

RESEARCH SEMINAR IN INTERNATIONAL ECONOMICS

Gerald R. Ford School of Public Policy
The University of Michigan
Ann Arbor, Michigan 48109-1220

Discussion Paper No. 502

**Storage, Slow Transport, and the Law of One Price:
Evidence from the Nineteenth Century U.S. Corn Market**

Andrew Coleman

University of Michigan

February, 2004

Recent RSIE Discussion Papers are available on the World Wide Web at:
<http://www.fordschool.umich.edu/rsie/workingpapers/wp.html>

Storage, slow transport, and the law of one price: evidence from the nineteenth century U.S. corn market.

Andrew Coleman*

February 2004

Abstract

This paper develops a rational expectations model of physical arbitrage incorporating storage and trade to explain how markets are integrated when trade is costly and non-instantaneous. The paper finds a striking empirical verification of the model from an analysis of the late nineteenth century corn markets in Chicago and New York. The dataset is particularly high quality and includes weekly data on spot and future prices, storage quantities and the cost of three modes of transport for a fourteen year period. In keeping with the model, it is shown that the New York spot price frequently exceeded both the New York futures price and the Chicago spot price plus the transport cost by several percent when inventories in New York were low, but not when they were high. The paper also derives a supply of storage curve for New York corn and argues it can be explained as the outcome of rational arbitrage when transport is slow.

*Department of Economics, University of Michigan. coldman@umich.edu

The model in the paper was formulated as part of my thesis at Princeton University, and I am grateful to Professors Angus Deaton, Kenneth Rogoff and especially Chris Paxson for helpful comments and suggestions. An early version of the paper was circulated under the title “Storage, Arbitrage, and the Law of One Price” and benefited from comments by participants at seminars at Boston University, UCLA, the Kennedy School, the University of Michigan, the University of Otago and Princeton University. The data was collected subsequently, and I would like to thank Ed Knotek for usefully assisting in this task. I would also like to thank Stephen Salant, Linda Tesar and Jeffrey Williams for insightful comments on the full draft.

1. Introduction

In the late nineteenth century, complex marketing infrastructures were developed in Chicago and New York to facilitate the export of grain from the Great Plains to Europe. Each year, millions of bushels of corn were sent to Chicago for sale to shipping agents, who transported the grain to New York, where it was resold, transferred to elevators, and then shipped to European markets. Both cities had elaborate facilities for receiving, storing, and forwarding grain, and financial exchanges that offered an array of liquid spot and futures markets. In such a setting it would be imagined that prices always obeyed the law of one price, i.e. that the price of corn in New York was always equal to the price of corn in Chicago plus the cost of shipping. However, this was not the case. Although the law of one price held most of the time, on numerous occasions — approximately ten percent of the weeks in the period under analysis — the New York spot price spiked upwards and reached a level considerably higher than the Chicago price plus the transport cost. On these occasions, the spot price also exceeded the price for delivery in New York the following month, typically by the same amount¹. These price spikes were temporary and almost always occurred when New York inventories of corn were extremely low. Although infrequent, the fact that these spikes occurred throughout the period in one of the world's most organised markets suggests that they were not incidental deviations from the law of one price but an essential feature of the arbitrage process.

This paper develops a simple model of price arbitrage that explains these price patterns. The model builds upon a recent series of papers that examine how prices in separate competitive markets would be determined if agents optimally used storage and transportation systems to arbitrage prices. In the key paper in this series, Williams and Wright (1991) developed a model to investigate how commodity prices in two different locations would be determined if forward looking, rational, and risk neutral arbitrageurs could either store goods or transport them instantaneously from one location to the other. They showed that storage should have an important role in the arbitrage process, as intertemporal arbitrage could be used as a cheap substitute for transportation to ensure that prices in the two locations never exceeded each other by more than the cost of transport. Their model captures one aspect of the data as they showed that if neither centre had inventories then the spot price in each centre would exceed the future price, and that such “stock-outs” would occur regularly in equilibrium. However, the price

¹ If the spot price exceeds the future price, prices are said to be in backwardation.

difference between centres would never exceed the transport cost in their model, since transport was instantaneous.

This paper argues that the Williams and Wright model can explain all of the observed price relationships in the corn market if the assumption that transport is instantaneous is relaxed. To make this argument, a rational expectations model of commodity price arbitrage is specified in which it takes one period to ship goods. It is shown that if inventories in both centres are positive, the difference between the spot prices will be less than or equal to the difference between the transport cost and the storage cost. However, if inventories in a centre fall to zero, the spot price will spike upwards and exceed the spot price in the other centre plus the difference between the transport cost and the storage cost. Since the price spikes only last until supplies are replenished, the spot price also exceeds the local future price on these occasions. Furthermore, by examining the condition for profitable shipments, it is shown that storage in one centre should regularly fall to zero. Therefore, this model has the implication that there should be regular occasions when the spatial price difference exceeds the difference between the transport cost and the storage cost and when commodity prices are in backwardation.

The intuition of the model is straightforward. Rational, risk neutral arbitrageurs export from one centre to the other if the expected future price in one centre exceeds the price in the other by the cost of transport. If there are sufficient inventories in the importing centre when the goods are sent, arbitrageurs there run down their inventories until the spot price equals the expected future price minus storage costs, and thus the spot price difference between centres equals the transport cost minus the storage cost. If inventories in the importing centre are zero, however, the spot price cannot be arbitrated down and the price difference between the centres exceeds the difference between transport costs and storage costs until the shipment arrives. This explains why the spatial price difference can exceed the transport cost². The spatial price difference should exceed the transport cost on occasion because exporters limit the quantities they send to ensure that random fluctuations in demand regularly cause inventories in the importing centre to fall to zero. Unless they do this, the exporters make negative profits because the price in the importing centre does not cover the cost of exporting if the exports arrive when supplies are plentiful.

² More precisely, it explains why the spatial price difference exceeds the difference between the transport cost and the storage cost. In practice the cost of storage is much smaller than the transport cost, so if the price spikes upwards the spatial price difference will typically exceed the transport cost as well. See section 3 for the precise formulation of the result.

The late nineteenth century Chicago and New York corn markets provide an ideal setting in which to detect the price relationships implied by the model as unusually detailed data are available including weekly spot prices, future prices, storage quantities, and transportation costs for three modes of transport. The data provide a striking verification of the model. During the period most corn was shipped from Chicago to New York via the Great Lakes and the Erie Canal in a trip that took approximately three weeks, although faster but more expensive rail transport was available. The paper shows that while the New York price for future delivery was normally equal to the Chicago spot price plus the cost of the lake and canal transportation, the difference between the New York and Chicago spot prices depended on the amount of storage in New York. When New York inventories were high, the difference between the New York price and the Chicago price plus the transport cost was normally less than four cents. However, when inventories in New York were low, the New York price spiked upwards and the price difference ranged as high as ten cents. These price spikes were temporary, and followed by declines in the New York price. The relationship between New York storage quantities and the spatial price difference can be summarised by plotting a “spatial arbitrage-storage” curve that links the spatial price difference adjusted for transport costs to the quantity of storage in New York.

The “spatial arbitrage-storage” curve examined in this paper has the same form as a supply of storage curve, a graph of the difference between the spot price and the future price of a commodity versus the quantity of storage. A supply of storage curve has a characteristic form: when storage volumes are low, the spot price is typically higher than the future price, but when storage quantities are high, the future price exceeds the spot price. The standard explanation for this curve, dating back to Kaldor (1939), Working (1949), and Brennan (1958), was that the holders of the inventories gained a “convenience yield” from their stocks and thus held them even though the spot price was higher than the future price. Recently, however, it has been argued that the supply of storage curve might simply be an artifact of an inappropriate method of aggregating inventory levels. In particular, Wright and Williams (1989) argued that a supply of storage curve exists because the spot-future price spread in one location is compared against inventories held at a wide range of locations. These other inventories are not immediately shipped to the central location to take advantage of the high spot prices, however, because to do so at short notice would incur extremely high transport costs.

This paper offers a related explanation for the supply of storage curve: that because transport takes time, inventories held elsewhere cannot be transported to the central market in time to take

advantage of temporarily high prices, and thus total inventories are positive even though the spot price exceeds the future price in one location. The theoretical model suggests that the spatial arbitrage-storage curve and the supply of storage curve should be closely related because when a centre has zero storage the spot price exceeds both the local future price and the other spot price plus the transport cost. To examine this relationship, a supply of storage curve for New York corn is calculated, by plotting the difference between the spot price and the one-month future price against the quantity of inventories in New York. It proves that the supply of storage and the spatial arbitrage storage curves are related in the suggested manner.

The similarity between the two curves suggests that the model is a useful framework for analyzing the supply of storage curve phenomena. Both the spatial arbitrage-storage curve and the supply of storage curve exist for two reasons: first, because New York prices frequently spiked upwards when stocks were low as grain could not be imported immediately; and, secondly, because storage quantities did not literally fall to zero when the New York price spiked upwards. It is likely that storage was never zero because there was always some corn in the grain elevators as they were used to both store grain and to transfer it from railcars and canal boats to ocean vessels. Furthermore, these data provide some support for Wright and Williams' hypothesis that a supply of storage curve can exist even when there is no convenience yield from holding the commodity. New York prices were in backwardation not because the owners of the grain gained convenience yield but because purchasers in New York were willing to pay high prices for grain to be delivered immediately rather than wait for new supplies to arrive from Chicago.

This paper is organised as follows. In section 2, a rational expectations model of storage and trade with non-instantaneous transport is specified and the solution technique is outlined. The key aspects of the solution that can be derived analytically are presented in the first half of section 3, while some of the results of numerical simulations are presented in the second half. In section 4, several conditional moments of the price distribution are derived and contrasted to the conditional moments usually estimated in cointegration models of spatial arbitrage. In section 5 key features of the Chicago-New York corn trade are described and it is established that grain was predominantly shipped by the slowest and cheapest mode of transport, namely by ship across the Great Lakes and then by canal to New York. The relationship between New York and Chicago prices, the cost of shipping, and storage quantities is examined in section 6, and a discussion of the results is offered in section 7.

2. A Model of Storage and Commodity Arbitrage

Structural models of spatial price arbitrage have antecedents that can be traced to Cournot (1838). The basic modeling approach, exemplified by Samuelson (1952), has been to specify a set of equations representing demand and supply curves in different locations and a set of no-arbitrage conditions that must hold if excess profits cannot be made by transporting goods instantaneously from one location to another. A series of prices is found that ensures aggregate demand equals aggregate supply and no profitable arbitrage opportunities exist:

$$P_t^i - P_t^j \leq K^T \quad (P_t^i - P_t^j - K^T) T_t^j = 0 \quad (1)$$

where P_t^i is the price at centre i at time t , K^T is the transport cost, and T_t^j are the exports from j to i at time t .

Mathematical models of rational storage have a similar structure (see for example Williams (1936), Williams and Wright (1991), or Deaton and Laroque (1992, 1996).) Risk neutral arbitrageurs are assumed to make a forecast of the future price and purchase and hold inventories until the expected price increase just offsets the costs of storage; conversely, inventories will be zero if the expected appreciation is less than the cost of storage. There are three possible storage costs. First, there can be an elevator charge K^S per unit to store goods each period. Secondly, the commodity depreciates at rate δ so if S_t is stored in period t , $(1-\delta)S_t$ will be available in period $t+1$. Thirdly, there is an interest cost r foregone when storage is undertaken. These relationships are represented by the following equations:

$$\left(\frac{1-\delta}{1+r}\right) E_t[P_{t+1}^i] - P_t^i \leq K^S \quad \left[\left(\frac{1-\delta}{1+r}\right) E_t[P_{t+1}^i] - P_t^i - K^S\right] S_t^i = 0 \quad (2)$$

where E_t is the expectations operator conditioning on information known at t , and S_t^i is the quantity of inventories held at time t .

These models of spatial arbitrage and rational storage were combined and solved by Williams and Wright (1991). They investigated how prices in two locations would be determined if rational, forward looking, fully informed, and risk neutral arbitrageurs could either store goods or transport them instantaneously from one location to the other. Their model has two centres with random production and independent demand functions, no capacity constraints to transport or storage technology, and is solved using numerical techniques to find the set of optimal storage and trade

functions that generate a stationary rational expectations equilibrium. The model developed in this paper copies their approach but relaxes the assumption that trade is instantaneous.

Because transport is modeled to take time, the model is also related to those used in the logistics management literature. This literature has examined the optimal ways for a company with sales in multiple locations to minimise the sum of procurement, transport, and storage costs³. A general theme of the literature is that careful inventory management enables a firm to substitute low cost but slow transport systems for faster but higher priced systems. Indeed, a cost minimizing firm will use fast transport systems only when inventories unexpectedly fall to such low levels that slow transport systems cannot be used to replenish them before they run out. The model in this paper extends the logistics management literature by endogenising the penalty that occurs when inventories to fall to zero, but at the cost of assuming there is only one type of transport system.

The model

There are two centres, A and B, each with a separate inverse demand function for a commodity:

$$P_t^{D,i} = P^{D,i}(D_t^i) : \quad P^{D,i}(0) < \infty, \quad \lim_{D \rightarrow \infty} P^{D,i}(D) = 0 \quad (3)$$

where D_t^i is the amount purchased for final use at time t and $i = A, B$. All production, consumption, storage and trade activity takes place at the beginning of the period, and the length of a period is the time that it takes to ship goods from one centre to another⁴. To be consistent with the data envisaged being used, the period length is typically taken to be between one and four weeks.

Output is assumed to be price inelastic but stochastic in the period under consideration, because it has a long gestation period. Output in each centre follows an independent first order autoregressive process around a constant mean:

$$(X_t^i - \bar{X}^i) = \rho^i (X_{t-1}^i - \bar{X}^i) + e_t^i \quad i = A, B \quad (4)$$

³ See Baumol and Vinod (1970) for an original statement, or Tyworth (1991), McGinnis (1989) or de Jong (2000) for a review.

⁴ This assumption is made for analytic convenience, for decisions to store and trade could be made at higher frequencies than the time it takes to transport goods. However, if decisions were made at higher frequencies, each decision would be a function of additional state variables representing quantities at different stages of transit, which would substantially complicate the model.

where e_t^i is a white noise process and $|\rho^i| < 1$ ⁵. It is assumed that unlimited amounts of the good can be stored, and that goods produced in two different periods are indistinguishable from each other so that they trade at the same price. The good can be stored in either centre, or shipped from one centre to the other. It is assumed that it is more expensive to transport goods from one centre to the other than it is to keep them in the same centre, so $K^S < K^T$. Note that an arbitrageur will store the commodity until it needs to be transported to minimise interest costs.

Let product availability, M_t^A and M_t^B , be the total quantity of stored and imported goods available in each centre at the beginning of the period,

$$M_t^i = (1 - \delta)(S_{t-1}^i + T_{t-1}^j) \quad (5)$$

where S_{t-1}^i is the non-negative quantity stored in centre i and T_{t-1}^j is the non-negative quantity exported from centre j . The quantities stored and exported are such that $S_t^i + T_t^i \leq X_t^i + M_t^i$.

It is assumed that risk neutral, profit maximizing, and rational speculators in both cities undertake a mixture of trade and storage to take advantage of expected price differences. The speculators have rationally determined expectations about future prices that incorporate all information about output, storage, and trade in both centres. The behavior of risk neutral speculators can be represented by four inequalities. Let $y_t = [M_t^A, M_t^B, X_t^A, X_t^B]$ be the vector of state variables.

