The elasticity of substitution: evidence from a UK firm-level data set*

Sebastian Barnes[†] Simon Price[‡] Maria Sebastia-Barriel[§]

November 28, 2006

Abstract

Using a panel of UK firms spanning three decades, we provide estimates of the long-run elasticity of substitution between capital and labour, the (negative of the) elasticity of capital and investment with respect to the user cost. The parameter is estimated with long-period differenced data and pooled mean group panel methods. The robust result is that it lies below 0.5, confirming previous results obtained using aggregate UK data, and consistent with some recent results using US data. The estimated returns to scale exceed unity, but when constant returns are imposed the estimated elasticity of substitution is substantially unchanged.

Keywords: investment, firm level data, elasticity of substitution, panel.
JEL codes: E200, E220.

1 Introduction

It hardly seems necessary to justify the importance of accurate estimates of the elasticity of substitution between capital and labour (σ), as it is a

*This paper represents the views and analysis of the authors and should not be thought necessarily to represent those of the Bank of England or Monetary Policy Committee members. We are grateful for comments by Robert Chirinko, Neal Hatch and Ron Smith.

[†]OECD. E-mail: Sebastian.BARNES@oecd.org
[‡]Bank of England. E-mail: simon.price@bankofengland.co.uk
[§]Bank of England. E-mail: maria.barriel@bankofengland.co.uk

fundamental parameter in many economic models.[1] From the point of view of monetary policy, it is of interest not least because investment is a major part of the monetary transmission mechanism, and the elasticity precisely determines the responsiveness of capital and investment to the user cost. Yet there is controversy about its size. Some researchers believe that the elasticity is around unity: others that it is substantially less, perhaps below 0.5. The evidence might be thought to favour the second camp. For example, Hamermesh (1993), Nadiri (1970) and Nerlove (1967) survey a range of early estimates of σ, which are generally between 0.3 and 0.7. Krusell et al. (2000) has a model with skilled and unskilled labour: the elasticity between skilled labour and capital is 0.67 although for unskilled labour it is 1.67. Antràs (2004) reports a range of estimates for the elasticty, generally significantly below one. One reason why views are so strongly held may be that many macroeconomists assume it must be unity to account for the widely-believed stylised fact that factor shares are constant over time. In fact, while Cobb-Douglas technology does deliver this, labour augmenting technical progress is also sufficient in a constant elasticity of substitution (CES) world.[2]

In this paper we bring to bear evidence from the first order conditions determining the demand for capital. Here much of the debate is about the proper treatment of dynamics in estimation. When Kiyotaki and West (1996) examined Japanese investment, they found that investment was inelastic with respect to the user cost. They argued that this reflects firm behaviour. If firms expect shocks to the user cost to be quickly reversed, they will not react to them, and this explained the empirical results emerging from their estimated VARs. This is not a bias to a structural prarameter - more a reflection of the general equilibrium nature of the problem, and that VARs are reduced forms. But Caballero (1994) used an econometric argument that if variation in the user-cost is dominated by short-term transitory movements, then there is a substantial downward bias in estimates of the long-run parameter. Thus both Caballero and Schaller (2006) used Stock and Watson's Dynamic OLS (DOLS) method of estimating cointegrating relationships for the United States and Canada respectively. This method should be robust to transitory dynamics, and they both found long-run estimates close to unity. Goolsbee (1998) emphasised another potential bias induced by the short- to medium-run supply elasticity, and obtained similar results. Caballero et al. (1995) use plant-level data. They found widely dispersed estimates of the

[1] See Chirinko and Mallick (2006) for a wideranging discussion.

[2] In the 2004 version of the Bank of England Quarterly Model (Harrison et al. (2005)), which is characterised by balanced growth, the parameter was set at 0.317. There is some further discussion of the elasticity of substitution and factor shares in Section 3.2.1 below.

elasticity, but these averaged close to unity.

By contrast, Chirinko et al. (2004), using a large US panel, found a precisely estimated elasticity of approximately 0.40. Their method is designed to be unaffected by the dynamic issues mentioned above. By essentially estimating the cross-sectional relationship using long-period changes, it is both simple and robust, and in this paper we apply it to a UK data set. We find similar values. We also find that dynamic panel methods designed to capture long-run parameters in the presence of heterogeneity produce similar results. This confirms results on aggregate data using systems cointegration methods reported in Ellis and Price (2004).[3]

2 The demand for capital

The firm's optimisation problem is to maximise the discounted sum of expected future profits. Following the literature, we assume a CES production function,

$$Y = \left\{ \omega K^{\frac{\sigma-1}{\sigma}} + (1-\omega) X^{\frac{\sigma-1}{\sigma}} \right\}^{\frac{\eta\sigma}{\sigma-1}} A \tag{1}$$

where Y is output, K is capital, and X is a composite of other inputs (such as labour). σ is interpretable as the negative of the elasticity of substitution between capital and other factors. ω affects the share of the two factors, and η is the scale parameter. A is the stock of technology. The capital stock evolves as

$$K_t = I_t + (1-\delta_t)K_{t-1} \tag{2}$$

where I is gross investment and δ is depreciation. The Jorgensonian user cost of capital (Jorgensen (1963)) C is defined as

$$C_t = \frac{P_{k_t}}{P_t} \left(r_t + \delta_t - \Delta p_{kt+1} \right) (1 - nt) \tag{3}$$

where r is the cost of finance, P_k (p_k) is the price (log price) of a unit of the capital good, P_t is the price of output and nt represent net taxes (including subsidies and capital depreciation allowances in addition to income taxes). In the absence of adjustment costs or other sources of lags the dynamic aspect of the problem appears only in the user cost, and the optimal equilibrium

[3] Systems methods should be robust to transitory short-run dynamics. And in that data set single equation methods appear not to be biased: the results were unaffected by the use of DOLS.

capital stock for a given level of output is obtained by setting the marginal product of capital equal to the user cost:

$$\frac{\delta Y}{\delta K} = \eta\omega Y^{1+\frac{1-\sigma}{\eta\sigma}} K^{-\frac{1}{\sigma}} A^{\frac{\sigma-1}{\eta\sigma}} = C \tag{4}$$

Hence the equilibrium capital stock is given by

$$K_t = (\eta\omega)^{\sigma} C^{-\sigma} Y^{\sigma+\frac{1-\sigma}{\eta}} A^{\frac{\sigma-1}{\eta}} \tag{5}$$

It is normally assumed that there are costs of adjustment or other sources of lags which enter the dynamic optimisation problem. Under some circumstances (for example, as in Chapter VI of Sargent (1979), where there are constant relative prices and quadratic costs of adjustment) the static first order condition defines the optimal long-run value which serves as an attractor for the capital stock. In other cases the static solution may serve as an approximation to the dynamic long-run solution. This is the *de facto* assumption in the empirical literature that examines long-run estimates of the user-cost elasticity.[4] Such an approach may justify a simple dynamic equation explaining the log of the capital stock, k, such as the following:

$$k_t = \rho k_t^* + (1-\rho) k_{t-1} \tag{6}$$

where ρ is a parameter representing the speed of adjustment of the capital stock to its optimal level, k_t^*.

