April 3, 2004

A Conceptual Model for Demands at the Pool Level

Contract No: DACW72-99-D-0005
Task Order # 0140: U.S. Army Corps
Of Engineers Revealed Choice
Estimate of the Demand for Barges
On the Mississippi River

Kenneth D. Boyer
Michigan State University

Introduction and Motivation

This research project is designed to provide information on the willingness to pay for navigation improvements at the pool level. It will produce estimates of the elasticities of demand for navigation services that can be used to calculate the gross benefits to users of increases in the speed of navigation at each of the pools.

Credible estimates of demand elasticities must have at least the following properties:

1) They must be based on plausible and realistic decision settings. These settings must be consistent with economic theory and must represent believable alternatives from the user public.

2) The models must be spatially motivated. It is widely recognized that transportation economics is motivated by spatial interactions, and these interactions must be at the center of the model to be estimated.

3) The model must be identified and thus capable of separating the effects of shifts in the supply of transportation services from changes in the demand for those services.

4) The model must be estimated at a level of aggregation that allows us to use the results to make useful estimates of the benefits provided by facilities expansion.

There are no published transportation demand estimations for any mode or facility that meet all of the four criteria listed above. Freight transportation demand elasticities are notoriously difficult to calculate, and thus few have been attempted.[1] Four primary difficulties have been cited as reasons for the unsatisfactory state of freight transportation demand elasticities: 1) data availability; 2) econometric identifiability; 3) the complexity of the choice setting of the problem; 4) Transferability—meaning the ability to use demand estimates from one transportation system to infer similar information about other systems. The last element on this list should not concern us since we are not intending to make inferences about the demand for shipment on other waterways or for other modes based on the estimations made here. The first three, however, need comment.

The unique and exciting aspect of this research project is that an opportunity is present to solve each of the other three problems that have held back freight transportation demand estimation. If this project is successful, we will have made the first fully-identified freight demand estimations that are directly usable for economic valuation of the benefits of transportation facilities improvement. This is possible for two basis reasons: first, data are available at a level of disaggregation that is unique in the

[1] Tae Hoon Oum, W.G. Waters II, and Jong Say Yong, "Concepts of Price Elasticities of Transport Demand and Recent Empirical Estimates," *Journal of Transport Economics and Policy*, Vol. 26, No.2 (May 1992), pp. 139-154.

transportation sector. And second, the linearity of the system and limited number of commodities handled give us the hope that the complexity of system can be simplified to something that is estimatable. While it handles a number of commodities and handles shipments of coal, petroleum, fertilizer and aggregates, the Upper Mississippi Waterway system specializes predominantly in handling grain, particularly soybeans and corn, with grain overwhelmingly moving from north to south. This makes modeling the system much easier than the multi-commodity highway or rail systems that move goods in all directions.

Choice Setting

This initial focus of this project will be on grain movements, with shipments of corn being the first grain to be modeled. The same basic method can be used for all grains and seasonal commodities that have a uni-directional commodity flow. As is the tradition in transportation, demand analysis, we will model freight transportation demand as the result of a sequence of decisions. For purposes of exposition, we will characterize the decision maker as the farmer, though we recognize that the authority to divert grain may be passed to elevators in the system.

1) A farmer decides on the number of acres to devote to corn. This decision is based on the forecast of market conditions for corn at harvest time as well as the forecast of market conditions for other commodities. Farmers in different locations may have different alternative crops that they might plant if market conditions for corn are projected to be poor. A reasonable first estimate of the decision to plant corn may be based on the net price that the farmer received for corn at the end of the last growing season. I hope to have an estimate of the number of acres planted to corn in areas in

rough proximity to each pool and to see whether this can be predicted by the market price minus a forecast of shipping costs. The decision on the number of acres of corn to plan is characterized as long run and beyond the basic structure of demand modeling attempted here. It is not clear that there will be enough independent variation in the data to allow us to estimate this effect in the short ten-year data series available to us, but it is worthwhile investigating after the other estimations have been performed.