Then, at each point y_t :

$$\left(\frac{1-\delta}{1+r}\right)E[P_{t+1}^A | y_t] - P^A(y_t) \leq K^S; \left[\left(\frac{1-\delta}{1+r}\right)E[P_{t+1}^A | y_t] - P^A(y_t) - K^S\right] \cdot S^A(y_t) = 0 \quad 6a$$

$$\left(\frac{1-\delta}{1+r}\right)E[P_{t+1}^B | y_t] - P^B(y_t) \leq K^S; \left[\left(\frac{1-\delta}{1+r}\right)E[P_{t+1}^B | y_t] - P^B(y_t) - K^S\right] \cdot S^B(y_t) = 0 \quad 6b$$

$$\left(\frac{1-\delta}{1+r}\right)E[P_{t+1}^B | y_t] - P^A(y_t) \leq K^T; \left[\left(\frac{1-\delta}{1+r}\right)E[P_{t+1}^B | y_t] - P^A(y_t) - K^T\right] \cdot T^A(y_t) = 0 \quad 6c$$

$$\left(\frac{1-\delta}{1+r}\right)E[P_{t+1}^A | y_t] - P^B(y_t) \leq K^T; \left[\left(\frac{1-\delta}{1+r}\right)E[P_{t+1}^A | y_t] - P^B(y_t) - K^T\right] \cdot T^B(y_t) = 0 \quad 6d$$

where

⁵ The assumption that mean output is predetermined and price inelastic can in principle be relaxed; see Williams and Wright (1991). Bailey and Chambers (1996) and Deaton and Laroque (1996) argue that if demand is linear there is an inherent lack of identification between changes in the mean and variance of

$$P^i(y_t) = P^{D,i}(X_t^i + M_t^i - S^i(y_t) - T^i(y_t)),$$

$$M_{t+1}^i(y_t) = (1 - \delta)(S^i(y_t) + T^j(y_t)), \text{ and}$$

$$E[P_{t+1}^i | y_t] = \iint_X P^{D,i}(X_{t+1}^i + M_{t+1}^i(y_t) - S^i(y_{t+1}) - T^i(y_{t+1}))f(X_{t+1}^i, X_{t+1}^j) dX_{t+1}^i dX_{t+1}^j \quad (7)$$

The first two of these inequalities are the conditions for profitable storage in either centre, while the second two are the conditions for profitable trade between centres. Each of the four inequalities holds with equality if the control variables (storage or trade) are non-zero. Since $K^S < K^T$, the centres will not export simultaneously.

The model solution, which is found numerically, comprises two parts. The first part is the set of optimal storage and trade functions [$S^A(\cdot)$, $S^B(\cdot)$, $T^A(\cdot)$, $T^B(\cdot)$] that satisfy the four inequalities 6a - 6d. Each function depends on the vector of four state variables. The second part of the solution is the distribution of the state variables that occurs in equilibrium, which depends on the assumed stochastic process determining output and the optimal storage and trade functions. The solution fulfils two conditions: first, that storage and trading decisions are profit maximizing conditional on expectations of future prices; and, secondly, that price expectations are consistent with the storage and trading decisions and expectations of future output quantities.

The model solution technique

The numerical solution to the model is calculated over a discrete four dimensional grid corresponding to the four state variables. The solution technique has four key steps. First, a discrete joint probability distribution over the grid values for the stochastic variables X^A and X^B is chosen, and the double integral formula in equation 7 is replaced by the equivalent summation formula. The joint probability density for X is chosen to mimic an autocorrelated process with normal innovations, and is represented by a $m^2 \times m^2$ Markov transition matrix Π specifying the probability of going from one point (X_{i1}^A, X_{j1}^B) to a second point (X_{i2}^A, X_{j2}^B) .

Secondly, a solution to the optimal storage problem for the limiting case that transport costs equal zero and trade is instantaneous — the “combined centre” case — is found using solution

output and changes in the demand function, so the random shocks can be considered either demand shocks or supply shocks.

techniques similar to those documented in Deaton and Laroque (1995, 1996)⁶. The solution is a function linking the optimal values of storage to combined centre output and combined centre availability. A linear demand function for each centre was specified:

$$P^{D,i}(D_t^i) = \begin{cases} \alpha^i & \text{if } D_t^i = 0 \\ \alpha^i - \beta^i D_t^i & \text{if } 0 < D_t^i \leq \alpha^i / \beta^i \\ 0 & \text{if } D_t^i > \alpha^i / \beta^i \end{cases} \quad (8)$$

The third step is an algorithm that calculates the optimal amounts of storage and trade in the two centres. The model is solved by finding a series of successive approximations to the optimal storage and trade functions, $S_k^i(y_t)$ and $T_k^i(y_t)$, where k refers to the k^{th} approximation. The algorithm is briefly described in Appendix 1. The starting value of the algorithm is based upon the combined centre solution, and the algorithm is repeated until the difference between successive values of the control values is small.

The fourth step, once the optimal storage and trade functions are calculated, is the calculation of the invariant probability distribution of the model solution. The invariant distribution of the model is the unconditional probability of being at a particular grid point, which is used to calculate various statistics about the price distributions in each centre. The method used to calculate the invariant distribution is also described in Appendix 1.

The model with capacity constraints or variable transport costs.

In practice, most logistics systems have capacity constraints that occasionally bind, so that the set of equations 6a-6d does not always hold. Rather, when the constraints bind the spot price spikes down, because the good must be consumed immediately or thrown away. Coleman (1998) contains an outline of a model in which there are storage capacity constraints: in particular, it is a model of optimal storage and trade decisions for a good that can be stored for only one period, so that the maximum amount that can be stored is the previous period's production. Other than the occasional downward price spikes that occur when then there is a large quantity of the perishable

⁶ The centres are combined to have a single demand function and a single production function, and the optimal storage function is calculated using equation 6a applied to the combined centres. The solution technique is slightly different to that of Deaton and Laroque. Their model is solved at annual frequencies, so the maximum possible level of storage is typically 20 times maximum annual production, assuming an annual depreciation rate of 5 percent. When the model is solved at weekly frequencies, maximum storage can be several hundred times as large as maximum weekly output, and several thousand times as much as the variation in output. Since it is not practical to choose a grid with several million points to describe output in both centres, three interlocking grids were used, with the finer grids focussed around the area where storage is zero and a coarse grid used to model the situation when availability is very large.

good available, the solution to this model is qualitatively similar to the case when the good can be stored indefinitely. The similarity of these solutions gives grounds for believing that the model solution is robust to a variety of assumptions about the nature of a good's storability.

The model is more difficult to solve if transport costs are both time varying and predictable, because the solution depends on whether there are transport or storage capacity constraints. Suppose there are no capacity constraints and one centre primarily exports to the other. If transport costs are expected to increase substantially, the optimal solution is for arbitrageurs in the exporting centre to transport all their goods at the low cost time and store them in the importing centre until they are needed. If capacity constraints limit the amount that can be shipped, however, the solution depends on whether transport capacity is allocated on a “first-come, first-served” basis or whether transport prices are allowed to adjust to ensure that demand for transport just equals the available capacity. If transport prices are flexible, they adjust to ensure the set of inequalities 6a-6d still hold. If transport prices are fixed and capacity is rationed, the inequalities do not hold. Rather, under these circumstances the expected future price in one centre exceeds the spot price in the other plus the transport cost, and the agents who obtain the transport capacity make windfall profits.

The late nineteenth century corn market had transport capacity constraints, but transport prices were flexible. Indeed, there is evidence that transport prices adjusted to equate supply with demand, as shipping prices typically increased towards the end of the season. Consequently, the model described by equations 6a-6d should adequately describe the behavior of prices during this period.

3. Properties of the Model

Analytical Results

The key results of the paper can be derived analytically by considering various combinations of the complementary conditions associated with equations 6a - 6d. The possible combinations are summarised in Table 1. Numerical simulations are used to calculate the relative importance of each combination; the probabilities in Table 1 correspond to the parameters described in footnote 8 below, but are representative of those pertaining to a wide variety of parameters. The probabilities suggest that the most common combinations are positive storage in both centres and zero trade (row 1), positive storage in both centres with trade from one centre to the other (rows 2 and 3), and positive storage in the exporting centre and zero storage in the importing centre (rows

6 and 8). The other combinations appear to occur rarely⁷. The set of conditions 6a - 6d indicates how prices adjust in each of these cases.

First, consider a point y_t when inventories are positive in each centre. Because inventories are positive, the price in each centre is expected to increase and therefore prices are expected to diverge; more precisely, by equations 6a and 6b

$$E[P_{t+1}^i | y_t] = \left(\frac{1+r}{1-\delta} \right) (P^i(y_t) + K^S) \quad i = A, B \quad (9)$$

and consequently

$$E[P_{t+1}^A - P_{t+1}^B | y_t] = \left(\frac{1+r}{1-\delta} \right) (P^A(y_t) - P^B(y_t)) \quad (10)$$

In addition, equations 6c and 6d imply the spatial price difference lies within the range

$$-(K^T - K^S) \leq P_t^A - P_t^B \leq +(K^T - K^S) \text{ with the inequality holding if trade is positive.}$$

(Otherwise it would be possible to make profits by reducing storage in the low priced centre and exporting to the high priced centre.) Therefore, when inventories are positive in both centres and trade is zero, prices lie in the range $|P_t^A - P_t^B| \leq (K^T - K^S)$ and are expected to diverge.

Secondly, consider a point y_t at which inventories in centre B are positive and arbitrageurs in centre B export to centre A. If centre A also has positive inventories, equations 6a, 6b, and 6d hold with equality and imply

$$P^A(y_t) = P^B(y_t) + K^T - K^S \quad (11a)$$

$$E[P_{t+1}^A - P_{t+1}^B | y_t] = \frac{1+r}{1-\delta} (K^T - K^S) \quad (11b)$$

In this case the spatial price differential is exactly $K^T - K^S$, but again prices are expected to diverge. Alternatively, if centre A has zero inventories, equations 6b and 6d hold and imply

$$P^A(y_t) > P^B(y_t) + K^T - K^S \quad (12a)$$

⁷ As the stochastic process determining output is changed, the optimal storage and trade functions change but the solution still fulfills the set of conditions 6a-6d. Consequently, analytic statements about the solution will hold irrespective of the assumed stochastic process, but the distribution of the state variables in equilibrium will differ. With other modeling assumptions, other combinations of the state variables may become much more important.

$$E[P_{t+1}^A - P_{t+1}^B | y_t] = \frac{1+r}{1-\delta} (K^T - K^S) \quad (12b)$$

In this case the spatial price difference exceeds the difference between the transport cost and the inventory holding cost, as prices are unusually high in the importing centre. However, for reasonable values of δ and r the centre A price is expected to fall when the goods arrive at $t+1$ and thus prices are expected to converge.

The last two sets of equations determine how goods price arbitrage occurs when the exporting centre has large inventories and shipping takes time. Equations 11b and 12b imply that the price difference at time $t+1$ must exceed $K^T - K^S$ under some circumstances. Since the price difference at $t+1$ cannot exceed $K^T - K^S$ when the imports arriving in centre A are so plentiful that there is a surplus that is held over to the following period, on some occasions the imports must be sufficiently small relative to demand that inventories fall to zero and temporarily high prices occur (equation 12a, applied at $t+1$). For the zero profit condition to hold, exporters must export a sufficiently small quantity that on regular occasions the importing centre inventories are zero. Here the non-negativity of storage causes an important asymmetry. If output is higher than expected in the destination centre when the goods arrive, the surplus can be stored and prices fall but little; but if output is very low there will be a shortage and prices will increase sharply. For the zero profit condition to hold, a sufficiently large quantity must be shipped that there are only a small number of sharp price increases offset by a large number of small price decreases. Consequently, if the output shocks are symmetric, the median trade will be unprofitable. Furthermore, the zero-profit condition for trade means that an importing centre will normally have positive inventories, because larger shipments arrive on the typical occasion than are needed for immediate consumption.

Numerical Results

In the rest of this section numerical results corresponding to two different situations are presented. In the first case the two centres are identical, so that trade occurs as frequently one direction as the other. In the second case, both centres have the same demand functions but mean production in centre B is higher than in centre A so that exports typically flow from B to A. The parameters are chosen so that transport costs are approximately 5 percent of the average price, the

period is one week, and the annual interest and depreciation rates are approximately 5 percent⁸. The solution is used to calculate the distribution of the spot price difference, $P_t^A - P_t^B$, and various statistics about the distribution of storage and trade.

Distribution of Prices

Figure 1 shows the unconditional distribution of the spot price difference $P_t^A - P_t^B$ in the two cases. Mean output is assumed to be 100 when the centres are equal; otherwise, mean output in centre A is 95 and mean output in centre B is 105. In both cases the dominant features of the distribution are the spikes at $\pm (K^T - K^S)$ corresponding to the occasions when one centre exports to the other, and the scattering of density outside these spikes corresponding to the price difference that occurs when at least one centre has zero inventories. When the centres are identical, the distribution is symmetric and appears to have a local maximum at $(P_t^A - P_t^B) = 0$; otherwise the spike at $P_t^A - P_t^B = (K^T - K^S)$ is larger than the spike at $P_t^A - P_t^B = -(K^T - K^S)$, and the density in the region $-(K^T - K^S) < P_t^A - P_t^B < (K^T - K^S)$ is increasing in the spatial price difference⁹. As the production asymmetry is exacerbated, more and more of the density is located at $P_t^A - P_t^B = (K^T - K^S)$.

The Two Centre Model with Equal Centres

Table 2 presents selected statistics of the optimal price, storage, and trade variables under the baseline parameterization when mean production in both centres equals 100. The statistics are compared with the case when trade is costless and timeless (i.e. the combined centre solution) and the case when trade is not possible. Results are also presented for two variations of the base parameters: when transport costs are increased to $K^T = 10$; and when they are reduced to $K^T = 2.5$.

Three features of the results should be noted. First, prices are smoothed primarily through the adjustment of storage quantities within a centre, not the transport of goods between centres. For each of the three transport costs, exports were zero more than 80 percent of the time and the mean

⁸ In the baseline case, the following model parameters are used: the demand function $P^{D,i}(X) = \alpha - \beta X = 200 - X$; the production conditional variance $\sigma^2 = 100$; the production autocorrelation $\rho = 0.9$; the weekly interest rate $r = 0.001$; the weekly depreciation rate $\delta = 0.001$; $K^S = 0$; and $K^T = 5$.

⁹ The use of a discrete grid to calculate the solution means that confidence in the exact shape of the density is limited. The problem is exacerbated in the “equal centre” case as when output and storage are exactly the same in each centre (an event with zero probability in the continuous case, but with positive probability in the discrete case) the spatial price difference equals zero. This is the cause of the sharp spike at $(P_t^A - P_t^B) = 0$ in Figure 1.

export value was only between 2 and 4 units or 2 – 4 per cent of average weekly output. In contrast, storage in each centre was positive 97 percent of the time, and mean storage in each centre was two to three times weekly output.

Secondly, even though trade volumes were small, trade had a large effect on prices because storage behaviour was altered so that the spatial price difference between the centres was nearly always less than the cost of transporting goods. Irrespective of the transport cost, the spatial price difference exceeded the difference between the transport cost and the storage cost only 3 percent of the time, so as the cost of transport was reduced price dispersion between the centres declined.

Thirdly, the possibility of trade reduced average storage by 25 to 50 percent of the level that would have prevailed were no trade possible. The average quantity stored in either centre fell because of access to supplies from the other centre when output fell below normal levels.

The effect of various changes to parameters on the distribution of prices and storage can be easily calculated. For instance, when the autocorrelation of output is low, there are few extended periods when output is very high and consequently large stocks are not built up. Moreover, inventories are used up immediately when output is unusually low because it is expected to rapidly return to average levels. This means the fraction of time that one or other centre has zero storage and when the spatial price difference exceeds the transport cost increases as the autocorrelation of output decreases. A more interesting variation concerns the effect of adjusting the length of the transport period, keeping the cost of transport the same. This can be done by changing the mean, variance, and autocorrelation of output as well as the demand curve parameters (and the interest rate and depreciation rate) in line with the changing period length: when the transportation period is halved, for instance, mean output is halved and the slope of the demand curve is doubled¹⁰. The simulations show that as the transport time lengthens the average amount of storage increases, the fraction of time that one centre has zero inventories increases, and thus the variance of the spatial price difference increases.

