UNIVERSITY OF BRISTOL

DEPARTMENT OF ECONOMICS

A Grain of Truth in Medieval Interest Rates? Re-examining the McCloskey-Nash Hypothesis

Liam Brunt & Edmund Cannon*

February 1999

Abstract

Using modern econometric techniques, we show that it is feasible to accurately extract the interest rate from the appreciation of asset prices within the year. On the basis of movements in the price of wheat, we estimate that the interest rate in medieval England was around 12 per cent per annum. This enables us to reject two hypotheses which have featured prominently in the debate on the medieval economy. First, the interest rate is significantly above zero; this runs contrary to the Komlos-Landes view that medieval agents were not economically rational profit-maximisers. Second, the interest rate is significantly below 30 per cent; this undermines the McCloskey-Nash argument that the cost of capital prevented the year-on-year storage of foodstuffs. A discount rate of 12 per cent could also generate output-sharing cooperatives in Kimball's model of farmer behaviour and explain the puzzling pattern of strip scattering in medieval agriculture.

Keywords: Interest rates, wheat prices, medieval agriculture.

JEL Classification: N53, O13, Q14

---

* Brunt: Nuffield College, Oxford; Cannon: University of Bristol. Correspondence to Cannon at Dept. of Economics, 8, Woodland Road, Bristol. BS8 1TN. (Edmund.Cannon@bristol.ac.uk) Brunt received financial support from the Economic & Social Research Council, the Fulbright Commission and Nuffield College; this paper was written whilst he was visiting the Economics Dept. at Harvard University. The data entry for this paper was partly funded by the Royal Economic Society under the small grant scheme. We should also like to thank Cliff Attfield, Andrew Chesher, David Demery, Kees Jan van Gardaren and Lucy White for helpful suggestions; the Warden & Fellows of Merton College for access to the account rolls; Rachel Hands for advice on reading the account rolls; Paul Harvey for providing us his unpublished notes on Cuxham grain prices. Any remaining errors are our own.

Introduction

A puzzling feature of the medieval European economy is the custom of strip scattering. Landholdings were divided into small strips (just a fraction of an acre) and each farmer held multiple strips which were geographically scattered around the village. This was productively inefficient because it prevented the farmer from gaining economies of scale. Yet the transaction cost of reorganising the strips into larger units would have been small (Stefano Fenoaltea, 1988), so there must have been a rational reason to prefer scattering. Unfortunately, medieval scholars left us no explanation for this apparently strange behaviour. The mystery is compounded by the fact that strip scattering did not occur in every locality and its popularity fluctuated over time. Any plausible explanation of strip scattering therefore needs to explain its spatial pattern and evolution. The strip scattering mystery is probably the most prominent question in medieval economic history – given both the importance of agriculture in the overall economy and the low standard of living endured by the population.

Donald N. McCloskey (1976) suggested that strip scattering was a form of insurance against very localised harvest failure. McCloskey and John Nash (1984) bolstered this argument by presenting evidence that other forms of insurance would have been prohibitively expensive. Notably, grain could be stored from one year to the next - but the stored grain would have had to appreciate fast enough cover the cost of capital, storage losses, depreciation and inflation (which together comprise the gross rate of return on grain). Since McCloskey and Nash calculated that the medieval gross rate of return was between 30 and 40 per cent per annum (and most of this was due to the rate of interest), year-on-year storage was very unlikely to be an optimal insurance strategy.

The claim by McCloskey and Nash that high interest rates made strip scattering the optimal insurance strategy has generated considerable debate. Miles S. Kimball (1988) shows that a farmers' cooperative could be a better insurance strategy under some levels of harvest uncertainty. The advantage of cooperatives grows as the interest rate falls below the 30 per cent postulated by McCloskey and Nash (or farmers' discount rates were lower). It seems likely that discount rates were indeed lower, given the success of medieval merchants in forming self-enforcing cooperative trade organisations (Avner Greif, 1993, 1994). John Komlos and Richard Landes (1991) note that medieval peasants had few alternative investments to holding grain and argue that, in any case, they were simply not rational profit-maximisers and would not arbitrage between holding grain and other investments. If this were true then presumably the rate of interest (or even the gross return) on grain holdings would be close to zero. Of course, this issue clearly goes far beyond the question of grain storage: such a high interest rate would have enormous implications for any kind of investment or accumulation decision in the medieval economy.

Remarkably, no one has disputed the claim that medieval interest rates were in excess of 30 per cent per annum. We show that McCloskey and Nash massively over-estimate the interest rate. Using exactly the same data, we show that the rate of interest was actually around 12 per cent.

Our result is significantly below 30 per cent – enabling us to reject the McCloskey and Nash idea that interest rates were prohibitively high. It is also significantly above zero – enabling us to reject the Komlos and Landes argument that medieval farmers were not rational. Finally, an interest rate of 12 per cent would ensure cooperation in Kimball's farmer cooperatives where the coefficient of variation in harvests was high (0.6) but not low (0.3). This could explain both the spatial pattern of strip scattering and its tendency to die out over time.

The reason McCloskey and Nash over-estimate the interest rate is that they systematically ignore evidence suggesting that the interest rate was low. If grain prices follow the sawtooth pattern which McCloskey and Nash suppose, then the size of the price rise over the year and the size of the price fall both contain useful information about how much the grain appreciated from the beginning to the end of the year. McCloskey and Nash ignore the evidence on price falls, and this leads them to an inefficient estimate that turns out to be very inaccurate.

Our argument will proceed as follows. In section 1 we replicate the McCloskey and Nash method and show that is flawed – and that correcting their method would lead to a much lower estimate of the interest rate. In section 2 we formulate a simple econometric model of grain prices; and in section 3 we use our model to estimate the gross return. In section 4 we address the issue of storage costs. In section 5 we discuss the plausibility and importance of our results. We conclude in section 6. The peculiarities of the data set are discussed in some detail in the Appendix.

1. Replication of McCloskey & Nash

The conclusions of McCloskey and Nash are based upon the estimates of the gross return to holding wheat. In this section we replicate their results and explain why they are misleading.

By far the best documented physical asset in the medieval economy is wheat[1], so McCloskey and Nash use wheat prices for their data series (taken from James E. Thorold Rogers, 1866). They note that the stylised path of grain prices is a sawtooth pattern. Grain is harvested in August and placed in storage; then it is gradually consumed through the year until stocks become exhausted on the eve of the next harvest. During the storage period, the unit value of grain is rising by the gross rate of return. When the new harvest begins the price falls sharply to its old low point (in the absence of shocks to supply or demand). Then the cycle of appreciation begins again.

In order to estimate the gross return it is sufficient to measure the appreciation of asset prices between two dates in the same village and year (hence removing village and year-on-year effects).[2] The gross return between month *t* and month *s* is then the mean value of $(\ln P_t - \ln P_s)$ /

[1] Other than the source we use, data is readily available in, for example, Nicholas W. Posthumus (1964), William H. Beveridge (1939) and Earl J. Hamilton (1965).

[2] Our convention is to define the year to run from August 1 to July 31. This corresponds very closely with the majority of accounts with usable data, which begin on 27 July. McCloskey and Nash conducted a much

($t-s$). We calculated these figures using the data of Thorold Rogers and report our results in Table 1, which corresponds to Table 2 on p.179 of McCloskey and Nash. A simple comparison can be obtained by using only the numbers on the leading diagonal.[3] Our results suggest a return from September to July of 56.5 per cent; and our unweighted measure 53.3 per cent (both on an annual basis). The analogous figure from McCloskey and Nash is 59.9 per cent.[4] Thus we agree with the analysis of McCloskey and Nash as it stands.

