Integration in the English wheat market 1770-1820

Liam Brunt[1] & Edmund Cannon[2]

August 2012

Abstract

Cointegration analysis has been used widely to quantify market integration through price arbitrage. We show that price variations are caused in this framework by: (i) the magnitude of price shocks; (ii) the correlation of price shocks; (iii) within-period arbitrage; (iv) between-period arbitrage. We show formally that identification of each component depends upon sampling frequency. We measure the variation of these components across time and space using English weekly wheat price data, 1770-1820. We show that conclusions about arbitrage are sensitive to the precise form of cointegration model used; that the different components behave differently; and that different explanations – in terms of transport and information – are needed to explain the components. Previous analyses need to be interpreted with caution.

[1] Department of Economics, NHH – Norwegian School of Economics, Helleveien 30, 5045 Bergen, Norway. liam.brunt@nhh.no

[2] School of EFM, University of Bristol, 8 Woodland Road, Bristol BS8 1TN, England. edmund.cannon@bristol.ac.uk

1. Introduction

The last decade has seen many papers analysing market integration using data sets for many countries and time periods and a variety of econometric methods: Federico (2012) provides a comprehensive survey. The simplest way to analyse integration is to look at a measure of contemporaneous price dispersion (typically the coefficient of variation) and then see how that changes over time (see, for example, Jacks, 2011; he uses the same data as we use in this paper). Alternatively one can use more sophisticated econometric techniques to analyse the prices jointly (i.e. a vector-auto-regressive approach such as that of Ejrnæs, Persson and Rich, 2008, or Studer, 2008; these build on the seminal work of Ravallion, 1986).

Federico (2012) notes that there are two requirements for distinct markets to be integrated: (i) the long-run equilibrium should have similar prices in the different markets (the Law of One Price, or "LOOP"); and (ii) in the short run price differences should correct relatively quickly if the equilibrium is disturbed (which Federico refers to as "efficiency"). Econometric time-series models using cointegration techniques potentially have the ability to confirm that LOOP characterises prices in the long run and to estimate the speed with which price differentials are arbitraged away.

Phillips (1991) shows that estimating and testing long-run relationships – which is LOOP for our purposes – is typically unaffected by data problems, so long as there is a sufficiently long span of data. Since many economic historical studies use very long spans of data, this is generally not a problematic aspect of the literature. Therefore most of our analysis here is devoted to the second issue – namely the speed of adjustment towards equilibrium – and we take LOOP as given.

In our analysis we make two methodological points. First, we quantify the significance of using high-frequency (weekly) data to estimate the dynamics of price movements. Although it is well understood that higher frequency data are better, we have seen no attempt to put a number on just how much empirical results depend upon the frequency of observation.[3] Since we have weekly data, we are able to say how our results would look if we used weekly, monthly or annual data.

Second, much of the literature measures market efficiency using the "half life" (that is, the time for half of a price disequilibrium to be corrected). It does this using a Vector Error Correction Model (VECM) with no additional lags. We show that this model typically does not fit the data and hence conclusions about market efficiency using this model are misleading. In particular, it is not possible to construct a simple estimate of the half life and it is better to estimate the full impulse response function.

Stepping back from these econometric concerns, we turn to the determinants of market efficiency during the Revolutionary and Napoleonic Wars. Throughout this period market integration would be expected to have improved due to better canals and road transport (Bogart, 2005). The implicit model in much of the market integration literature is that prices are hit by random shocks and then arbitrage, dependent upon transport links, results in prices moving back to equilibrium. This seems to ignore another important ingredient – namely the role of information transmission. If information about changes in demand and supply reaches all markets at roughly the same time, then prices need never move out of equilibrium at all. This is true even if transportation of goods is to take place in the future, since, unless credit markets are completely dysfunctional, agents can calculate that such arbitrage will occur and immediately alter their prices accordingly.

To test the importance of information transmission, we estimate models of market efficiency with both transport and communication variables as possible explanators. We show that market efficiency increased in England during the period 1770-1820, even during the Revolutionary and Napoleonic Wars. During that conflict the magnitude of shocks to the grain market increased, but there was an underlying improvement in

---

[3] Taylor (2001) shows that parameter estimates will be inconsistent if the data consist of prices averaged within observation periods (for example, if we use monthly data consisting of the average of daily prices within the month); this phenomenon is called *temporal aggregation* and direction of bias means that the efficiency of the market (i.e. estimated speed of convergence to equilibrium) will be under-estimated. The data that we shall use in this paper are weekly data and in many instances markets only traded once or twice per week, so temporal aggregation is not our primary concern.

market efficiency due to improved roads, canal building and increased newspaper circulation.

Ironically, none of the transport or communication variables were responsible for shortening the measure most widely used in the economic history literature – namely, the half life. Half lives did shorten throughout the period 1770-1820, and even fell during just the period of the Revolutionary and Napoleonic Wars. However, we are unable to find any variables that correlate with our half life estimates. Instead we find that, conditional on the general state of the economy, improved transport and communication resulted in *smaller shocks to prices* and that prices *moving more closely together*. These are the routes by which better transport and communication generated arbitrage.

The rest of our paper is organised as follows. In section 2 we describe our data quite carefully and summarise it with a series of measures that have been used elsewhere in the literature. This approach suggests that very little changed in the economy, but that is due to a large number of effects which have a tendency to cancel each other out (better transport increased market efficiency, while the Revolutionary and Napoleonic Wars increased market turbulence). In section 3 we describe the econometric issues formally. Section 4 presents examples of our econometric procedures together with our estimates of half lives and other measures of market efficiency. In section 5 we correlate our measures of market efficiency with transport and communication variables. Section 6 concludes.

2. Wheat Prices 1770-1820

In this section we provide an overview of the English wheat market in the period 1770 to 1820: our data are described in detail in Brunt and Cannon (2012). We have weekly data for the average price within each county in England,[4] where the weekly average is calculated from prices from known towns within a county. There is good reason for confining our attention to wheat prices. First, the markets for oats, beans and peas have more missing observations. Second, the quantities traded between 1770 and 1820 were probably relatively unimportant (data on the quantities traded are not available for this

---

[4] Data for London are only available until 1794 and we do not use them because the London market is not a representative area. Monmouthshire was treated as an English county at this time.

earlier period, so we cannot test this hypothesis directly, but we know that the quantities were relatively small from 1820 onwards). Third, there are even more problems with barley prices: although barley sales were relatively large, they were concentrated in a relatively small part of the year (September-November) and the market for the rest of the year is so thin that the prices are unlikely to be informative.

This leaves wheat as our focus of analysis. Given data from 1818 and later periods, we know that this grain was traded steadily throughout the year, as would be expected from the crop that was fairly easy to store and which provided the main foodstuff in the UK at this time (Petersen, 1995). There are also relatively few missing observations: out of 2604 weekly observations many counties are missing only a few data points (the worst county, Hereford, lacks just 61).

Figure 1 about here (wheat price)

Figure 1 illustrates the movement of grain prices over the whole period, plotting the minimum and maximum price in each week. The range of prices in each week is large – on average about 33 pence, or a third of the price. An alternative measure of the range of prices observed at any point in time is the standard deviation of logged prices

$$(1) \qquad \sigma(p)_t = \sqrt{N^{-1}\sum_i \left(\ln P_{it} - \overline{\ln P_t}\right)^2} = \sqrt{N^{-1}\sum_i \left(p_{it} - \overline{p}_t\right)^2}$$

where we use the convention that lower case p_{it} refers to the log-price in county i at time t. With 2604 time-series observations, we can calculate this statistic 2604 times: we plot the annual averages in Figure 2, using the year October-September as the approximate harvest year. As a robustness check we also calculate the within-year average price for each county and then the standard deviations of these annual averages; as can be seen from Figure 2, this alternative measure is similar to the first.[5] The standard deviation in any week is typically about 0.08, which we can interpret by saying that the standard deviation of prices was consistently about 8 per cent of the price.

[5] There is no significant seasonal pattern in the standard deviation, so most of the difference between the two measures in the graph appears to be due to Jensen's inequality.

Figure 2 about here (dispersion of prices between counties)

From both Figures 1 and 2, and the related calculations, there appears to be no systematic change over the fifty year period. That is, the range of prices does not trend down (which we might expect to be a consequence of greater transport links) and there is no obvious increase during the Napoleonic Wars. So, if these two effects were important, then they must have cancelled each other on this measure.[6]

Of course the fact that the standard deviation of prices is fairly constant tells us relatively little about how the prices were interacting with each other. Simple correlations of the price series are uninformative because there is considerable variation in prices and prices move sufficiently closely together that the underlying trend will dwarf any other effects (typical correlations are 0.98-0.99).

A more interesting question is to ask how relative prices changed over time. To do this we take annual cross sections of the average within-harvest-year prices at the beginning of the harvest year and calculate the correlation with the corresponding prices at the beginning of the following year:

$$\text{corr}\left(p_{i.y}, p_{i.y+1}\right) = \frac{\sum_i \left(p_{i.y} - \overline{p}_y\right)\left(p_{i.y+1} - \overline{p}_{y+1}\right)}{\sqrt{\sum_i \left(p_{i.y+1} - \overline{p}_{y+1}\right)^2 \sum_i \left(p_{i.y} - \overline{p}_y\right)^2}} \tag{2}$$

If the pattern of relative prices in the different counties were to stay the same, then we should expect this statistic to be high. Figure 3 shows the value of the statistic for consecutive pairs of years over the whole period: given the sample size, these correlations are statistically significantly positive if bigger than 0.29. The only years when the pattern of prices changed much from the previous year are 1772, 1779, 1800 and 1808, suggesting that relative prices changed fairly slowly. In Figure 4 we plot all of the correlations $corr\left(p_{i.y}, p_{i.y+j}\right)$: where the correlation is high this is shaded red and where it is negative it is shaded blue.

---

[6] There is no evidence that price dispersion was different in the Napoleonic Wars: the first series averages 8.0 per cent up to 1792 and then 7.9 per cent for 1793-1815, while the second falls from 6.9 per cent to 6.6 per cent (neither of these changes are statistically significant).

Figure 3 about here (year-on-year correlations of cross sectional prices)

Figure 4 about here (correlations of cross sectional prices)

The pattern of relative prices is remarkably stable throughout the whole period: for example, the correlation of prices in 1788 with 1813 is 0.78. The bottom right-hand corner of the graph contains the correlations for years which are furthest apart and even here the correlations tend to be high.

Our final summary characterisation of the data is to see how price differences depend upon proximity of counties. There are a variety of spatial descriptive statistics and here we confine ourselves to Moran's *I* statistic[7] calculated in each week *t*:

(3) $$I_t = \frac{N\sum_j \sum_i a_{ij}\left(p_{jt} - \overline{p}_{jt}\right)\left(p_{it} - \overline{p}_{it}\right)}{\left(\sum_j \sum_i a_{ij}\right)\left(\sum_k \left(p_{kt} - \overline{p}_{kt}\right)^2\right)}$$

where a_{ij} is an indicator variable taking the value one if counties *i* and *j* are adjacent and is zero otherwise. Under the null hypothesis of no spatial correlation, the expected value of this statistic is –0.024: our calculated *I* statistics average 0.41, typically in the range from 0.2 to 0.6, and are almost invariably statistically significant with (standard Normal) Z-statistics averaging 4.34. Yet again, there appears to be no systematic variation over time (there is no seasonal pattern in the Moran statistic: a regression on seasonal dummies yields a test statistic of F(52, 2550) = 0.72 [p = 0.94]).

Figure 5 about here (Moran's I)

We summarise our analysis so far by noting that, although prices were very different in the various counties for which we have data, these prices all moved closely together over the entire period. Moreover, from an analysis of summary statistics, their behaviour does not appear to have changed much over the period 1771-1820. There was a high degree of spatial correlation, which did not change much, and the relative prices in different counties was also roughly the same at the beginning of the period as it was at the end. Very similar results can be obtained whether using end-of-year prices or within-year averages, so the frequency of measurement is not a major determinant of our conclusions.

---

[7] More sophisticated measures of spatial correlation would yield similar conclusions.

This evidence provides only the most minor support for the idea that improved transport significantly affected grain prices. The most important prediction of any model of falling transport costs would be some form of convergence, some change in relative prices or some change in the relevance of distance. From our analysis of summary statistics in this section, none of these things happened. Therefore we can only conclude that we need to model price behaviour of the individual series much more closely. In the following section we consider a framework for discriminating between different determinants of price movements.