2) A farmer decides on the level of attention to devote to the crop based on contemporaneous market conditions as well as growing conditions. I think it is unlikely that we will have sufficient data to estimate this effect, which is likely to be small in any case, but it is worth listing for the sake of completeness.

3) A farmer decides how much of his crop to harvest and how much to plow under based on the net price available to him at the elevator that is in a position to offer the highest net price. I doubt that data will be available to estimate this effect, and again the effect is likely to be minor. I expect, however, that an index of total harvest at different points along the waterway will be available.

4) A farmer decides on the timing of the movement to and release of harvest from storage elevators depending on current conditions and anticipated future market conditions. This is one of the three key decisions affecting demand for transportation that we should be able to deal with. I do not expect to being able to model decisions made using speculative motivations. I do, however, believe that we will be able to measure the extent to which a harvest is being held back from normal shipping patterns in response to system congestion or whether the harvest is being unusually rushed to the river to take advantage of exceptionally high prices.

5) A farmer decides whether to deliver the harvest (or to allow the release of the harvest from the country elevator where it is stored) to the river for export by water. There are many alternatives to river transportation. For example, the farmer could deliver the harvest to an elevator that would send the harvest by rail or truck to an alternative port (e.g., Duluth or Portland), deliver the grain to a local processor, or delver the grain to an elevator that will load it on a train for delivery to a point that bypasses the lock system. We will give the name "leakage" to the loss of harvest to modes, uses, or destinations or than the delivery to the river for export.

Leakages represent one of the two basic short- to medium-term sources of flexibility that farmers will have that will determine the elasticity of demand for transportation. We should not presuppose what the best alternative is for each farmer—whether it is rail to a location off of the UMR/IW system, local consumption, or even trucking to a distant port. For this reason, we will not try to model and estimate the geographic structure of alternatives available to each shipper. It is clear that each farmer or elevator will have a different options available and a different ordering of alternatives. Rather, we will infer the extent of leakage of grain from the river system by the difference between the amount of grain that would normally be delivered to the River based upon the harvest in each year, and the amount that ultimately is delivered to the rivers.

6) A farmer, having decided to deliver the harvest to the river, decides which pool to deliver the harvest to. Farmers can reduce their exposure to congestion by delivering grain to a pool farther south on the system. This decision will be based on the pool whose elevators pay the highest price for the harvest as well as the cost of trucking the

corn to the river. It should be noted that the cost of road transportation declined precipitously during the period covered by this study and thus is should be possible to observe in the data a trend of increased flexibility in the port of delivery[2]. Beyond the inventorying of grain and the price-induced leakages of harvest in decision 5) above, this decision about the river location to which to deliver grain is the third source of flexibility in response to prices that we will measure as we estimate the elasticity of demand for navigation services. We will give the name "lock bypass" to this loss of distance shipped in response to congestion. Ultimately, if the system is extremely congested, delivery of grain to a point below the last lock might be economically justified, in which case the harvest is lost to the system and the lock bypass becomes a leakage.

Measurement of the price charged for navigation

Elasticities of demand measure the extent to which farmers (or their agents) will purchase fewer ton-miles of navigation services on the River as the price of those services rise. To estimate demand elasticities, it is critical to measure either prices charged or a proxy for the prices charged for navigation services. In this study, three different measurements of prices will be used, each of which has certain advantages and disadvantages.

1) <u>A direct measurement of prices charged</u>.

This is apparently the most accurate measurement of the charges that grain shippers will face. It should be noted, however, that in the case of other modes of transportation, these figures are not available—at best, list prices are published for most

---

[2] See, for example, Kenneth D. Boyer and Stephen V. Burks, "Productivity and Cost Trends in the Trucking Industry, 1977-1997," Working paper at Michigan State University, April 2004.

modes of transportation, with the actual price paid being the result of a secret negotiation between the buyer and seller of transportation services. Where there is public posting of prices (air fares, for example) the price structure is stunningly complex, making it impossible for an outside observer to even guess at the price paid by the marginal unit. Given the simplicity of the River system, it is possible that the same data problems that have prevented researchers from observing the price paid by the marginal shipper may be less extreme here. It is clearly prudent, however, to think about other proxies for prices charged.