The one piece of information that McCloskey and Nash do not report is the price fall from August to September. If we denote the average gross return per month to be γ, then the average fall from August to September must be 11γ, since wheat prices have no trend. But in fact, we calculate a weighted average price fall of 0.13 per cent and an unweighted fall of 0.03 per cent (again on an annual basis). This is clearly very worrying.

We can see that McCloskey and Nash discard important information about the return to holding grain by ignoring the price fall. But it is not immediately clear how much importance should be attached to it. There are two reasons to suppose that the price fall information should be given a low weighting. First, there are relatively few price pairs in those months. Second, Taub (1987) argues that the information in the price drop might be contaminated by the arrival of information about the next harvest (such as might be revealed by weather shocks, for example).[5]

However, there are good statistical grounds for attaching a very high weight to the price fall. We can illustrate this point most easily by considering the case of Gaussian disturbances. Suppose that the gross monthly return for the first eleven months is $N(\gamma,\sigma^2)$ and that fall at the end of the year is $N(-11\gamma,\omega\sigma^2)$. The log-likelihood would be

$$L = \sum_{i=1}^{M_1} \ln f\left(\frac{\Delta \ln P_i - g}{s}\right) + \sum_{j=1}^{M_2} \ln f\left(\frac{\Delta \ln P_j + 11g}{s\sqrt{w}}\right)$$

where ϕ is the density function, M_1 is the number of price rises for the first 11 months and M_2 the number of price falls at the end of the year and $\Delta \ln P_i$ and $\Delta \ln P_j$ are the price rises and price falls respectively. The parameter ω measures the relative variance of price falls at harvest to price rises throughout the rest of the year and will thus be greater than one. The reason for

more extensive re-basing of the accounts so that they run from September 1 to August 31. We have made all the calculations both ways and found that it makes no difference to the results.

[3] The figures above the leading diagonal suggest a rather lower gross rate of return, both in our analysis and that of McCloskey and Nash. To use all of the information in the matrix, we would need to take an appropriately weighted average (McCloskey & Nash use a simple average, which results in a degree of double counting). The issue of weighting is implicitly solved in our analysis in Section 4.

[4] We use very slightly fewer observations than McCloskey and Nash because we reject price observations on transactions which have no specified quantity. This is partly because they seem less reliable and partly because we also use the quantities to conduct Weighted Least Squares estimation in our later analysis.

[5] In fact, very little useful information about the weather arrives in advance of the harvest – because the crucial time period is immediately before and during the harvest itself (Brunt, 1998). Moreover, contemporaries found it very difficult to get an accurate estimate of the amount of grain produced even after it was harvested. They had to wait until a sample had been threshed and the grain was separated from the straw.

attaching a high weight to price falls can be seen from the first order condition for the maximum likelihood estimate of γ:

$$\sum_{i=1}^{M_1}\frac{f'\left(\left(\Delta\ln P_i - g\right)/s\right)}{f\left(\left(\Delta\ln P_i - g\right)/s\right)} - 11\sum_{j=1}^{M_2}\frac{f'\left(\left(\Delta\ln P_j + 11g\right)/s\sqrt{w}\right)}{f\left(\left(\Delta\ln P_j + 11g\right)/s\sqrt{w}\right)} = 0$$

so, *ceteris paribus*, the price falls are given 11 times as much weight as the price rises. If we solve for the maximum likelihood estimator of γ, we obtain

$$\hat{g} = \frac{w\sum_{i=1}^{M_1}\Delta\ln P_i - 11\sum_{j=1}^{M_2}\Delta\ln P_j}{wM_1 + 121M_2}$$

The estimate of McCloskey & Nash corresponds to either the assumption that $M_2 = 0$ or that $\omega = \infty$, neither of these assumptions being correct.

An estimate of γ would be straightforward to calculate if we knew ω, but, given the number of observations we have for price falls in August-September, we would have little confidence in a maximum likelihood estimate of ω. Instead, we calculate the maximum likelihood estimate of γ and the corresponding gross return at an annual rate for a range of plausible estimates of ω (still using just the data from the leading diagonal of Table 1):

ω		1	2	3	4	5	6	7
$\hat{g}$	.006	.009	.011	.013	.015	.016	0.16	
Gross return (per cent)	7.41	11.5	14.6	17.0	19.0	20.5	21.8	

Thus even if the relative variance ω is very large, the maximum likelihood estimate of the gross return would not be much in excess of 20 per cent, dramatically less than the estimate of McCloskey & Nash.

It would be possible at this stage to refine our procedure to take account of observations above the leading diagonal. But using observations above the leading diagonal involves some double counting, since, if we already have observations for April-May and May-June for the same village and year, then we also have an April-June observation (this difficulty is ignored by McCloskey & Nash). The optimal weighting will depend upon the dynamic process generating the data: this will be complicated unless there is a stochastic trend within the year. Since we cannot be sure that there is a stochastic trend within the year, we prefer to model the level of the grain price and deal with econometric issues such as auto-correlation and heteroskedasticity directly.

2. An Econometric Model

2.1 Seasonal Grain Prices

In this section we build a simple econometric model of the level of grain prices. Denote the logarithm of the market price in any month m, village v and year t as $p_{v.t.m}$. We decompose this into

$$p_{v.t.m} = \alpha\, p^*_t + \kappa_v + f(m) + u_{t.v.m}$$

where p^*_t is the general price level for the year, κ_v is a village dummy, f(m) determines the deviation of the price in any given month from the annual price, $u_{t.v.m}$ is a random error with mean zero denoting the deviation of the market price from its trend. For most of the time we shall consider f(m) = γm, where m is zero in the first month of the agricultural year: the gross return on an annualised basis is then $e^{12\gamma} - 1$. An analogous model expressed in terms of the day x, is

$$p_{v.t.x} = \alpha\, p^*_t + \kappa_v + \gamma_{(x)}x + u_{t.v.x}$$

where the gross return is $e^{365\gamma(x)} - 1$. Since these two models are structurally identical we shall frame our discussion entirely in terms of the former.

Our data is not "the" market price, but prices of individual trades, which may deviate from the market price due to unobserved differences in the quality of grain or other factors. So our model of observed prices, $p_{v.t.m.i}$, is

$$p_{v.t.m.i} = p_{v.t.m} + e_i = \alpha\, p^*_t + \kappa_v + f(m) + u_{t.v.m} + e_i,$$

where e_i is a random error with mean zero. Where we have several price observations in the same month and take an average of these price observations, such an average will tend to have a smaller variance than a price based on only one price observation, so this will cause heteroskedasticity. The prices of very large trades are more representative of the market price than for small trades, which will also cause heteroskedasticity. We control for these two effects individually by using appropriate weighted least squares (WLS).