The Two Centre Model with a Dominant Exporter

There are two cases to note when one centre is a dominant exporter: first, when the asymmetry between the centres is small, so that both centres still regularly export; and secondly, when

¹⁰ It is assumed that there is only one shipment per period no matter the length of the transport period. This method of varying the time period has been chosen as it requires minimal programming changes. The exercise is not harmless, however. It implicitly assumes that output is evenly distributed through time during the period and that there is no depreciation within a period.

production is sufficiently specialised that trade almost always flows from one centre to other. These cases are modeled by lowering the mean production level in centre A progressively from 100 to 95 to 70 while maintaining total production in the two centres at a mean level of 200. Various moments of the solution are presented in Table 3. Note that the dominant exporter model best describes the U.S. corn market analyzed in section 5, as corn was transported from Chicago to New York but never in reverse.

When $\bar{X}_A = 95$ and $\bar{X}_B = 105$, there is only a weak tendency for net exports to flow from one centre to the other, and the solution is similar to the symmetric case. Storage adjustment remains the dominant means by which price fluctuations are smoothed, and while storage is higher in centre B than A it is positive in each centre over 96 percent of the time. Centre B is the dominant exporter, exporting 27 percent of the time, but centre A exports 8 percent of the time. (In contrast, when the centres are equal each centre exports 14 percent of the time.) The mean spatial price difference is 1.9 or less than the transport cost because cheap storage technology is used much more frequently than expensive transport technology to arbitrage prices.

The asymmetry in the solution becomes more pronounced as output in centre A is reduced to 70. Exports from B to A occur more than 80 percent of the time, while exports in the reverse direction take place less than 0.1 percent of the time. Mean trade volumes from B to A are 27 units, or close to the level that would prevail if there were no uncertainty. Although total storage declines as the production asymmetry becomes larger, it is increasingly held in the exporting centre. When $\bar{X}_A = 70$, the total of average storage plus trade in both centres is only 76 percent of average storage plus trade when the centres were equal, but centre B has storage 98 percent of the time. As production in centre A declines further, total storage declines to the level of storage in the combined centre model.

4. Conditional Moments of the Prices

In this section several conditional moments that are commonly estimated in cointegration models of spatial prices are calculated numerically¹¹. The first of these conditional moments is

$E[z_{t+1} | z_t]$, where $z_t = P_t^A - P_t^B$ is the spatial price difference between the two centres.

Typically when two regional price series are tested for cointegration, the difference in prices z_t is

¹¹ Of course, since output in the model is stationary, prices are stationary unless there is general inflation.

calculated, and the first order unconditional autocorrelation coefficient ρ in the regression $z_{t+1} = \rho z_t + e_{t+1}$ is estimated and tested to see whether or not it equals 1. Implicitly this formulation implies $E[z_{t+1} | z_t] = \rho z_t$. When the conditional mean associated with this model is calculated from the set of equations 6a - 6d, however, it proves not to be linear (see Figure 2). The basic shape reflects the pattern of prices when either both centres store (with or without one exporting) or one centre stores and exports but the other centre has zero inventories, for one of these combinations occurs most of the time. When there is storage in each centre, $-(K^T - K^S) \leq z_t \leq +(K^T - K^S)$ and according to equation 10 the price in each centre is expected to increase by a factor of $\frac{1+r}{1-\delta}$. Because z_t is primarily in the range $-(K^T - K^S) \leq z_t \leq +(K^T - K^S)$ when storage in each centre is positive¹², $E[z_{t+1} | z_t \leq K^T - K^S] \approx \frac{1+r}{1-\delta} z_t$. In contrast, at a point y_t where centre B has positive inventories and is exporting, but centre A has zero inventories, $z_t > (K^T - K^S)$ and equation 11b holds: $E[z_{t+1} | y_t] = \frac{1+r}{1-\delta} (K^T - K^S)$. Since this is the dominant reason why $z_t > K^T - K^S$, $E[z_{t+1} | z_t > K^T - K^S] \approx \frac{1+r}{1-\delta} (K^T - K^S)$ ¹³. Similar considerations mean that $E[z_{t+1} | z_t < -(K^T - K^S)] \approx -\frac{1+r}{1-\delta} (K^T - K^S)$. The model therefore provides a structural justification for the use of threshold autoregression models to estimate the process of dynamic spatial price adjustment. Unlike the model of Obstfeld and Taylor (1997), however, the speed of adjustment parameters are known.

The conditional mean has the same form irrespective of the parameters of the model. This is not true for the conditional variance $Var[z_{t+1} | z_t]$, however. In particular the shape of the conditional variance depends on whether trade is symmetric or not. If each centre is the same the conditional variance is symmetric and “U” shaped, and much higher outside the range $-(K^T - K^S) \leq z_t \leq$

¹² In some rare circumstances stocks in each centre will be zero and output will be very low but similar so that prices are high in each centre but no storage or trade takes place.

¹³ When there are no stocks in either centre it is possible that both $|z_t| > K^T - K^S$ and $E[z_{t+1} | z_t] > (K^T - K^S)$: in Figure 3, these occasions are the cause of the spikes in the region $|z_t| > K^T - K^S$. On these occasions, one centre will export to the other but not keep any stocks itself; the spatial price difference in the next period is expected to be greater than the transport cost as prices are expected to fall by different amounts in both centres in the next period.

$+(K^T - K^S)$ than inside it. The conditional variance outside this range is high because at least one centre has zero inventories and there is a reasonably high probability of having zero stocks and thus high prices in the subsequent period¹⁴. If trade is not symmetric, the conditional variance, while basically “U” shaped, is rather more complex. If centre B normally exports to centre A, the conditional variance is high outside the range $-(K^T - K^S) \leq z_t \leq +(K^T - K^S)$ and “tick” shaped within the bands, reaching a minimum close to $z_t = K^T - K^S$. The conditional variance of z_{t+1} is high when z_t is small or negative because output in centre B is unusually low in these circumstances and prices in centre B are volatile.

In addition to testing spatial price differences for cointegration, an “error correction” regression is normally estimated to find out how prices change through time: $\Delta P_{t+1}^i = \alpha^i z_t + u_{t+1}^i$, $i=A,B$.

This formulation also implies the conditional mean $E[\Delta P_{t+1}^i | z_t]$ is a linear function of z_t .

However, when the conditional means are calculated they are not linear either (see Figure 3).

When both centres have storage, the price difference is in the range $-(K^T - K^S) \leq z_t \leq +(K^T - K^S)$ and in each centre the price is expected to increase to compensate for the costs of storage:

$$E[\Delta P_{t+1}^i] = \frac{r + \delta}{1 - \delta} P_t^i + \frac{1 + r}{1 - \delta} K_t^S \quad i = A, B.$$

In the diagram $E[\Delta P_{t+1}^i | z_t] \approx 0.2$ as the average price is 100 and the weekly interest and depreciation rates are both 0.01¹⁵.

Outside this range there is an asymmetry. If centre A has no inventories but centre B has inventories and exports to A, $P_t^A - P_t^B > (K_t^T - K_t^S)$ and

$$E[\Delta P_{t+1}^A] = -(P_t^A - P_t^B - K_t^T) + \frac{r + \delta}{1 - \delta} K_t^T + \frac{r + \delta}{1 - \delta} P_t^B$$

¹⁴ In Dumas(1988), the conditional variance is also “U” shaped when the centres are symmetric. He does not calculate the non-symmetric case, however.

¹⁵ It should be emphasised that the price change is not directly related to the spatial price difference but to the level of prices in each centre. It is not necessarily the case that for two points in the state space y_1 and y_2 with $P^A(y_1) - P^B(y_1) = P^A(y_2) - P^B(y_2)$ that $E[\Delta P_{t+1}^i | y_1] = E[\Delta P_{t+1}^i | y_2]$, $i=A,B$ because the level of prices associated with y_1 and y_2 can be quite different.

$$E[\Delta P_{t+1}^B] = \frac{r + \delta}{1 - \delta} P_t^B + \frac{1 + r}{1 - \delta} K_t^S$$

The converse relationship holds if A is the exporting centre and B the importing centre.

In combination, therefore, a threshold regression of ΔP_{t+1}^A against $P_t^A - P_t^B$ should have a value of approximately zero if $P_t^A - P_t^B$ is less than $(K_t^T - K_t^S)$, and the value of minus one if it exceeds $(K_t^T - K_t^S)$, while a threshold regression of ΔP_{t+1}^B against $P_t^A - P_t^B$ should have a value of one if $P_t^A - P_t^B$ is less than $-(K_t^T - K_t^S)$ and zero if it exceeds $-(K_t^T - K_t^S)$.

These results stem from the interplay of both storage and trade. P_t^A is unusually high when there is a stock-out in centre A, and in this case prices are expected to fall in the subsequent period because of supplies arriving from centre B. The price in centre B will already have increased in period t , however, when the goods are removed from the local market.

While the conditional variances of P_{t+1}^A and P_{t+1}^B (and hence ΔP_{t+1}^A and ΔP_{t+1}^B) can easily be calculated, they are difficult to summarise. When the centres are symmetric, the conditional variances are “U” shaped, with $Var[P_{t+1}^B | z_t] > Var[P_{t+1}^A | z_t]$ when $P_t^B > P_t^A$ and $Var[P_{t+1}^B | z_t] < Var[P_{t+1}^A | z_t]$ when $P_t^B < P_t^A$. When B mainly exports to A, the conditional variances also cross, but at a point closer to $K^T - K^S$.

The shapes of the conditional moments $E[\Delta P_{t+1}^i | z_t]$ stem directly from the structural assumptions of the model, and can be contrasted with the ad-hoc price adjustment moments usually postulated in empirical models. The implication that price adjustment should be dominated by a reduction in the price in the higher priced centre is examined in section 6.

5. The Late Nineteenth Century U.S. Corn Trade.

Trade between Chicago and New York, 1878-1890

In the late nineteenth century, corn was predominantly consumed where it was produced, primarily as animal feed. Only a small fraction of the total crop was transported any distance,

most of which was shipped east and south from the main producing area, the Great Plains states¹⁶. Chicago was the preeminent inland shipping centre, receiving and shipping an average of 67 million bushels per year between 1878 and 1890. Much of the grain was sent to east coast ports prior to export. New York was the most important of these ports, receiving an average of 34 million bushels per year and exporting as much as Boston, Baltimore, and Philadelphia combined. Both Chicago and New York had active markets and elaborate physical infrastructure to handle the large volumes of corn and other grains.

Corn was first sent by rail and canal to Chicago, and then sent to New York using one of three transportation methods. The slowest and least expensive method was to send grain to Buffalo by ship via the Great Lakes, and then to forward it to New York along the Erie Canal. This method, which took approximately three weeks, was not available between November and late April, however, as both the canal and the Great Lakes froze. A faster and more expensive method, taking 10 days, was to ship grain over the Great Lakes to Buffalo and then send it by rail to New York. The fastest and most expensive method was to send grain to New York by rail, a trip that took 3 or 4 days. Between 1881 and 1891, when average annual costs were reasonably stable, the average cost of shipping a bushel of corn from Chicago to New York was 7.7 cents by lake and canal, 10.3 cents by lake and rail, and 14.6 cents by rail¹⁷. The average price of a bushel of corn in Chicago during this time was 45 cents.

Between 1878 and 1890, an average of 45 million bushels of corn per year was shipped from Chicago by lake, and 22 million was transported by rail. Of the latter, however, 15 million bushels were “through-shipments” — shipments that started west of Chicago, that were routed through Chicago, but that were never sold or handled in Chicago¹⁸. Unfortunately, the through-shipments were included in all of the annual and weekly shipping statistics, creating a misleading picture of the true extent of rail shipments of grain that was sold in Chicago. When the through shipments are excluded, the fraction of corn transported from Chicago that was shipped across the Great Lakes increases from 67 per cent to 86 per cent. Lake transport was even more dominant during the open water season. Of the 7 million bushels of corn shipped by rail that were processed

¹⁶Nebraska, Iowa, and Illinois produced a third of the total U.S. crop, with Illinois alone producing some 225 million bushels annually. In contrast, combined production in New York, New Jersey, and Pennsylvania was 75 million bushels per year.

¹⁷ Chicago Board of Trade, 1892, p122.

in Chicago, only 2.7 million bushels were shipped between April and November. Consequently, 95 per cent of corn that was transported in the open water season was shipped by lake¹⁹. Figure 4 indicates the seasonal pattern of the three types of shipments.

While it is straightforward to establish that 95 per cent of the grain leaving Chicago in the summer went across the Great Lakes, it is more difficult to calculate how much of it went to New York, and by which transport mode. Most Chicago corn was shipped through Buffalo, however, and since most corn arriving in New York came from Buffalo it is possible to establish the link indirectly. Unfortunately detailed shipping statistics for New York are not available for the whole period, although annual data is available for the years 1878-1881 and 1888- 1890. During these years, an average of 39 million bushels of corn were shipped by lake from Chicago to Buffalo, and another 10 million bushels were sent by rail along the Michigan Central and the Lake Shore and Michigan railroads to the same city²⁰. From Buffalo, most of the corn was sent to New York either by canal boat or by one of four railroads. On average New York received 40 million bushels of corn, of which 21 million arrived by canal boat, 16 million arrived on trains coming from Buffalo, and 3 million came from elsewhere. It is established below that almost all of the grain sent by rail from Chicago to Buffalo was forwarded to New York by rail, meaning that at most 6 million bushels of the grain arriving in New York by rail could have arrived in Buffalo by ship. Consequently, of the 27 million bushels of corn that arrived in Buffalo by water and were sent to New York, 21 million (or 78 per cent) were forwarded by canal. The dominance of the lake and canal route over the lake and rail route is even greater than this fraction suggests, however, because the lake and canal season was typically 3 weeks shorter than the lake and rail season. Since the season typically lasted 30 weeks, the lake and canal route accounted for some

¹⁸ The through shipments are calculated as the total of the monthly through shipments on the Chicago and Northwestern, Illinois Central, Chicago, Burlington, & Quincy, Chicago, Rock Island, & Pacific and Chicago and Alton railroads that are reported in the Chicago Board of Trade Annual Reports each year.

¹⁹ Unless it is recognised that most rail shipments were through shipments two puzzles arise. First, rail shipments and lake shipments took place simultaneously even though rail shipment was considerably more expensive than lake shipment. There are some reasons why shippers would prefer to send grain by rail — it was faster, and there was less risk of additional heat damage if the grain were already damaged, for instance — but these do not seem to have justified the additional cost. Secondly, rail shipments during winter took place when the price difference between New York and Chicago was lower than the rail cost.

²⁰ The data are from the annual reports of the Chicago Board of Trade. Only 70 per cent of corn sent by ship from Chicago went to Buffalo. Since comparable New York data for 1882-1887 are unavailable, the averages are calculated for the years 1878-1881 and 1888-1890. For the full period 1878-1890, an average of 32 million bushels was sent from Chicago to Buffalo by lake, and 9 million by rail.

87 per cent of the grain arriving in New York during the period that the lake and canal route was open²¹.

It can be demonstrated that grain sent by rail from Chicago to Buffalo went to New York by comparing the monthly shipments of corn along the railroads connecting Chicago and Buffalo with the monthly shipments along the railroads connecting Buffalo and New York during the years that such data exist, 1877 to 1881²². New York receipts along these lines were fifty percent higher on average than rail shipments from Chicago to Buffalo, because the Hudson and Erie lines were used to rail some of the corn shipped to Buffalo on the Great Lakes²³. Nonetheless, as indicated in Figure 5, there was an almost one-for-one correspondence between the variation in rail flows from Chicago to Buffalo and flows from Buffalo to New York. Formally, a regression of New York rail receipts from Buffalo with Chicago rail shipments to Buffalo has a slope estimate of 1.20, with a ninety five percent confidence interval of 0.93 to 1.46²⁴.

$$(\text{NY rail receipts})_t = -125000 + 1.20 (\text{Chicago rail shipments})_t + 0.087 (\text{Lake shipments})_t + e_t$$

$$(130000) \quad (0.131) \qquad \qquad \qquad (0.020)$$

$R^2 = 0.67$; 60 observations.

Similar data can be used to show that most of this corn was part of a through shipment. Figure 5 also suggests a near one-for-one correspondence between the variation in rail flows from Chicago to Buffalo and the through-shipments through Chicago. A regression of Chicago rail shipments to Buffalo against Chicago through-shipments has a slope estimate of 0.85, with a ninety five percent confidence interval of 0.72 to 0.99²⁵.