3. Cointegrated prices: explanation and example

It is obvious from the previous section that individual price series show both large and persistent variation over time. This is true not just for our data but for many other data sets. In fact it is almost universally the case that one cannot reject the null hypothesis that any given price series has a unit root, meaning that standard statistical theory will not apply to some estimation and testing procedures. It is also common for price series to move closely together, so that the difference in – or the ratio of – two price series is much less variable and more persistent. This suggests that there is a simple equilibrium relationship between the two price series; if the difference or ratio of prices does not have a unit root then the series are cointegrated. A good introduction to this approach is provided in the appendix by Ejrnæs in Persson (1999) and we shall build on that analysis here.

We start our analysis with a fairly general cointegrating model, which illustrates the main points of this paper and can encompass many, but not all, of the other issues which may be relevant. Our exposition is confined to a situation with just two markets, i and j, and we write the logarithm of prices in period t as p_t. Then our starting point is the Data Generating Process (DGP):

(4)
$$\begin{bmatrix}\Delta p_t^i \\ \Delta p_t^j\end{bmatrix} = \begin{bmatrix}\alpha_i \\ \alpha_j\end{bmatrix}\left(p_{t-1}^i - p_{t-1}^j\right) + \begin{bmatrix}\mu_i \\ \mu_j\end{bmatrix} + \begin{bmatrix}\pi_{i.i}^{(1)} & \pi_{i.j}^{(1)} \\ \pi_{j.i}^{(1)} & \pi_{j.j}^{(1)}\end{bmatrix}\begin{bmatrix}\Delta p_{t-1}^i \\ \Delta p_{t-1}^j\end{bmatrix} + \cdots + \begin{bmatrix}\pi_{i.i}^{(K)} & \pi_{i.j}^{(K)} \\ \pi_{j.i}^{(K)} & \pi_{j.j}^{(K)}\end{bmatrix}\begin{bmatrix}\Delta p_{t-K}^i \\ \Delta p_{t-K}^j\end{bmatrix} + \begin{bmatrix}\varepsilon_t^i \\ \varepsilon_t^j\end{bmatrix};$$
$$\begin{bmatrix}\varepsilon_t^i \\ \varepsilon_t^j\end{bmatrix} \sim N\left(\begin{bmatrix}0 \\ 0\end{bmatrix}, \begin{bmatrix}\omega_{i.i} & \omega_{i.j} \\ \omega_{i.j} & \omega_{j.j}\end{bmatrix}\right); \qquad \alpha_i \leq 0; \quad \alpha_j \geq 0; \quad \alpha_j - \alpha_i < 1$$

where $\Delta p_t^i \equiv p_t^i - p_{t-1}^i$ which we refer to as the price change.[8] The cointegration equation can be written more compactly in vector notation as[9]

$$\text{(5)} \qquad \Delta \mathbf{p}_t = \boldsymbol{\alpha\gamma}\mathbf{p}_{t-1} + \boldsymbol{\mu} + \sum_{k=1}^{K} \boldsymbol{\pi}^{(k)} \Delta \mathbf{p}_{t-k} + \boldsymbol{\varepsilon}_t; \quad \boldsymbol{\varepsilon}_t \sim N\left(\mathbf{0}, \boldsymbol{\Omega}\right); \quad \boldsymbol{\gamma} \equiv \begin{bmatrix} 1 & -1 \end{bmatrix}$$

One potential objection to this model is that it requires prices to be cointegrated regardless of how far apart the prices actually are. One view of the market efficiency is that prices will only move back towards equilibrium when they are further apart than a set of commodity points (perhaps due to transport costs); in such a case, it would be appropriate to use a Threshold ECM. While we are not unsympathetic to this view, it must be remembered that our prices are not market prices but averages within a geographical area: even if the average price gap between two markets is less than the transport costs, individual markets' prices may have been further apart and so prices would still have been adjusting. For this reason we confine ourselves to the VECM model.

If the lagged dependent variables are not needed then the DGP can be re-written:[10]

---

[8] To avoid ambiguity, we avoid using the phrase *price difference*, which might imply either the change in price $\Delta p_t^i \equiv p_t^i - p_{t-1}^i$ or the market disequilibrium $\boldsymbol{\gamma}\mathbf{p}_t \equiv p_t^i - p_t^j$.

[9] In our specification we follow Marks (2010) and place no restriction on the constant terms in the vector $\boldsymbol{\mu}$. In some published research the constant is restricted. In some studies, it is constrained to lie in the cointegrating space (for example, in equation 5A.4 in Ejrnæs' appendix in Persson, 1999); in others, it appears to be omitted altogether (such as in Buyst, Dercon and Van Campenhout, 2006). In calculations that we do not report here, restricting the constant has negligible impact on estimates of the half life; but omitting the constant leads to it being biased up by a factor of as much as two.

[10] If it is correct that the variables are cointegrated (with LOOP), then the variable $\left(p_t^i - p_t^j - \lambda\right)$ is stationary and it is not necessary to use the Dickey-Fuller distribution to test the parameter in this regression: we are not sure why Studer (2008, p. 409) tests this using the DF test rather than a conventional test.

(6) $$\left(\boldsymbol{\gamma}\mathbf{p}_t - \lambda\right) = \left(\boldsymbol{\gamma\alpha} + 1\right)\left(\boldsymbol{\gamma}\mathbf{p}_{t-1} - \lambda\right) + \boldsymbol{\gamma\varepsilon}_t \qquad \lambda \equiv \frac{\mu_i - \mu_j}{1 - \left(\alpha_i + \alpha_j\right)}$$

so that the difference between the two prices is a first-order auto-regressive process; the effect of a shock dies away geometrically; and it is possible to describe this decay by a single statistic. The time taken for half of the magnitude of a price difference to die away is referred to as the *half life*, defined as

(7) $$HL \equiv \frac{\ln\left(0.5\right)}{\ln\left(1 + \alpha_i - \alpha_j\right)} = \frac{\ln\left(0.5\right)}{\ln\left(1 + \boldsymbol{\gamma\alpha}\right)}$$

Most authors simply substitute $\widehat{\alpha_i - \alpha_j}$ into this formula to estimate the half life. In principle this is wrong: the half life is an increasing function of $\alpha_i - \alpha_j$: when this quantity is less than about -0.57 the function is concave; thereafter it is convex. This suggests that the expected value of the half life will not be the same as the half life evaluated at the expected value of the parameters. However, in nearly all cases the standard error of $\hat{\alpha}_i - \hat{\alpha}_j$ is sufficiently small that it makes no difference (we discuss details of this in the appendix).

If we re-arrange equation (4) we can obtain the squared price difference: this has the advantage that it looks at the absolute distance between prices and so it does not matter which price is higher[11]

(8) $$\mathsf{E}\left[\left(p_t^i - p_t^j\right)^2\right] = \underbrace{\frac{\exp\left\{-2 \times \text{Half life}\right\}}{4}}_{\text{a function of the Half life}} \mathsf{E}\left[\left(p_{t-1}^i - p_{t-1}^j\right)^2\right] + \underbrace{\left(\mu_i - \mu_j\right)^2}_{\substack{\text{constant}\\ \text{terms}}} + \underbrace{\omega_{i.i} + \omega_{j.j}}_{\substack{\text{variance of}\\ \text{the disturbances}}} - \underbrace{2\omega_{i.j}}_{\substack{\text{covariance of}\\ \text{the disturbances}}}$$

This allows us to decompose the squared price difference into four components, all of which are intuitive:

First, obviously, the average difference in prices, determined by the constant term $\boldsymbol{\gamma\mu\mu}'\boldsymbol{\gamma}'$ as discussed above. Subtracting this term from the RHS of the equation leaves us with the variance of the price difference;

[11] Since $HL = \ln\left(0.5\right)/\ln\left(1 + \boldsymbol{\gamma\alpha}\right)$ it follows that $1 + \boldsymbol{\gamma\alpha} = 0.5\exp\left\{-HL\right\}$.

Secondly, the variances of the disturbances $\omega_{i.i} + \omega_{j.j}$ which cause prices to change: larger disturbances result in greater dispersion of prices;

Thirdly, the covariance of the disturbances $\omega_{i.j}$, which reduces the dispersion in prices. There are two ways of interpreting this variable: it is a measure of how much the markets respond instantaneously in the same way (which is a measure of integration); it is also the adjustment of prices which takes place within the period of observation (i.e. within-week price adjustment in our data). It is not possible to identify these two different effects.

Fourthly, the speed of adjustment (which Federico refers to as efficiency): the shorter the half life, the less the price dispersion this period is due to prices being out of equilibrium in the previous period.

The first component is really a measure of LOOP, rather than market efficiency. From our data set we are able to obtain estimates of the last three of these components of price efficiency. The advantage of this is that we are able to ask, not just what correlates with reduced or increased price dispersion, but also the *mechanism* through which changes determine the price dispersion. For example, does improved transportation reduce the variance of domestic price shocks (by allowing imports to flood in to the domestic market more cheaply or quickly, and therefore offset changes in the domestic harvest)? Or does improved transportation increase the covariance of shocks (by making linking regional markets more strongly to the London market)? Or does improved transport increase the speed of adjustment of one market to another? Before we analyse these effects in more detail, however, we address some issues of estimation.

4. Cointegration estimation

In this section we estimate the VECM models for both the whole period and for sub-periods to allow for parameter instability. To clarify our procedures we illustrate many of our models using prices from Bedfordshire and Buckinghamshire for 1770-1820. Alphabetically, these counties are the first adjacent-county pair our data set (i.e. we chose them randomly, not because their prices series have any special characteristics). Geographically, they are large, adjacent counties; both of them are agricultural and both have almost complete data (just one missing observation each). Using data of different frequencies, we estimate the following models (regardless of the data frequency, the half lives are measured in weeks)[12]:

(9)

$$\text{annual data} \quad \widehat{\begin{bmatrix} \Delta p_t^i \\ \Delta p_t^j \end{bmatrix}} = \begin{bmatrix} -0.603 \\ 0.272 \end{bmatrix} \left(p_{t-1}^i - p_{t-1}^j \right) + \begin{bmatrix} -0.010 \\ 0.015 \end{bmatrix} \quad \begin{matrix} HL = 17.4 \\ EHL = 19.0 \end{matrix}$$

$$\text{monthly data} \quad \widehat{\begin{bmatrix} \Delta p_t^i \\ \Delta p_t^j \end{bmatrix}} = \begin{bmatrix} -0.256 \\ 0.249 \end{bmatrix} \left(p_{t-1}^i - p_{t-1}^j \right) + \begin{bmatrix} -0.006 \\ 0.009 \end{bmatrix} \quad \begin{matrix} HL = 4.3 \\ EHL = 4.3 \end{matrix}$$

$$\text{weekly data} \quad \widehat{\begin{bmatrix} \Delta p_t^i \\ \Delta p_t^j \end{bmatrix}} = \begin{bmatrix} -0.127 \\ 0.133 \end{bmatrix} \left(p_{t-1}^i - p_{t-1}^j \right) + \begin{bmatrix} -0.004 \\ 0.004 \end{bmatrix} \quad \begin{matrix} HL = 2.3 \\ EHL = 2.3 \end{matrix}$$

Including seasonal dummies for the monthly and weekly data makes no quantitative difference. It is notable that using higher frequency data results in much shorter estimated half lives: using weekly data, rather than monthly, results in a half life of two weeks, rather than four.[13]

Figure 6 and Table 1 about here

The results from Bedfordshire and Buckinghamshire are fairly representative of other pairs of counties. We estimate the half life for each county pair and illustrate our results in Figure 6 and Table 1. On average, the half life estimated from data measured at an annual frequency is twenty weeks, skewed heavily to the right; whereas the half life estimated using weekly data is about eight weeks, with much less skew. The second panel of Figure 6 uses half lives estimated from weekly data to compare the distribution of half lives for all counties and for adjacent counties. Markets appear more efficient

---

[12] One problem with constructing monthly data is that our data are "week-ending" data: sometimes there are five observations in a month and sometimes only four. So the monthly data are not equally spaced. We also estimated a model using every fourth observation so that the data were approximately monthly but equally spaced: the results were almost identical to those from monthly data.

[13] This is not due to temporal aggregation (i.e. averaging observations within period, as discussed in Taylor, 2001): the annual data are taken from the first observation in October and the monthly data from the first observation of the month.

(have a shorter half life) when counties are adjacent: the average half life is only four weeks, instead of eight.