We consider two ways of modelling the year effect. Our first method is to consider a dummy for each year, which is also differs across villages, which we shall call a village-year dummy. This approximates to the method of McCloskey & Nash, differing only in the implicit weights given to different observations. An alternative strategy, which loses fewer degrees of freedom is to proxy p* by the average annual price series from Thorold Rogers (1866, vol.1 pp.227.ff), which we shall denote p^{TR}. The variable p^{TR} is effectively measured with error, so, to prevent mis-estimation of the parameters of interest, we consider the regression

$$p_{v.t.m} = \alpha\, p_t^{TR} + \kappa_v + f(m) + u_{t.v.m},$$

where we do not restrict $\alpha = 1$. So long as p^{TR} is not correlated with m or κ, the errors in variable problem will lead to biased estimates of α but not biased estimates of the parameters of interest. The observed correlation is sufficiently small[6] for us to be confident that this method should also be quite reliable.

The final consideration is the error term, $u_{t.v.m} + e_i$. We have already noted reasons why e_i may contribute to heteroskedasticity. The component $u_{t.v.m}$ may be both heteroskedastic and autocorrelated.

The effect of autocorrelation is most important in conventional short time series, because the initial random error persists and biases the estimate of the parameters. In our data, however, we have many sets of observations which are effectively independent: the residuals in data from Cuxham in 1344 will be unrelated to the residuals in Letherhead in 1314, so they are effectively repeated drawings of the "initial" random error, which will tend to cancel each other out and reduce the variance of our estimates of the parameters in f. But although the variance of our estimates of these parameters is reduced, the variance of our standard errors will be increased. Furthermore our data is too piece-meal to enable us to remove autocorrelation through inclusion of lagged dependent and explanatory variables or estimate using auto-regressive least squares.

We therefore estimate our model using OLS or WLS, but estimate the standard errors using the Heteroskedastic and Autocorrelation Consistent Standard Error (HACSE) proposed by Whitney K. Newey & Kenneth D. West (1987). Andrew Chesher & Ian Jewitt (1987) observe that even these standard errors may be seriously underestimated when there are points of high leverage in a regression if heteroskedasticity is "moderate". Since we cannot rule out "moderate" heteroskedasticity, and will even remove some sources of heteroskedasticity by the use of Weighted Least Squares (WLS), we both calculate the likely downward bias of the standard errors and consider regressions using only subsets of the data, where high leverage is less likely.

3. Results

In this section we discuss our empirical results. In Section 3.1, we present the estimates of the gross return when we use the larger month-dated data set. In Section 3.2, we analyse the small data set and use the technique suggested by Chesher (1998) to remove any bias arising from errors in variables.

The data was analysed using OX 2.00 (see Jurgen A. Doornik, 1998). Our notation is as follows: HCSE is the heteroskedastic consistent standard error due to Halbert White (1980); HACSE is the heteroskedastic and autocorrelation consistent standard error due Newey & West (1987); N is the total number of observations; σ^2 is the equation standard error; DW is the Durbin-Watson test for first order residual correlation; C-J bias is the lower bound on the

[6] In many of the regressions corr(p^{TR},m) = – 0.05 and the average correlation between village dummies and p^{TR} is 0.005.

bias of the HCSE, when we assume that the ratio of the smallest to largest error variance is two (which is Chesher & Jewitt's (1987) definition of "moderate heteroskedasticity").

3.1 Results from Wheat Data

We consider initially four types of regression using month-dated data, in each case regressing dated grain prices on various explanatory variables.

1. OLS regression where the explained variable is the set of all the dated grain prices: observations within the same village month are not pooled. For these regressions we provide HCSEs: because some observations are contemporaneous, the HACSEs are not always well defined, since the HAC Moment Estimator need not be positive definite.

2. OLS regressions where contemporaneous observations from the same village have been combined in a weighted average, using the quantities as weights (we also analysed simple averages, but found no substantial difference). The HACSEs are always well defined in this case.

3. WLS estimates, where the weights are respectively the number of prices making each observation and the squared logarithm of the total quantity of the grain transacted. Again the HACSEs are always well defined.

Our first set of regressions have p^{TR} and twelve month dummies for explanatory variables. Descriptions of the regressions are shown in Table 2. The parameter estimates for the 12 month dummies are shown in Figures 1 and 2, with 5 per cent confidence intervals taken from either the HCSEs or HACSEs, depending on whether the latter are available. The graphs are normalised so that January's dummy is zero.

It is clear that all four regressions lead to the same conclusions. There is a substantial drop in prices between July and August and a steady rise thereafter, although the drop is not statistically significant. These findings contradict McCloskey & Nash in three ways:

1. The phase of the wheat cycle is different, since prices fall in August, not September. Since harvesting may have started as early as late July and was well under way in August, wheat from the new harvest would be available in August, albeit at the cost of diverting labour from the relatively urgent harvest to threshing. There is very weak evidence that prices fell further in September when the opportunity cost of labour would be lower.

2. There is little if any evidence for concavity in the curve. Since McCloskey & Nash mis-measured the phase of the cycle, however, this is not entirely surprising. The estimates in their Tables 2 and 3 are averages of price rises and price falls. The relative weight of the price fall in August becomes larger as the period under consideration becomes smaller, therefore biasing the result towards concavity.

3. The gross return is only in the region of 10 – 15 per cent (and is not statistically significant when estimated via 12 dummies). This lower estimate is unsurprising given the discussion in Section 1.

We re-estimated all of the equations with dummies for each village, with virtually no change in the pattern of our point estimates, although the HACSEs fell to less than 0.07. This latter result, however, was due to the imbalanced nature of the regressions: using the Chesher-Jewitt formula we found that the HACSEs could be biased down by as much as 76 per cent. We then dispensed with p^{TR} and controlled for village and year differences using dummies for each village year. Again the gross return was no more than 17 per cent: inclusion of the village dummies resulted in unreliable HACSEs.

We then simplified the analysis by replacing the month dummies with a trend taking the value 0 in August and incremented by 1 each month. We report the implied gross returns and their confidence intervals in Table 3 together with the equation diagnostics. The analysis is by OLS with contemporary data averaged and HACSEs for the confidence intervals: the gross return is in the region 17 – 18 per cent and is statistically significantly different from both zero and 30 per cent. We favour the wider confidence intervals of the regression using p^{TR} instead of the ones obtained from using village-year dummies because there is evidence that the latter are biased.

To check the robustness of our findings we duplicated the analysis using WLS: estimates were very similar. Attempts to include quadratic or cubic trends to test for concavity were unsuccessful, since the quadratic and cubic terms were always insignificant.

Our final analysis using month dated data was to look at three subsets of the data, reported in Table 4: Cuxham is that village on its own, XX is all of the other villages for which the number of price observations is greater than 20 and "Other" is the remainder. The evidence from Cuxham and the XX villages suggests a slightly higher gross return, whereas the "Other" villages suggest a lower return. The latter sub-set of the data consists of a large number of villages for which there are only a few observations: this sub-set is much more heterogeneous and the confidence intervals are correspondingly larger.

4.2 Daily-dated Data and Errors in Variables

In this section we discuss the extent to which measurement error in the date will result in biases in the parameter of interest. First, we show why the bias will be small. Second, we re-estimate the regression for the data set with daily dates using the technique suggested by Chesher to remove the bias.

