W_{kwn} is largely a function of route length. But the pre-processer which figures out W_{kmn} for each ship category and route can handle all sorts of non-linearities. If told enough about the ship and the route's terminals, it can figure out the time loading and discharging, correctly determining that ships spend relatively more time in port on short routes than long. If it is told the terminal and draft restrictions on each route, it can throw out ships that are simply too large for the route, part-load others. It can figure out deadweight lost to bunkers. If it is told the crude density on each route, it can figure out whether the ship is weight or volume limited. If the route involves a canal, it can figure out the time lost in canal transit and throw the canal toll into the marginal cost.[3]

In short, many non-linearities within the transport sector can be handled in the pre-processing step. In the case of the tanker network, the pre-processor which figures out the unit tanker intensity for each route and ship category is essentially a full fledged voyage estimator

[3] In the real tanker world, there is not a one-to-one correspondence between source-sink combination mn and route. In some cases, there is no route between a particular producing area and market. In many cases, there are multiple routes. From the PG to Europe, you can go around the Cape both ways, through the Suez one way, and through the Suez both ways. The correct choice depends on bunker price, Canal tolls, and the momentary scarcity of tankers. Models of real world networks usually have to maintain a separate route list which includes the route's starting and ending node; and then figure out which route goes into which constaints in the above linear program.

This allows backhauls — routes with more than one loaded leg. The only difference between a backhaul and a normal route is that a backhaul route shows up in multiple supply/demand constraints.

And then there are transhipment nodes, some of which are serviced by pipelines. For each such transhipment node, such as Gulf of Mexico lightering, there is a simple mass balance constraint: what goes in, must come out.

5 A Still More Real Network

The tanker market is still more interesting than we have indicated. In the short run, a tanker owner must decide not only which route to put his ship on, but also the ship's speed. A ship's fuel consumption curve is highly non-linear. Figure 3 shows a typical fuel consumption curve for a 2000 built diesel VLCC. Fuel consumption is very roughly cubic in speed.

This means that the supply curve of tankship services is critically dependent on bunker price. The red line in Figure 4 shows the *speed up curve* for our typical VLCC when bunkers are \$50 per ton, about what they were in 1999. The speed-up curve displays optimal steaming speed as a function of spot rate. At any given bunker price, a tanker owner must balance the extra fuel cost of speeding up versus the additional revenue earned. This in known as *slow-steaming.* If bunkers are \$50 per ton, the owner comes out of lay up as soon as the sport rate reaches Worldscale 8, and it only takes another 7 WS points before he should be steaming at full speed. This curve is very J-shaped. A little reflection will reveal that the ton-mile supply curve of the fleet as a whole is simply the sum of all the fleet's speed up curves, and it too will be very J-shaped.[4]

If on the other hand, ther bunker price is \$450 per ton as it was in end-2007, then the VLCC speed up curve looks like the blue in Figure 4. Any any rate below, WS29 the owner is better off laid up. At about WS30 he should come out of lay up, but at this rate his optimal speed is only 11 knots. As rates improve, he slowly speeds up, but he does not reach full speed until the spot rate is WS 100. This speed up curve is far less J-shaped, far more elastic than the \$50 per ton bunkers curve. The same thing is true of the supply curve of the fleet as a whole. In a Worlscale 50 market, 15% more ton-miles will be supplied at a bunker price of \$50 per ton than at a bunker price of \$450 per ton. In short, any model of the tanker market that hopes to be realistic must model slow-steaming correctly; and that means modelling the highly non-linear fuel consumption curves for each tanker category.