From this we conclude that it is possible to get dramatically different estimates of the half life by using data of different frequencies. Although there is probably parameter instability in all of the models, this is unlikely to explain the differences in half life estimates entirely, since all three regressions are based on the same span of data (i.e. 1770-1820) and would suffer the same instability. The more important problem is that the weekly models display significant serial correlation, suggesting that the VAR of equations (4) and (5) does not fit the data if we impose the restriction $\boldsymbol{\pi}^{(k)}$. All of the models in (9) have this restriction; all have biased parameter estimates and the degree of bias depends upon the frequency of the data used in estimation.

An important consequence of this is that we cannot compare half lives from models estimated on data of different frequencies. So research based on annual data (Studer, 2008), is not comparable to research using monthly data (Bateman, 2011; Buyst, Dercon and Van Campenhout, 2006; Goodwin and Grennes, 1998; Goodwin, Grennes and Craig, 2002; Jacks, 2005; Marks, 2010; Trenkler and Wolf, 2005), which is not comparable to research using data with two observations per month (Ejrnæs and Persson, 2000), which is not comparable to research using weekly data (Ejrnæs and Persson, 2010; Federico, 2007; Hynes, Jacks and O'Rourke, 2012).

Interestingly, estimated half lives from annual data appear slightly longer for adjacent counties than for all counties, suggesting that attempts to correlate market efficiency with distance may be ineffective or misleading when data are measured at low frequencies. If this result could be generalised then it might explain why Studer (2008, Table 5) finds only weak or ambiguous correlation between market efficiency and distance.[14]

The presence of serial correlation suggests that more lags are needed in the VAR, but this complicates half life estimation. It is convenient to re-write the DGP as

[14] For example, Studer's average estimate of $\alpha_i - \alpha_j$ for 1870-1914 is −0.46 when the distance is 150-300 km (a half life of 1.12 years); when the distance is 600-1000 km, the adjustment is −0.60 (a half life of 0.76 years), which is considerably faster. But Studer is using annual averages (p. 396), so the half lives are all biased up (Taylor, 2001).

$$\mathbf{p}_t = \mathbf{\Phi}_1 \mathbf{p}_{t-1} + \mathbf{\Phi}_2 \mathbf{p}_{t-2} + \cdots + \mathbf{\Phi}_{K+1} \mathbf{p}_{t-K-1} + \boldsymbol{\mu} + \boldsymbol{\varepsilon}_t$$

(10)

$$\mathbf{\Phi}_1 \equiv \left(\mathbf{I} + \boldsymbol{\alpha\gamma} + \boldsymbol{\pi}_1\right); \quad \mathbf{\Phi}_2 \equiv \left(\boldsymbol{\pi}_2 - \boldsymbol{\pi}_1\right); \ldots \quad \mathbf{\Phi}_K \equiv \left(\boldsymbol{\pi}_K - \boldsymbol{\pi}_{K-1}\right); \quad \mathbf{\Phi}_{K+1} \equiv -\boldsymbol{\pi}_K$$

Pesaran and Shin (1996) show that the speed of adjustment towards equilibrium is the same regardless of which random shock causes the initial price difference. Defining

(11)

$$\mathbf{B}_s \equiv \mathbf{0} \quad \text{for} \quad s < 0$$

$$\mathbf{B}_0 \equiv \mathbf{I}$$

$$\mathbf{B}_s \equiv \sum_{k=1}^{K} \Phi_k \mathbf{B}_{s-k} \quad \text{for} \quad s > 0$$

we can calculate the squared effect of a shock in period zero on the error correction (price difference) in period *s* to be $\boldsymbol{\gamma}'\mathbf{B}_s\boldsymbol{\Omega}\mathbf{B}_s'\boldsymbol{\gamma}$, which can be normalised by adjusting for the variance of the original shock to obtain the impulse response function

(12)

$$i(s) = \sqrt{\frac{\gamma \mathbf{B}_s \Omega \mathbf{B}_s' \gamma'}{\gamma \Omega \gamma'}}$$

Where the model is includes no lagged dependent variables (e.g. Buyst, Dercon and Van Campenhout, 2006; Hynes, Jacks and O'Rourke, 2012; Studer, 2008) equation (12) reduces to the same geometric decay implied by equation (7), although this will still be a misleading estimate of market efficiency if the omission is inappropriate. Several authors include additional lags in the VAR (Persson, 1999, ch.5; Bateman, 2011, whose procedure is explained in Bateman, 2007; Marks, 2010); but they do not plot the full impulse response function and appear to measure market efficiency using the loadings alone, despite the fact that this gives no meaningful description of the response of prices to market disequilibrium. Trenkler and Wolf (2005) estimate a VAR with more lags, but then re-estimate the model with just one lag to get a half life. The only papers that we have found that illustrate the full dynamic response are Goodwin and Grennes (1998) and Goodwin, Grennes and Craig (2002).

To see the effect of this, we start by returning to the Bedfordshire and Buckinghamshire data, using weekly data. Two sample models that we estimated are (omitting the constant and seasonal dummies):

(13)

$$\widehat{\begin{bmatrix} \Delta p_t^i \\ \Delta p_t^j \end{bmatrix}} = \begin{bmatrix} -0.095 \\ 0.099 \end{bmatrix} \left(p_{t-1}^i - p_{t-1}^j\right) + \begin{bmatrix} -0.005 & 0.257 \\ 0.211 & -0.034 \end{bmatrix} \begin{bmatrix} \Delta p_{t-1}^i \\ \Delta p_{t-1}^j \end{bmatrix}$$

$$(14)\quad \widehat{\begin{bmatrix}\Delta p_t^i \\ \Delta p_t^j\end{bmatrix}} = \begin{bmatrix}-0.082 \\ 0.081\end{bmatrix}\left(p_{t-1}^i - p_{t-1}^j\right) + \begin{bmatrix}-0.039 & 0.289 \\ 0.242 & -0.077\end{bmatrix}\begin{bmatrix}\Delta p_{t-1}^i \\ \Delta p_{t-1}^j\end{bmatrix} + \begin{bmatrix}-0.071 & 0.078 \\ 0.112 & -0.062\end{bmatrix}\begin{bmatrix}\Delta p_{t-2}^i \\ \Delta p_{t-2}^j\end{bmatrix}$$

If we look at these regressions in conjunction with the weekly regression in equation (9), it is clear that the loadings get smaller as more lags are included. If an attempt were made to estimate the half life just from the loadings from equations (13) and (14) and equation (7), regressions with more lags would suggest longer half lives, illustrated in the first row of Table 2. This table also contains results for more lags, going up to 53 weeks, to take account of any additional seasonal effects. The disadvantage of including so many lags is that the confidence intervals (not reported here) are much wider.[15]

Table 2 about here

However, in this instance the loadings now under-state market efficiency. Consider an hypothetical situation where, from market equilibrium, prices diverge due to a disturbance causing a rise in Buckinghamshire prices in period 1 while Bedfordshire prices are constant. From equation (13) the prices move towards each other in period 2, not just due to the error correction term, but also because – in that period – the Bedfordshire price rises by $0.257\Delta p_1^{Buck}$ and the Buckinghamshire price changes by $-0.034\Delta p_1^{Buck}$. Thus, in addition to the decrease in the disequilibrium of 0.194 from the error correction, there is an additional 0.291 from the effect of the lagged price changes. These effects are illustrated in the second row of Table 2.

Figure 7 about here

We illustrate this graphically in Figure 7, which shows impulse response functions for Bedfordshire and Buckinghamshire from models with differing lags of price changes: since the data are weekly, we consider up to 53 lags to allow for seasonal effects

---

[15] We do not address the issue of optimal lag length in this paper. Conventional criteria, such as information criteria, typically choose a compromise to maximise goodness of fit subject to minimising the number of explanatory variables. Such criteria sometimes choose models where there is still residual correlation. Which criterion is optimal is sensitive to the objectives of the research (so estimation, testing and forecasting might all yield different answers); for the purpose of this paper more lags will generally be better than fewer.

(although the model also includes seasonal dummy variables, which make little difference). The solid black line shows the model estimated with no lags and demonstrates geometric decay, i.e the disequilibrium decays at the same speed regardless of the length of time since the disturbance. So long as two or more lags are included the impulse response function is more-or-less the same: the shape is quasi-hyperbolic, with relatively fast decay for the first few weeks and thereafter relatively slow decay. There is no simple formula for calculating the half life and we do so by simple linear interpolation (where the graph cuts the horizontal line). Linear interpolation is a slightly crude method, but the half life from this method, 2.33 weeks, is very close to the value of 2.30 using equation (7). When more than two lags are included in the estimation, the half life ranges from 1.23 to 1.04.

The third and fourth rows of Table 2 show that the example cannot be generalised: if we look at all 780 county pairs then the mean average half life is actually *higher* when several lags are included (the same is true for the median). However, it is the case for adjacent-county pairs that including more lags results in shorter average half lives. Figure 8 shows the distributions of the half lives for all 780 county pairs for differing lags. Adding a few lags results in longer half lives and adding very large numbers of lags results in slightly shorter half lives.

Our models hitherto have all been estimated on the whole sample from 1770 to 1820. Obviously this is inappropriate if there is parameter instability, especially since the economic issue is potential changes in efficiency. To address this, we divide the data set into 4650 sets of weekly data for a given harvest year for each adjacent county pair: so, for example, one data set would be Bedfordshire and Buckinghamshire for the harvest year 1780-81.

The problem with this approach is that there is a strong seasonal pattern in prices that is variable over time, and we are unable to model seasonal effects when using data within a single year. However, Brunt and Cannon (2002) show that the seasonal pattern is approximately saw tooth: in about the 33rd week of the year, at harvest time, prices fall dramatically until about the 45th week: thereafter they rise approximately exponentially (so log prices rise linearly). From this stylised fact we use the forty observations from the 45th week of year t to the 33rd week of year $t+1$ and ignore seasonal effects (the stochastic trend is modelled through the constant term, which is not restricted to lie in the cointegrating space). We refer to the parameter estimates for this year as belonging to year t. Forty data points is a relatively small number of observations, and we lose observations because the need for lagged variables and due to missing data: where there are fewer than thirty observations we do not estimate parameters at all. For each of the data sets we estimate the root-mean-square price difference, the magnitude of the disturbances (measured as the standard deviation of the two disturbances), the relative size of the disturbances (measured as the ratio of the larger standard deviation to the smaller), the correlation between the disturbances and the half life. The averages of the statistics across all 93 adjacent-county pairs are illustrated in Figures 9-13.

Figures 9-13 about here.

Figure 9 shows the root mean square price difference between price pairs averaged across all adjacent-county pairs: this is the empirical measure of the left-hand-side variable of equation (8). Since this is just a description of price dispersion it is independent of any econometric modelling issues. The vertical axis is measured in percentages, so for the first part of the period this measure of price dispersion was about 4 per cent. From about 1793 onwards it rose, coinciding with the Revolutionary and Napoleonic Wars. This is consistent with Jacks (2011), but not consistent with Figure 2, which showed that the standard deviation of all prices did not rise. The apparent contradiction is resolved by the observation that relatively close markets became less integrated while overall dispersion of prices did not rise. This suggests that any effect of higher volatility from the Revolutionary and Napoleonic Wars was masked or even offset by other factors.

We now turn to our modelling results, which decompose the price variation. We do this along the lines of equation (8) but, since we allow for lagged price changes in the VECM model, this is only an approximation to the relationship between the different components.

Figure 10 confirms that the one of the major causes of the greater price dispersion illustrated in Figure 9 was that the disturbances were larger: the peaks in price disturbance in Figure 10 coincide with peaks in Figure 9, although the magnitudes are not necessarily the same. This is prima facie evidence that prices became more dispersed, not due to declining efficiency of the market, but due to the shocks hitting the economy. However, Figure 11 shows that over time the disturbances to markets became more correlated and this attenuates the effects of larger shocks on price dispersion. Higher correlation does not mean that the disturbances became more similar in size, and so we look also at the ratio of the magnitude of the shocks in Figure 12: this is the ratio of the standard deviation of the more variable disturbance to that of the less variable disturbance. If the Revolutionary and Napoleonic Wars resulted in more similar shocks (i.e. a source of additional shocks that was the same for all markets) then the variances of the shocks should have become more similar. Over the whole period, when the magnitude of shocks is high, they are both more correlated (correlation of 0.73) and the relative size falls (correlation of -0.46).[16]

Figures 10 - 12 report four separate estimates: one for each set of VECM models (with zero, one, two or three lags). Especially in Figures 10 and 11, there is little difference between our estimates, as long as we include at least one lag (although the estimates from models with zero lags have a small, but noticeable, differences).