2) The difference in the bid price for grain at different pools.

In a perfectly competitive market with identical products, grain buyers would be willing to pay prices for grain at different pools that exactly match the difference in the cost to them of transporting grain for longer or shorter distances. Thus the difference in the price of corn in two different pools should be equal to the cost of transportation between those two pools. We believe that we will be able to observe at least some of these prices and we hope to analyze whether they fit well as the price of transportation. However, the inference from price differences in different locations to the cost of transportation between them depends critically on the competitive nature of markets. Grain marketing is, however, dominated by large firms and the relationship between grain exporters, grain elevators, and shipping companies is complex and not characterized by the sort of arms-length dealings that is assumed by the competitive model. It remains to be seen whether in fact this is a good measure of the price of transportation.

3) A constructed price based on the speed of transportation. The logic of this measurement is that towboat costs are quoted by the hour, and thus the cost of using the system should be proportional to the length of time it takes to navigate the river system. The more congested is the system, the longer will each voyage take, and thus the higher will be the price charged for navigation services. Similarly, shipments from locations farther upstream will take longer and thus necessarily will have higher shipping costs per ton for delivery to the exporter.

The cost per ton will also depend on the hourly charge of using a towboat and this logically should depend on the overall usage of the entire river system, not merely the part that is the focus of this study. To the extent that towboats can be shifted from service on other rivers when the demand for navigation on the Upper Mississippi/Illinois River Waterway is high, the prices charged for towboat services should reflect the higher opportunity costs of using them.

In addition, transportation economics tells us that the price of making a run should depend on whether a backhaul is available. It is predictably more expensive to hire a towboat and barge flotilla that comes upstream empty to pick up a load than to use the service if the upstream and downstream movements of commodities and power units is more balanced. Data are available on the frequency of empty movements into a pool as well as on the overall level of equipment usage in the river system, thus suggesting that this proxy for the price of transportation may be most feasible.

Part of this project will be a study of the relationship between all three of these measurements of prices. Under ideal circumstances, it may be possible to use relationships among them to verify data validity and to fill in data gaps and to identify implausible data items in each series.

Spatial Aspects of Demand

The geographic structure of the navigation system allows us to identify some regularities of the demand system. With the basic movement of grain from north to south, grain that is put on the river in the north traverses some of the same stretches of river as grain that is brought to the river lower down. If congestion slows the system, or if the hourly opportunity cost of using a boat rises, the per-ton price increases must be greater for shippers upstream than downstream. In short, supply shifts for each pool will have a similar pattern over time, but the effect will be exaggerated for shippers located farther north. This expectation, which is derived from spatial aspects of the problem, can be used to improve the estimation of demands at the pool level.

A second spatial aspect of demand is the prediction of the location of shifts in locations at which grain is put on the River in response to the higher price of shipping. If the cost of navigation rises, upstream elevators will not be able to pay as much for the crop. The price paid for the crop by riverside elevators should fall faster at upstream locations than at downstream ports. Profit maximizing farmers, deciding on which river location to access river transportation, will shift the location of the point at which they access the river, moving south in response to rises in navigation costs. Thus, an unexpectely low shipping level at an upstream pool should be balanced by an unusually high level of shipping from a pool below. Since we will wish to estimate demands at the pool level, knowing the spatial sequencing of pools allows us to specify an error structure to the individual pool-level demands. This will help to improve the estimating efficiency of the model.

A structural model

At the beginning of this paper, we identified three major short term responses that farmers might have to higher shipping prices on the River: they can delay shipment until prices drop, they can divert their harvest into local consumption or an alternative off-river or below-St. Louis outlet, and they can deliver their crops to a river location farther south, but still above- St. Louis. In addition, in the long run, they can change their planting and cultivation patterns to economize on river transportation.   The goal of this project is to measure the extent to which, in the aggregate, at each pool, shippers use each of these strategies in response to rising shipping prices. The result can be quoted as an elasticity of demand and can be used to estimate the marginal value that shippers would place on facilities improvements that reduced the cost of river transportation.