²¹ i.e. 21 million out of the 24 million bushels that arrived in New York during the lake and canal season.

²² Data for these years was published in the annual reports of the Chicago Board of Trade and the New York Produce Exchange, but I have been unable to find it for other years. The Chicago data is the monthly shipments along the Michigan Central and the Lake Shore and Michigan Railroads. The New York data is the monthly shipments along the Erie and the New York Central and Hudson railroads. For information on through shipments, see footnote 17.

²³ The shipments to New York were also higher than those from Chicago in winter, when there was no water transport. Since little corn was stored in Buffalo, at least in the years for which I could find data, 1882 – 1886, the winter shipments to New York must have originally come from elsewhere.

²⁴ OLS regression with standard errors calculated using the Newey-West method using 5 lags. The regression was also estimated using a feasible generalised least squares estimate that corrected for first order serial correlation and assumed the variance of errors was a linear function of Chicago shipments. In this case the slope estimate was 1.16 with standard error of 0.13.

²⁵ OLS regression with standard errors calculated using the Newey-West method using 5 lags. The feasible generalised least squares estimate (see footnote 23) was 0.95 with standard error of 0.10.

$$\begin{aligned} (\text{Chicago rail shipments})_t &= 205000 + 0.85 (\text{Chicago through shipments})_t + e_t \\ &\quad (86000) \quad (0.067) \end{aligned}$$

$R^2 = 0.65$; 60 observations.

These two regressions strongly suggest the corn sent by rail to Buffalo was forwarded to New York and that the corn originating in Chicago was part of a through shipment. A third regression linking New York rail receipts to Chicago through shipments confirms this link²⁶:

$$\begin{aligned} (\text{NY rail receipts})_t &= 215000 + 1.127 (\text{through rail shipments})_t + 0.038 (\text{Lake shipments})_t + e_t \\ &\quad (180000) \quad (0.131) \qquad \qquad \qquad (0.024) \end{aligned}$$

$R^2 = 0.57$; 60 observations.

New York Receipts, Export Shipments, and Storage

Inward and outward shipping activity occurred at New York throughout the year. When grain arrived in New York, it was first transferred to an elevator or a lighter and then either sold or delivered in fulfillment of a futures contract²⁷. Grain was often stored temporarily, but the storage capacity was rarely fully utilised, even in winter²⁸. Inventory levels never fell to zero, because the elevators always contained some grain in transit as they were used to transfer grain from arriving canal boats and rail cars to departing ocean vessels. The average of the five lowest inventory levels each year for the entire period was 260 000 bushels. Corn inventories typically fell to their low points each year in the middle of May, prior to the opening of the summer transport season, and in August. Storage quantities between 1881 and 1891 are shown in Figure 7.

Corn inventories in Chicago had a marked seasonal pattern. Receipts in Chicago were continuous throughout the year, but were higher than shipments between December and March, when the lakes were closed, and between August and September. Consequently, inventories built up to a late March peak averaging 4 million bushels on shore or 6.2 million bushels if corn loaded on board ships is included. Inventories reached seasonal lows in late July and November averaging 1.4 million bushels.

²⁶ OLS regression with standard errors calculated using the Newey-West method using 5 lags. The feasible generalised least squares estimate (see footnote 23) was 1.10 with a standard error of 0.15.

²⁷ Grain from canal boats could be unloaded to an elevator or be sold "afloat", whereupon it could be transferred directly to a ship using a lighter.

²⁸ In 1888 there were 24 million bushels storage capacity in New York and Brooklyn, and a further 3 million in New Jersey. However, peak grain storage in New York and Brooklyn between 1887 and 1889 was just over 16 million bushels, of which 11 million bushels were wheat and 4 million bushels were corn.

Storage charges varied little during the period (Goldstein (1928); Ulen (1982)). In 1888 it cost $\frac{5}{8}$ cents per bushel to deposit grain in an elevator, including the cost of 10 days storage; thereafter, storage cost $\frac{1}{4}$ cents per bushel per ten days. There were additional charges for trimming from canal boats and for trimming into ocean boats. Charges in Chicago were similar. In addition to these direct charges, the cost of holding inventories included insurance and the opportunity cost of holding the grain. Working (1929) estimated that in 1913 these costs were approximately 1.4 cents a bushel per month.

Transport prices between Chicago and New York.

There was a marked seasonal pattern in shipping costs (see Figure 6). Lake and canal and lake and rail prices were typically high at the opening of the season, but they then declined during the summer before rising towards the close of the season. On average, lake and canal rates increased by 0.2 cents per week between July and the end of the shipping season. Rail rates varied seasonally between winter and summer, particularly prior to 1886 when railroads competed aggressively with each other and with shipping lines for the grain business. The competition was sufficiently fierce to divert substantial quantities of the grain trade from the water route to rail, essentially by inducing “through shipments” from shipping agents west of Chicago (Tunell (1897), United States Treasury (1898)). This price competition is understated in the rail price data collected by the Chicago Board of Trade, as much of the business was transacted at lower, unrecorded prices, particularly during periods when the rail cartel broke down²⁹. The seasonal pattern in rail prices declined after the passing of the Interstate Commerce Act 1887, which regulated rail transport and substantially reduced price competition between the rail lines.

Market Prices

Both New York and Chicago had well developed grain markets in which both spot and futures contracts were bought and sold. A standard contract was settled by the delivery of grain to a warehouse or elevator. The main contracts were for immediate delivery (the spot contract) or for delivery at any time within the current month, the next month, two months’ time, or in May of a particular year (the futures contracts). The seller had the option as to the date in a particular month the grain was delivered, so spot prices normally exceeded or were equal to the zero-month

²⁹See the discussion by Nimmo in his reports on the internal commerce of the United States. (United States Bureau of Statistics, 1877, 1881, 1884). See Porter (1983, 1985) for a discussion of the price cutting wars that occurred prior to 1886.

future price. This paper uses the Wednesday closing price for all of the analysis. (See Appendix 2 for the precise definition and source of the data.)

6. Spatial Corn Price Arbitrage and Transport Costs, 1878-1891.

Spot Prices, Spatial Arbitrage and New York Storage Volumes

According to the model, the difference between the New York and Chicago spot prices on the day the corn was sent should have been equal to the transport cost minus the storage cost when there were positive inventories in New York, but it should have exceeded the transport cost minus the storage cost when inventories in New York were zero or very small. To examine this implication of the model, in Figure 8 the spatial price difference $P_t^{NY} - P_t^{CH}$ is plotted against the cost of lake and canal transport for each week that transport cost data is available, 1878 – 1891.

Observations for which storage quantities in New York were less than 300,000 bushels are distinguished from those for which storage quantities exceeded 300,000 bushels.

Three features of the graph stand out. First, most of the observations lie above the 45 degree line, indicating that spatial price difference normally exceeded the transport cost. While the spatial price difference need not have exceeded the transport cost each week, as corn may not have been shipped each week, it did so on 93 per cent of all weeks in the sample and 98 per cent of the weeks when New York storage was less than 600,000 bushels³⁰. On only one occasion (not shown) was the spatial price difference negative, when there was a corner in the Chicago market.

Secondly, the spatial price difference was far more likely to exceed the transport cost by a large amount when storage was less than 300 000 bushels than when storage exceeded that amount. In Table 4 the data are grouped by the level of storage and the distribution of the spatial price difference minus the transport cost is calculated. On 62 per cent of the 21 occasions when there was less than 150 000 bushels of corn in store, and 31 per cent of the 45 occasions storage was between 150 000 and 300000 bushels, the New York price exceeded the Chicago price plus transport cost by 5 cents or more. In contrast, the New York price was more than 5 cents above the Chicago price plus transport costs only 4 percent of the 224 occasions when storage exceeded 600 000 bushels. The relationship between the spot price difference and inventory levels can be summarised by constructing a “spatial arbitrage-storage” curve that plots the difference between

³⁰ In contrast, the spatial price difference exceeded the cost of lake and rail shipment only 63 per cent of the time.

the New York spot price and the Chicago spot price adjusted for transport costs against inventory levels (see Figure 10). The scatter-plot of points in Figure 10 is accompanied by a non-parametric kernel regression showing the average relationship between the future premium and the storage quantity³¹.

To test whether these differences between groups were random, the data was split into six different storage categories, 0 – 150000 bu, 150001 – 300000 bu, 300001 – 600000 bu, 600001 – 1000000 bu, 1000001 – 2000000 bu, and more than 2000000 bu, and Wilcoxon-Mann-Whitney test statistics were calculated. These statistics are used to test the hypotheses that the cumulative distribution functions of the spatial price difference minus the transport cost were the same for each storage category, against the alternative that one distribution lay above the other. The hypothesis that the distribution of the spatial price difference minus the transport cost was equal to that of the group that had the lowest storage was rejected for every storage category. In addition, it is possible to reject the hypotheses that the cumulative distribution functions of the first four storage categories were the same, although for the categories above 1000000 bushels the distributions were similar. Consequently, it is possible to conclude that the large spatial price differences that occurred when storage in New York was low were not due to simple random variation.

Thirdly, the average spatial price difference increased at an almost one for one rate with the transport cost. The best regression line, estimated using feasible generalised least squares to take into account first order autocorrelation, is³²:

$$(P_t^{NY} - P_t^{CH}) = 1.75 + 0.85 \text{ Transport Cost}_t + 2.05 \text{ 1(Storage}_t < 300,000) + u_t$$

$$(0.29) \quad (0.06) \qquad (0.29)$$

$$u_t = 0.45u_{t-1} + e_t \qquad R^2 = 0.69 \qquad N = 367$$

where $\text{1(Storage} < 300,000)$ is an indicator function that takes on the value one if inventories were less than 300,000 and zero otherwise. The slope of the line, 0.85, suggests that the average spatial price difference did not increase at an exactly one-for-one rate with transport costs, although as is shown below, when the New York spot price is replaced by the New York future price the slope coefficient is very close to and insignificantly different from 1. The positive and significant

³¹ The kernel regression is estimated with an Epanechnikov kernel with bandwidth 300 000 bushels.

³² The observation for 23 September 1884, when there was a corner in Chicago and the Chicago price was 16 cents higher than the New York price, was omitted. When included, the autocorrelation coefficient was 0.33, and the coefficient on the transport variable was 0.89 with a standard error of 0.061.

coefficient on the low inventory variable confirms that the New York spot price spiked upwards relative to the Chicago price and the transport cost when inventories were low.

Future - Spot Prices Differences and New York Storage Volumes

It is of interest to compare the above relationship with the relationship between New York inventories and the difference between the New York future price and the Chicago spot price plus the transport cost. According to equation 6c, the expected spot price in New York on the date when the corn was expected to arrive should have been equal to the spot price in Chicago three weeks earlier plus the cost of transport, independently of New York inventory levels. Although the expected spot price in New York on the date that the corn was expected to arrive is not known, it is possible to use the price for future delivery in New York as a proxy. Since the seller had the option of delivering any time during the month, and the trip took three weeks, the price for delivery “this month” was used if the date of the month was before the eleventh, and the price for delivery in the subsequent month was used if the date occurred on or after the eleventh.

The results for the New York future — Chicago spot price gap and the New York spot — Chicago spot price gap are quite different. This can be seen immediately by comparing Figure 8 with a graph of the difference between the New York future price and the Chicago spot price plotted against the transport cost, Figure 9. While there is a similar one-for-one relationship between the spot-future price difference and the transport cost, the spikes in the New York spot price that occurred when New York inventories were low are conspicuously missing. The difference can be confirmed statistically in one of two ways. First, the best fitting regression line corresponding to the graph is:

$$(F_t^{NY} - P_t^{CH}) = 1.11 + 0.95 \text{ Transport Cost}_t + 0.22 \text{ 1}(\text{Storage}_t < 300,000) + u_t$$

(0.23) (0.044) (0.22)

$$u_t = 0.41u_{t-1} + e_t \quad R^2 = 0.56 \quad N = 357$$

The coefficient on the low storage variable is small and insignificantly different from zero, indicating that there was no systematic tendency for the future-spot price gap to be high when inventories were low. (Also note that the slope of the coefficient on the transport cost variable is very close to, and insignificantly different to 1, indicating a one-for-one relationship between the New York future – Chicago spot price gap and the transport cost.) Secondly, a set of Wilcoxon-Mann-Whitney statistics can be calculated to test the hypothesis that the distribution of the spot-future spatial price difference was independent of inventory levels; these are reported in Table 5. These also show that the New York future price did not spike upwards when inventories were low; indeed, it is not possible to reject the hypothesis that the cumulative distribution of the spot-

future price difference minus the transport cost was different to that of any other group for any of the six inventory categories. These results clearly show that only the New York spot price spiked upwards when inventories were low, not the New York future price.

Threshold Regressions

Further support for the model is obtained by examining the relationship between the change in the price in each centre and the price difference between centres. According to the model, the spatial price difference would have been higher than the transport cost when there were low inventories in New York, and thus when the New York price was temporarily high. Consequently a regression of the price change in New York against the spatial price difference (adjusted for transport costs) should have a negative slope, while there should be no relationship between the spatial price difference and the change in Chicago price. This proves to be the case when the equations are estimated³³:

$$(P_{t+1}^{NY} - P_t^{NY}) = 0.75 - 0.29(P_t^{NY} - P_t^{CH} - K_t^T) + e_{t+1}^{NY} \quad N = 368$$

$$(0.23) \quad (0.089) \quad R^2 = 0.08 \quad \sum e^2 = 2073$$

$$(P_{t+1}^{CH} - P_t^{CH}) = -0.10 + 0.07(P_t^{NY} - P_t^{CH} - K_t^T) + e_{t+1}^{CH} \quad N = 368$$

$$(0.21) \quad (0.08) \quad R^2 = 0.005 \quad \sum e^2 = 1789$$

In section 4 it was argued that the relationship between the change in the spot price and the spatial price difference ought to depend on whether or not the spatial price difference exceeded the transport cost. Let $z_t = P_t^{NY} - P_t^{CH} - K_t^T$. The following threshold regressions were estimated:

$$\Delta P_{t+1}^{NY} = \alpha_0 + \alpha_1 z_t + [\alpha_2 + \alpha_3(z_t - \gamma)] 1(z_t > \gamma) + e_{t+1}^{NY}$$

$$\Delta P_{t+1}^{CH} = \beta_0 + \beta_1 z_t + [\beta_2 + \beta_3(z_t - \gamma)] 1(z_t > \gamma) + e_{t+1}^{CH}$$

where $1(z_t > \gamma)$ is an indicator function equal to one if $z_t > \gamma$ and zero otherwise. If all transport costs were known exactly and output were described by a first order autoregressive process, the threshold for the first equation would occur at $\gamma=0$, and the coefficients $(\alpha_1, \alpha_1 + \alpha_3)$ would equal $(0, -1)$ respectively. However, because New York never exported to Chicago, there would not be a

³³ New West-corrected standard errors (with 3 lags) in parenthesis. The observation for 23 September 1884 was omitted. If included, the coefficient for New York increased to -0.21 (0.11) and the coefficient for Chicago decreased to -0.005 (0.09) .

threshold in the relevant data range for the second equation, so $(\beta_1, \beta_1+\beta_3)$ would equal $(0,0)$ ³⁴. In practice not all components of the transport cost, such as the cost of transferring grain to an elevator, are known exactly and consequently the threshold parameters need estimation.

The equations were estimated by conducting a grid search on the value of γ , with the parameters $\alpha(\gamma)$ or $\beta(\gamma)$ estimated using OLS for each value of γ . The value of γ that minimised the sum of square errors, $\hat{\gamma}$, was selected; the estimates of α and β corresponding to $\hat{\gamma}$, and the associated Newey-West standard errors are reported below. Following Hansen (1996), separate Wald tests of the hypotheses that $(\alpha_2, \alpha_3) = (0,0)$ and $(\beta_2, \beta_3) = (0,0)$ were calculated. The distribution of the statistics are non-standard, so the sample statistic was compared to a simulated distribution. The grid search was initially conducted on a range that ensured that at least 15 percent of the observations were on each side of the threshold, using an increment of 0.02 cents. Since the threshold was estimated to be close to the endpoints, they were re-estimated on a larger range that only required that at least 10 per cent of the observations were on each side of the threshold. The results are qualitatively similar, although the standard errors are larger.