The remaining measure of market efficiency is the half lives, illustrated in Figure 13. Unlike the previous three graphs, the estimated half lives vary considerably depending on the number of lags used in the VECM. In all cases, however, the half lives tend to fall over the period: regressing the average half life on a trend results in a coefficient of about -1 per cent: this is true even if the estimation is only for the period 1792-1815.[17] This suggests that market efficiency was improving throughout the period including the Revolutionary and Napoleonic Wars. The reason that this does not show up in measures of price dispersion is that the shocks to the economy were simultaneously increasing.

To summarise this section: estimates of market efficiency are highly sensitive to both the frequency of the data and the number of lags included in time series models. The differences are sufficiently large that they suggest comparison of research using different methods or frequencies is hazardous. We have shown informally how the dispersion in prices can be decomposed into the magnitude of the shocks, the correlation of the shocks and the speed to convergence (half life).

Our data confirm that the Revolutionary and Napoleonic Wars saw increased price dispersion, but we show that this was not due to less efficient markets. The evidence suggests that market efficiency continued to increase, even while the magnitude of

[16] Results are almost identical regardless of the number of lags in the VECM.

[17] This is only a crude calculation: using Newey-West standard errors to compensate for the obvious serial correlation suggests that the relationship is statistically significant at the 5% level when the half lives are calculated from VECMs with zero, two or three lags and at the 10% level for one lag.

shocks grew larger: the reason for greater price dispersion was that the latter predominated.

5. The effect of transport on market efficiency

In the previous section we showed that there was evidence that market efficiency improved, but that this failed to reduce dispersion in prices because the magnitude of the shocks hitting the economy were simultaneously increasing. This raises the question of whether we can find any effect of transport and communication variables on market efficiency.

Our procedure is similar to that of Jacks (2011), but for three differences. First, we only consider adjacent county pairs. This is mainly because transport – such as roads or canals – is only conceptually easy to measure for adjacent counties: where counties are not adjacent it is not obvious how they would be linked for arbitrage purposes.[18] We are also concerned about the statistical properties of using all 780 county pairs: since these are based on only 40 price series, they are not independent.

Secondly, Jacks looks at a single measure of price dispersion (albeit a slightly different one to us), but we look at three of the components of price dispersion.

Finally, we increase the number of controls by the using both year and county-pair fixed effects instead of the alternative of time-series variables (such as severity of war, measured by battle casualties). The reason for this is that we are primarily interested in the effect of transport variables (we take it as read that warfare disrupted markets) and so are content to use a relatively large set of control variables.

We use two transport variables. The first is a dummy variable indicating that the two counties were linked by a canal. The second is a measure of turnpiked roads in the two counties defined as

$$\text{Road}_{(i,j).t} \equiv \frac{M_{i.t} + M_{j.t}}{A_i + A_j} \qquad (15)$$

where $M_{i.t}$ is the mileage of turnpiked road in county *i* in year *t* and A_i is the area of the county in hundreds of square miles. At first it might seem strange to measure road

[18] One possibility is coastal traffic. To the extent that this is constant, it is modelled in the fixed effect.

linkages by the average road density, since transport links are typically thought of as *between* two markets. But recall that our data are average prices within counties and therefore for we would only expect one county's average price to converge to the other's if all markets were connected within the two counties. Given the price and road data that are available, this measure seems appropriate.

The final variable that we use is a measure of communication. The timing of price changes is determined less by the existence of transport than by the arrival of news. We use data on newspaper circulation in the towns from which our wheat prices were collected. Underlying data on newspaper circulation were taken from Gibson (1991) and from this we calculated the proportion of towns in adjacent-county pairs that had at least one newspaper.[19]

Table 3 about here (regression results)

Our first set of regression results are reported in Table 3. The first column reports the regression for price dispersion, which we know to have increased during the Napoleonic Wars. Given the huge variability in price dispersion – and the fact that road and canal and newspaper networks only evolved relatively slowly – it would be unsurprising if none of the variables were statistically significant. However, the Canal indicator is statistically significant at conventional levels and it suggests that a canal reduced the root-mean squared price difference by one-quarter of one per cent. The effect for newspapers is statistically significant at the ten per cent level, but the effect is relatively small: in a county pair with a total of ten towns, the presence of one additional newspaper in a town previously without a newspaper would reduce the root-mean-squared price difference by only 0.7 per cent.

The remaining four columns of Table 3 report the regression results for the components of market integration (based upon estimates of the VECM with no lags). Both Roads and Canals appear to reduce the variance of the shocks: to make the coefficients easier to interpret we use the standard deviation rather than the variance. An extra ten miles of turnpike per hundred square miles would reduce the standard deviation of the shocks by about one-third of one per cent, while the presence of a canal would reduce the standard deviation by about one-sixth of a per cent. Theoretically, the effect of

---

[19] As a robustness check we also considered the average number of newspapers in circulation: the results were quantitatively very similar.

transport on the variance of price changes is ambiguous (depending on elasticities of supply and demand); but it appears in this instance that the lower transport costs allowed risk-sharing through pooling of risks in separate locations. The effect is large, as evidenced by the R-squared.

Our other two measures of the shocks are the ratio of the magnitudes and the correlation. Roads and Newspapers appear to reduce the ratio of the variance of the shocks: in other words, if a shock hits one market then the size of the shock hitting the other market is more likely to be the same size. This is prima facie evidence that both Roads and Newspapers increase market efficiency, as the disturbances in the two markets have a more similar magnitude. Surprisingly the Roads variable has only a minimal effect on the correlation of the disturbances, suggesting that it does not increase efficiency, but Canals and Newspapers do.

Both Roads and Canals have a positive and statistically significant effect on the half life, which suggests that they reduce market efficiency. The effects appear to be large: an extra ten miles of road per hundred square miles apparently increases the half life by two-thirds of a week (i.e. four to five days) and a canal by one-third (two days). These results are counter-intuitive.

Our analysis so far is based on measures of market integration from models where there are no lagged price changes in the VECMs. Table 4 provides comparative analysis for components estimated from VECMs with different numbers of lags (the first row reproduces the information in Table 3).

Table 4 about here (comparative regression results)

We have already seen from Figures 9-12 that our estimates of most components of market integration are similar regardless of the number of lags used in the VECM estimation. So it is unsurprising that the first four panels of Table 4 provide very similar conclusions regardless of which version we use. The first panel shows consistently that the magnitude of shocks (measured by the standard deviation of the disturbances) is reduced by the transport variables. The effect of Roads is marginally statistically significant, but the Canal indicator is significant at the 1 per cent level.

The statistical significance of newspapers on the correlation of the disturbances is very strong. The correlation is measured conventionally (i.e. it takes a value between plus and minus one), so if one extra town got a newspaper (out of two counties with a total of ten towns) then the correlation would change by 0.02, which is quite a large effect

As expected, the biggest effect of including different lags is felt in models of the half life. The parameter estimate for Newspapers is very sensitive to the half-life measure, but is never statistically significant. Parameter estimates for the transport variables continue to have the “wrong sign”, suggesting mis-specification or omitted variables in the panel regression, but are nowhere near statistically significant. These results confirm the methodological point that estimating half lives from VECMs with insufficient lags not only biases the half lives up, but results in measures which cannot reliably be correlated with other explanatory variables.

In terms of our understanding of market efficiency in the Revolutionary and Napoleonic Wars, we conclude that – although prices converged to equilibrium more quickly – we are unable to explain why. Our transport and communication variables seem to have had more effect on price changes at a frequency of less than a week.

6. Summary and Discussion

We have analysed the comprehensive data set of London Gazette English grain prices for 1771-1820. In the spirit of Federico (2012), who notes that different authors have used different techniques, we have reported a variety of measures. Summary statistics of the data set suggest that the market was remarkably stable over this fifty year period, despite the expansion of transport networks and the shocks of wars. Looking at the graphs in section 2, it is difficult to see anything that has changed over the period, other than prices all moving up during the Napoleonic wars.

An increasingly popular tool for measuring market efficiency is the use of VECM models. Whilst it is well understood that the conclusions of these models depend upon the data frequency, we have – until now – had little idea of the magnitude of this effect. Since we have a complete set of weekly data, we have been able to estimate the speed of convergence to equilibrium not only on high frequency data, but also on lower frequency data, and thus quantify the importance of this issue. We show (in Table 1) that, on data where the “half life” appears to be about eight weeks, the use of monthly data would raise our estimate to eleven weeks; and the use of annual data would raise our estimate to twenty-two weeks. Furthermore, on our data the upward bias seems to be higher for those county pairs that are adjacent and for which the true half life is lowest.

A further issue that we estimate is whether the underlying assumption of geometric decay in price dispersion is correct. Using models with richer dynamic structures, we find that the convergence to equilibrium is quasi-hyperbolic, rather than geometric; that estimates of the half life may differ significantly; and that this may change the ordering of which markets we believe to be most efficient.

These two points taken together suggest that it may be difficult to compare reliably previous studies that use different frequency data or omit additional lags in time series estimation.

Our analysis supports the work of Jacks (2011), which finds that prices became more dispersed during the Napoleonic Wars, although price dispersion also remained high immediately after the conflict was over in 1816-17. The major reason for the increase in price dispersion was disturbances in the price dynamics: shocks from abroad mattered more, and so the disturbances to prices became larger, more highly correlated and more similar in size. There was an increase in market efficiency, as measured by the half life, but the effect of this was relatively small.

Our final contribution is to see whether transport and communication variables can explain either the overall behaviour of prices or the underlying components. The transport variables, but not our measure of newspapers, reduce the magnitude of random changes in prices, suggesting that arbitrage acts as a form of risk-pooling and reduces overall price variation. The primary importance of the transport variables appeared to be on the magnitude of the shocks, although this was not the only mechanism. Newspapers sped up the transmission of information, so that shocks to prices were more correlated: information arrived in different places at the same time. Market efficiency (moving towards equilibrium) occurs both within the period of observation (i.e. within the week) and over longer periods: the latter is measured through the half life. We know that the half life fell over the period 1770-1820, but estimates of this variable are sensitive to the model used: regardless of this, we are unable to explain the decline in the half life with the transport and communication variables that we have used here.

References

Bateman, V.N. (2007) "The evolution of markets in early modern Europe, 1350-1800: a case study of wheat prices" University of Oxford Department of Economics Discussion Paper No. 350.

Bateman, V.N. (2011) "The evolution of markets in early modern Europe, 1350-1800: a case study of wheat prices" *Economic History Review 64(2)*, pp. 447-471.

Bogart, D. (2005) "Turnpike trusts and the transportation revolution in 18th century England" *Explorations in Economic History 42,* pp. 479–508.

Buyst, E., Dercon, S., and Van Campenhout, B. (2006) "Road Expansion and Market Integration in the Austrian Low Countries during the Second Half of the 18th Century" *Histoire & Mesure, 21(1),* pp. 185-219.

Brunt, L. and Cannon E.S. (2002) "Do Banks Improve Financial Market Integration?" University of Bristol, mimeo..

Brunt, L. and Cannon, E.S. (2012) "The truth, the whole truth, and nothing but the truth: The English Corn Returns as a data source in economic history, 1770-1914" University of Bristol, mimeo. .

Ejrnæs, M., and Persson, K.G. (2000) "Market Integration and Transport Costs in France 1825–1903: A Threshold Error Correction Approach to the Law of One Price" *Explorations in Economic History, 37*, pp. 149–173.

Ejrnæs, M., and Persson, K.G. (2010) "The gains from improved market efficiency: trade before and after the transatlantic telegraph" *European Review of Economic History, 14(3),* pp. 361–381.

Ejrnæs, M., Persson, K.G., and Rich, S. (2008) "Feeding the British: convergence and market efficiency in the nineteenth-century grain trade" *Economic History Review, 61(S1)*, pp. 140–171. .

Federico, G. (2007) "Market integration and market efficiency: the case of 19th century Italy" *Explorations in Economic History, 44,* pp. 293-316. .

Federico, G. (2012) "How much do we know about market integration in Europe?" *Economic History Review, 65(2),* pp. 470–497.

Gibson, Jeremy S.W. (1991) *Local newspapers, 1750-1920: England and Wales, the Channel Islands and the Isle of Man: a select location list* (Birmingham: Federation of Family History Societies)..

Goodwin, B.K., and Grennes, T.J. (1998) "Tsarist Russia and the World Wheat Market" *Explorations in Economic History, 35,* pp. 405-430.

Goodwin, B.K., Grennes, T.J., and Craig, L.A. (2002) "Mechanical Refrigeration and the Integration of Perishable Commodity Markets" *Explorations in Economic History, 39*, pp. 154–182..

Hynes, W., Jacks, D.S., and O'Rourke, K.H. (2012) "Commodity market disintegration in the interwar period" *European Review of Economic History, 16,* pp. 119–143.