A simple calculation shows that the bias in the estimate of the gross return will be small when we use the data set with monthly data. Suppose that the data had been dispersed in a fairly extreme fashion, so that all of the observations in the first half of the year were actually from the month before that in which they had been recorded and all of the data in the second half were actually from the month after that in which they had been recorded. Assuming that dates are never mis-measured by more than one month, this will lead to the maximum downward bias in our estimate of the gross return. We can obtain the true gross return by re-allocating price observations to the hypothesised true date. In the figure below we show the bias under the hypothesis that the dependent variable is non-stochastic (we use the empirical distribution of the wheat price dates):

Figures are percentages				
Estimated gross return	13	27	43	62
Actual gross return	14	30	48	68

When the estimated gross return is low, the bias is small, despite the fact that the degree of dispersion we have assumed might be considered exaggerated.

A more sophisticated attempt to correct for this problem is to use the procedure suggested by Chesher (1998). This procedure involves assuming the functional form for the relationship between the variables and then uses non-parametric methods to correct for measurement error by looking for departures from this functional form. In the circumstances, we do not consider the imposition of a linear functional form *a priori* to be too onerous. This approach, however, relies upon the explanatory variable being (approximately) continuous, so we cannot use data grouped in months and use just the daily-dated data.

Chesher's technique involves creating a new variable, g(x), which will be a function of the mismeasured explanatory variable of interest, in this case x, defined as follows:

$$g(x) \equiv \frac{d \ln q(x)}{dx}$$

where θ is the marginal density function of x. We obtain $\hat{g}_X(x)$ by use of a kernel instead of the method proposed by Chesher, who suggests approximating the log-density using a sum of Legendre polynomials. The reason for our decision can be seen from Figure 3, which is the empirical density: this is characterised by six very large spikes which would have only been approximated by a very large number of polynomials. We used a wide bandwidth binomial kernel and discarded from our analysis the first 45 days of the year: these were where the kernel estimate was least satisfactory and also where the noise in the price series and departures from our chosen functional form would be most serious.

Having estimated the log-density, we estimate $\gamma_{(x)}$ through the regression:

$$p_{v.t.x.i} = \alpha\, p_t^{TR} + \kappa_v + \gamma_{(x)}x + \lambda\, \hat{g}_X(x) + u_{t.v.x} + e_i$$

In the regression, λ is an estimate of the variance of the unobserved error in the measurement of x, so, if it is insignificant, we can reject the hypothesis that measurement error matters. The distribution of λ is not fully understood, however, so we do not place much weight on this result.

Estimates of this equation are shown in Table 5. The notable feature is that the estimate of the gross return on wheat is almost exactly the same as in our estimates in the previous section, suggesting that errors in variables is not causing any downward bias. Using the HCSE, λ is insignificant, which is consistent with measurement error being unimportant.

4.3 Barley

Barley was the most important crop after wheat in the medieval period and barley prices are therefore the second most abundant source of price data. Barley prices are potentially interesting for two further reasons. First, it has been shown by B. Taub (1988) that under certain conditions it is possible to use the prices of two grains together to separate the cost of capital from the cost of storage and depreciation (See Appendix A.2). Since McCloskey and Nash do not give any direct evidence on storage costs or depreciation in their original analysis, the Taub approach could prove useful. Second, barley prices provide an independent check on our method because we can replicate our wheat analysis using barley prices and we would hope to get similar results (with barley showing a slightly higher gross rate of return).

Typical results of our analysis of barley prices are presented in Table 6. Two points are worth noting.

First, there are too few observations in August, September and October to estimate these month dummies and we decided to discard these observations. This means that it is very difficult to be sure at which point barley prices actually fell: our preference was to assume that they were flat during November and December and then rose, so we used a trend with two initial zero points to reflect this. Using alternative specifications did not lead to dramatically different results.

Second, our estimates of the gross return to barley were much more sensitive to the precise way that we modelled year-on-year and village effects, perhaps unsurprisingly given the small samples involved. Thus the two regressions that we report suggest returns of 26 per cent and 38 per cent respectively. We also divided the data into sub-sets based on village (Elham (62 observations), Oxford (44 observations) and all the rest), with results shown in Table 7. Again there is wide variation in our estimates of the gross return.

We conclude that the gross return on barley does appear to be higher than that for wheat, but, compared with wheat, our estimates are both considerably less precise and less robust to alternative ways of modelling the other causes of price variation.

Although the application of our econometric model to barley prices is instructive, the precision with which we can estimate the gross return is too low to implement Taub's approach with any degree of confidence. It is true that the gross return on barley does appear to be higher than for wheat, consistent with Taub's model and this is prima facie evidence for storage costs being properly reflected in barley and wheat prices.[7] But the confidence intervals around the appreciation rates of wheat and barley are too large compared to the difference between them to isolate the storage cost with any degree of precision.

In fact, we will not pursue the issue of barley prices and storage costs for two further reasons. First, Taub's model is predicated on the assumption that wheat and barley were in competition for storage. But the different production patterns for wheat and barley means that they were probably not in competition for much of the year.[8] Second, and most important, historians have recently presented direct evidence on storage costs and wheat depreciation – thereby allowing us to separate out directly the cost of capital. We now know that the rate of depreciation of wheat was around 3 per cent per annum; and storage costs were in the region of 2.5 per cent per annum.

5. Discussion

We have now established that the gross rate of return on wheat in medieval England was about 18 per cent. Campbell, Galloway, Keene & Murphy (1993) provide direct evidence to suggest that storage costs were in the order of 2.5 per cent of the value of wheat; and depreciation losses were in the region of 3 per cent: this depreciation figure is also consistent with the contemporaneous figure implicit in Walter of Henley (see *e.g.*, Dorothea Oschinsky, 1971). Hence the cost of capital was actually around 12 per cent per annum. We now consider the plausibility of our estimates and discuss the implications for grain storage.

Whilst 12 per cent is substantially above the levels currently experienced in the OECD, it is not exceptionally high. In fact, 12 per cent may seem surprisingly low. After all, the investible surplus generated by the medieval economy was probably small and farmers in LDCs today show widespread evidence of being capital-constrained. But if we consider the supply and demand for capital in the medieval economy then we will see that 12 per cent is a plausible figure for two reasons.

First, it may be correct that the medieval economy generated only a small investible surplus (although it is difficult to be sure, given the current state of knowledge). This would obviously reduce the supply of capital, ceteris paribus. The demand for capital, however, was probably also low. Farmers in LDCs currently have a high marginal productivity of capital because they can import technology from developed countries (and most technology imports are capital-intensive). But it is far from obvious that medieval farmers had any investments open to them

[7] The absolute cost of storage was the same for a given volume of wheat or barley – so, given wheat generally cost twice as much as barley, the percentage increase in the barley price would have to be twice as high as wheat to make the same absolute increase.

[8] Barley was planted in the spring, whereas wheat was planted in October: hence barley was harvested rather later than wheat in the autumn. Also barley was processed at a later date, because it was a lower value crop and animals preferred to eat the straw from wheat before they ate the straw from barley (which meant the wheat had to be threshed first). Finally, it is not clear to what extent barley was stored as grain rather than as malt or beer.

which had a high marginal productivity. We conclude that both the demand and supply of capital were low, which is obviously consistent with a low interest rate.