The New York equation is not particularly well behaved, as the value of the threshold that minimises the sum of squared residuals is not the same as the value that maximises the Wald statistic. This may be due to heteroskedasticity in the error process, or it could indicate more than one threshold. The equation that minimises the sum of squared residuals is:

$$\Delta P_{t+1}^{NY} = 0.43 - 0.08z_t + [0.89 - 0.96(z_t - 4.46)]1(z_t > 4.46) + e_{t+1}^{NY} \quad n = 368$$

(0.23) (0.11) (0.83) (0.39)* $R^2 = 0.13 \quad \sum e^2 = 1954$

The estimated threshold is 4.46 cents, although this appears to be imprecisely estimated³⁵. The spot estimates of α_1 and $\alpha_1+\alpha_3$, $(-0.08, -1.04)$, are very close to and insignificantly different to the

³⁴ The threshold for the second equation would occur when $P_t^{CH} = P_t^{NY} + K_t^T$ or $z_t = -2K_t^T$. There was only one observation for which $z_t < -2K_t^T$.

³⁵ A * indicates the coefficient is statistically different to zero at the five percent level. The thresholds were restricted to lie between 0.44 and 4.52, (i.e, so at least 15% of observations were on each side of the threshold), so this estimate is very close to the maximum permitted threshold. The estimate when the thresholds were restricted to lie between 0.18 and 5.50 (i.e, so at least 10% of observations were on each side of the threshold) is

$$\Delta P_{t+1}^{NY} = 0.48 - 0.13 z_t + [1.99 - 1.57(z_t - 5.50)]1(z_t > 5.50) + e_{t+1}^{NY} \quad n = 368$$

(0.22) (0.10) (1.10) (0.58)* $R^2 = 0.15 \quad \sum e^2 = 1902$

When the threshold region was widened further, the best estimate of γ was 5.68.

predicted values of (0, -1). To test whether a threshold model was appropriate, an average Wald statistic was calculated (see Hansen (1996, p415)); the value of the statistic was 5.45, corresponding to the 95.3 percentile of the appropriate simulated distribution, and thus marginally statistically significant³⁶.

The best equation for Chicago prices is³⁷:

$$\Delta P_{t+1}^{CH} = 0.24 + 1.34 z_t + [-0.72 - 1.35(z_t - 0.46)]1(z_t > 0.46) + e_{t+1}^{CH} \quad n = 368$$

(0.38) (0.37)* (0.53) (0.38)* $R^2 = 0.07 \quad \sum e^2 = 1670$

The estimated threshold, 0.46 cents, is very close to zero, indicating that there was a different price dynamic in Chicago when the price was greater rather than less than the New York price minus transport costs. The spot estimate of $\beta_1 + \beta_3$, -0.01, is very close to and insignificantly different from the predicted value of zero. In contrast, the spot estimate of β_1 is +1.34, close to and insignificantly different from 1. It appears, therefore, that when prices were high in New York relative to Chicago, the Chicago price was independent of the spatial price difference, but when prices were relatively high in Chicago and the spatial price difference was less than the transport cost, Chicago prices were temporarily high and prices subsequently fell. The value of the average Wald test was 11.5, corresponding to the 99.8 percentile of the appropriate simulated distribution, and thus statistically significant at conventional levels.

It is worth emphasizing that the Chicago threshold regression is not consistent with the model, for according to the model on most occasions when the spatial price difference was less than the transport cost (and the Chicago price was not so high that corn was imported from New York) both centres should have had storage, no trade should have taken place, and prices should have been expected to gradually increase in both centres to cover the cost of storage. This appears not to have been the case, however, as the estimated regression is consistent with a model in which

³⁶ One thousand series 368 observations long were generated, and for each series a threshold regression was estimated using a grid search on γ with 0.02 increments, with minimum and maximum values of γ selected so that at least 15 percent of the sample was on each side of the threshold region. In each case the average Wald statistic was calculated as a simple arithmetic average of the Wald statistic associated with each value of γ . The simulated series were generated to have the same properties as the sample properties of the real data. In particular, the data were generated by the following processes:

$$P_{t+1}^{CH} = 0.97P_t^{CH} + e_{t+1} \quad e_{t+1} \sim N(0, 2.56) \quad K_t^T \sim N(8.15, 2.62)$$

$$\Delta P_{t+1}^{NY} = 0.75 - 0.29(P_t^{NY} - P_t^{CH} - K_t^T) + u_{t+1} \quad u_{t+1} \sim N(0, 2.38)$$

³⁷ When the thresholds were restricted to lie between 0.18 and 5.50 rather than 0.44 and 4.52 the estimated equation was

the spatial price difference is normally less than transport cost because of a temporary local shortfall in Chicago, and a corresponding temporary price spike in Chicago. These results suggest that the stochastic process used in the model is not an appropriate description of output in this market.

The Spatial Arbitrage – Storage Curve

It is well known that if the difference between the spot price and the future price of a commodity is plotted against the quantity of storage, a “supply of storage” curve is generated. When storage volumes are low, the spot price is typically higher than the future price, but when storage quantities are high, the future price exceeds the spot price. It is possible to estimate a supply of storage curve for New York corn prices, using the difference between the spot price and the price for delivery the next month as the future premium. When this is done, it has the standard shape (see Figure 11).

It is apparent that the supply of storage curve and the “spatial arbitrage-storage” curves have similar shapes. By plotting the future price premium $F_t^{NY,t+1} - P_t^{NY}$ against the spatial price difference adjusted for transport cost, $P_t^{CH} + K_t^T - P_t^{NY}$, it is possible to show that when the New York spot price exceeded the future price, the New York spot price also exceeded the Chicago price plus the transport cost (see Figure 12). The data in Figure 12 indicate that there was a strong positive correlation between the two variables when storage was less than 600 000 bushels, and a somewhat weaker correlation when storage was higher than this level. More formally, when the variables are regressed against each other, the best fitting line is:

$$(F_t^{NY,t+1} - P_{NY,t}) = [1.21 + 0.93(P_t^{CH} + K_t^T - P_t^{NY})] \mathbb{1}(Storage \leq 600000) \\ (0.23) \quad (0.05) \\ + [0.78 + 0.49(P_t^{CH} + K_t^T - P_t^{NY})] \mathbb{1}(Storage > 600000) + u_t \\ (0.23) \quad (0.05)$$

$$u_t = 0.49u_{t-1} + e_t \quad R^2 = 0.75 \quad N = 367$$

$$\Delta P_{t+1}^{CH} = 0.88 + 1.64z_t + [1.06 - 1.65(z_t - 0.18)] \mathbb{1}(z_t > 0.18) + e_{t+1}^{CH} \quad n = 368 \\ (0.49) \quad (0.35)^* \quad (0.55) \quad (0.36)^* \quad R^2 = 0.08 \quad \sum e^2 = 1647$$

This regression confirms that in this case the supply of storage curve is closely related to the spatial arbitrage-storage curve³⁸. The reason is that when storage was low, the spot price in New York temporarily increased and exceeded both the future price in New York and the spot price in Chicago plus the transport cost. While the set of equations 6a-6d suggests that these price spikes should have only occurred when storage was literally zero, in this case they occurred when storage was positive but small because the elevators were used to transfer grain to ocean going ships and thus always had some grain in transit.

7. Conclusions

This paper has attempted to enhance economists' understanding of how spatial arbitrage occurs by examining how the interaction of storage and trade affects prices in different locations. Its theoretical contribution has been to link the economics literature analysing spatial price arbitrage with the logistics management literature analysing how the speed of transport determines inventory holdings. It has made the link by relaxing the standard assumption in models of spatial arbitrage that transport is instantaneous; in doing so it has emphasised the manner in which inventories are used to smooth price fluctuations. This role had long been recognised in the logistics management literature, but only at the level of the firm, rather than across competitive markets. Its empirical contribution has been to assemble a set of price, transport cost, and storage data that is detailed enough to detect how logistics issues have affected commodity prices in one specific market. It has shown that the spatial price difference frequently exceeded the cost of shipping goods, in a heavily traded commodity market in which there were large investments in logistics infrastructure and thick financial markets. In doing so, it has established that prices exceeded the transport cost when there were low supplies in the importing city, causing a temporary price spike followed by a price decrease when new supplies arrived from the exporting centre. The result can perhaps be best summarised by plotting a spatial arbitrage-storage curve that shows how the spatial price difference depends on inventory levels in the importing centre.

The theoretical model suggests that the process of physical arbitrage is significantly more complex than has previously been recognised. The central feature of the model is that the

³⁸A similar relationship can be estimated between the Chicago future premium and the spatial price difference adjusted for transport costs. There were some occasions when the Chicago future price was less than the Chicago spot price at the same time that the New York spot price was higher than the Chicago spot price plus transport cost. Most of these occasions occurred in 1891 when storage quantities in both Chicago and New York were low. The corresponding regression has much less explanatory power, suggesting the relationship is specifically between the New York future premium and the spatial price difference.

quantity of goods arbitrageurs ship each period will not be the amount necessary to ensure that the spatial price difference is always equal to the cost of transport. Rather, in order to make normal profits on average, arbitrageurs will ship an amount such that there will be insufficient supplies and high prices in the destination centre on a regular but infrequent basis. The profits they make on these occasions offset the small losses they make on the more frequent occasions that the exports arrive when local supplies are adequate. Put more starkly, traders make normal profits on average only by obtaining high prices and extraordinary profits on occasions that supplies in the importing market are low. If this phenomenon is generally true in practice, it provides an argument against governments attempting to stabilise prices in times of shortage.

Three other implications of the model are of interest. First, prices in different centres should diverge when the spatial price difference is less than or equal to the cost of transportation (adjusted for the storage cost), for in these circumstances inventories in each centre are positive and prices in each centre should rise when inventories are positive. Secondly, the difference between prices in different centres should on occasion be greater than the cost of transportation because local price spikes occur when local inventories are exhausted. Thirdly, importing centres will typically but not always have small inventories of goods available, simply because when transport takes time the exporting centre cannot export precisely the correct amount to keep inventories equal to zero.

If the price behavior evident in the New York and Chicago corn markets is representative of price behavior in other markets, there are five interesting implications. First, even though Deaton and Laroque (1992, 1996) showed that rational expectations models of storage do not explain annual frequency commodity price behavior at all well, their models and those solved by Williams and Wright (1991) may be consistent with micro-economic data. The model in this paper is a straightforward extension of these storage models, and appears to satisfactorily explain the relationship between New York and Chicago prices. Consequently, it may be the case that the empirical failure of rational expectations storage models to explain long term commodity price series is because the data is collected at a too aggregated level, as Wright and Williams (1989) suggested.

Secondly, studies of market integration that assume that the spatial price difference is always equal to the transport cost need modification. For example, Spiller and Wood (1988) developed a

popular methodology to estimate spatial transactions costs from price data that is explicitly predicated on the assumption that the price difference is always equal to the cost of transport, but that the cost of transport is variable. If that methodology were used on this data set, it would overestimate the mean and variability of transport costs because it wrongly assumes that the large price differences that occurred when one centre had low storage were due to temporarily high transport costs.

Thirdly, if cheap transport is slow, studies of market integration need to focus more on logistics issues — the way arbitrageurs use transport systems and storage to integrate markets — rather than just transport issues. The logistics literature has shown that when companies ship goods between different plants, storage choices are as important as transport choices. The data analysed here, and the models of Williams and Wright, suggest that the interaction of storage and transport is key to understanding the short term behavior of commodity markets as well.

Fourthly, the similarity of the supply of storage curve in New York and the spatial arbitrage curve between New York and Chicago suggests that spot prices exceed future prices in a centre when storage is low and new supplies cannot be quickly and cheaply reordered. If the supply of storage and spatial arbitrage curves are closely related in general, it suggests that the supply of storage curve exists because it is most profitable to use slow, cheap transport systems to ship goods. If so, a rational expectations model of storage and trade incorporating slow transport may be the appropriate framework for investigating the phenomena of a supply of storage curve. Further research should be directed at establishing whether the spatial arbitrage-storage curve is as ubiquitous as the supply of storage curve.

Lastly, this data supports Wright and Williams' contention that a supply of storage curve can exist even when a commodity has no convenience yield. According to the theoretical model, the spot price can temporarily exceed the future price in an importing centre if inventories fall to zero. The model further suggests that such occasions should occur regularly, explaining the frequency of occasions when prices are in backwardation. While the model does not explain why the spot price spikes upwards when inventories are low but not literally zero, the New York example suggests that inventories may never fall to zero because the elevators and warehouses in which they are held are dual use, being used in this case to transfer grain as well as to store it.

Appendix 1: Solving the Rational Expectation Model

This appendix contains more details about the solution technique used to solve the model in Section 2. A solution is found by constructing an algorithm that calculates a series of successive approximations to the optimal storage and trade functions, $S_k^i(y_t)$ and $T_k^i(y_t)$, where k refers to the k^{th} approximation. The iteration process is as follows. First, given the k^{th} value of these functions, the price functions for each centre are calculated at all grid points, using the inverse demand function $P_k^i(y_t) = P^{D,i}(D_k^i(y_t))$, where $D_k^i = X_t^i + M_t^i - S_k^i(y_t) - T_k^i(y_t)$.

Secondly, given the k^{th} storage and trade rules $S_k^i(y_t)$ and $T_k^i(y_t)$, a schedule of the expected future price conditional on being at M_{t+1}^A and M_{t+1}^B in the subsequent period is calculated. At each point $(M_{t+1}^A, M_{t+1}^B, X_t^A, X_t^B)$ the expected future price is calculated by multiplying the price functions by the transition matrix probabilities Π :

$$E_X \left[P_{k,t+1}^i \mid M_{t+1}^A, M_{t+1}^B, X_t^A, X_t^B \right] = \sum_{X_{i,t+1}^A, X_{j,t+1}^B} P^{D,i}(D_k^i(M_{t+1}^A, M_{t+1}^B, X_{i,t+1}^A, X_{j,t+1}^B)) \pi(X_{i,t+1}^A, X_{j,t+1}^B \mid X_t^A, X_t^B) \quad (13)$$

Thirdly, the values of the four complementary conditions associated with each of the equations 6a-6d are calculated at each point y_t . The value of the expected future price $E[P_{k,t+1}^i \mid y_t]$ is equal to $E_X [P_{k,t+1}^i \mid ((1-\delta)(S_k^A(y_t) + T_k^B(y_t)), (1-\delta)(S_k^B(y_t) + T_k^A(y_t)), X_t^A, X_t^B)]$ and is calculated using linear interpolation of equation 13. Linear interpolation is used because it ensures that if a linear inequality restriction holds at contiguous grid points, it will also hold in between the grid points. If the value of the complementary condition is inconsistent with the value of the associated control variable, — for example, if in inequality 6a, $S_k^A > 0$ but

$\frac{1-\delta}{1+r} E[P_{t+1}^A \mid y_t] - P_t^A(y_t) - K^S \neq 0$, or $S_k^A = 0$ but $\frac{1-\delta}{1+r} E[P_{t+1}^A \mid y_t] - P_t^A(y_t) - K^S > 0$ — then the

value of the control variable is recalculated at the grid point; otherwise, it remains unchanged.

The new $k+1^{\text{th}}$ values of the control variables associated with the inconsistent inequalities are simultaneously calculated at the grid point using an optimising routine such as the Newton Rhapsod or the secant method. The process is repeated at each grid point until S_{k+1}^i and T_{k+1}^i are calculated over the whole domain.

The whole algorithm is repeated until the difference between successive values of the control values is small. It proved necessary to sacrifice accuracy at some points of the domain of y_t simply to construct a four dimensional grid small enough to modeled by a computer; in the end, the grid structure consisted of some 20000 to 40000 points. The algorithm was solved on a personal computer using Gauss software. Typically convergence took 300 iterations to achieve acceptable accuracy.

The starting value for the algorithm was calculated by splitting the combined centre solution into two, that is by estimating storage functions S^A and S^B such that $S^A(M^A, M^B, X^A, X^A) + S^B(M^A, M^B, X^A, X^A) = S(M^A, M^B, X^A, X^A)$, where $S(\cdot)$ is the solution for storage in the combined centre problem. The initial value for the trade variables was zero.