The basic estimating form to be employed is a multiplicative form defined at the pool level for a week of grain shipments:

$$Q_{yip} = A_p(H_{yp})(W_{yip})^{\beta 1}(P_{yip})^{\beta 2}e_{yip} \quad (1)$$

Where Q_{yip} is the quantity of grain shipped southbound from pool p, in week i of year y.

A_p is a constant term for pool p.

H_{yp} is an index of the harvest level in the year y-1 in the hinterland of pool p.

W_{ip} represents a set of i weekly dummy variables to capture the seasonality of grain shipments from pool p.

P_{yip} is the price of shipping from pool p to pool 30 in week i of year y

e_{yip} is the error term associated with the southbound movement of grain from pool p in week i of year y. In practice, equation (1) will be estimated in the log form.

Equation (1) says that the amount of grain shipped out of pool p (all of which is assumed to be ultimately bound for delivery beyond pool 30) in week i, is proportional to the harvest in the hinterland to the pool, distributed according to the weekly seasonal pattern of shipment. Changes in the price of shipping from each pool then has a constant proportional effect on the quantity moved based on the seasonal distribution of the harvest. The parameter β2 is the first-pass estimate of the elasticity of demand for shipping from pool p.

To take into account the fact that all pools are affected simultaneously by changes in world market conditions, we will estimate equation 1 as a system of simultaneous equations in which the error terms are contemporaneously correlated. This will allow shipping spikes that result from sudden shifts in export markets to be treated as affecting all pools rather than being the effect of a change in the price of shipping.

Equation (1) can be improved by recognizing that deviations from the normal seasonal shipping pattern early in the shipping season may have an effect later in the season as grain stocks pile up or are drawn down faster than normal. To take this into account, equations (1) can be modified below as shown in equation (2)

$$Q_{yip} = A_p(H_{yp})(W_{yip})^{\beta 1}(P_{yip})^{\beta 2}(S_{yp})^{\beta 3}e_{yip} \qquad (2)$$

Where S_{yp} is a measure of the ratio of the weekly storage to the normal storage level in pool p.

The parameter β3 represents the elasticity of shipping with respect to unusual inventory levels. If β3 is insignificantly different from zero, we can conclude that price-induced deviations in normal shipping levels earlier in the season do not affect shipping levels later and thus the reduced and induced traffic are gains or leakages from the river system, rather than simply a shifting in the timing of shipments. If β3 is significant, we can conclude that inventory levels have an effect on shipments independent of the level and seasonal pattern of the harvest. In this case, shippers react to high transport prices by holding back commodities, storing them until congestion eases. The seasonal pattern of grain shipments from a pool will already include this effect, and thus the parameter β3 represents the transportation demand response to unusual inventory build-ups or draw-downs.

Inventory levels may be available at the pool level or they can be constructed as the accumulated deviations in any pool of the difference between the forecasted grain flows out of a pool and the observed flows since the last harvest. If they are constructed, then equation (2) must be estimated recursively since it depends on the difference between forecasted and actual flows and the forecast is based on parameter estimates. The equation will be estimated repeatedly with parameter estimates from one pass used to create forecasts of inventory holdings in the next, until the parameter values and variable values stabilize.

The introduction of inventory holding as a response to congestion changes the calculation of demand elasticities since high prices during one week may cause transporters to hold grain for later shipment; a direct measure of the price elasticity during that week would falsely lead us to conclude that the shipment had been lost to the system. Instead, the elasticity will be constructed through a simulation in which the transport price from one pool is arbitrarily raised by 1%, with the week-by-week reduction in shipments from a pool calculated along with the implied inventory build-up. The parameter value on the inventory measure then tells us the rate at which the unusually high inventory levels cause shipments from the pool to be accelerated. At the end of the shipping season, we can calculate the extent to which reduced shipments from the pool due to the direct effect of higher shipping prices have been offset by higher than normal shipments due to increased inventory levels later in the season. The difference will be assumed to be leaked from the system and lost to the river as a result of higher transport prices.