Second, even if there were a shortage of capital in the medieval economy it is unlikely that this would be reflected in high interest rates. It was more likely to be reflected in low interest rates and capital-rationing. This is partly because there were strict laws against usury (charging either excessive rates of interest or any interest at all). And it is partly because capital-rationing is a well-documented response to excess demand for credit, particularly where farmers have few assets to guarantee a loan such as in developing countries.[9]

A lower interest rate has several implications for grain storage. Farmers' cooperatives would be a feasible insurance strategy for some but not all levels of harvest variability (Kimball, 1988). With a discount rate of 12 per cent, cooperatives could emerge where harvest variation was high (0.6); but strip scattering would be used where variability was too low to support cooperation (0.3). This is an important result because strip scattering was not adopted universally (Fenoaltea, 1988). Given a discount rate of 12 per cent, it is plausible to suggest that temporal and spatial variations in the pattern of strip scattering were caused by different degrees of harvest variability.

McCloskey and Nash imply that storage might also become a feasible insurance possibility at a lower interest rate. Ironically, without a more detailed model of storage, it is rather unclear whether the interest rate was actually the most important factor regulating the level of storage. The decision to carry grain over into the next harvest year does not depend solely on this year's rate of interest, but also on the likely price at which one can sell next year compared to this year. If this year's harvest was very good (*i.e.,* the circumstances under which storage is most likely), price levels would be low throughout the year – even if the rate of appreciation (interest rate) was the same as usual. If next year were bad or average then price levels might return to normal and there would be no price drop in August-September. On the other hand, when prices were high and will therefore be likely to fall next year, the shortage of grain would mean that there would be very little scope to carry over any grain anyway. It follows that grain storage across harvest years could be perfectly rational, even with risk-neutrality and high gross returns within the year.

---

[9] See, for example, Mark R. Rosenzweig & Kenneth I. Wolpin (1993).

6. Conclusion

There are five conclusions to be drawn from our analysis:

First, an important contribution of this paper lies in deriving a more efficient means of estimating the interest rate from grain price observations within the year. Grain prices are one of the most widely available data series in both historical sources and developing countries. So where direct evidence on interest rates is missing, our method offers the potential for estimating both long time series and broad cross-sections.

Second, the gross rate of return on grain in medieval England was around 18 per cent. Direct evidence suggests that storage costs were in the order of 2.5 per cent of the value of wheat; and depreciation losses were in the region of 3 per cent. Hence the cost of capital was actually around 12 per cent per annum. Whilst this is substantially above the levels currently experienced in the OECD, it is not exceptionally high.

Third, we can reject the Komlos and Landes hypothesis that medieval farmers were not rational profit-maximisers. They argue that there were few alternative investments to holding grain and that, in any case, medieval farmers would not arbitrage between holding grain and other investments. If this were true then presumably the rate of interest on grain holdings would be close to zero. Since the 18 per cent gross rate of interest is significantly different from zero, we can reject the Komlos and Landes hypothesis. This is not very surprising because Komlos and Landes implicitly take the view that peasants were responsible for most grain storage in the medieval economy. In fact, medieval agriculture was dominated by religious institutions and the whole debate has revolved around data drawn from the accounts of manors owned by the Oxford colleges. These manors were supervised by the most educated minority in the country (the Cambridge colleges had yet to be founded) who also provided the civil service for the monarch. It seems likely that they were capable of recognising and exploiting seasonal patterns in grain prices.

Fourth, year-on-year grain storage was not prohibitively expensive in the medieval period due to the cost of capital, contrary to the claims of McCloskey and Nash. The 18 per cent gross rate of interest which we have estimated is significantly below the 30 per cent or more postulated by McCloskey and Nash. The cost of year-on-year storage no longer seems to have been excessive, although it is difficult to be precise because McCloskey and Nash do not postulate a threshold value below which storage would have become economical.

Fifth, Kimball shows that a 12 per cent discount rate was low enough for farmers to want to cooperate and provide mutual insurance under high levels of crop variability. But under low levels of crop variability, strip scattering would be a less costly method of insurance. This could explain the puzzling fact that strip scattering was restricted to particular times and places. This generates an interesting and testable hypothesis which is the subject of ongoing research.

Appendices

A.1 Discussion of the Data

Following McCloskey & Nash, we use the data compiled by Thorold Rogers (1866). In this section we begin by discussing the nature and reliability of the data, including its transcription and translation. Of greatest concern is the accuracy and precision of the data. We then discuss several other troublesome econometric characteristics of the data.

The data were compiled by analysing accounts of manors to their owners. Most of these were Oxford Colleges: the majority of the data that we used come from manors owned by just one college, Merton.

A typical piece of data is that a quantity of 2 quarters of wheat was sold by the manor in Marlborough for 5 shillings per quarter in October 1282. Combined with the fact that 1 quarter and 5 bushels of wheat was sold in the same place at 5s. 8d. in December, one might infer that prices rose by about 12.5 per cent in those three months.[10] In isolation these two pieces of data would suggest that the gross return on holding grain is 50 per cent per year.

For most of the transactions in Thorold Rogers, the month is not recorded, so it is impossible to make any inferences about seasonal price changes. To make a valid price change observation we need at least two dated prices within the same village (variation between villages appears substantial and we know that variation between years is massive from the average annual price series) and this reduces the number of usable prices in Thorold Rogers considerably.

Some of the data have a relatively vague date: we delete these data. Some data have a date of the form: Feb 2 – Mar 31. We would delete this observation. However, if the range of the date were the last or first date of the month, we might retain it: for example we assign "Nov.30–Xtmas" to December. We always assume that "Xtmas" is Dec 25, although both "Xtmas" and "Dec 25" appear in Thorold Rogers.

If we continue to look at the data for Marlborough, we may note that these are the only usable observations in the agricultural year 1282–3: there are five usable prices[11] (in April, Nov, Dec, Mar, May) for 1281–2 and no usable prices at all for 1280–81 and 1283–4. Thus we are estimating time series properties of data which in no way resembles a normal time series: some

[10] To avoid creating an upward bias in estimates of price changes, the estimate of the interest rate should be based upon the difference in the logarithm of the prices and not the percentage change. We use the log method throughout this paper.

[11] It is apparent that the April price is entered out of order in Thorold Rogers, almost certainly indicative of being entered out of order in the original accounts. In some instance we might be concerned that this indicated that the accounts ran for a period longer than a year and that the April figure was thus for the previous year. This is unlikely to be so here, since Thorold Rogers does provide the accounts for Marlborough in 1280–81, although none of the prices are dated.

of the price pairs are so far separated both spatially and temporally that we could effectively treat them as cross section observations.

There are 203 village years for which two or more wheat prices are recorded. Of these the most significant Cuxham (pronounced to rhyme with “books ‘em”), a village in Oxfordshire owned by Merton College, Oxford. The relative importance of Cuxham is larger than this, however, because it frequently has quite a large number of prices within a year: note that Cuxham wheat data alone (382 observations from 260 distinct months) is bigger than our entire database on barley. The detailed nature of the accounts and other records also provide considerable background information, discussed in detail in Harvey (1965).

We now turn to the quality of the data, especially Cuxham. Thorold Rogers believed that the Cuxham data were superior to those of other accounts.[12] During the time in which we are most interested, Merton’s holdings were run by two reeves, Robert Beneyt (1288–1311) and Robert Oldman (1311–1349), whose length of service is *prima facie* evidence of their reliability. After the Black Death, however, Merton had great difficulty in obtaining reliable managers and eventually ran the estate directly. Thus the quality and quantity of accounts diminishes after 1349 and we have conducted tests without these data on the assumption that they may be less reliable.