Jacks, D.G. (2011) "Foreign wars, domestic markets: England, 1793–1815" *European Review of Economic History, 15*, pp. 277–311..

Marks, D. (2010) "Unity or diversity? On the integration and efficiency of rice markets in Indonesia, c. 1920–2006" *Explorations in Economic History, 47,* pp. 310-324.

Moran, P.A.P. (1950) "Notes on Continuous Stochastic Phenomena" *Biometrika 37(1),* pp. 17–23..

Ó Gráda, C., and Chevet, J-M. (2002) "Famine and Market in Ancien Régime France" *Journal of Economic History, 62(3),* pp. 706-733.

Persson, K.G. (1999) *Grain Markets in Europe, 1500-1900* (Cambridge: Cambridge University Press).

Petersen, C., edited by Jenkins, A. (1995) *Bread and the British Economy c. 1770-1870* (Aldershot: Scholar Press).

Phillips, P.C.B. (1991) "Error correction and long-run equilibrium in continuous time" *Econometrica, 59(4)*, pp. 967-980..

Ravallion, M. (1986) "Testing Market Integration" *American Journal of Agricultural Economics, 68,* pp. 102-9..

Studer, R. (2008) "India and the Great Divergence: Assessing the Efficiency of Grain Markets in Eighteenth- and Nineteenth-Century India" *Journal of Economic History, 68(2),* pp. 393-437.

Taylor, A.M. (2001) "Potential Pitfalls for the Purchasing-Power-Parity Puzzle? Sampling and Specification Biases in Mean-Reversion Tests of the Law of One Price" *Econometrica, 69 (2)*, pp. 473-498..

Trenkler, C., and Wolf, N. (2005) "Economic integration across borders: The Polish interwar economy 1921–1937" *European Review of Economic History, 9(2)*, pp. 199-231.

Van Bochove, C. (2008) *The Economic Consequences of the Dutch* (Amsterdam: Aksant Academic Publishers)..

Appendix: The Expected Half Life

As noted above the formula for the half life is

$$HL \equiv \frac{\ln(0.5)}{\ln(1+\alpha_i-\alpha_j)} = \frac{\ln(0.5)}{\ln(1+\boldsymbol{\gamma\alpha})}$$

Where only estimates of the parameters are available, the issue is how to obtain the expected value of this quantity:

$$\mathsf{E}[HL] = \mathsf{E}\left[\frac{\ln(0.5)}{\ln(1+\boldsymbol{\gamma\alpha})}\right] \neq \frac{\ln(0.5)}{\ln(1+\boldsymbol{\gamma}\mathsf{E}[\boldsymbol{\alpha}])}$$

To check the importance of this, we tried a Monte Carlo procedure to calculate the *expected half life*. So long as the sample is sufficiently large, or the disturbances have a distribution which is not too far from the Normal, it follows that

$$\widehat{\alpha_i-\alpha_j} \sim \mathsf{N}\left(\alpha_i-\alpha_j, \mathsf{var}(\alpha_i-\alpha_j)\right)$$

Using this as a basis, we simulate 10,000 values $\alpha_i-\alpha_j$ from a Normal distribution $N\left(\widehat{\alpha_i-\alpha_j}, \mathrm{var}\left(\widehat{\alpha_i-\alpha_j}\right)\right)$ and calculate the corresponding 10,000 half lives (in a very small number of cases the draw of $\alpha_i-\alpha_j$ is negative and these are discarded). The mean of these then corresponds to the expected half life. We found that in nearly all cases the standard error of $\widehat{\alpha_i-\alpha_j}$ is so small that the HL and EHL are almost the same: the reason for this is that the correlation of $\hat{\alpha}_i$ and $\hat{\alpha}_j$ is very high.

Figures and Tables

Figure 1: Wheat prices 1770-1820

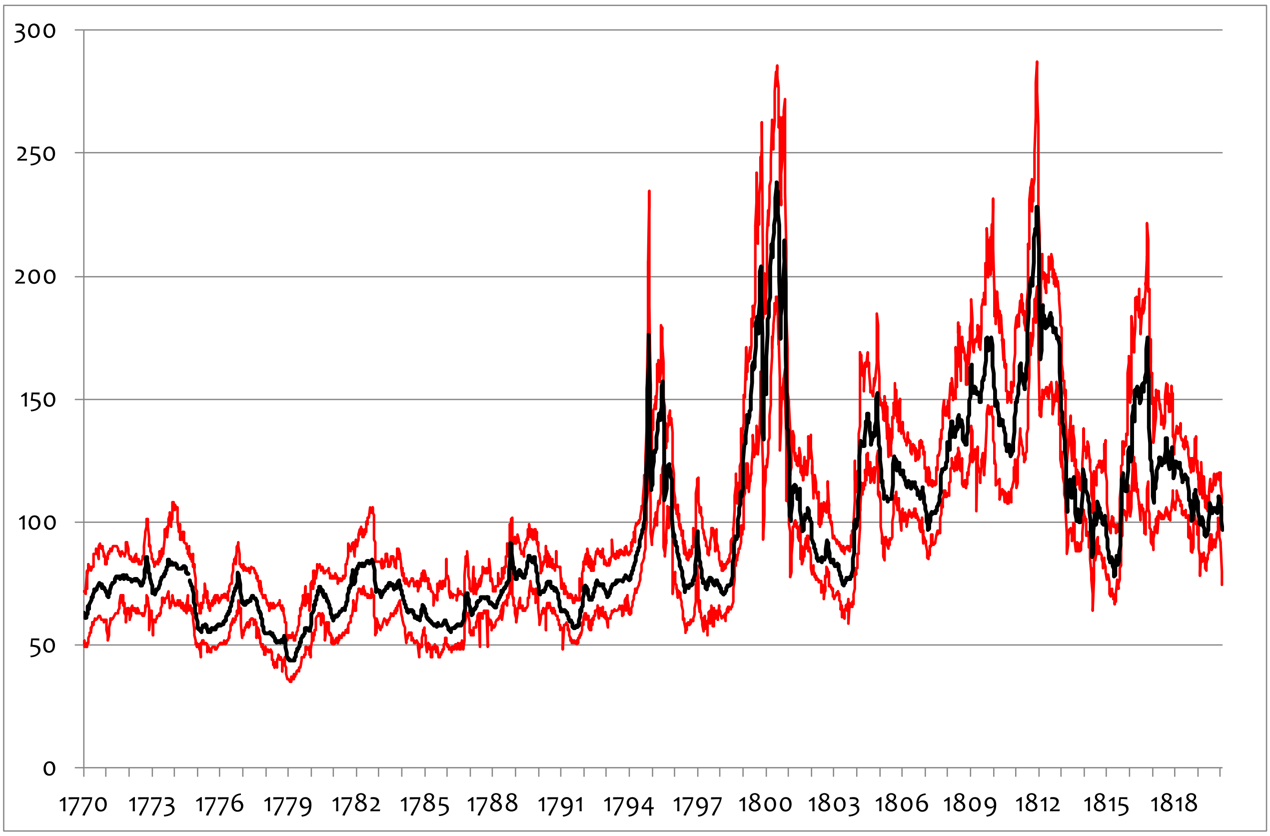

Figure shows the minimum, maximum and average London Gazette wheat price in each week from November 1770 to September 1820.

Figure 2: Dispersion of prices between counties

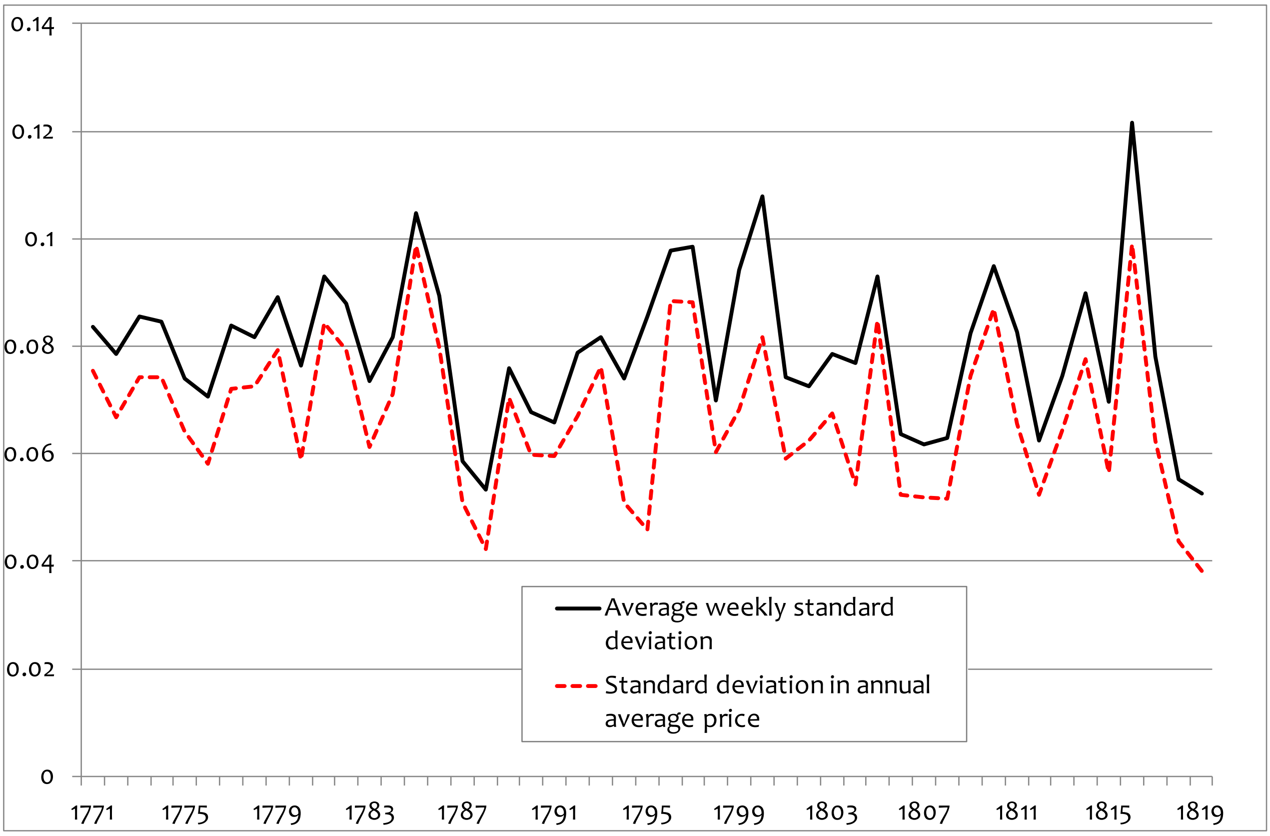

Average weekly standard deviation: the standard deviation of log prices is calculated for each week of the sample and then the 52 standard deviations are averaged for a harvest year (Oct-Sep); standard deviation in annual average: the harvest-year mean price is calculated for each county and then the standard deviation is calculated of the forty mean prices.