The Invariant Distribution.

The invariant distribution of the model is the unconditional probability of being at a particular grid point. Deaton and Laroque (1995) suggest a method for finding the invariant distribution as follows. Suppose Y is a $n \times 1$ vector of all possible grid points corresponding to the four state variables. Let the function $\omega: (Y \times Y) \rightarrow \mathbf{R}$ be the conditional transition probabilities of moving from one point one period to another in the next period:

$$\omega(y_1, y_2) = \text{Prob}(Y_{t+1} \in B(y_2) | Y_t = y_1)$$

where $B(y_2)$ is a region around y_2 defined so that the regions form a partition of the space and only include one grid point. Let Ω be the $n \times n$ matrix of these probabilities, with $\Omega_{ij} = \omega(y_i, y_j)$. Ω has at least one unit eigenvalue, since all the rows of Ω sum to 1 and all elements of Ω are strictly less than 1 by construction. The eigenvector corresponding to this eigenvalue will be the invariant probability distribution of Y .

Deaton and Laroque suggest that the eigenvalue and its corresponding eigenvector can be found by inverting the $n \times n$ matrix Ω . This is not practical in this case as n typically exceeds 20000. The alternative method for finding the invariant distribution is simply to multiply an arbitrary initial distribution on Y by the transition matrix Ω' until some successive values of the product meet some convergence criteria. This method proves to be fast, for even though Ω has several hundred million elements, the vast majority of these are zeros and it is straightforward to devise

an algorithm that only uses the several hundred thousand positive elements in the iterative procedure.

Appendix 2: Data Sources

Six kinds of data have been assembled for this project: the spot price of corn in New York and Chicago; the future price of corn in New York and Chicago; transport costs between Chicago and New York; transport volumes between Chicago and New York; storage prices in Chicago and New York; and storage volumes in Chicago and New York.

Spot Price of Corn.

Prices were collected for Number 2 Yellow corn. Number 2 corn was the primary future grade and comprised a large fraction of the spot market. Grades were defined as follows.

New York: “YELLOW CORN shall be sound, dry, plump and well cleaned; an occasional white or red grain shall not deprive it of this grade. No.1 CORN shall be mixed corn of choice quality, sound, dry and reasonably clean. No.2 CORN shall be mixed corn, sound, dry and reasonably clean. ” New York Produce Exchange (1882) p207

Chicago: “No. 1 YELLOW CORN shall be yellow, sound, dry, plump and well cleaned. No. 2 CORN shall be dry, reasonably clean, but not plump enough for No. 1” Chicago Board of Trade (1882) p 79-80

Spot prices for both cities were collected in the Thursday edition of the New York Times, 1878-1891. The prices were for the preceding Wednesday. If the Wednesday were a public holiday, the Tuesday price was collected. If the markets were closed on both Wednesday and Tuesday, the data was skipped for that week.

Daily spot prices for New York are also available in some years in the Annual Report of the New York Produce Exchange. However, since The New York Times had to be used to collect the Chicago spot price and the New York future price, as well as the New York spot price in years where it was not reported in the Annual Report, the weekly New York Times data was used.

Future Price of Corn.

Prices were collected for Number 2 Yellow corn. The Chicago future price was collected from the Annual Report of the Chicago Board of Trade. The quotes is for seller delivery: the seller could choose any day to deliver within the said month. Wednesday quotes were collected.

The New York Wednesday future prices were collected from the Thursday edition of the New York Times. The seller also had the option as to the delivery date.

Corn Trade and Storage Data.

Storage and trade data for Chicago was sourced from the Chicago Board of Trade Annual Reports. The New York data came from a variety of sources. Where possible, it came from the New York Produce Exchange Annual Reports, but these documents had little data between 1882 and 1887. Storage data for these years came from the weekly Commercial and Financial Chronicle. Export data for several years came from the Chicago Board of Trade Annual Reports. Storage cost data come from the Chicago Board of Trade and New York Produce Exchange Annual Reports, and from Goldstein (1928).

Transport Data

The transport cost data were published by the Chicago Board of Trade and New York Produce Exchange Annual Reports. They are similar not identical to the data published in the Aldrich Report, (United States 52nd Congress 2nd Session (1893) *Senate Report 1394: Wholesale Prices, Wages, and Transportation. Report by Mr Aldrich from the Committee on Finance March 3 1893 Part 1.* (Washington: Government Printing Office).

Table 1: Analytical results corresponding to equations 6a-6d.

S_t^A	S_t^B	T_t^A	T_t^B	Prob(.)	$ P_t^A - P_t^B $	$E[P_{t+1}^A - P_{t+1}^B y_t]$
>0	>0	=0	=0	68.1%	$< (K^T - K^S)$	$\frac{1+r}{1-\delta}(P_t^A - P_t^B)$
>0	>0	=0	>0	13.4%	$= (K^T - K^S)$	$\frac{1+r}{1-\delta}(K^T - K^S)$
>0	>0	>0	=0	13.4%	$= (K^T - K^S)$	$\frac{1+r}{1-\delta}(K^T - K^S)$
>0	=0	=0	=0	0.05%	Uncertain	Uncertain
>0	=0	=0	>0	0.5%	$= (K^T - K^S)$	Uncertain
>0	=0	>0	=0	1.4%	$> (K^T - K^S)$	$\frac{1+r}{1-\delta}(K^T - K^S)$
=0	>0	=0	=0	0.05%	Uncertain	Uncertain
=0	>0	=0	>0	1.4%	$> (K^T - K^S)$	$\frac{1+r}{1-\delta}(K^T - K^S)$
=0	>0	>0	=0	0.5%	$= (K^T - K^S)$	Uncertain
=0	=0	=0	=0	0.1%	Uncertain	Uncertain
=0	=0	=0	>0	0.5%	$> (K^T - K^S)$	$> \frac{1+r}{1-\delta}(K^T - K^S)$
=0	=0	>0	=0	0.5%	$> (K^T - K^S)$	$> \frac{1+r}{1-\delta}(K^T - K^S)$

The table gives the absolute value of the price differential $P_t^A - P_t^B$ and the expected future price differential $E[|P_{t+1}^A - P_{t+1}^B| | y_t]$ where they can be determined exactly by equations 6a - 6d, the arbitrage conditions describing storage and trade. The probabilities of each set of conditions occurring pertain to the baseline simulation for the two centre model with symmetric centres.

**Table 2: Price, storage, and trade statistics corresponding to the model.
(Identical centres, changing transport costs.)**

Statistic	Combined Centre	$K^T = 2.5$	$K^T = 5$	$K^T = 10$	No Trade
Prices					
Mean(P^A)	100.2	100.2	100.3	100.3	100.4
S. Dev.(P^A)	8.5	8.6	8.5	8.6	10.4
Mean($P^A - P^B$)	—	0	0	0	0
S. Dev.($P^A - P^B$)	—	3.1	4.6	7.3	7.5
% ($ P^A - P^B > K^T - K^S$)	—	2.9%	3.1%	3.1%	—
Storage					
Mean(S^A)	186	227	261	315	422
S. Dev.(S^A)	157	232	251	285	291
% ($S^A = 0$)	—	3.5%	3.1%	2.6%	1.0%
% ($S^A, S^B = 0$)	3.1%	1.7%	1.1%	0.6%	0.0%
Trade					
Mean(T^A)	—	3.9	3.2	2.3	—
S. Dev.(T^A)	—	10.6	10.5	10.2	—
% ($T^A = 0$)	—	80%	84%	90%	—
% ($T^A, T^B = 0$)	—	61%	68%	81%	—

P^A : the price in centre A. S^A : storage in centre A. T^A : trade from centre A to centre B.

% ($|P^A - P^B| > K^T - K^S$): the fraction of time the price difference exceeds the difference between the trade cost and the storage cost. Note $K^T = 5$ and $K^S = 0$ in these simulations.

% ($S^A[T^A] = 0$): the fraction of time storage [exports] = 0.

**Table 3: Price, storage, and trade statistics corresponding to the model.
(Different centres, mean output changes across columns.)**

Statistic $\bar{X}^A =$ $\bar{X}^B =$	100 100	95 105	90 110	80 120	70 130
Prices					
Mean(P^A)	100.3	101.2	101.9	102.5	102.8
S. Dev.(P^A)	8.5	8.4	8.5	9.0	9.3
Mean(P^B)	100.3	99.3	98.6	97.9	97.7
S. Dev.(P^B)	8.5	8.7	8.8	8.8	8.7
Mean($P^A - P^B$)	—	1.9	3.3	4.7	5.1
S. Dev.($P^A - P^B$)	4.6	4.4	3.8	2.7	2.0
% ($ P^A - P^B > K^T - K^S$)	3.7%	3.8%	3.2%	3.0%	2.1%
Storage					
Mean(S^A)	261	215	166	110	91
S. Dev.(S^A)	251	216	171	107	78
% ($S^A = 0$)	2.8%	3.1%	2.5%	2.4%	2.1%
Mean(S^B)	261	286	296	288	285
S. Dev.(S^B)	251	273	284	286	294
% ($S^B = 0$)	2.8%	3.8%	4.2%	7.7%	9.4%
% ($S^A, S^B = 0$)	0.9%	1.3%	1.3%	1.8%	1.9%
Trade					
Mean(T^A)	3.2	1.7	0.8	0.1	0.0
S. Dev.(T^A)	10.5	8.0	5.4	2.2	0.5
% ($T^A = 0$)	86%	92%	96%	99%	99.9%
Mean(T^B)	3.2	5.7	9.1	17.7	27.4
S. Dev.(T^B)	10.5	13.0	15.1	18.8	23.1
% ($T^B = 0$)	86%	73%	61%	32%	19%
% ($T^A, T^B = 0$)	72%	65%	56%	31%	18%

P^A : the price in centre A. S^A : storage in centre A. T^A : trade from centre A to centre B.

% ($|P^A - P^B| > K^T - K^S$): the fraction of time the price difference exceeds the difference between the trade cost and the storage cost. Note $K^S = 0$ in these simulations.

% ($S^A[T^A] = 0$): the fraction of time storage [exports] = 0.

Table 4: Distribution of spot price difference adjusted for transport costs ($P^{NY} - P^{CH} - K^T$) by storage level

Storage level bu	<150000	<300000	<600000	<1000000	<2000000	2000000 +
N	21	45	79	62	87	75
% obs < 0	0%	2%	2.5%	3%	10%	15%
% obs $0 \leq x < 5$	38%	67%	85%	90%	87%	83%
% obs ≥ 5	62%	31%	12.5%	6.5%	2%	3%
Mean	6.01	3.84	2.70	1.73	1.45	1.10
Std	2.93	2.57	1.89	2.82	1.61	1.81
WMW test 1		-2.77*	-4.44*	-5.23*	-5.95*	-5.96*
WMW test 2		-2.77*	-2.43*	-2.52*	-1.94	-1.03

The table shows the fraction of observations in each group which are less than zero, between 0 and 5 cents, and more than 5 cents.

The first Wilcoxon-Mann-Whitney test, WMW test 1, tests whether the distribution is the same as the distribution of the group for which storage is less than 150000 bushels.

The second Wilcoxon-Mann-Whitney test, WMW test 2, tests whether the distribution is the same as the distribution of the group immediately to the left.

The WMW test is asymptotically distributed as $N(0,1)$ and a * indicates significance at the 5% critical level i.e. that the two cumulative distributions lie above each other.

Table 5: Distribution of spot-future price difference adjusted for transport costs ($F^{NY} - P^{CH} - K^T$) by storage level

Storage level bu	<150000	<300000	<600000	<1000000	<2000000	2000000 +
N	19	44	77	60	84	74
% obs < 0	21%	7%	10%	8%	7%	11%
% obs $0 \leq x < 5$	74%	91%	86%	92%	92%	85%
% obs ≥ 5	5%	2%	4%	0%	1%	4%
Mean	1.04	1.60	1.60	1.29	1.45	1.55
Std	2.12	1.33	1.76	3.03	1.12	2.09
WMW test 1		1.00	1.07	0.92	0.73	0.90
WMW test 2		1.00	0.07	-0.31	-0.49	0.56

See Table 4. The table shows the fraction of observations in each group which are less than zero, between 0 and 5 cents, and more than 5 cents.

Figure 1

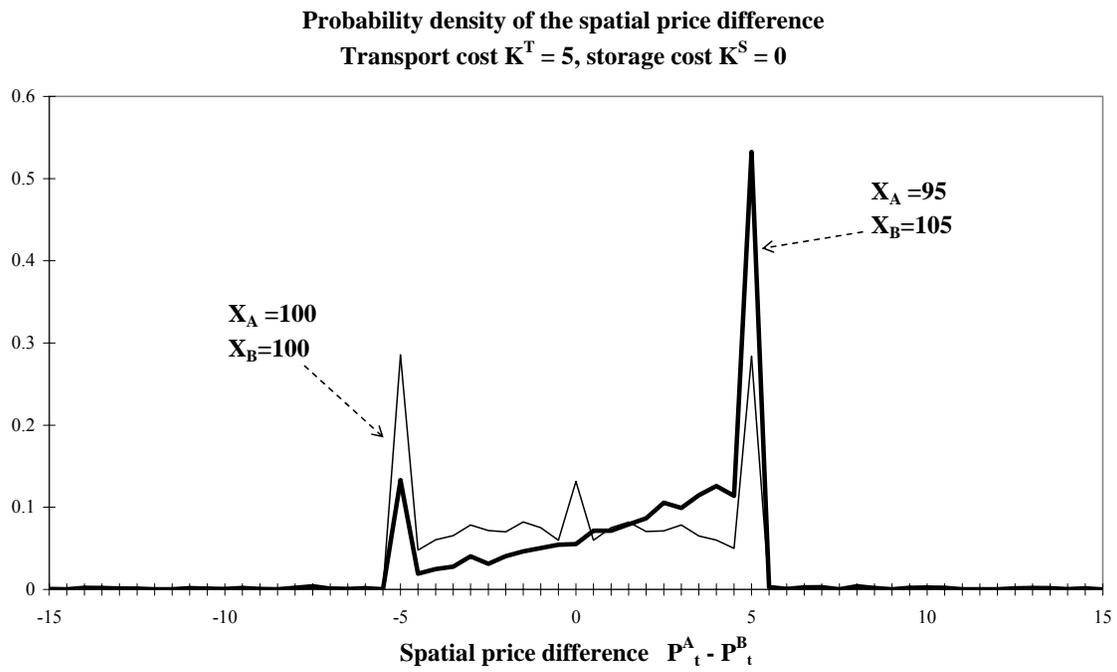


Figure 2

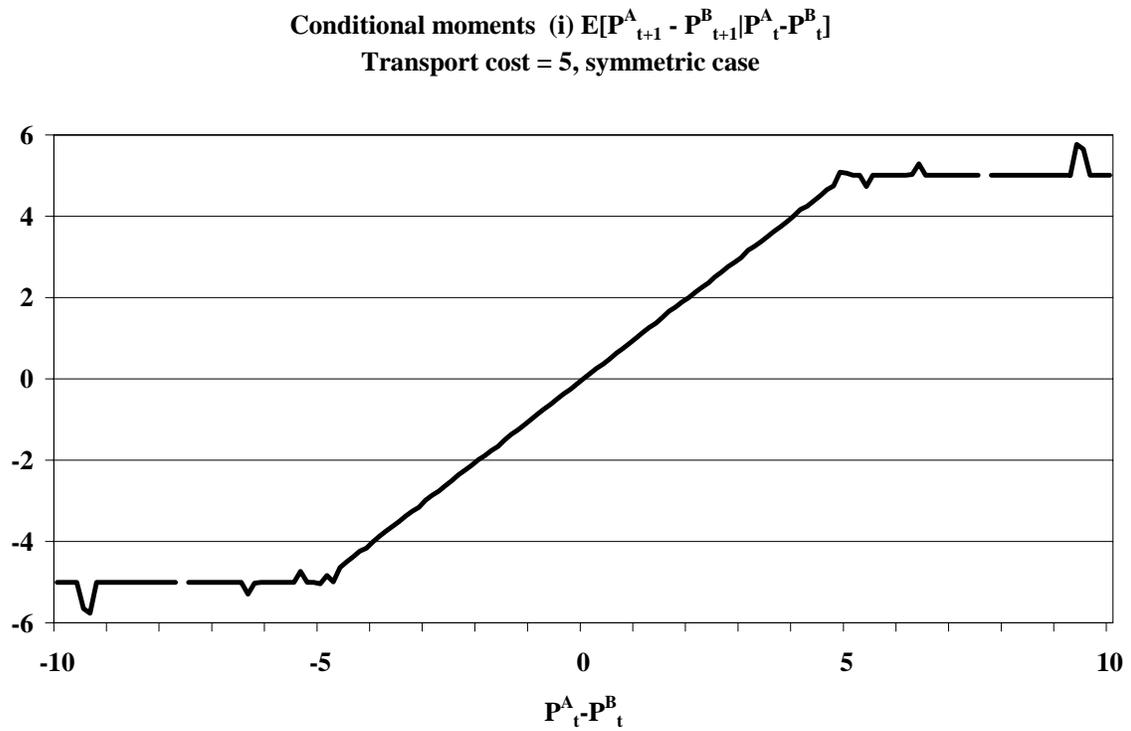


Figure 3

Conditional moments (ii) $E[P^A_{t+1} - P^A_t | P^A_t - P^B_t]$ and $E[P^B_{t+1} - P^B_t | P^A_t - P^B_t]$.
Transport cost = 5, symmetric case