3 Econometric methods

We adopt three approaches to estimating the parameter of interest, which we will discuss in turn. In the first, we follow Chirinko et al. (2004) by setting to one side the panel nature of the data set and estimating a cross-section using time averages. Estimated in long differences, this allows us to abstract from time series properties. In the second and third, we use panel methods on the levels of the variables. As the data are non-stationary, this necessitates some prior tests for cointegration. Then we use a pooled mean group dynamic panel method due to Pesaran et al. (1999), allowing for heterogeneity in the dynamics. Finally, we estimate a static regression allowing for heterogeneity in the long-run parameters and cross-unit correlations due to Pesaran (2006), where static means a regression in the levels of the variables with no lagged terms.

[4] A discussion of various approaches to investment and the demand for capital can be found in Chirinko (1993): see also Caballero (1999).

3.1 Time averaging

Our discussion of the investment decision is not explicit about the treatment of either adjustment costs or signal extraction (relevant as firms need to form expectations about the future path of conditioning variables), but as pointed out above the debate has emphasised the importance of correctly treating transitory and persistent shocks. Chirinko et al. (2004) suggest a method that, as they put it, 'avoids' rather than 'overcomes' this difficulty. In particular, they propose the following formulation. (5) is rewritten in logs

$$k_{f,t} = \psi_{f,t} - \sigma c_{f,t} + \beta y_{f,t} + a_{f,t} \tag{7}$$

where lower case indicates the log of a variable, $\psi_{f,t} = \ln\left((\eta\omega)^{\sigma}\right)$ and $\beta = (\sigma\eta + 1 - \sigma)/\eta$. f subscripts the firm and t time.

In this formulation, the technology parameter is specified as:

$$a_{f,t} = \zeta\left[v_f + v_i + w_t + w_{f,t} + w_{i,t}\right] \tag{8}$$

where f denotes firm-specific variables, i denotes industry-specific variables and $\zeta = (\sigma - 1)/\eta$. v are firm- and industry-specific fixed productivity terms. w are firm- and industry-specific time-varying productivity terms.[5] The productivity terms in (8) are unobservable but the fixed (v) terms can be eliminated by differencing. Furthermore, taking time averages over two intervals allows the changes in firm- and industry-specific productivity $(w_{f,t} + w_{i,t})$ to be captured by a firm-specific intercept term $\gamma = \zeta\Delta w_t$ and an industry dummy variable $\lambda_j = \zeta\Delta w_{j,t}$. They then examine the cross-section regression using changes in time-averaged variables

$$\Delta k_f = -\sigma(\Delta c_f) + \beta(\Delta y_f) + \gamma + \lambda_i + \varepsilon_f \tag{9}$$

where the error term ε_f is assumed to be iid normally distributed. By estimating the model in first differences, this approach is also robust to the presence of I(1) variables.

Chirinko et al. (2004) argue for this specification on a number of grounds. It abstracts from any assumptions about the short-run dynamics of investment while retaining the better-founded theoretical assumptions about the long-run determination of the capital stock. In addition, it avoids the potential downward bias in the estimates of the long-run parameters of the model if a large proportion of the shocks are transitory rather than persistent. And this

[5] This specification does not specify economy-wide productivity growth, but of course this is subsumed within the components.

specification will not suffer from upward bias if the short-run supply curve of investment is upwards sloping, which tends to accentuate the response to investment in the short-run. The authors present a formal frequentist interpretation of the interval-difference estimator, where the use of intervals of several years is shown to give much greater weight to low frequency variation than simply using annual data.

They also argue that OLS is a robust estimation method. It is unlikely that there is a relationship between the stochastic element ε_f (which includes firm-specific changes in productivity Δw_f) and the regressors because (i) it is likely that most of the variation in productivity is specific to industries rather than firms[6] and, (ii) the output term may also absorb some of the effect of the productivity shock (see Shapiro (1986)). Moreover, the model is robust to some types of mis-specification of the firms' underlying optimisation problems. Mis-specification that generates levels differences, such as heterogeneity in mark-ups across firms, is eliminated by differencing. Some mis-specification of the model in differences, such as biased technological change, will be captured in the constants. And time-averages imply that there is less of a distinction between actual and potential output, addressing the criticism of models that fail to account for this.[7] And the model is robust to classic measure error, although not systematic measurement error.

Furthermore, unlike some other specifications, time averaging does not rely on identifying as 'persistent' only shocks that are permanent. This more supple definition of 'persistent', which includes shocks that are not permanent but may be very long-lived, is more in the spirit of the firm's forward-looking profit maximisation problem. One further advantage of this approach is that it eliminates some of the difficulties introduced by the 'lumpiness' in investment, proposed elsewhere in the literature. It is argued that plant-level investment is likely to be 'lumpy' from year to year due to the discrete nature of some capital purchases such as new structures or large pieces of equipment. These plant-level effects may translate into 'lumpiness' at the firm-level if the plant is large relative to the firm. Time-averaging under these circumstances will help to smooth the uneveness in investment and present a better picture of the relationship between investment and its determinants.

Chirinko et al. (2004) perform a number of robustness checks that support this approach in their dataset.

[6]For example, there is a widely held view that there have been major productivity advances in some particular industries, notably telecommunications and computing: see Basu et al. (2003).

[7]See references in Chirinko et al. (2004) to papers by Coen (1969) and Hall (1995).