Figure 3: Year-on-year correlations of cross-sections of prices

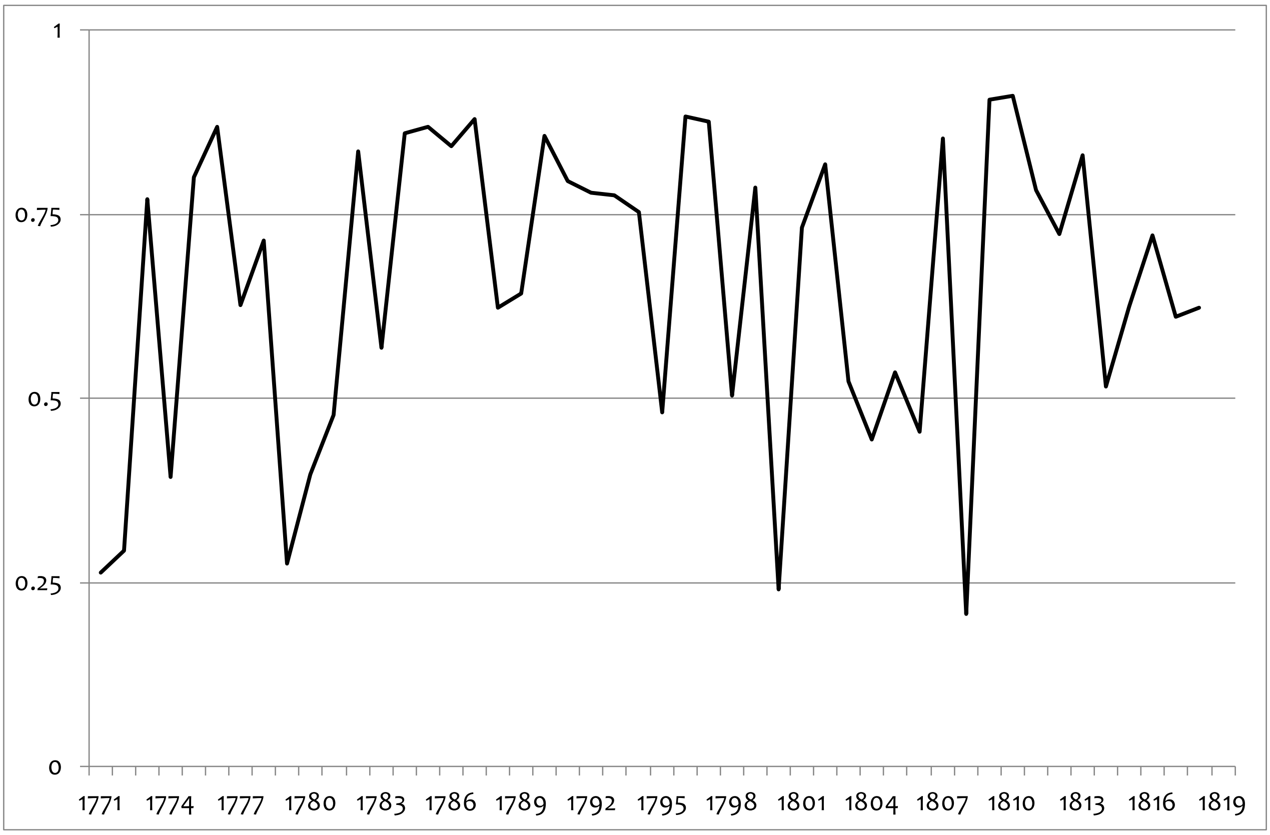

The graph plots the correlations of county prices in each year with prices in the following year (equation 2).

Figure 4: Correlations of cross-sectional price series for all year-pairs

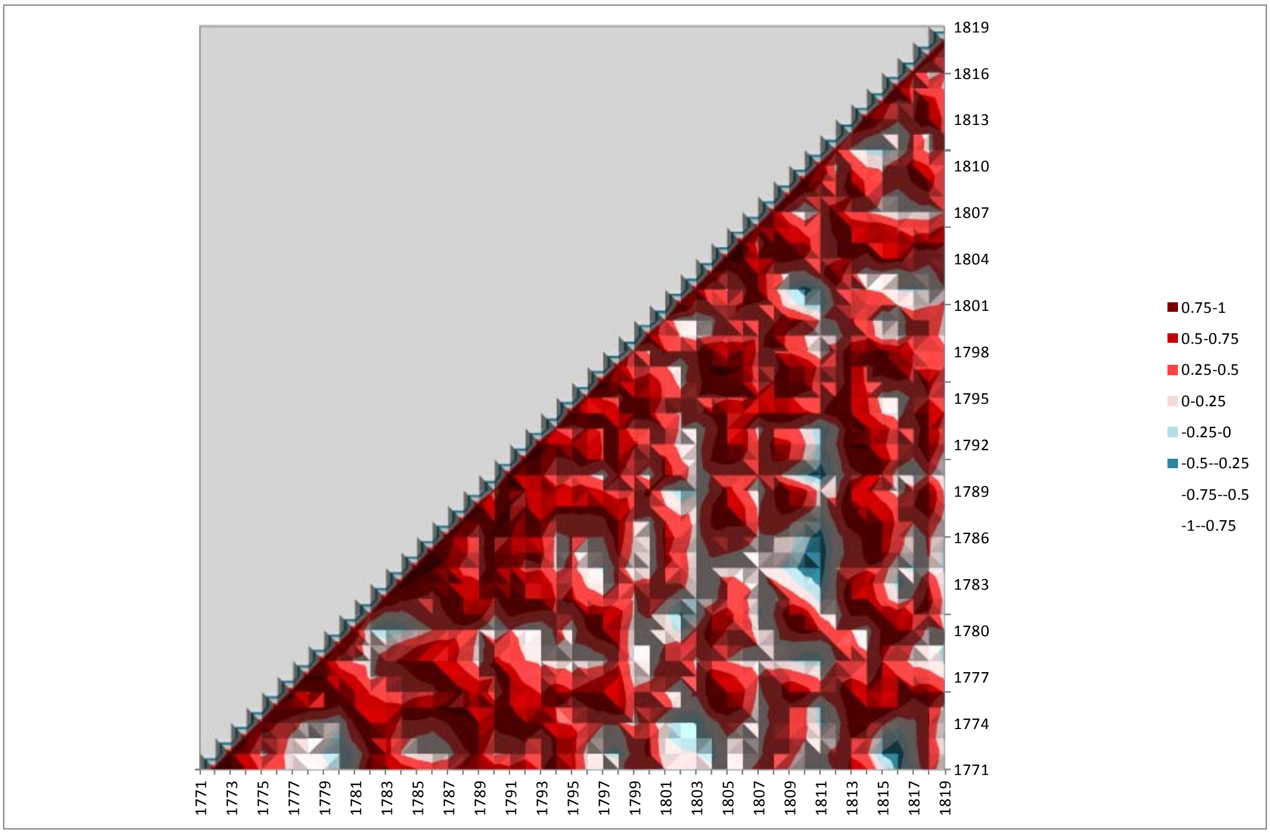

The graph plots the correlation of county prices in each year with all other years.

Figure 5: Moran's I Statistics

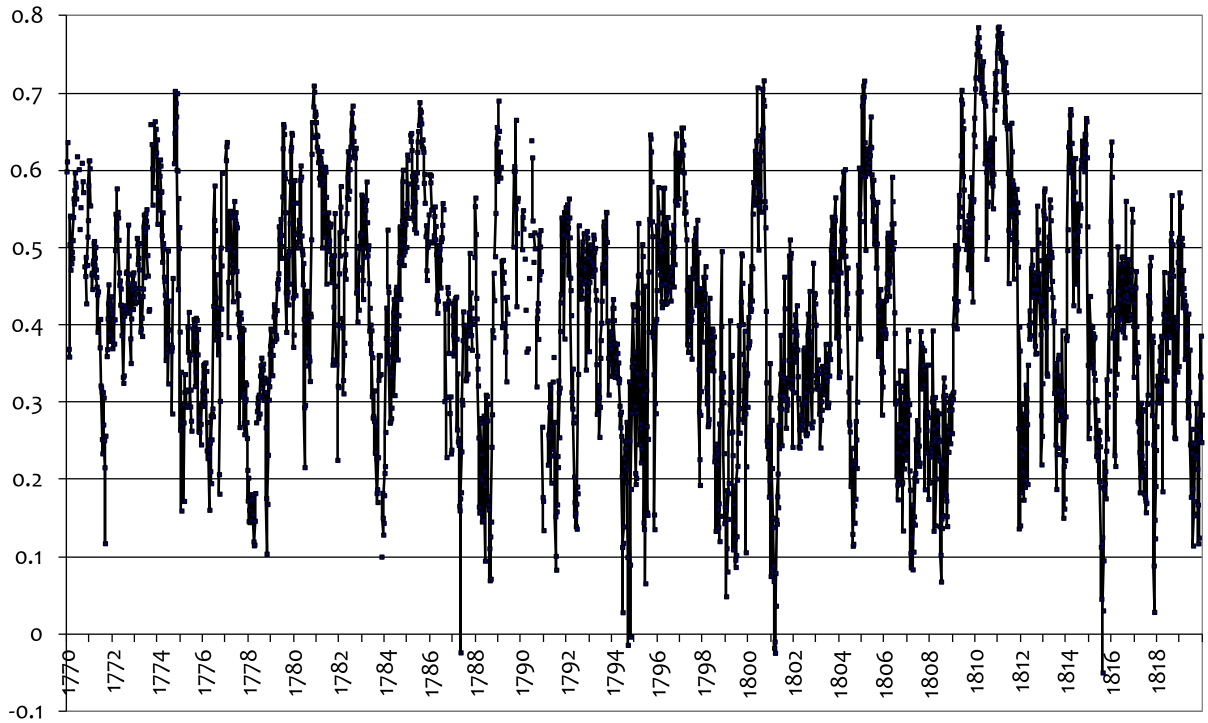

Each point plotted in the figure is a Moran I statistic calculated from a separate cross-section of weekly wheat prices using the formula in equation (3).

Figure 6: Distribution of half lives from models estimated on 1770-1820 data

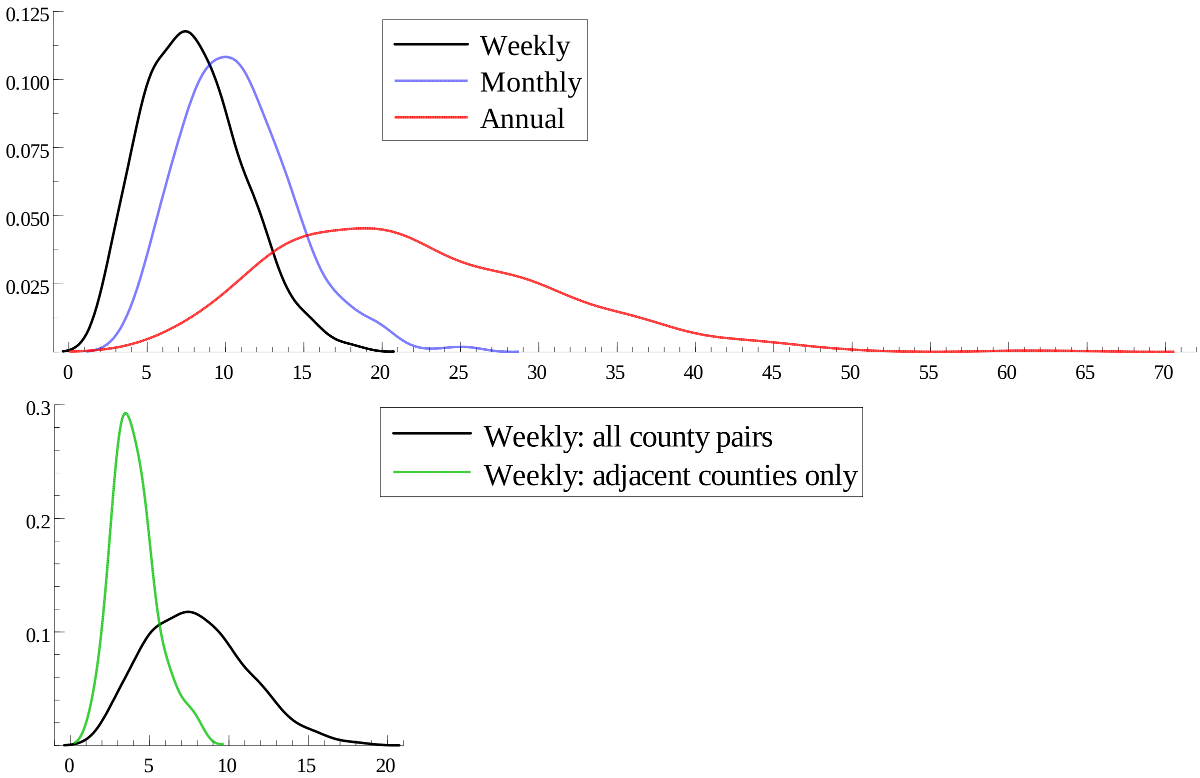

Each distribution in the top panel is based on 780 half lives (slightly fewer for annual data, where some half lives could not be calculated). Each half life is estimated using a model of the form in in equation (9) using data from the entire period 1770-1820, except where one of the prices is from London when it is 1770-1793. The bottom panel reproduces the distribution of all 780 half lives from the top panel and compares it to the distribution of the 103 half lives where the counties are adjacent.