We were also concerned with the quality of Thorold Roger’s transcription, due to a contradiction in his data. Thorold Rogers provides the bailiff’s account for Cuxham for 1316–17 as a specimen of the accounts that he used as sources. On vol 2, page 617 we can find the paragraph about wheat sales, which should correspond to the entry on page 74 in his table of grain prices. The relevant entries are reproduced in Table 8. It transpires that this particular year has examples of every sort of problem in the data.

Consider first the precision of the date. Many dates consist of a month alone, *e.g.*, in Table 8, the reference “July”. As can be seen from the original accounts, the actual date was Thursday before 1 Aug, which we know is 29 July.[13] Since we do not have access to the original accounts, we take Thorold Rogers’ data at face value and consider two ways of using the data: first we use all of the data dated to a month; second, we use the subset of data for which the day is available.

Second, there is some mismeasurement in the date: for example, the entry for Nov 1, clearly reads *before* All Saints’ Day, suggesting that the sale took place in *October* and not November. This problem is acute because some of the commonest dates are near the beginning or end of the month: 2 Feb (Candlemass); 1 May (SS Philip & James); 29 June (SS Peter & Paul); 29 Sep (Michaelmas); 1 Aug (Lammas/St Peter “in chains”); 1 Nov (All Saints’); 30 Nov (St Andrew).

[12] “Out of the many thousand accounts which I have investigated none equal those of Cuxham for intelligence, accuracy and order.” Thorold Rogers (1860, vol.1, p.508.n).
[13]We have obtained relevant dates from the Perpetual Calendar in Harvey (1932), Appendix IV.

Finally there is the concern that Thorold Rogers made major errors in transcription, especially of the date: since this is an explanatory variable, mismeasurement may cause biased estimates of the parameter of interest. We were particularly concerned with the third entry in Table 8, since it appears to confuse St. Nicholas with St. Matthias. We have checked the original account roll and the relevant entry is unclear. We also checked all of Thorold Rogers' figures with Harvey's independent reading of the same account rolls: some of these are published in Harvey (1976) and Harvey kindly made his notes for all of the other years available to us. With the exception of the discrepancy already noted, the two readings agree exactly and we decided to proceed on the assumption that there were no further errors of this kind.

We now outline the main features of the data. The logarithm of annual wheat and barley prices are shown in Figure 4: this data is taken from Thorold Rogers (1866, vol.I, pp226.ff). It can be seen that there is no trend in the prices over the period concerned. Barley is always less valuable than wheat. The logarithm of the relative price is plotted in Figure 5 and the relative price of barley to wheat is plotted against the wheat price in Figure 6. There appears to be little change in the relative price when wheat prices are high or *vice versa* and the years when wheat prices are very high do not display a significant change in the relative price. We accordingly treat the two series as separate.

We now turn to the prices which are dated within the year. It is always difficult to display such data visually. Figure 7 shows these data for wheat where the ordering of the data is: year; within year, village; within village, month of the year.

We now turn to the pattern of sales within the year. The proportion of wheat sales for the whole data set and several subsamples is shown in Figure 8. January, August and September are poorly represented except from the village of Cuxham. For this reason, we sometimes dispense with these months' observations.

The locations of all the villages that we were able to identify with certainty are shown in Figure 9. Most of the villages are from the South of England. The omission of the North of England is unfortunate, but the proximity of the villages does mean that we might expect the data to be more homogeneous. We are less concerned about the absence of the West Country, since the two and three field systems were much less prevalent there at this time.

We consider several sub-sets of data, which are detailed in the main text. This results in more homogeneous data; correspondingly it provides an idea of whether parameter non-constancy is a problem. Consider the wheat data set, for which there are 822 observations drawn from 682 distinct months. Of these 122 come from 40 villages from which the total number of observations is less than ten, while Cuxham alone has 382 observations and Cheddington, Elham, Farley, Ibstone, Letherhead, Oxford and Wolford between them have a further 270 observations: note that all of these villages are Merton owned except Wolford.[14] Any regression with all the data will have almost half of its observations from Cuxham: these observations might have a lower variance, thus introducing heteroskedasticity.

Another consideration is the definition of the year. The English agricultural year began on 29th September (Michaelmas), but not all accounts follow this convention. Where accounts do follow this convention, it is sometimes unclear whether "Mich" refers to an observation at the beginning or the end of the year, and we therefore discard it. Cuxham, Elham, Farley, Ibstone, Letherhead, Oxford and Stratton all usually began their accounting years on 25th July (St James' Day) (but sometimes on the 22nd of July) and Basingstoke, Cambridge, Cheddington and Wolford begin on 7th July. These villages (nearly all owned by Merton) account for the largest part of the data. Because of this, we follow a convention of beginning all years on 1st August: we check every July figure for the above accounts to see whether it is at the end or the beginning of the year: for all of the other accounts we move the August/September observations to the preceding year. This occasionally results in the loss of data for these months, which is unfortunate, since trades at this time of the year were relatively scarce.

The effect of using an accounting year based on September to August (the method used by McCloskey & Nash) is equally problematic, since it loses a lot of observations for the manors named above, where the data are most comprehensive and probably most accurate. The observations which are not lost by this method tend to come from villages on which we have very little data.

To see the effects of this, consider barley:

There are a little less than 500 barley prices that are dated within the year. Of these only 446 satisfy the requirement that there is more than one price within the year. As with the wheat prices, some of the dates are too vague to be of much use. Furthermore, barley prices consist of barley, drage, barley malt, drage malt; there are also prices of old barley. We discarded the latter four types of data to retain comparability, reducing the data set to 302 observations. Of these, we discarded one September observation (Michaelmas, Bosham 1308) because it was unclear which end of the year it belonged to, one August observation (Dengmarsh, 1369), because it belonged to August 1370 instead of 1369 and the data from Oxford 1294/5 because the July was in 1294 instead of 1295. There remained three potentially ambiguous July observations (Elham 1343/4, 1344/5; Wolford 1346/7); however, these were all from accounts where all figures are always in date order, so that we were satisfied that they belonged to the July at the end of the year and not the beginning.

The result was that we had 296 prices drawn from 22 villages, with 57 village-years. 100 prices were from Elham, 76 from Oxford; no other village had more than 14 prices. Some of these

[14] An overview of Merton's interaction with its manors can be found in available in Roger L. Highfield & Geoffrey E. Martin (1997).

prices were from the same month, however, so we averaged the prices within month (using both a simple average and a quantity-weighted average). This reduced the number of observations to 66 prices from Elham, 42 from Oxford and 205 in total.

A.2 Barley and Wheat Prices

Taub (1988) has noted that it may be possible to interest rates from gross returns. We find Taub's terminology confusing and so distinguish here two factors determining the gross return:

1. Depreciation. We use this term to cover all forms of cost which can be related directly to the value of the grain. Thus if the grain rots or is eaten by mice then the value of the loss per initial quarter of grain is δp.

2. Storage Costs. We use this term to denote costs which relate directly to the volume (and hence, ignoring differences in density, the weight) of grain. Thus let S be the storage cost of storing one quarter of any grain.

Now suppose that two grains are in competition for storage, so that they pay the same price S (if the grains were not in competition, then S would change over the year as more or less capacity was available). Ignoring any risk premia, the return to holding one quarter of grain of type k is then

$$r_k = \frac{\Delta P_k - dP_k - S}{P_k}$$

where ΔP is the change in price. Assuming arbitrage, the return will be equal for all grains j, k, so one can always solve for S.