1. Introducing industry-specific dummies (λ_i) has little impact on the estimated user cost elasticity coefficient. This increases the standard error of the estimate four-fold, due to collinearity with the industry-specific components that explain much of the variation in *firm-level* usercost.

2. To test for correlation between the error term (including firm-specific changes in productivity Δw_f) and firm output growth, output growth is dropped as regressor (by imposing constant returns to scale ($\eta = 1$), which implies $\beta = 1$). This has little impact on the estimate of the usercost elasticity, suggesting that simultaneity is not a problem (the paper does not test this formally).

3. The model rejects the restriction of a Cobb-Douglas production function ($\sigma = 1$,$\eta = 1$).

4. There is a potential bias from systematic measurement error in constructing the capital stock (*e.g.*, a change in technology that increases growth and depreciation, which the model assumes is fixed in equation 2). The paper assesses the potential bias from mismeasurement of either explanatory variable. Given the use of differencing and industry variables, variation from measurement would have to be very large to bias the estimate of σ. The test in Rao (1973)) shows that the likely bias from measurement error in output is trivial.

5. Instrumental variable (IV) estimates of the model yield almost identical results to the main model, suggesting that the endogenous variables are independent from the regressors as assumed. This is confirmed using Hausman tests. The instrument set comprises variables from an earlier interval than the two used in the model, in some sense predicting the variables used in the model but independent from its errors.

6. The model is robust to changes in the sub-sample, which variously split the sample by a cashflow variable, size and a measure of Brainard-Tobin Q.

We adopt this method in the current paper, although data limitations prevent us from carrying out all the robustness checks that Chirinko et al. (2004) carry out.

3.2 Panel estimates

The time-averaging method is arguably robust, but may have limitations. By construction, it ignores all short-run information, although there is no presumption that this leads to inefficient or biased estimates. But it does ignore information about the speed of adjustment to the long-run, which is known to be helpful in estimating long-run relationships in a cointegrating context. And it also maintains the hypothesis that the long-run parameters (returns to scale and the elasticity of substitution) are equal across firms.

3.2.1 Mean group estimates

An alternative method, without these disadvantages, identifies the long-run responses of the capital shocks using the Pooled Mean Group (PMG) estimator of Pesaran et al. (1999). This is estimated in a dynamic panel model, where the long-run parameters of interest are restricted across the panel but the short-run dynamics are estimated without restriction for each member of the panel. In principle, this might be more efficient.

Allowing the dynamics to be estimated freely is important. Static models are rarely adequate for time series. But there are well-known problems that arise from estimating pooled dynamic models. The small-T problems with dynamic panels[8] are not relevant here as the fixed-effects problem from the initial conditions declines rapidly as T rises. But instead, there are profound problems that result from heterogeneity in the model parameters that emerge as soon as a lagged dependent variable is introduced. This problem was forcefully addressed by Pesaran and Smith (1995). Unlike in static models, estimates are inconsistent even in large samples, essentially because the heterogeneity is not reduced by increasing the size of the cross-section. Happily, in an increasing number of data sets T is sufficiently large to allow individual cross-sectional - in our case, firm-level - estimation. Pesaran and Smith observe that while it is implausible that the dynamic specification is common to all units, it is at least conceivable that the long-run parameters of the model may be common. We can then exploit the cross-sectional dimension to gain more precise estimates of these average long-run parameters. They then propose estimation by either averaging the individual unit estimates (the Mean Group method: MG), or in their later paper with Shin (Pesaran et al. (1999)) by pooling the long-run parameters with the PMG method, if the data allow. And even if the common parameters restriction

[8] Arellano and Bond (1991).

is rejected, there may still be benefits from pooling, including robustness to outliers.

The dynamic model is formulated by taking logs of (5), combining with (6) and including lags in all variables.

$$k_{f,t} = \sum_{1}^{j=p} \gamma_{1fj} k_{f,t-j} + \sum_{1}^{j=p} \gamma_{2fj} y_{f,t-j} + \sum_{1}^{j=p} \gamma_{3fj} c_{f,t-j} \tag{10}$$

The role of technology requires some discussion here. (10) suppresses the constant and technical progress terms. While (1) contains a term in technology, A, in the time averaged case the effects are differenced out. However, it is normal not to include technological shift terms in expressions such as (5). The usual justification is that technological progress is assumed labour augmenting. While it may not be immediately obvious, this is a reasonable assumption, as balanced growth requires this assumption (or the closely related assumption that technology is Cobb Douglas): see Barro and Sala-i-Martin (1995), pages 54-55, and Jones (2003). The case for balanced growth is in turn powerful, given the relative constancy of factor shares in many countries, and perhaps especially for the United States and United Kingdom. There is some evidence on this issue. Antràs (2004) estimates a model in which biased technological growth for both labour and cappital are estimated. As well as reporting estimates for σ that are less than one, he finds that in his sample labour augmenting efficiency grew about 3% faster than capital augmenting efficiency. In fact, the latter is estimated to trend downwards, which may not seem too plausible, but may suggest zero growth is not unlikely. In the amusingly entitled Klump and Willman (2007) similarly allow for differently biased growth, with a flexible deterministic trend. They find that for both the United States and Euro area there is no current evidence for capital augmenting technical progress.

(10) can be reparametrised as an error correction mechanism (ECM):

$$\Delta k_{f,t} = \lambda_{f1} k_{f,t-1} + \beta_{1f} y_{f,t-1} + \beta_{2f} c_{f,t-1} + \sum_{1}^{j=p-1} \omega_{1fj} \Delta k_{f,t-j}$$

$$\sum_{1}^{j=p-1} \omega_{2fj} \Delta y_{f,t-j} + \sum_{1}^{j=p-1} \omega_{3fj} \Delta c_{f,t-j} \tag{11}$$

where in practice the lag legnth may vary. With the PMG estimator, this model is estimated subject to the restriction $\beta_{1f} = \beta_1$ and $\beta_{2f} = \beta_2$ for all i.