Figure 7: Impulse response functions for Bedfordshire-Buckinghamshire

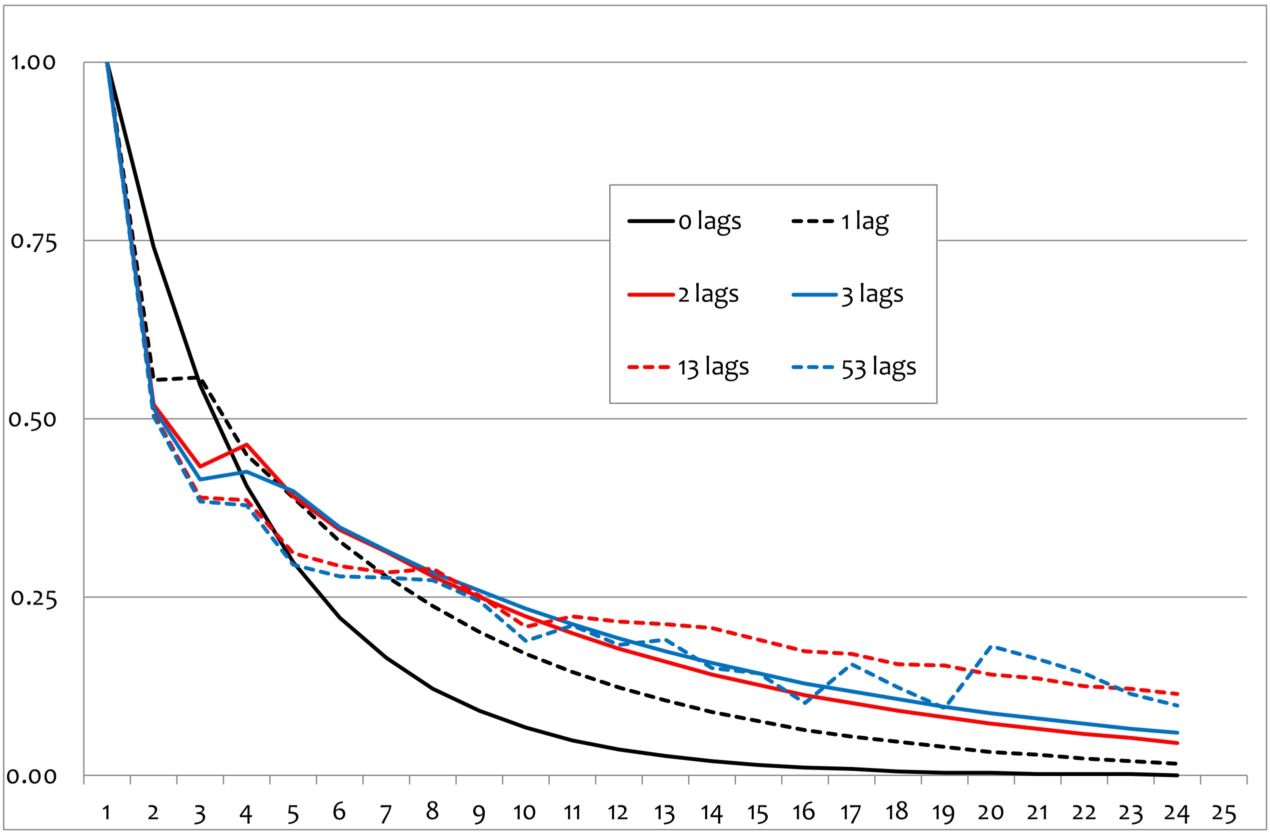

The graphs illustrate the speed with which a log-price difference dies away over time (the horizontal axis is measured in weeks). Each impulse response function is estimated using equations (10) to (12). The underlying models are estimated on the full sample of weekly data from 1770-1820 and differ only in the number of lagged dependent variables (the parameter K).

Figure 8: Distributions of half life estimates

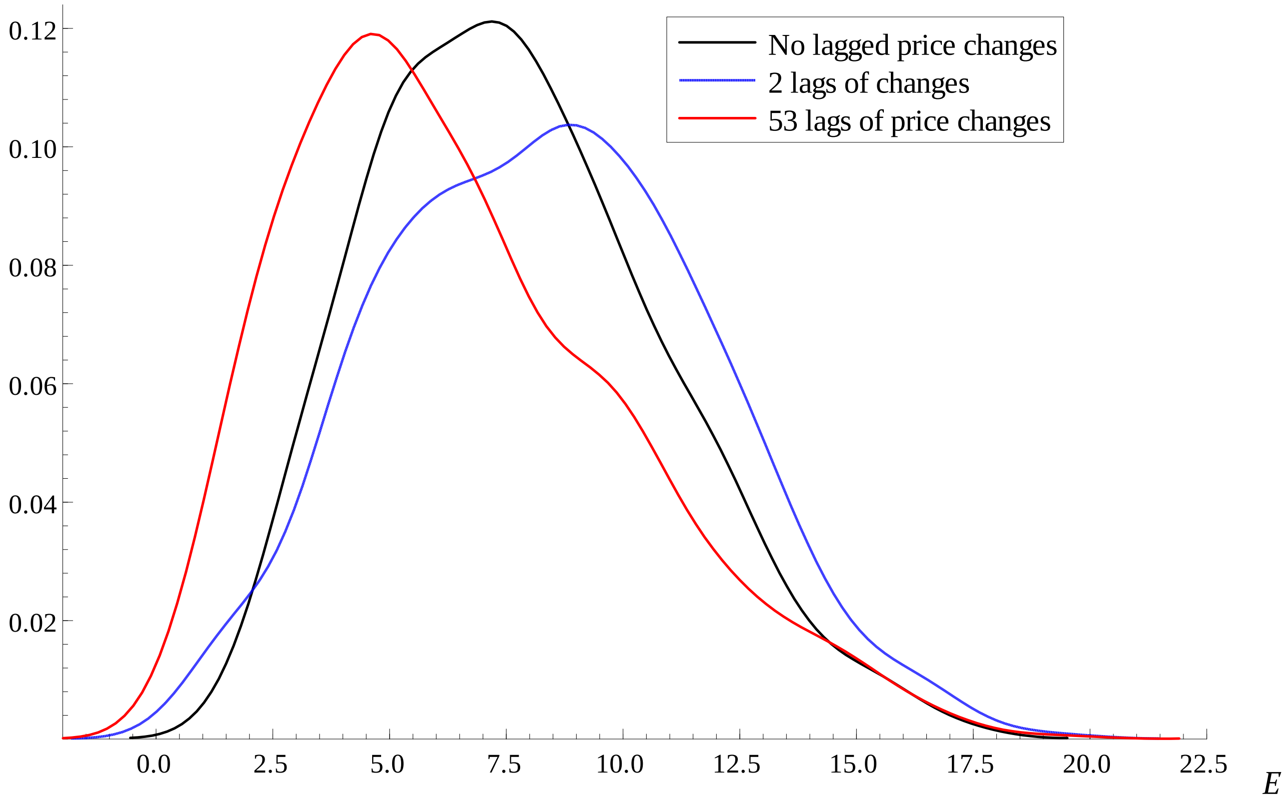

ach of these distributions summarises the half lives from 780 regressions, each of which is estimated on weekly data for a county pair for the entire period 1770-1820. The only difference between the distributions is the number of lagged price changes used in the regression. These distributions are based on the same information as the third row of Table 2. Note that all the half lives were positive but an artefact of the kernel smoothing method used to estimate the density was that the curves appear to extend to the left of the origin.

Figure 9: Dispersion of prices (average rmspd of adjacent-county pairs, per cent)

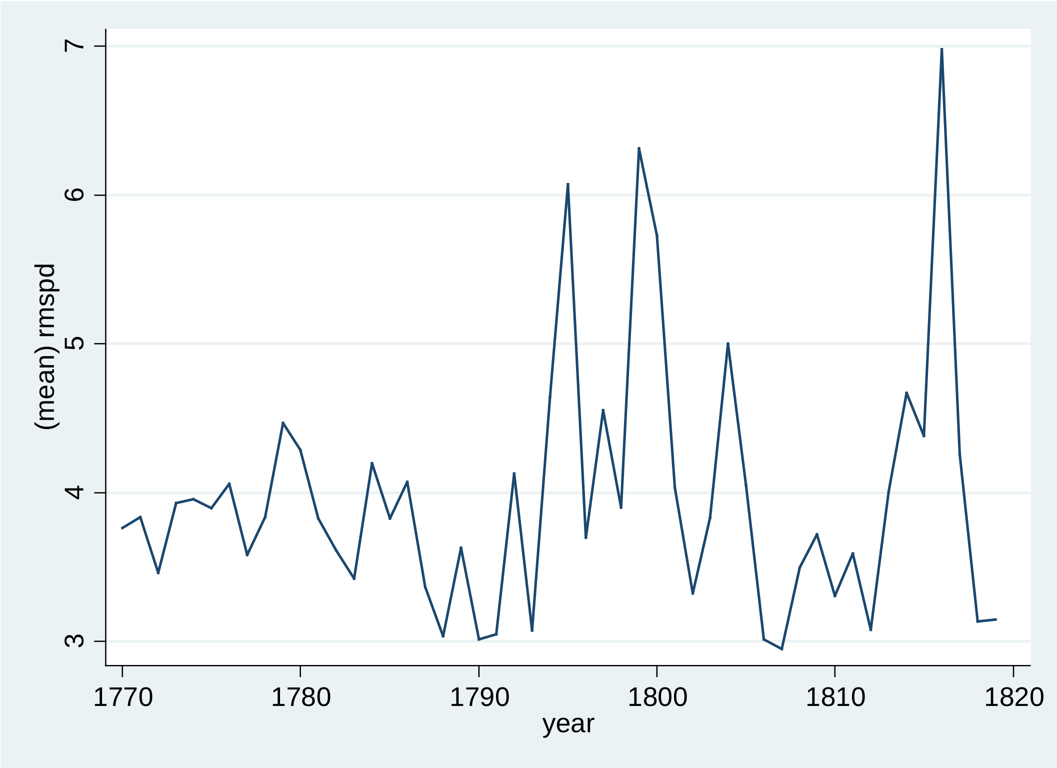

Figure 10: Mean standard deviation of disturbances for each year (per cent)