Note that $\Delta P_k / P_k = \Delta P_j / P_j$ will only occur if

$$\delta_k - SP_k = \delta_j - SP_j$$

Since P_j and P_k always differ, this would occur only if $S = \delta_k - \delta_j = 0$ or the difference between δ_k and δ_j exactly offset the difference in prices. Neither of these conditions can be imposed *a priori*, so it is unreasonable to assume that the gross return on different grains will be the same. For this reason we do not pool different grain prices.

A.3 A Critique of Clark

As noted above, we only obtained Clark's undated mimeo (which we believe to date from mid 1998) as we finished this paper. In the main body of the paper and the Appendices we resolve several problems which arise in obtaining an unbiased estimate of both the gross return on holding grain and the precision with which it is estimated. Clark does not address these problems, confining himself to OLS estimation, so, although his results are superficially similar to ours, we suggest that his results are less robust and less comprehensive. We use this brief appendix to highlight his paper's flaws, with which we have dealt satisfactorily.

1. Estimation of confidence intervals. Clark does not say how his confidence intervals are obtained, but his estimates are very close to the simple estimates that we obtain by OLS. We find comprehensive evidence of both heteroskedasticity and autocorrelation and we appropriately correct for both. Furthermore, Clark is in error when he says that contemporaneous observations should not be given additional weight: means of contemporaneous or near contemporaneous observations have a smaller variance about the unobserved average market price and should be given greater weight through WLS to ensure that additional heteroskedasticity is not introduced.

2. Clark is wrong to pool wheat and barley prices. Following Taub's (1988) analysis, we demonstrate in Appendix A.2 that imposing the same gross return on different grains is equivalent to making assumptions which are unmerited *a priori*. In fact our empirical analysis shows that wheat and barley prices differ not just in their average level but both their phase and amplitude of cycle; Clark only deals with the first difference. From additional analysis that we conducted, it appears that the differences in phase and amplitude nearly cancel out, so that regressing barley on a linear trend starting in August instead of December provides a slope estimate similar to the true slope for wheat.

3. Clark's method of adjusting for year-to-year fluctuations, i.e., by using $p_{i.v.t.m}/p^{TR}$ is equivalent to imposing a unit elasticity on p^{TR} in his regressions. Although the elasticity of $p_{i.v.t.m}$ with respect to the true market price, p*, will be unity, our measures of p* are only imperfect proxies (*i.e*. they contain measurement error), so such a procedure risks biasing the estimates of the gross return.

4. Clark does not discuss the issue of downward bias in estimates of the gross return arising from mismeasurement in the dates. Since the issue is whether McCloskey & Nash have over-estimated the gross return, it is necessary to show that contrary results are not due to such bias: we have done this in Section 4.2

References

Beveridge, William, H. *Prices and wages in England from the twelfth to the nineteenth century*. (London, 1939).

Campbell, B M.S.; Galloway, J; Keene, D and Murphy, M. "A Medieval Capital and its Grain Supply: Agrarian production and Distribution in the London Region c.1300." Historical Geography Research Series, 30, 1993.

Chesher, Andrew “Measurement error bias reduction” University of Bristol Discussion Paper 98/449, 1998.

Chesher, Andrew & Jewitt, Ian. “The bias of a heteroskedasticity consistent covariance matrix estimator” *Econometrica, 55,5* September, 1987, pp.1217–1222.

Clark, Gregory “The cost of capital and medieval agricultural technique” *Explorations in Economic History, 25*, July, 1988, pp.265 – 94.

Clark, Gregory “A precocious infant? The evolution of the English grain market, 1208 – 1770” mimeo (undated), University of California in Davis.

Doornik, Jurgen A. *OX: An Object-Oriented Matrix Programming Language* (Kent: Timberlake Consultants Ltd, 1988) `http://www.nuff.ox.ac.uk/Users/Doornik/`.

Fenoaltea, Stefano. “Transaction Costs, Whig History and the Common Fields.” *Politics & Society, 16*, 1988.

Greif, Avner. “Contract Enforceability and Economic Institutions in Early Trade: The Maghribi Traders’ Coalition.” *American Economic Review, 83, 1* (1993), pp.525 – 548.

Greif, Avner. “Cultural Beliefs and the Organization of Society: a Historical and Theoretical Reflection on Collectivist and Individualist Societies.” *Journal of Political Economy, 102, 5* (1994), pp.912 – 950.

Hamilton, Earl J. *American treasure and the price revolution in Spain, 1501-1650* (New York : Octagon Books, 1965).

Harvey, P.D.A. *A Medieval Oxfordshire Village: Cuxham 1240 to 1400* (Oxford: Oxford University Press, 1965).

Harvey, P.D.A. (ed.) *Manorial Records of Cuxham, Oxfordshire, circa 1200 – 1359* Historical Manuscripts Commission JP 23 (London: Her Majesty’s Stationery Office, 1976).

Harvey, Sir Paul. *The Oxford Companion to English Literature* (Oxford: Oxford University Press, 1932).

Highfield, Roger L. & Martin, Geoffrey E. *The History of Merton College* (1997)

Kimball, Miles S. “Farmers’ cooperatives as behavior toward risk” *American Economic Review, 78, 1* March, 1988, pp.224 – 232.

McCloskey, Donald, N. "The prudent peasant: new findings on open fields" *Journal of Economic History, 51,* 1976, pp. 343 – 355.

McCloskey, Donald, N. & Nash, John "Corn at interest: the extent and cost of grain storage in medieval England" *American Economic Review, 74, 1* March, 1984, pp.174 – 187.

Newey, Whitney K. & West, Kenneth D. "A simple, positive definite, heteroskedasticity and autocorrelation consistent covariance matrix" *Econometrica, 55,3* May, 1987, pp.703 – 708.

Oschinsky, Dorothea (ed.) *Walter of Henley and other treatises on estate management and accounting* (Oxford : Clarendon Press, 1971).

Posthumus, Nicholas W. *Inquiry into the history of prices in Holland*, (Leiden, 1964).

Rosenzweig, Mark R. and Wolpin, Kenneth I. "Credit Market Constraints, Consumption Smooting, and the Accumulation of Durable Production Assets in Low-Income Countries: Investment in Bullocks in India." *Journal of Political Economy, 101, 2,* 1993, pp.223 – 244.

Taub, B. "A model of medieval grain prices: comment" *American Economic Review, 77, 5* December, 1987, pp.1048 – 1053.

Thorold Rogers, James E. *A History of Agriculture and Prices in England* (Oxford: Oxford University Press, 1866).

White, Halbert "A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity" *Econometrica, 48,4* May, 1980, pp.817 – 838.