This model could be estimated by iterated least squares, imposing and testing the cross-firm restrictions on β_{1f} and β_{2f}. However, this will be inefficient as it ignores the contemporaneous residual covariances. A natural estimator is Zellner's SUR method,[9] which is a form of feasible GLS. But SUR estimation is only possible if N is (usually much) smaller than T. Thus Pesaran et al. (1999) suggest a maximum likelihood estimator.[10]

There are also potential problems with inference. Arguably, omitted group specific factors or measurement errors are likely to severely bias the unit estimates. Perhaps because of this, it is a commonplace in empirical panel studies to report a failure of the 'poolability' tests based on the group parameter restrictions.[11] Pesaran et al. (1999) suggest that a Hausman test be used as an alternative to (*e.g.*) a Wald test of the hypothesis that the long-run coefficients are common. The test is based on the result that an estimate of the mean long-run parameters in the model can be derived from the average (mean group) of the unit regressions. This is consistent even under heterogeneity. However, if the parameters are in fact homogeneous, the mean and the individual parameters coincide and the PMG estimates are more efficient. Thus we can form the test statistic

$$H = \hat{q}'[var(\hat{q})]^{-1}\hat{q} \sim \chi^2_k$$

where $\hat{q}$ is a $(k \times 1)$ vector of the difference between the mean group and PMG estimates and $var(\hat{q})$ is the corresponding covariance matrix. Under the null that the two estimators are consistent but one is efficient, $var(\hat{q})$ is easily calculated as the difference between the covariance matrices for the two underlying parameter vectors. If the poolability assumption is invalid then the PMG estimates are no longer consistent and the test rejects.

The Hausman should be seen as a misspecification test rather than a test of the restrictions on the parameters. Holly (1982) discusses the Hausman null. In this case we should interpret that null as not that the parameters are equal, but that the mean (i.e., MG) estimate of the parameters is not significantly different from the PMG estimate. As an empirical issue, it is this average value with which we are concerned, rather than the hypothesis of homogeneity, so this does not present a problem. Indeed, it might be considered implausible that the parameters are homogeneous; which is not

[9] Zellner (1962).

[10] Implemented in a GAUSS program available on Hashem Pesaran's website: we are grateful to the authors for making this available.

[11] For example, Baltagi and Griffin (1997) state that although the poolability test is decisively failed (F(102,396) = 10.99; critical value 1.3), 'like most researchers we proceed to estimate pooled models.'

to say that pooling cannot increase efficiency. It should also be noted that the test applies only where there are both lagged dependent variables (LDVs) and exogenous regressors. If there are no LDVs the Plims of both are the same and there is no heterogeneity bias, given random parameters. If there are no exogenous variables the asymptotic covariance matrices are the same for the pooled and MG estimates. But neither of these apply in our specification, where the user cost is defined at an industry level and is surely exogenous to the firm's decisions.

As Pesaran et al. (1999) and others observe that standard static panel models may often provide good estimates of the long-run parameters in practice, we report these as well.

3.2.2 Cross-sectional correlation

Recent research has investigated the implications of heterogeneity and unmodelled cross-sectional correlations. A recent review of the relevant literature, focussing on non-stationary cases, is in Breitung and Pesaran (2005).

Pesaran (2006) presents a new and simple method to estimate relationships in large panels of the type we explore.[12] It presumes that there is a general multifactor error structure behind the data, which should be accounted for in estimation. This is highly plausible: firms are likely to be subjected to the same shocks. Were N small enough, this would suggest estimation *via* GLS, but in our case, as already observed, this is infeasible. It also allows for heterogeneity in the parameter estimates. The linear heterogeneous model is

$$y_{it} = \boldsymbol{\alpha}_i'\mathbf{d}_t + \boldsymbol{\beta}_i'\mathbf{x}_{it} + e_{it} \tag{12}$$

where $\mathbf{d}_t$ is a vector of observed common effects (including deterministics such as intercepts), $\mathbf{x}_{it}$ is a vector of observed regressors and the errors e_{it} have the multifactor structure,

$$e_{it} = \boldsymbol{\gamma}_i'\mathbf{f}_t + \varepsilon_{it} \tag{13}$$

where $\mathbf{f}_t$ is the vector of unobserved common effects. The mean-group method of Pesaran et al. (1999) estimates (12) by OLS. To estimate the $\boldsymbol{\beta}$ efficiently, we simply add as regressors the cross-sectional averages of y_{it} and the $\mathbf{x}_{it}$, $\mathbf{z}_t$. Thus we run the regression

$$y_{it} = \boldsymbol{\alpha}_i'\mathbf{d}_t + \mathbf{b}_i'\mathbf{x}_{it} + \boldsymbol{\gamma}_i'\mathbf{z}_t + u_{it} \tag{14}$$

[12] See Holly et al. (2006) for an application and Breitung and Pesaran (2005) for a recent survey of some related areas.

and $\hat{\mathbf{b}}$ is a consistent estimator of $\boldsymbol{\beta}$. Assuming the number of unobserved factors is sufficiently low, consistent estimates of $\boldsymbol{\beta}_i$ can be obtained, with variances given by standard Newey-West procedures. Even if the individual slope coefficients cannot be consistently estimated, the mean can be estimated under Pesaran's assumptions. The variance of the mean may be estimated very simply:

$$\sqrt{N}(\hat{\mathbf{b}}_{MG} - \boldsymbol{\beta}) \xrightarrow{d} \mathrm{N}(\mathbf{0}, \boldsymbol{\Sigma}_{MG}) \tag{15}$$

where

$$\boldsymbol{\Sigma}_{MG} = \frac{1}{N-1}\sum_{i=1}^{N}(\hat{\mathbf{b}}_i - \hat{\mathbf{b}}_{MG})(\hat{\mathbf{b}}_i - \hat{\mathbf{b}}_{MG}). \tag{16}$$

If any of the slope coefficients are the same, efficiency can be increased by pooling over them. To do this, prior regressions of y_{it} and $\mathbf{x}_{it}$ on the $\mathbf{z}_t$ are run and the residuals e_{1it} and $\mathbf{e}_{2it}$ retrieved. Then the regression

$$e_{1it} = \boldsymbol{\lambda}_i' \mathbf{e}_{2it} + \eta_{it} \tag{17}$$

provides consistent estimates of $\boldsymbol{\beta}$.

The case examined in Pesaran (2006) is for stationary variables, but Kapetanios et al. (2006) find that the presence of unit roots does not affect most theoretical results. And Monte Carlo experiments suggest that in small samples the method is robust to a wide variety of data generation processes and has lower biases than a range of alternative estimation methods they consider.