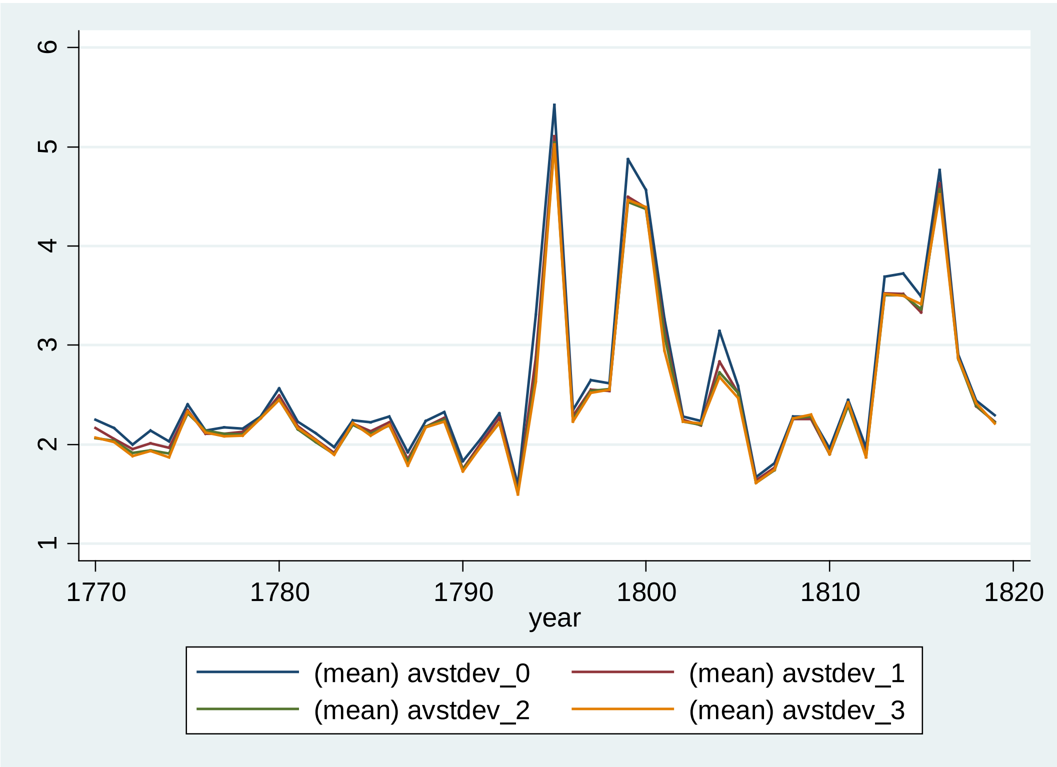

Figure 11: Mean correlation of disturbances for each year

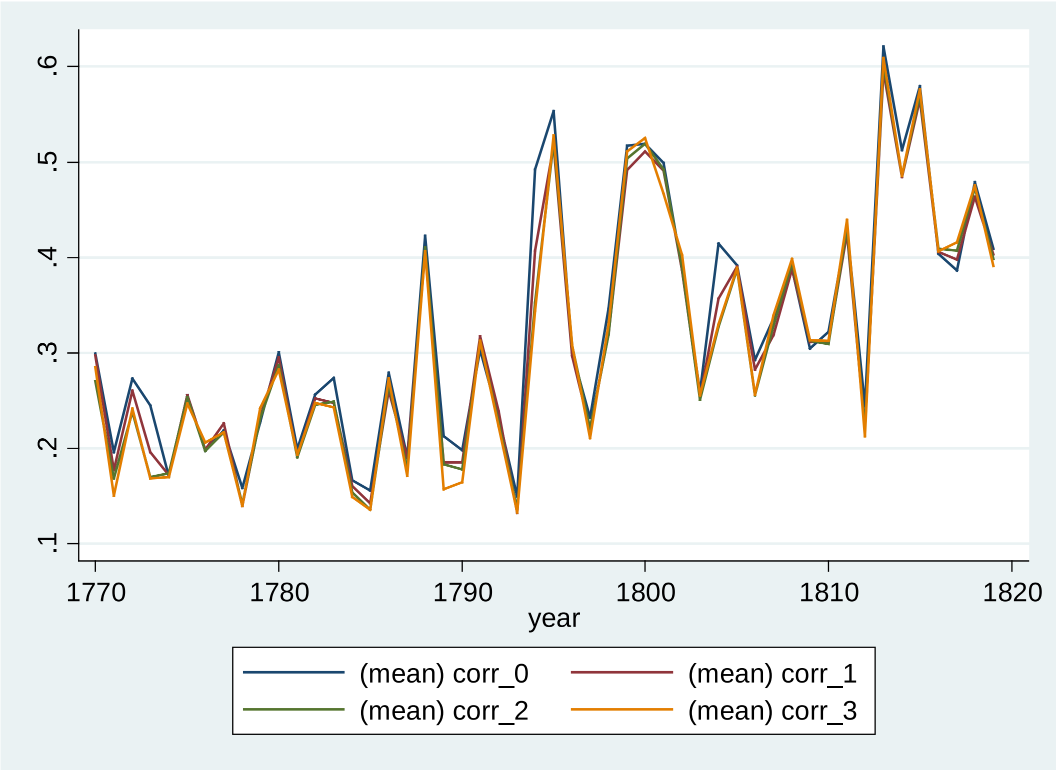

Figure 12: Mean ratio of standard deviations of disturbances

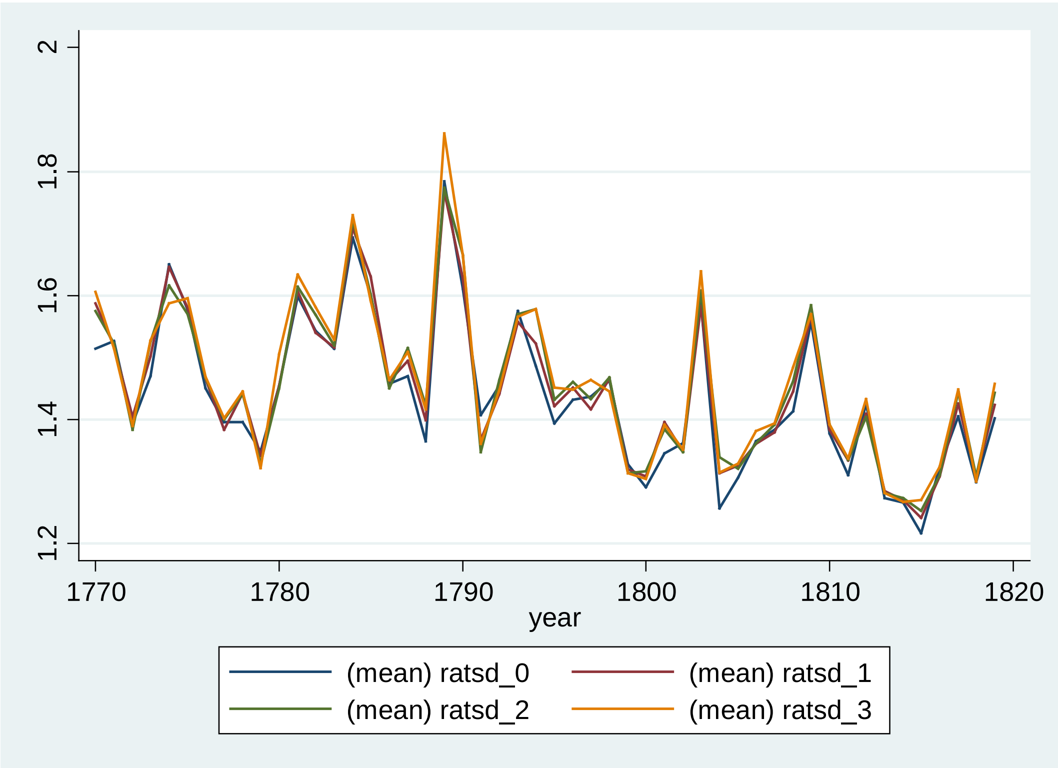

Figure 13: Mean of half lives for each year (weeks)

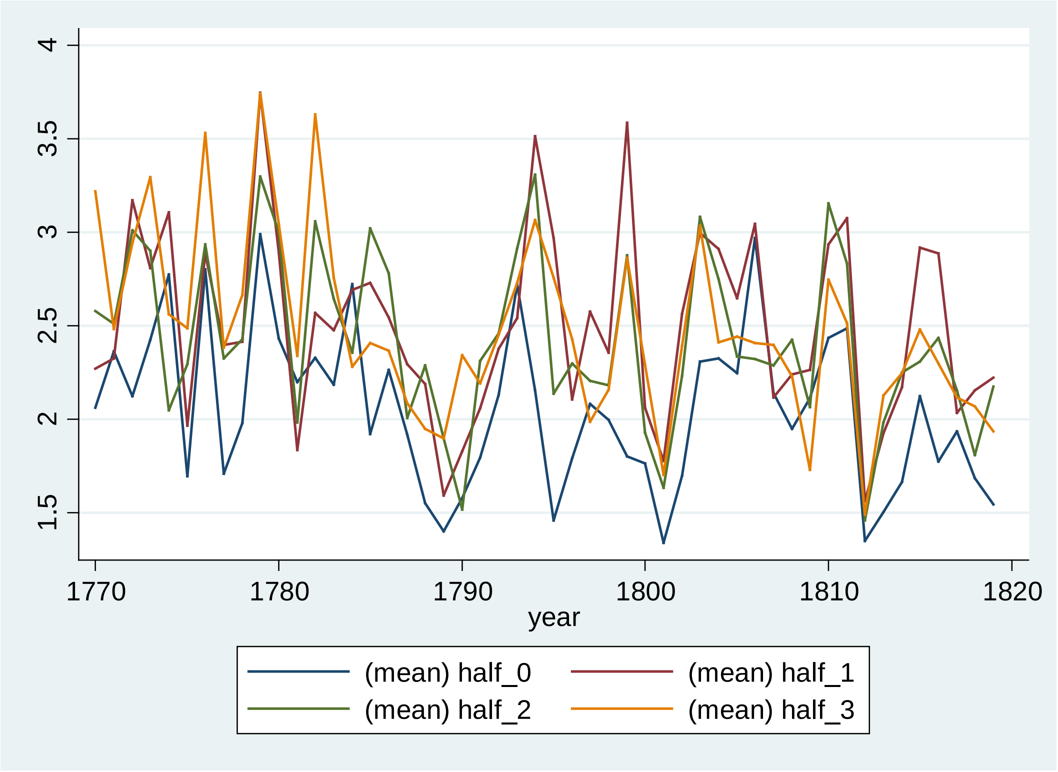

Table 1: Summary of half lives estimated for 1770-1820

	All county pairs			Adjacent-county pairs		
frequency of data	weekly	monthly	annual	weekly	monthly	annual
mean	8.0	10.8	22.1	4.1	6.9	22.9
median	7.7	10.4	20.7	3.7	6.4	21.2
st.dev.	3.2	3.7	9.0	1.4	2.2	10.9
minimum	1.5	3.3	5.5	1.5	3.3	5.5
maximum	18.7	25.8	65.0	8.0	14.0	60.9

The first three columns summarise the distribution of 780 half lives (slightly fewer for annual data, where some half lives could not be calculated). Each half life is estimated from a regression of the form in equation (9) using data from the entire period 1770-1820, except where one of the prices is from London when it is 1770-1793. The final three columns report analogous statistics for the 103 pairs where the counties are adjacent.

Table 2: Effect of including lagged price changes on half-life estimates

Number of lags	0	1	2	3	13	53
Bedfordshire and Buckinghamshire						
Half life based on loadings	2.30	3.18	3.86	4.07	4.66	6.73
Half life from impulse response function	2.33	2.53	1.23	1.14	1.06	1.04
Average for all 780 county pairs						
Half life from impulse response function	7.72	8.50	8.36	8.16	6.66	6.44
Average for the 93 adjacent-county pairs						
Half life from impulse response function	4.09	4.61	4.28	3.71	2.95	2.80

Results are based on regressions using weekly data for Bedfordshire and Buckinghamshire for the whole period 1770-1820: the only difference is the number of lagged price changes. The half life in row 1 is based upon equation (7). The half life in row 2 is based upon Figure 7: linear interpolation is used to see where the impulse response function crosses the line $h = \frac{1}{2}$. *Half lives in rows 3 and 4 are calculated analogously to those in row 2.*

Table 3: Regression analysis

	Root mean square price difference	Components of price dispersion (from VECM with no lagged price changes)			
		Average St. Dev.	Ratio St. Dev.	Correlation of disturbances	Half life
Roads	-0.027 (1.310)	-0.030 (2.517)	-0.015 (2.526)	0.002 (0.640)	0.067 (2.316)
Canals	-0.236 (1.967)	-0.171 (2.440)	0.011 (0.305)	0.031 (1.987)	0.307 (2.131)
Newspapers	-0.735 (1.708)	-0.073 (0.264)	-0.705 (3.874)	0.202 (3.044)	-0.296 (0.332)
N × T	4642	4642	4642	4642	4564
R-squared	0.356	0.722	0.090	0.340	0.051

Results are for five different regressions, each for a panel for fifty harvest years (1770-71 to 1819-20 and 93 adjacent-county pairs. Other than the Canals, Roads and Newspapers explanatory variable, there are adjacent-county-pair fixed effects and year fixed effects. T-statistics in parentheses are robust to heteroskedasticity and within-group correlation. The R-squared refers to within-group variation.

Table 4: Regressions using different measures of market integration

No. of lags	0	1	2	3
Dependent variable: Average standard deviation of disturbances				
Roads	-0.030 (2.517)	-0.024 (1.941)	-0.025 (1.965)	-0.026 (2.056)
Canals	-0.171 (2.440)	-0.196 (2.641)	-0.205 (2.720)	-0.204 (2.687)
Newspapers	-0.073 (0.264)	-0.034 (-0.117)	-0.124 (0.418)	-0.111 (0.366)
N × T	4642	4642	4642	4642
R-squared	0.722	0.688	0.680	0.674
Dependent variable: Ratio of standard deviations of disturbances				
Roads	-0.015 (2.526)	-0.014 (2.250)	-0.013 (2.079)	-0.013 (2.087)
Canals	0.011 (0.305)	0.015 (0.465)	0.005 (0.142)	0.005 (0.132)
Newspapers	-0.705 (3.874)	-0.709 (3.760)	-0.674 (3.524)	-0.601 (3.216)
N × T	4642	4642	4642	4642
R-squared	0.090	0.085	0.084	0.086
Dependent variable: Correlation of disturbances				
Roads	0.002 (0.640)	0.004 (1.236)	0.003 (1.039)	0.003 (1.122)
Canals	0.032 (1.987)	0.029 (1.717)	0.032 (1.851)	0.036 (2.060)
Newspapers	0.202 (3.044)	0.213 (3.320)	0.188 (2.939)	0.187 (2.953)
N × T	4642	4642	4642	4642
R-squared	0.340	0.316	0.313	0.302
Dependent variable: Half life				
Roads	0.067 (2.316)	0.063 (1.673)	0.039 (1.044)	0.064 (1.508)
Canals	0.307 (2.131)	0.199 (0.852)	0.397 (1.771)	0.281 (1.181)
Newspapers	-0.296 (0.332)	1.132 (1.084)	0.928 (1.051)	-0.243 (0.195)
N × T	4564	4384	4330	4308
R-squared	0.051	0.032	0.030	0.032

Results are for sixteen separate regressions. The explained variables are themselves estimated from regressions on weekly data for each county pair: the column headings refer to the number of lags in the first-stage time-series regressions.