How can we incorporate such non-linearity into a linear program? No prize for guessing that once again we approximate the fuel consumption curve by a series of steps. More precisely, we allow the owner only a discrete set of choices from the fuel consumption curve. For example, we might limit

[4] Tankers actually have two fuel consumption curves: loaded and ballast. The best way to mush these two together, as I have done in 4, is to assume that the optimal loaded consumption per day is the same as the optimal ballast consumption per day. This will result in a ballast speed 1 to 1.5 knots faster than a loaded speed.

him to picking only one of the five speeds shown as red dots in Figure 3.[5] We will index this discrete set of speeds by l. Associated with each combination of ship category k and steaming speed l, and route mn, there will be a marginal cost F_{klmn} and dwt intensity W_{klmn}, the deadweight required to move a ton of oil on route mn in ship category k steaming at speed l. It will be the job of a pre-processer to figure out all the F_{klmn}'s and W_{klmn}'s for the given bunker price.

The new optimization problem is a straightforward extrapolation of the last.

The new problem is:

$$max \sum_n \sum_i P_{ni} d_{ni} - \sum_m \sum_j C_{mj} s_{mj} - \sum_k \sum_l \sum_m \sum_n F_{klmn} t_{klmn}$$

subject to

$$\sum_l \sum_m \sum_n W_{klmn} t_{klmn} <= K_k \qquad for\ all\ ship\ categories\ k$$

$$\sum_j s_{mj} >= \sum_k \sum_l \sum_n t_{klmn} \qquad for\ all\ sources\ m$$

.

$$\sum_i d_{ni} <= \sum_k \sum_l \sum_m t_{klmn} \qquad for\ all\ markets\ n$$

$$0 <= s_{mj} <= S_{mj} \qquad for\ all\ j\ and\ for\ all\ m$$

$$0 <= d_{ni} <= D_{ni} \qquad for\ all\ i\ and\ for\ all\ n$$

where t_{klmn} is the amount shipped from m to n on tanker category k at speed l.

The magic of optimization guarantees that the program will pick that speed for each ship and route which maximizes the overall surplus. The program neither knows nor cares that the fuel consumption curve is highly

[5] This is not much of a restriction. Ship owners rarely attempt to specify steaming speed more accurately than a knot, in part because vagaries of weather, current, etc would render any such attempt futile.

non-linear, as long as it is an increasing function of speed.[6] It should be clear by now that you can play this trick for just about any non-linearity, converting non-linearity into a discrete set of variables.

Of course, each time we do this we increase the number of variables by a lot. In the case of slow steaming, the fuel consumption curves can be adequately modelled by about 5 steps, so we have roughly 5 times as many variables. But once again the efficiency of linear programming algorithms and modern day computer power are usually up to the task.

Take for example Martinet, one model of the world petroleum transportation network. It has about 30 producing areas, about 30 consuming regions, about 10 transhipment nodes, some 900 routes, 15 different tanker categories, each of which has about 5 steaming speeds. In total there are about 150 constaints (not including bounds), and 20,000 variables. It runs on a quarterly period. A 12 quarter run of this model on a standard PC takes less than 3 seconds.

Martinet assumes inelastic oil supply in each producing area and inelastic oil comsumption in each market. But, as you can see from the above formulation, assuming ten step supply and demand curves at each such node would not make much difference. Most of the variable are the t_{klmn} which dont change. We would turn 60 variables into 600, so the total number of variables would increase by less than 10%. The number of real contraints doesn't change. The number of bounds would go from 60 to 600, but modern LP algorithms are nearly insensitive to the number of bounds.

[6] To avoid multiple optima.