3.2.3 Cointegration

This raises the issue of non-stationarity. Given our trended data, the dynamic panel estimates require cointegration. Pedroni (1999, 2004) developed residual based cointegration tests. If there is cointegration, the residuals will be I(0). Panel tests have advantages over single time series tests, which are are well known to have non-standard distributions and low power. In panels the distributions tend to asymptotic normal as the cross-section dimension rises, and power usually increases. Nevertheless, it may be that the tests are powerful against an arguably uninteresting alternative. They examine the null of no cointegration - that is, that there are unit roots in the candidate cointegrating residuals in all of the N groups. Rejection does not, therefore, imply that all the series are I(0), but rather that some are I(0); they are not all I(1). Thus rejection means that there is evidence for cointegration in some of the series. While this does not rule out cointegration in all series, it does not imply it. The Pedroni tests which we use[13] allow for heterogeneity among the panel members. All are based on the residuals from the (most general) regressions

$$y_{it} = \alpha_i + \delta_i t + \beta_i x_{it} + e_{it}. \tag{18}$$

Pedroni constructs seven tests, four of which are based on pooling along the 'within-dimension' and three the 'between-dimension'. The former effectively pool the autoregressive coefficient in the residual based test and the latter take the average, allowing more heterogeneity. Pedroni refers to the within statistics as panel cointegration statistics, and to the between statistics as group mean panel cointegration statistics, a natural terminology given our discussion above. The panel tests constitue a panel v-statistic (a non-parametric variance bounds test), a panel ρ-statistic (analagous to the Phillips Perron ρ test) and nonparametric and parametric panel t-statistics (or more accurately, ADF statistics). The group tests are a group ρ-statistic and the two group t-statistics.

In each case the null is no-cointegration. The (one sided) test statistics are distributed asymptotic standard normal. The critical value for the panel v-statistic is positive, while the others are negative. In the spirit of the PMG approach heterogeneity in the autoregressive process is likely, so we prefer to be guided by the group statistics.

4 The dataset

We use a new firm-level dataset constructed by combining company accounts data from DataStream/Worldscope (DSW) and annual industry-level data from the Bank of England Industry Database (BEID), which covers 32 different industries for the period 1970-2005.

The firm-level DSW data cover UK publicly-listed firms only. There are around 3800 firms in the full sample, based on the list of currently listed companies and firms that no longer exist. Around 1200-1400 firms are active in any given year. The median turnover in 2003 was around £40 million in current prices, although turnover is distributed with a strong negative skew so that the mean turnover was around £700 million. The median number of employees was around 400 (the mean number was above 5000). Construction of the BEID is discussed in greater detail in Basu et al. (2003).

[13] We are grateful to Peter Pedroni for making his RATS procedures available to us.

The merged dataset has observations on around 1600 firms. The loss of observations relative to the full DSW dataset is largely due to missing industry identifiers in the DSW, although some observations are also lost for firms that belong to industries that are not covered by the BEID.[14] The dataset has fewer firms than comparable studies for the United States, due to the smaller number of listed companies, and a narrower sample of firms than some previous studies for the UK using the DSW data, primarily due to the use of BEID industry-level data rather than National Accounts aggregate data (see for example Bond et al. (2004)).[15]

4.1 Variable definitions

There is a consensus in the literature about variable definitions, which we largely follow (see Appendix A), although there are a number of alternative definitions of investment extant. We follow Bond et al. (1999) in defining investment as (gross) payments for fixed assets and excluding (net) sales of fixed assets. Bond et al. (1999) justify this on the grounds that asset disposals (*ds423*[16]) 'pre-1992 appear [...] to be contaminated by measurement error'.[17] We also ignore observations of payments of fixed assets by subsidiaries as do Blundell et al. (1992), although mainly on the pragmatic grounds that this series has been discontinued by Worldscope.[18] The main alternatives to this are to include part or all of investment by subsidiaries (*ds479*), and to subtract sales of fixed assets (*ds423*) as Bond et al. (2004). No method is *a priori* more satisfactory than the others; there is an underlying trade-off between the quality and availability of the data, and constructing a series with greater theoretical coherence. Our definition of investment is around 20 percent higher on average than the alternative used in Bond et al. (2004).[19]

The capital stock is constructed by the Perpetual Inventory Method (PIM),

[14]The allocation of firms to SIC industries is based on Bond et al. (2004).

[15]The main sample in that paper has 4263 observations on 703 firms for the period 1989-2000.

[16]Italicised codes of this form, *e.g. dsxxx* refer to the DSW identifiers.

[17]They also ignore negative values of *ds479* (FixedAssets(Subs)). See also Carpenter and Guariglia (2003).

[18]Bond et al. (1999) ignore negative values of *ds479* (FixedAssets(Subs)) on the same grounds of mis-measurement as for asset disposals more generally. This argument is also used in Carpenter and Guariglia (2003).

[19]Note that the choice of definition should not *a priori* have any implication for the investment to capital ratio, assuming there is a stable relationship between the two alternative definitions. But, the *level* of the capital stock is different in the two cases. This also, for example, also implies different level of Tobin's Q.

using firm-level investment and industry-level depreciation to project forward the capital stock from a firm-specific initial value $k_{0,f}$:

$$k_{t,f} = k_{t-1,f}\,(1 - d_{t,i}) + inv_{t,f}$$

where t is time, i is an industry, f is a firm, d is depreciation, inv is investment and k is the (net) capital stock. The industry-level measure of depreciation (d_i) is from the BEID and has some drawbacks as a measure of depreciation at the firm-level. It is an industry rather than firm-level concept, and is conceptually somewhat different to a direct measure of the rate of physical depreciation.[20] But using industry-level data may provide a better approximation to firm-level depreciation than the economy-wide measures used elsewhere.