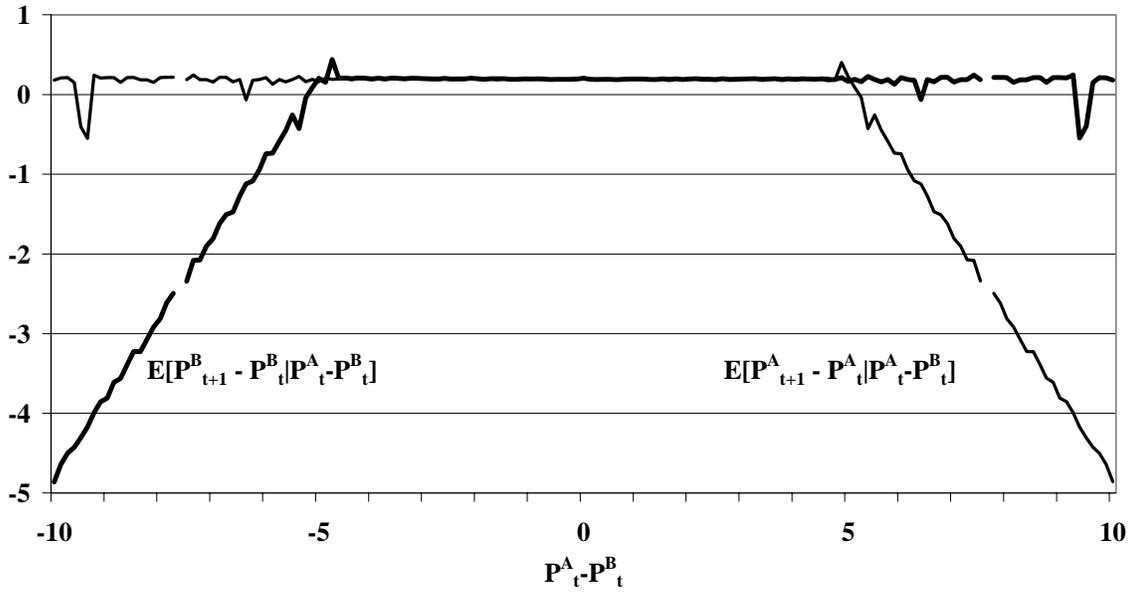


Figure 4

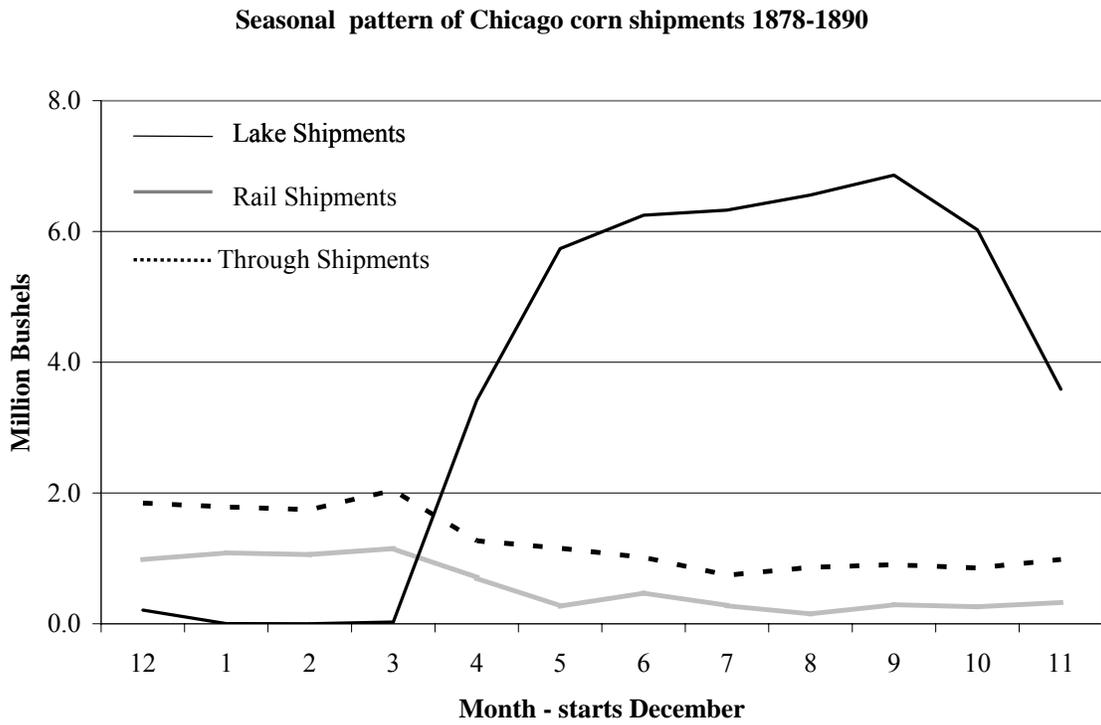


Figure 5

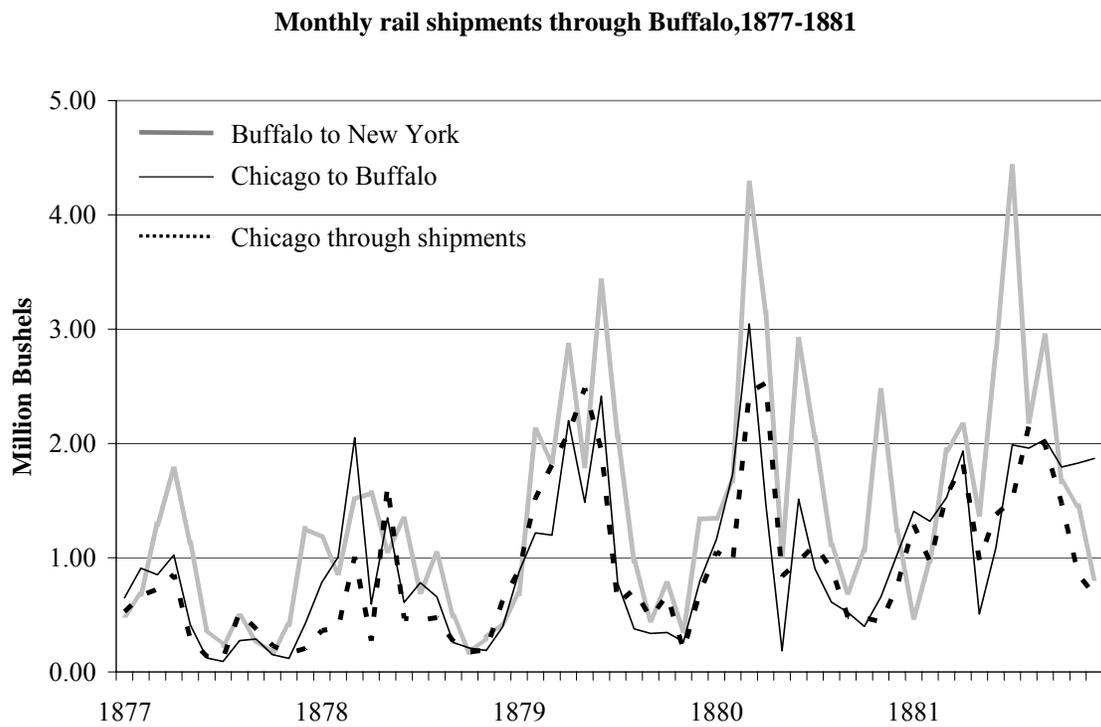


Figure 6

Average transport costs, 1880 - 1891, by week

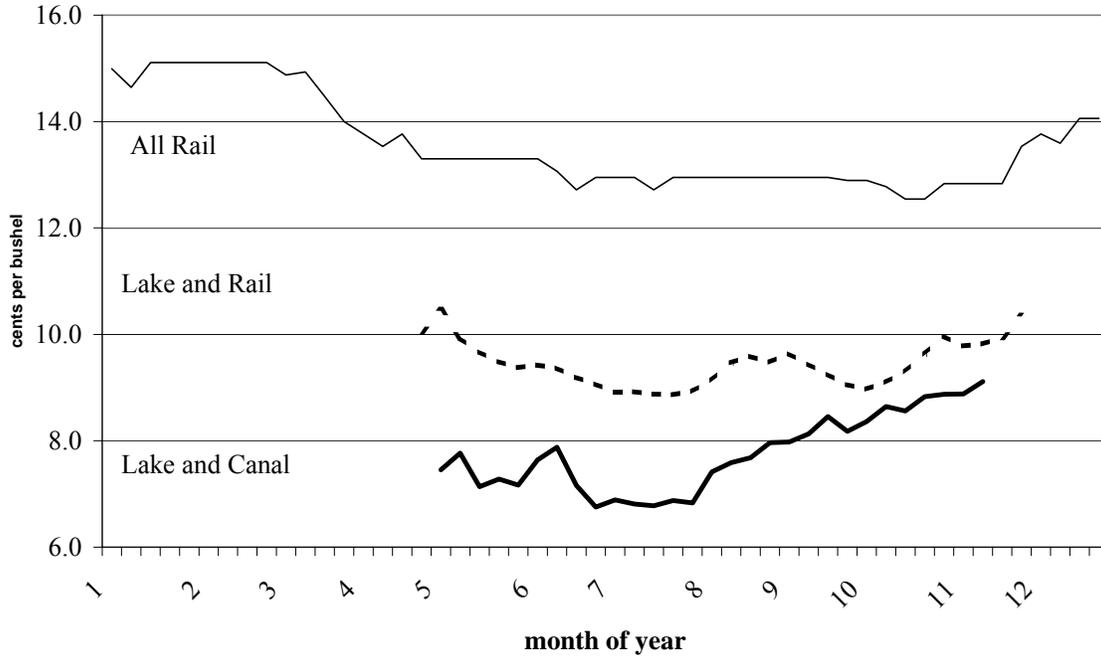


Figure 7

"Summer" Storage (May-November), New York, 1878-1891

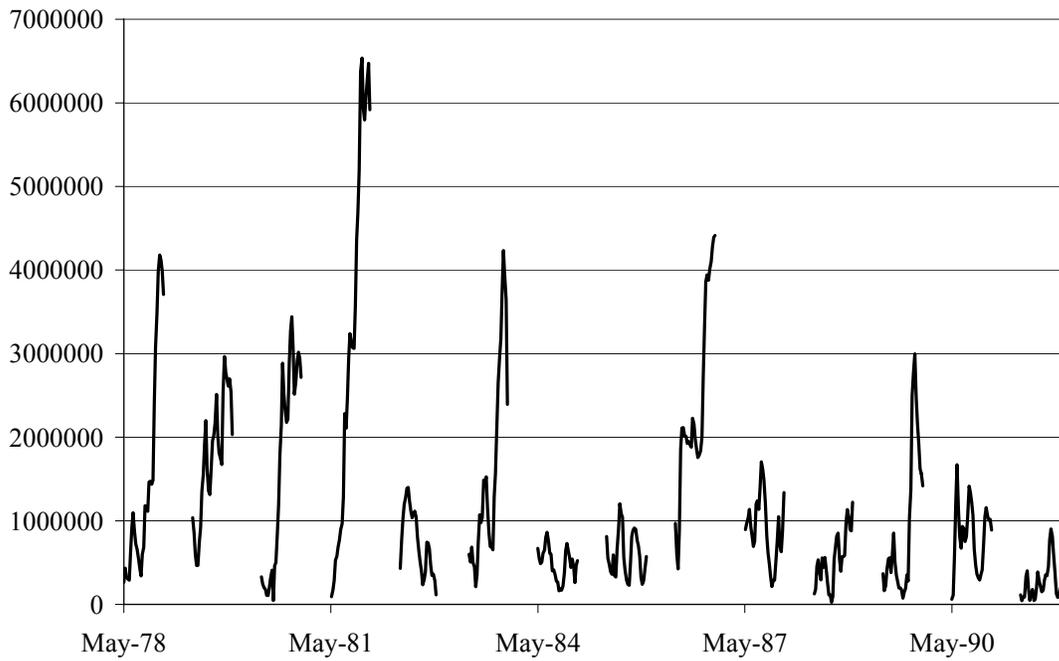


Figure 8

**NY spot and Chicago spot price difference versus transport cost
1878-1891**

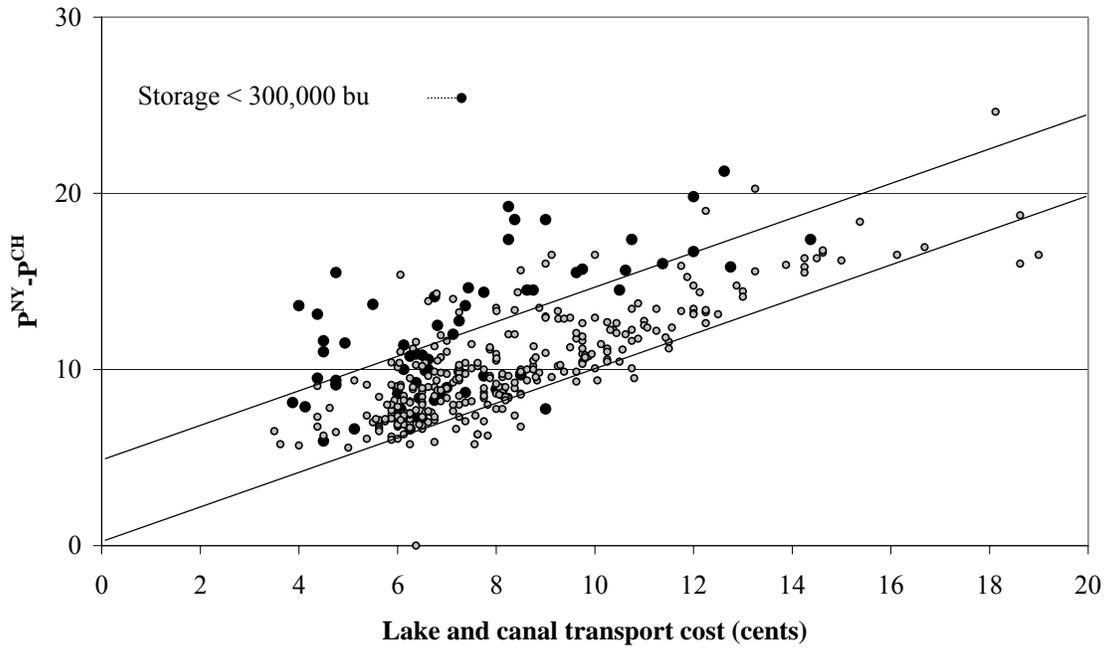


Figure 9

NY future minus Chicago spot price versus transport cost 1878-1891
New York future is "delivery this month" if date is before the 10th; otherwise "delivery next month"