Neither Equation (1) nor (2) have explicit spatial motivations. To enter such considerations it is necessary to identify linkages between small numbers of markets. Spatial considerations tell us that spillovers between markets are likely to be in close pools rather than in distant pools. The most direct way to do this is to introduce into the model two unobserved transfer variables, each of which depends on the price of transportation from the pool:

$$Q_{yip} = A_p(H_{yp})(W_{yip})^{\beta 1}(P_{yip})^{\beta 2}(S_{yp})^{\beta 3}e_{yip} + T^{+}{}_{vip-1}((H_{yp-1})^{Pyip-1}) - T^{-}{}_{vip+1}((H_{yp})^{Pyip}) \quad (3)$$

Where $T^{+}{}_{vip-1}((H_{yp-1})^{Pyip-1})$ is the amount of harvest transferred to a pool from the hinterland of the pool immediately upstream. This amount is assumed to be related to the level of harvest in the hinterland of the pool immediately above and to the price of transportation on the system.

$T^{-}_{vip+1}((H_{yp})^{Pyip})$ is similarly the amount transferred from the pool to the next pool down the river. This amount will appear as T^{+} in the regression for the next pool downriver.

Equation 3 allows for the inclusion of a cascade of transfers from the hinterland of the pool immediately above to the pool immediately below, with the level of transfers being determined by the height of the transportation price for river transportation. Since there has been a large secular decline in the price of trucking during the period covered by this study, it might be appropriate as well to include a trend variable in the function determining the transfers.

The drop in trucking costs also allows for the possibility that lock bypass may shift deliveries to more than one pool farther south. It is straightforward to allow for such transfer functions. However, allowing for such possibilities reduces the compactness of the estimating form, reduces the simple spatial motivation, and introduces the possibility that true randomness will be misinterpreted as a systematic bypass effect. My preference is to use a goodness of fit test to guide the expansion of bypass beyond one lock, and then confirm pool-skipping bypass through interviews. Similarly, if through interviews or freight waybill data, we can confirm the existence of a rail link with corn flows moving from an interior point to a river terminal in a different pool, we can introduce this bypass through the inclusion of an ad-hoc transfer.

The combination of equations 1, 2, and 3 will allow us to get an empirical estimate of the extent to which higher prices for river transportation cause grain shippers to respond in the three ways that would cause them to economize on river transportation: 1) delaying or accelerating transportation to avoid congested periods; 2) leaking grain to destinations that are off-river of below St. Louis; or 3) choosing a location to access river transportation that involves more trucking and less river transport.

The Supply Side

This study is intended to analyze the sensitivity of shipments to changes in the characteristics of the navigation system. The characteristics of locks and channels are considered to be exogenous, and not influenced by demand decisions. However, like all transportation systems, transportation supply in the Upper Mississippi River is the result of interactions between vehicles and infrastructure. The vehicle supply is shared between pools and across river systems.

The basic observation in this study is flows from an individual pool. As is typical in a demand study, the fundamental variable of interest is the price of transportation services. The price is the product of the hours of transit time to pool 30 and price per hour of navigation services. The latter is determined by the relationship between the demand for equipment on the river system as a whole. Transit times, however, are logically endogenous to traffic levels from each pool.

Identification

Both demand curves and supply curves are relationships between prices and quantities. The first task of a demand study is to insure that what is observed is a relationship that can be thought of as a demand curve rather than supply relationship. In this case, transportation prices and demands are both seasonal with high prices and high quantities transported occurring at the same time. A simple correlation would show a positive relationship between prices and quantity, in contradiction to the downward slope given to the demand curve by economic theory.

Identification of the demand part of the price/quantity relationship will be achieved by standard instrumental variables techniques. The key to identification is the fact that observations are at the pool level while the prices charged for transportation depend on demands not only from the pool from which there is a shipment, but from pools up and downstream as well, in addition to demands from other commodities and rivers that share the same equipment.