Setting the (unobserved) initial value of the capital stock $k_{0,m}$ is an important issue. We follow a method based on adjusted book-value, similar to that proposed by Chirinko and Schaller (2004a) (see Appendix B for details of the method).[21] They show that differences between the 'accountant's' and the 'economist's' definitions of capital arise through a 'price distortion', due to differing treatments of cost of acquiring capital and a 'depreciation distortion' due to the use of different depreciation rates. They propose initialisation of the capital stock with an adjustment to book-value which takes these factors and the price-level into account. Chirinko and Schaller (2004a) show the quantitative importance of these effects: the median difference in their study for all industries between book-value and the economic definition of the capital stock is 178%,[22] and greater than 60 percent for all but 2 of 46 industries. They find a number of cases where the difference is much larger. This has a number of implications. First, the convergence of PIM-based estimates of capital from book-value to the economist's definition may be large and very slow, taking around 15 years to reach an acceptable degree of convergence in the growth rate Chirinko and Schaller (2004a). This may induce spurious trends in the data. Furthermore, the common practice of merely dropping the initial three observations may not be sufficient to remove this. Second, it suggests that cleaning methods that drop observations where capital differs from book-value by a factor of 3 or 4 may be inappropriate

[20]It is the implied depreciation rated derived from the growth rate of a chain-weighted aggregate of capital less a chain-weighted aggregate of investment. As such, it partly captures shifts in relative prices between investment in different assets at the industry-level rather than simply the physical depreciation of capital (see Whelan (2003)).

[21]Bond et al. (2004) use an alternative method, where the 'price distortion' is addressed using average investment goods price inflation over the previous three years and the 'depreciation distortion' is allowed for by dropping the first three observations.

[22]That is, economic capital is almost three times greater than historic book-value.

as they remove valid observations (while retaining some invalid observations along the convergence path).[23] It is important in our sample to apply an appropriate correction as there are typically relatively few observations for each firm over time, so that any convergence-related mismeasurement would dominate in our sample.

The Jorgensonian user cost of capital (C) is defined at industry-level as the weighted sum of the user cost of 7 different assets in that industry. Depreciation and taxes differ across assets, but are the same across industries for a given type of asset. Prices for a given investment good are identical across industries, except for plant and machinery where they also vary by industry. There is a considerable advantage in this method relative to the common procedure of assuming a depreciation rate that is constant over time and common across assets, given the increasing importance of computer equipment in the capital stock and its relatively high rate of depreciation (see Schaller (2006)). The rate of return is the weighted-average cost of capital. This measure is forward looking in the sense that the cost of finance is an expected rather than *ex post* measure. This is assumed to be common across industries, which ignores sectoral differences in risk.[24] The assumption of a common cost of finance may be relatively innocuous if the capital stock does not depend strongly on the cost of finance relative to other elements of the user cost (see Schaller (2006)). Although this assumption does imply the loss of some firm or industry-level variation, the method mitigates a potential endogeneity problem in conventional measures of the cost of equity used in other studies. These are based on the ratio of equity prices to *current* earnings or dividends: they do not distinguish between the risk-adjusted cost of finance and expectations of future profitability.[25] A small number of observations of the user cost of capital are negative, primarily in the 1970s.[26]

Other variables are as defined in Appendix A and follow the definitions of Datastream/Worldscope.[27] Nominal sales deflated by industry-wide output

[23] See, for example, Bond et al. (2004) who drop observations that are more than 3 times out of line with book-value.

[24] By constrast, Chirinko and Schaller (2004b) use an industry-specific measure of risk derived from the covariance of equity prices using the Capital Asset Pricing Model (CAPM).

[25] This endogeneity problem still exists under this approach for economy-wide shocks, but the effect for firm- or industry-level shocks is mitigated by the use of a common measure.

[26] There are 51 negative observations, only 15 of which occur after 1980. These are most common in industries 3 and 17: see Table C1.

[27] We can construct a measure of Tobin's average Q, the ratio of the value of the firm to the capital stock. This approximates marginal only under the conditions set out in Hayashi (1982). No adjustment is made for taxation. But we do not use Q in this paper.

deflators proxy for firm-level output. Gross output is the appropriate conditioning variable here. Although economists tend to concentrate on value added, this is often because they are interested in value-added aggregates such as GDP: see Basu and Fernald (1997).[28] Nevertheless, we also construct a measure of value added using operating profit and employment costs, although this brings a loss of observations due to incomplete data.

4.2 Data cleaning

We drop implausible observations or those for which important information is missing. Missing, zero or negative observations of sales (*ds104*), market value (*dsMV*) and book-value (*ds339*) were dropped. Observations with missing or negative values of total loans (*x321*) were dropped together with cases where net assets *x390* are missing. Standard procedure was adopted whereby observations with increases in turnover greater than 200 percent in one year are excluded. For the reasons set out in Section 4.1, we do not follow the literature in dropping estimates of the capital stock more than 3 or 4 times greater than book-value.[29]

4.3 Mergers

One important issue is the treatment of firms where there is a merger or acquisition. Our broad definition of investment includes both acquisition of capital goods (*e.g.* a new machine) and other companies (as in Blundell et al. (1992)). We do not have an explicit identifier for mergers and acquisitions as has been exploited by some of the literature using the corresponding US data.[30]

This raises a number of difficulties. Some changes in the microdata will result from changes in ownership rather than changes in activity. This poses difficulties for estimating firm-level investment equations, as the motivation for the acquisition may lie more in the theory of the firm than of capital formation. Furthermore, such mergers make it difficult to analyse aggregate

[28]Under a Cobb-Douglas production function with constant returns to scale and productivity, where the industry output deflator is perfectly correlated with the economy-wide output price deflator, sales and value-added would be perfectly correlated over time.

[29]Our method is the same as Barnett and Sakelleris (1998), although they justify this on *a priori* grounds.

[30]Barnett and Sakelleris (1998) appear to use such identifiers to purge their dataset of post-merger observations.

Table 1: Descriptive statistics of firms in Large-N dataset

	mean	std. dev.	first quartile	median	third quartile
Total sales (£000)	1030982	4057262	44204	121949	511900
Total employment	8641	20129	501	1439	5746
Market value	859	3120	19	81	388
Pre-tax profits	78108	306600	1934	8076	35200

1996 data in current prices based on observations on 403 firms.

investment in a balanced panel if one of the original entities is dropped from the panel altogether as a result of the merger or acquisition. And mergers may create difficulties in using the PIM to estimate the capital stock. On the one hand, it requires the acquired entities' capital to be purchased at its economic value for investment to be correctly measured. On the other, the assumption about the constant ratio of book-value to economic capital used to initialise the PIM algorithm may also be violated in this case. Frequent mergers would also interfere with the convergence of the PIM to its 'true' value Chirinko and Schaller (2004a) if this value is frequently revised.