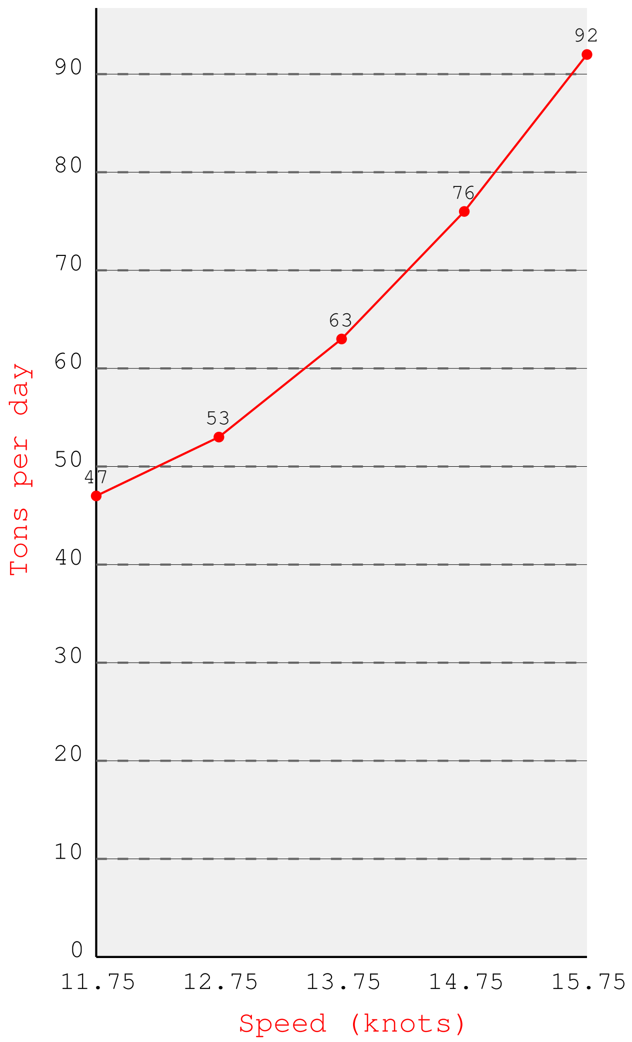

Figure 3: Typical Fuel Consumption Curve, VLCC Built 2000

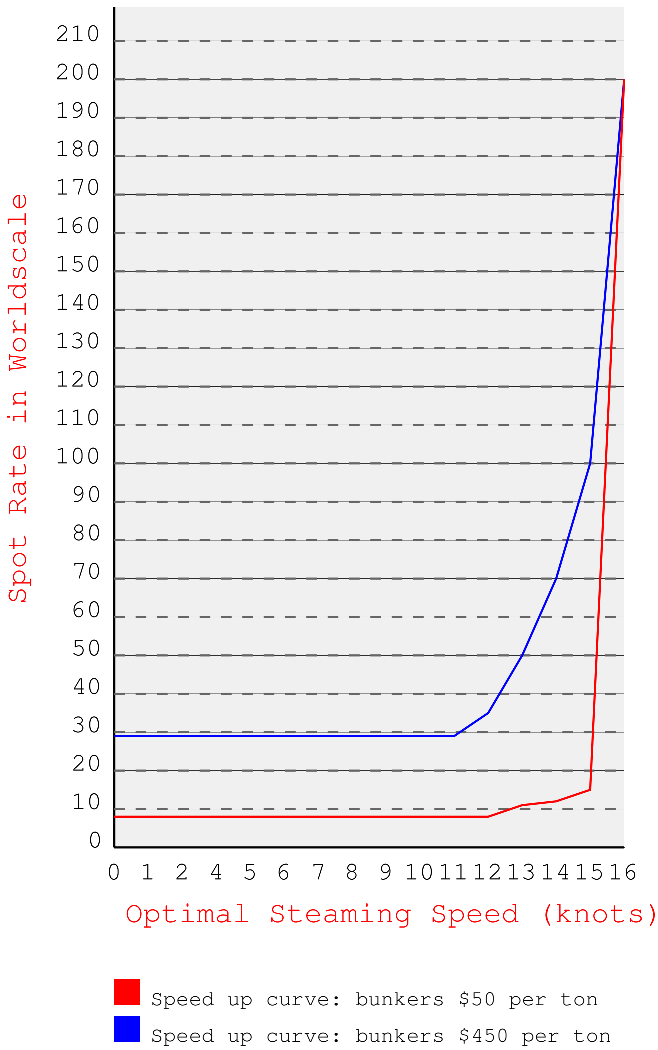

Figure 4: Hi/lo BFO Price VLCC Speed-up Curves

6 Conclusion

A wide range of non-linearities can be incorporated into transportation network problems and still be feasibly solved by linear programming techniques. These include source supply curves, sink demand curves, and non-linear fuel consumption curves.