The identification of pool level demands will be accomplished by including seasonal variables in the instrument set. In addition, the river system-level of congestion will be exogenous to the price at any pool. Barge availability in any pool, the extent to which movements represent backhauls as opposed to purposeful placement of barges, and the price of inputs to the barge industry can all be used as instruments.

Pooling tests

As currently envisioned, each pool will have a separate regression with no cross-equation restrictions on the size of price parameters. Since the motivation for the estimations is spatial rather than from the theory of the firm, there is no logical reason why the parameters should be related. The price coefficient represents the extent to which farmers have alternatives that become attractive as the price of river transportation rises. There are numerous factors that determine the attractiveness of alternatives—for example, the availability of local processing plants in an area or the presence of good rail lines, or the proximity of other ports.

It would nonetheless be gratifying if there were similar demand elasticities at individual pools, suggesting that there were fundamental spatial reasons for the demand elasticities rather than sensitivities being the result of the geographic idiosyncrasies of different regions. Pooling tests will be performed to determine whether there regularities in the demand responses of different regions. Since the price variation in upriver locations is logically much higher than for shippers closer to pool 30, it is reasonable to expect that demand elasticities will be better measured for pools closer to the Twin Cities than to St. Louis. It is not necessarily true, however, that demand elasticities should be higher in the north than farther south. Standard pooling tests will be performed to see whether such relationships exist.

Measuring multiple types of elasticities

The basic elasticities of interest in this study are those relating to the pool-level response to changes in the price of transportation. These will be useful for measuring the benefits of facilities improvement as well as analyzing factors that affect the overall usage of barges for moving cargoes on the River. All of these effects will be analyzed through simulation. For example, it will be interesting to use the results to simulate the effects of an increase in the cost of inputs to the barge industry (as, for example, might occur due to an increase in fuel or labor costs). The effect will be to change prices for service more at upstream locations than downstream and change the geographic composition of traffic. From this we can infer to what extent transportation is rationalized by shifts in the composition of traffic rather than from the responses from individual shippers.   We can then calculate the proportionate change in ton-miles that would be induced by a proportionate change in prices—a figure that is generally considered to be the elasticity of demand for a transportation service.

An alternative way of representing the results is to simulate the effect of improving an individual facility. This, too, will be the result of a simulation, in which the price of barge services decline in proportion to the saving in time, with different percentage price drops estimated for different pools. From this, it should be possible to estimate the reduction in leakages and the upstream shift in delivery locations that result from the lockage and from that derive a saving to shippers and another elasticity of demand.

Long-run feedbacks

The goal of this exercise is to measure the responses of grain shippers to changes in the cost of using the system.   The demand elasticity that we are focusing on is the short-to medium-run where planting is assumed to be exogenous to the cost of river transportation and where shippers decide on routes and destinations for their crops. As noted earlier, however there is also a long-run feedback loop from the profitability of using the river system to the decision on crops to plant and the level of intensity to apply to cultivation. Farmers who anticipate high transportation costs for getting their crops to market may not plant.

It would seem to be a simple matter to see if there is a connection between the level of transportation costs and the extent of planting. This project will try to estimate the final part of the feedback loop and find such a relationship, but in practice, however, it seems unlikely that this relationship can be uncovered. The primary reason is our inability to accurately measure expectations on transportation costs as well as on weather, the cost of other inputs, and world market conditions. For example, while transportation costs may have been high in one year, farmers may have considered these conditions to be unusual and not likely to be repeated the following year.

Nonetheless, this project will try to estimate the average annual cost of grain transportation from each pool and correlate that with planting in the hinterland of the pool the following growing season. The use of lagged variables to model expectations means that one of the data points will be lost, reducing our panel from ten years to nine. With such a small number of data points, the feedback relationship between transportation costs and planting will only be uncovered if there are large swings in the average cost of river transportation from year to year. It seems unlikely that these swings will be observed in the data set, and so we are pessimistic about the ability to estimate this final element of long run demand, but it is an exercise worth doing and we will attempt it.