The literature generally does not identify and treat mergers and acquisitions explicitly, by attempting to match firms as they merge or make acquisitions.[31] We follow the method of Bond et al. (2004) and others and drop extremely large changes in key variables as these are likely to indicate a merger/acquisition rather than an underlying change in activity.

4.4 Alternative datasets

We use our sample based on combining DSW and BEID data to construct three different datasets, which attempt to exploit the trade-off between the N (number of firms) and T (number of time periods) dimensions. All firm-level observations in the dataset are continuous to enable capital to be derived using the PIM.[32]

[31] An alternative method is developed in Chirinko and Schaller (2004b), where an algorithm is developed for dropping large changes in book value that do not correspond to proportionate changes in investment (where data on acqusitions and retirements of capital stock is used to measure investment precisely).

[32] Chirinko and Schaller (2004b) interpolate over some missing observations (including for some observations of investment), arguing that the increase in the number of observa-

We consider three different datasets, all based on the same underlying DSW and BEID data:

- Large-N, where there are *at least* 10 continuous observations on 403 firms over the period 1973-2003.
- Large-N-balanced panel of observations drawn from large-N on 261 different firms over the period 1986-1999.
- Large-T-balanced panel of observations for 1975-2000 on 142 firms (all of which belong to large-N-balanced).

Balanced panels may in principle allow aggregation across the series to produce an aggregate series that corresponds to measures of overall investment of all firms. This approach also provides some control for time effects, as no year is more represented in the sample than any other. These advantages are achieved at the cost of a loss of observations, where data are missing for some firms in some years.

Furthermore, there will be attenuation or survivorship bias if those firms with missing observations or that disappear from the dataset are not representative of the sample as whole. We avoid the bias associated with only including firms active at the end of the sample by choosing our firms from the list that includes the firms which are no longer active (Barnett and Sakelleris (1998)).[33] One source of bias is that our balanced firms may include disproportionately few firms that are likely to fail or merge in the sample. Some authors have taken the view that unbalanced panels are optimal: Chirinko and Schaller (2004b) argue that the gain in the number of observations and reduced attentuating survivorship bias outweigh the other difficulties.

More generally, the firms in the large-N subsample (defined below) are relatively large compared to the typical firm in the UK corporate sector. In 1999, the last year of the sample, median sales for firms in the sample were close to £200 million with a strong positive skew. The median number of employees was around 1,600 with the mean number of employees 10,000.[34]

---

tions is of greater value than the cost of measurement error introduced by interpolation.

[33] In the United States, the NBER maintains a dataset of different years of this type of data that provides a more comprehensive solution to this problem. Unfortunately, there are no corresponding data available for the United Kingdom.

[34] A description of the size distribution of firms in the corporate sector as whole can be found in Criscuolo et al. (1999).

Table 2: Summary statistics - whole sample, main variables

	mean	std.dev	first quartile	median	third quartile
I/K	0.22	0.28	0.09	0.15	0.25
Q	2.10	4.17	0.46	1.13	2.32
r	0.59	1.26	0.17	0.24	0.40
Δ real sales	0.04	0.23	-0.08	0.02	0.12
Δ employment	0.04	0.33	-0.06	0.01	0.09
Δ capital	0.11	0.28	-0.01	0.05	0.15
Δ value added	0.03	0.78	-0.68	0.04	0.15
δ	0.11	0.05	0.08	0.09	0.11

Based on 9505 observations covering 1973-2002 using 1995 prices.

5 Description of data

5.1 Summary statistics

Table 2 shows the summary statistics for the large-N dataset for the main variables. These are comparable to those from earlier work using related UK data (Bond et al. (2004)) and for studies for other countries (see, for example, Barnett and Sakelleris (1998)). The estimates of Tobin's Q averages 2.10, a little lower than in a Bond et al. (2004), who report a mean of 2.66; but this still exceeds unity by a large margin. However, the median value is not far above one. The numerator of the average Q expression is defined in the same way in both studies and our results are similar over a comparable sample (1987-2000), so this discrepancy is likely to be related either to the alternative methods of initialising the capital stock or measuring depreciation.[35]

The annual percentage change in the capital stock, following the approach of Doms and Dunne (1994) on firm-level data, has the advantage of being in principle a stationary variable. This rate is faster than for the corresponding aggregate measure in the National Accounts, but these figures are not comparable since the Table reports the unweighted mean of growth rates. Given that smaller firms are likely to growth faster, it is likely that the unweighted average growth rate will be faster than the aggregate growth rate.

[35]It is not meaningful to compare aggregate statistics based on chain-volumed indices, such as the investment/capital ratio, with their firm-level counterparts.