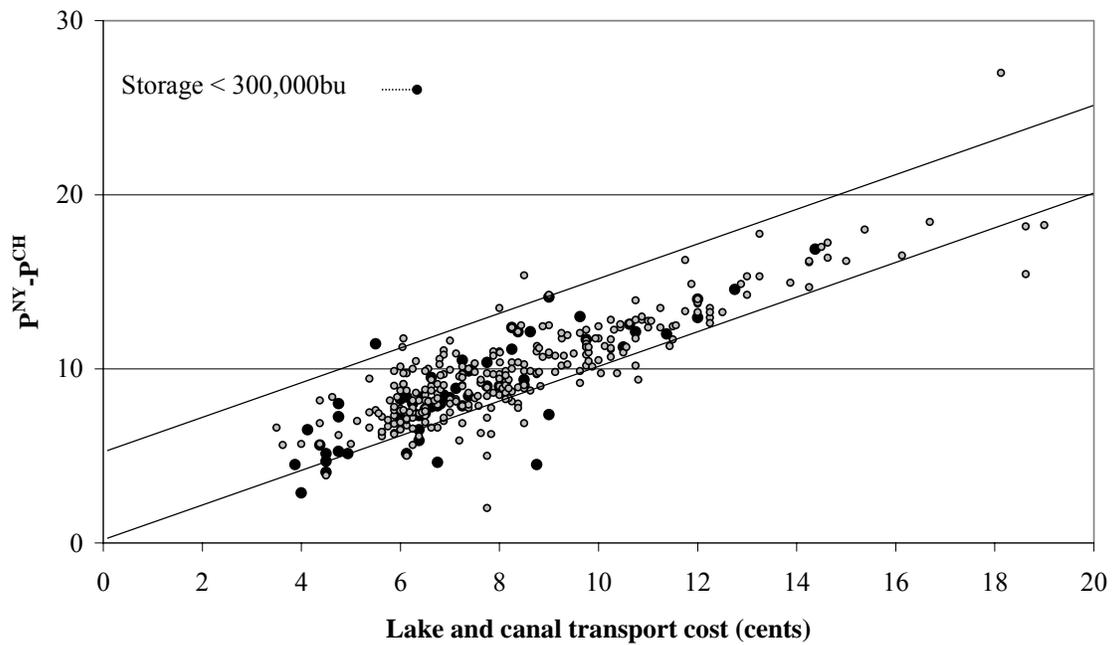


Figure 10

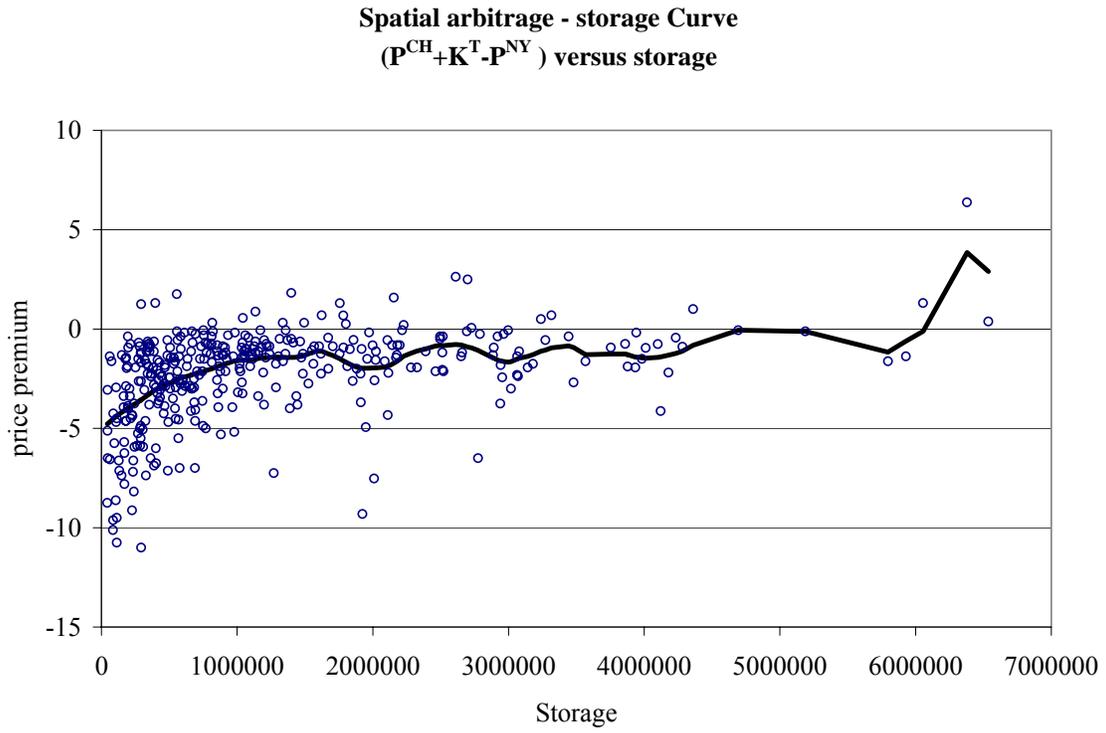


Figure 11

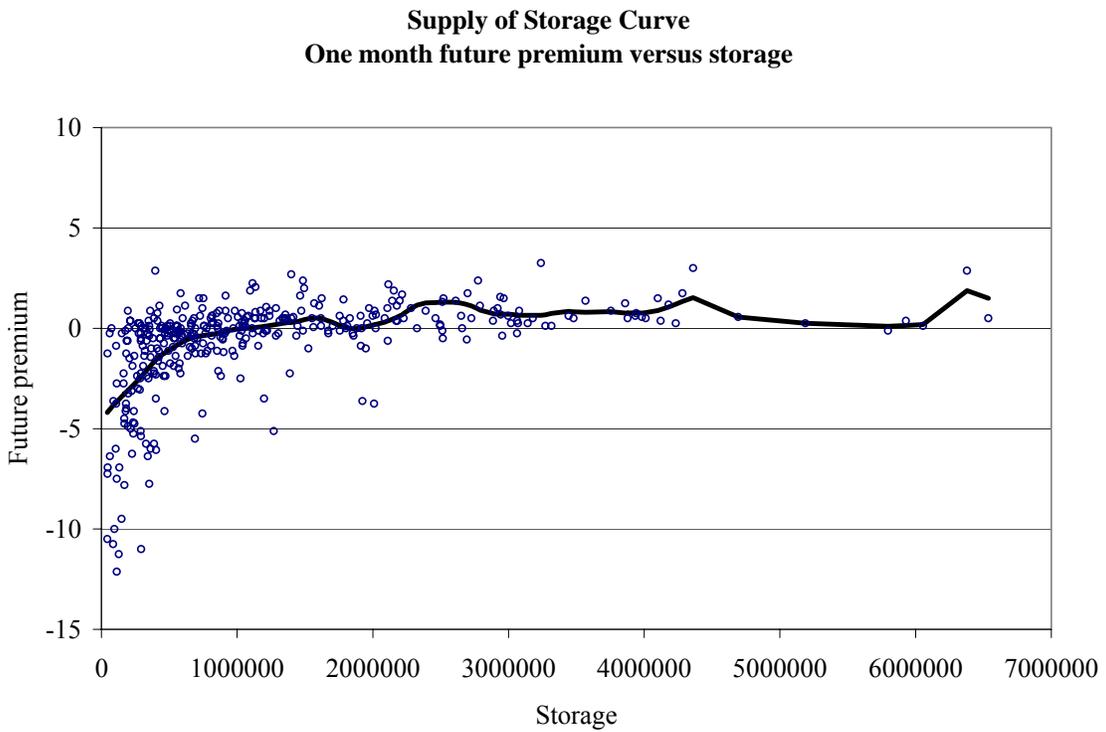
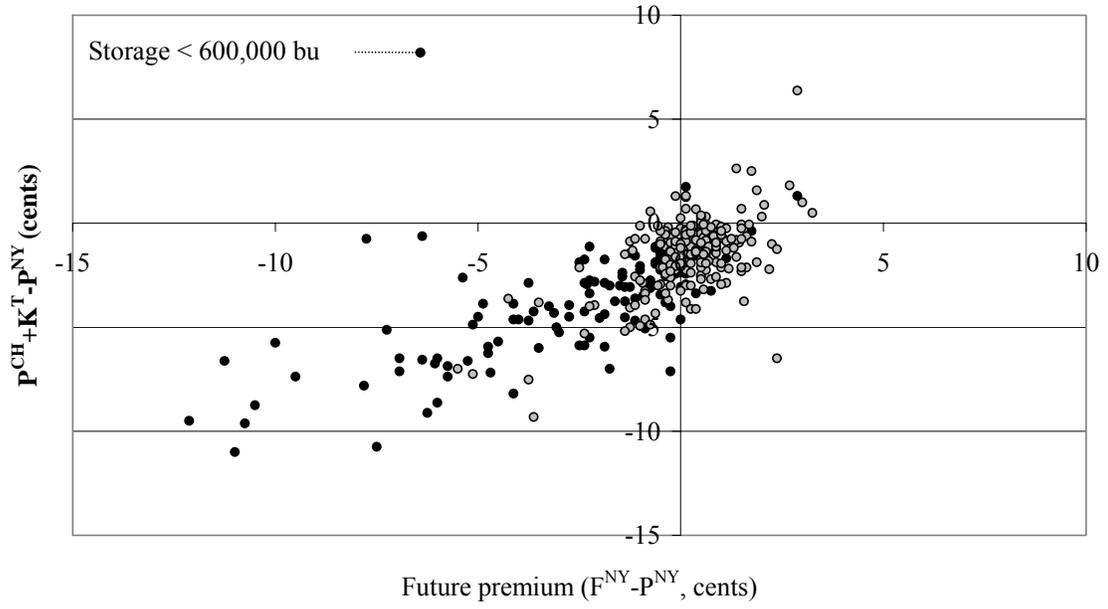


Figure 12

Future premium versus Spatial arbitrage-transport cost gap



Bibliography

- Bailey, Roy E. and Chambers, Marcus J. (1996) A Theory of Commodity Price Fluctuations. *Journal of Political Economy* 104 (5) 924 - 957
- Baumol, W. J. and H. D. Vinod (1970) An Inventory Theoretic Model of Freight Transport Demand. *Management Science* 16(7) 413-421
- Brennan, Michael J. (1958) The Supply of Storage *American Economic Review* 48(1) 50-72
- Chicago Board of Trade (1876 - 1892) *Annual Report of the Trade and Commerce of Chicago* (Chicago: JMW Jones Stationery and Printing Co.)
- Cournot, Augustin (1838) *Researches into the Mathematical Principles of the Theory of Wealth*. Translated by Nathaniel T Bacon, 1897. (New York: Augustus M. Kelly, Reprints of Economic Classics, 1960)
- The Commercial and Financial Chronicle* (New York: William B Dana & Co. (Weekly: 1882 - 1887; Volumes 35 - 43))
- Coleman, Andrew M.G. (1998) Storage, Arbitrage, and the Law of One Price (Princeton: Princeton University Thesis)
- Cournot, Augustin (1838(1960)) *Researches into the Mathematical Principles of the Theory of Wealth* [translated by Nathaniel T Bacon, 1897] (New York: Augustus M. Kelley, Reprints of Economic Classics)
- Deaton, Angus, and Guy Laroque (1992) On the Behavior of Commodity Prices *Review of Economic Studies* 59 1-23
- Deaton, Angus S. and Laroque, Guy (1995) Estimating a nonlinear rational expectations commodity price model with unobservable state variables *Journal of Applied Econometrics* 10 S9-S40
- Deaton, Angus, and Guy Laroque (1996) Competitive Storage and Commodity Price Dynamics *Journal of Political Economy* 104(5) 896-923
- de Jong, Gerard (2000) Value of Freight Travel-Time Savings pp 553 - 564 in Hensher, D.A. and Button, Kenneth J. (ed) *Handbook of Transport Modeling* (Oxford: Pergamon)
- Dumas, Bernard (1988) Pricing Physical Assets Internationally *NBER Working Paper* 2569
- Goldstein, Benjamin (1928) *Marketing: A Farmers' Problem* (New York: The MacMillian Company)
- Hansen, Bruce E. (1996) Inference When a Nuisance Parameter Is Not Identified Under the Null Hypothesis. *Econometrica* 64(2) 413-430
- Kaldor, Nicholas (1939) Speculation and Economic Stability *The Review of Economic Studies* 7(1) 1-27

- Lee, Guy A. (1937) The Historical Significance of the Chicago Grain Elevator System *Agricultural History* 11(1) 16-32
- McGinnis, Michael A. (1989) A Comparative Evaluation of Freight Transportation Models *Transportation Journal* 29(2) 36-46
- New York Produce Exchange (1876 - 1892) *Annual Report* (New York: Jones Printing Company)
- New York Times (New York, 1870 - 1892)
- Obstfeld, Maurice, and Alan M. Taylor (1997) Nonlinear Aspects of Goods-Market Arbitrage and Adjustment: Heckscher's Commodity Points Revisited. *Journal of Japanese and International Economics* 11 441-479
- Porter, Robert H. (1983) A Study of Cartel Stability: The Joint Executive Committee, 1880-1886 *Bell-Journal-of-Economics*. 14(2): 301-14
- Porter, Robert H. (1985) On the Incidence and Duration of Price Wars *Journal-of-Industrial-Economics*. 33(4): 415-26
- Samuelson, Paul A. (1952) Spatial Price Equilibrium and Linear Programming *American Economic Review* 283-303
- Spiller, Pablo T. and Wood, Robert O. (1988) "The Estimation of Transactions Costs in Arbitrage Models" *Journal of Econometrics* 39 309-326
- Tunell, George (1897) The Diversion of the Flour and Grain Traffic from the Great Lakes to the Railroads. *Journal of Political Economy* 5 339-375
- Tyworth, John (1991) The Inventory Theoretic Approach in Transportation Selection Models: A Critical Review. *Logistics and Transportation Review* 27(4) 299-318
- Ulen, Thomas S. (1982) The Regulation of Grain Warehousing and its Economic Effects: The Competitive Position of Chicago in the 1870s and 1880s. *Agricultural History* 56(1) 194-210
- United States Bureau of Statistics (1877) *First Annual Report on the Internal Commerce of the United States, by Joseph Nimmo, Jr. Part 2 of the Annual Report of the Chief of the Bureau of Statistics on the Commerce and navigation of the United States for the Fiscal Year ending June 30 1876*. (Washington: Government Printing Office)
- United States Bureau of Statistics (Treasury Department) (1881) *Report on the Internal Commerce of the United States, 1881, by Joseph Nimmo*. (Washington: Government Printing Office)
- United States Bureau of Statistics (Treasury Department) (1884) *Report on the Internal Commerce of the United States, 1881-1882, by Joseph Nimmo*. (Washington: Government Printing Office)
- United States 52nd Congress 2nd Session (1893) *Senate Report 1394: Wholesale Prices, Wages, and Transportation. Report by Mr Aldrich from the Committee on Finance March 3 1893 Part 1*. (Washington: Government Printing Office)

United States Treasury (1898) *Letter from the Secretary of the Treasury transmitting a report made to the Bureau of Statistics by Mr George G. Tunell of Chicago on Lake Commerce.* (Washington: Government Printing Office)

Williams, John Burr (1936) Speculation and the Carryover. *Quarterly Journal of Economics* 50(3) 436-455

Williams, Jeffrey (1986) *The Economic Function of Futures Markets* (Cambridge: Cambridge University Press)

Williams, Jeffrey C. and Wright, Brian D. (1991) *Storage and Commodity Markets* (Cambridge et al: Cambridge University Press)

Working, Holbrook (1929) The Post Harvest Depression of Wheat Prices. *Wheat Studies of the Food Research Institute* 6(1) 1-40

Working, Holbrook (1949) The Theory of Price of Storage *American Economic Review* 39(6) 1254-1262

Wright, Brian D and Jeffrey C. Williams (1989) A Theory of Negative Prices for Storage *Journal of Futures Markets* 9(1) 1-13