Table 1 – English Wheat Prices: Percentage Rates of Change per Month

Quantity Weighting

		Sep	Oct	Nov	Dec	Jan	Feb	Mar	Apr	May	June	July
	Aug	-1.5	-8.6	3.6	4.4	-0.6	0.7	1.0	1.2	2.7	1.1	0.6
	Sep		12.3	7.5	0.7	1.6	2.0	1.8	2.4	4.8	2.3	2.1
	Oct			5.1	1.1	0.1	0.2	-0.4	0.6	1.0	0.4	3.6
	Nov				5.0	1.9	1.3	1.1	1.6	0.3	0.4	1.2
From	Dec					3.7	3.4	2.0	2.2	1.4	2.2	0.5
	Jan						2.7	2.3	3.1	0.7	1.1	-0.7
	Feb							2.4	3.2	2.9	1.5	2.6
	Mar								2.0	2.8	1.7	2.2
	Apr									0.9	-0.3	1.8
	May										1.3	0.6
	June											3.1

Simple Weighting

		Sep	Oct	Nov	Dec	Jan	Feb	Mar	Apr	May	June	July
	Aug	-0.3	-8.4	3.7	4.4	-0.6	0.8	1.0	1.3	2.7	1.0	0.6
	Sep		10.7	7.5	0.9	1.7	1.7	1.7	2.5	4.7	2.0	2.0
	Oct			4.6	1.2	0.1	0.4	-0.5	0.6	1.0	0.3	3.6
	Nov				5.0	2.0	1.4	1.1	1.7	0.3	0.3	1.2
From	Dec					3.7	3.4	2.0	2.2	1.4	2.2	0.5
	Jan						2.7	2.4	3.4	0.7	1.0	-0.7
	Feb							2.3	3.3	2.9	1.4	2.6
	Mar								2.1	2.7	1.6	2.1
	Apr									0.9	-0.6	1.7
	May										1.1	0.6
	June											3.5

Number of Observations

		Sep	Oct	Nov	Dec	Jan	Feb	Mar	Apr	May	June	July
	Aug	9	4	10	6	10	9	8	9	7	11	8
	Sep		6	6	6	4	6	6	9	9	10	9
	Oct			7	9	5	8	8	8	11	10	5
	Nov				20	9	16	25	23	20	23	17
From	Dec					8	20	22	28	23	20	13
	Jan						16	13	12	13	15	9
	Feb							33	31	27	26	17
	Mar								45	31	41	25
	Apr									35	42	26
	May										47	25
	June											33

The total number of price pair observations is 1082, based upon 681 price observations. These 681 price observations are taken from 874 individual prices.

Source: Authors' Calculations based on data taken from Thorold Rogers.

Table 2

Statistics Describing the Regressions of lnP on lnP^{TR} and month dummies
(Accompanies Figures 1 and 2)

Wheat

Regression	OLS(i)	OLS(ii)	WLS(n)	WLS(q)
N	874	681	681	681
σ^2	0.18	0.19	0.21	0.38
DW	0.91	0.56	0.21	0.99
CJ	-0.084	-0.11	-0.32	-0.42

Barley

Regression	OLS(i)	OLS(ii)	WLS(n)	WLS(q)
N	286	204	204	204
σ^2	0.19	0.19	0.50	0.22
DW	0.83	0.78	0.72	0.69
CJ	-0.18	-0.26	-0.45	-0.51

Table 3

GROSS HOLDING RETURN FOR WHEAT

Regression is
$\ln P = \alpha \ln P^{TR} + \text{constant} + \gamma \text{Augtrend}$

Data is month-dated.

Variable	Parameter	HACSE
constant	-0.108	0.147
$\ln P^*$	0.994	0.0344
Augtrend	0.0143	0.0032

N	681
σ^2	0.19
CJ	-0.03

Implied gross return on annual basis is 18.76%
5% confidence interval is 9.51%% to 28.24%

Regression is
$\ln P$ = Village-Year Dummies (203) + γAugtrend

Data is month-dated.

Variable	Parameter	HACSE
Augtrend	0.0133	0.0033

N	681
σ^2	0.15
CJ	-1

Implied gross return on annual basis is 17.3%
5% confidence interval is 7.90%% to 23.93%

Table 4

GROSS HOLDING RETURN FOR WHEAT: SUBSETS OF DATA

Regression is
$\ln P = \alpha \ln P^* + \text{constant} + \gamma \text{Augtrend}$
Data is month-dated.

Results are gross return expressed at an annual rate

	Estimate	Confidence interval
Cuxham	20.2%	9.1% to 32.4%
XX	22.1%	12.3% to 32.6%
Other	13.0%	-0.4% to 33.5%

Table 5

RESULTS OBTAINED WHEN CORRECTING FOR BIAS USING CHESHER [1998]

Daily-dated Wheat data.

Without Chesher correction

Variable	Parameter	HCSE
constant	-0.150	0.131
ln*P**	1.01	0.031
Augtrend	0.00040	0.00011

N	424
σ^2	0.18
CJ	-0.06

Implied gross return on annual basis is 15.74%
5% confidence interval is 6.77% to 22.68%

With Chesher correction

Variable	Parameter	HCSE
constant	-0.183	0.148
ln*P**	1.01	0.031
ghat	-0.0052	0.011
Augtrend	0.00041	0.00011

N	424
σ^2	0.18
CJ	-0.063

Implied gross return on annual basis is 16.24%
5% confidence interval is 7.09% to 26.19%

ghat is insignificant.

Table 6

GROSS HOLDING RETURN FOR BARLEY

Regression is
$\ln P = \alpha \ln P^* + \text{constant} + \gamma\text{Dectrend}$

Data is month-dated, with all Aug-Oct observations missing

Variable	Parameter	HACSE
constant	0.0002	0.051
ln*P**	0.967	0.20
Dectrend	0.0193	0.0087

N	204
σ^2	0.19
CJ	-0.12

Implied gross return on annual basis is 26.10%
5% confidence interval is 2.24%% to 55.47%

Regression is
$\ln P = \text{Village-Year Dummies (62)} + \gamma\text{Dectrend}$

Data is month-dated.

Variable	Parameter	HACSE
Dectrend	0.027	0.0084

N	204
σ^2	0.15
CJ	-0.76

Implied gross return on annual basis is 38.38%
5% confidence interval is 12.34% to 69.25%

Table 7

GROSS HOLDING RETURN FOR BARLEY: SUBSETS OF DATA

Regression is
$\ln P = \alpha \ln P^* + \text{constant} + \gamma \text{Dectrend}$
Data is month-dated, with all Aug-Oct observations missing

Results are gross return expressed at an annual rate

	Estimate	Confidence interval
Elham	35.5%	-0.02% to 83.9%
Oxford	56.7%	0.0% to 145%
Other	19.8%	-1.04% to 70.3%

Table 8

Thorold Rogers' transcription of the 1316 Cuxham Accounts

(Vol 2., p.617)

Idem r: de iiij *li* xvj *s.* de vj quarteriis frumenti venditis die Jovis proxime ante gulam Augusti, pretium quarterii xvj s.

Et de vj *li* vj *s.* de ix quarteriis frumenti venditis ante festum omnium Sanctorum, pretium quarterii xiiij s.

Et de iiij *li.* xvj *s.* de vj quarteriis frumenti venditis die Jovis proxime ante festum Sancti Nicholai, pretium quarterii xvi s.

4/16/-, 6 quarters of wheat sold on the Thursday before Lammastide [1 August], price per quarter 16/-.

6/6/-, 9 quarters of wheat sold before the feast of All Saints, price per quarter 14/-.

4/16/-, 6 quarters of wheat sold on the Thursday next before the feast of St. Nicholas, price per quarter 16/-.

Thorold Rogers' entry in the Table of Grain Prices

(Vol 2., p.74)

6	July	16/-
9	Nov. 1	14/-
6	Feb. 24	16/-