Figure 1: Investment to capital ratio

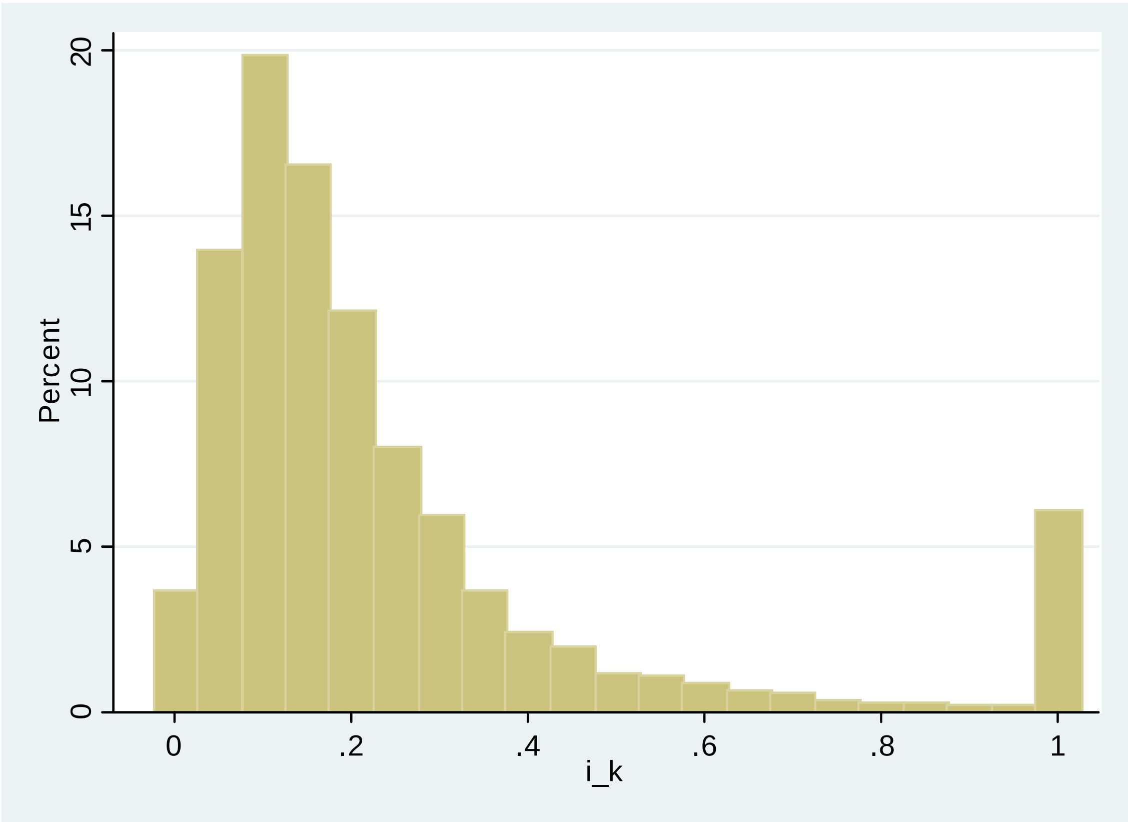

I/K ratio, where all observations greater than 1 have been allocated to the final bin. There are no negative observations for this sample.

Figure 2: Investment to sales ratio

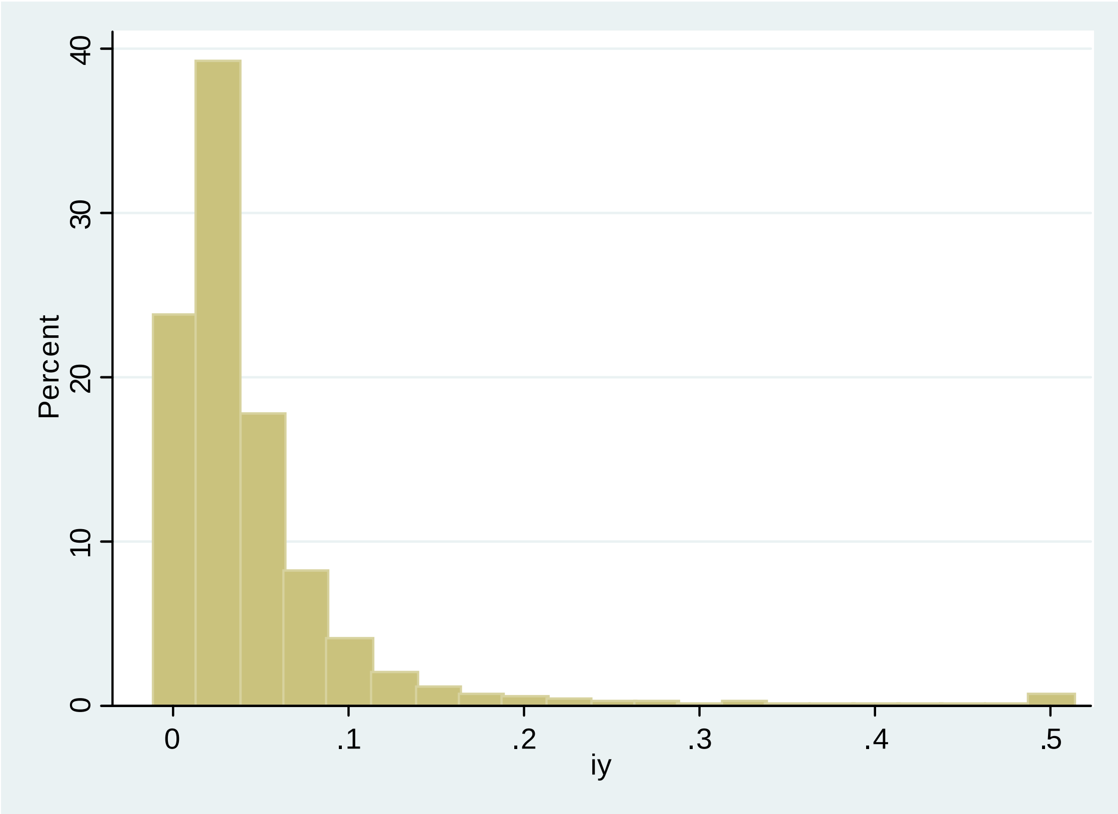

I/Y ratio, where all observations greater than 0.5 have been allocated to the final bin. There are no negative observations for this sample.

Figure 1 shows the investment-to-capital ratio, which has a mean of 0.22 (see above). The distribution has a negative skew, although with a substantial proportion of observations greater than 0.3. An alternative measure of investment is the investment-to-sales ratio (I/Y). Although it does not provide a very direct measure of the relationship between investment and capital, it has the advantage of being based entirely on DSW data rather than our estimates of the capital stock. Figure 2 shows that the I/Y ratio has a mean of around 0.04, although the distribution has a similar negative skew as that for the I/K ratio.[36]

[36]These numbers reveal that the capital to sales ratio is much smaller than the aggregate capital to output ratio, which was around 2.7 in 2005 (whole economy). However, the aggregate data is on a value-added basis, and in this sample sales are correspondingly larger than value added.

One useful analysis is to consider the number of 'spikes' and zero or extremely low investment episodes. Around 11% of observations have I/K ratios greater than 50%. These estimates of 'spikes' arguably provide a lower bound because, if they are due to single projects, some may span more than one financial year and so the 'spike' in investment recorded in one year may be less than the total amount of investment associated with that episode (Doms and Dunne (1994)).

We have observations of firms in 15 different industries, where just over half of firms are in the services sector with most of the rest in manufacturing. For industry-level variables, such as the user cost, we consider two methods of capturing variation in the data. The first method is to report unweighted statistics across the firms taken as a whole. This measure should respond more closely to the variation that is driving the econometric results (which are implicitly weighted by the *number* of firms). The second method considers variation in these variables at the industry-level (*i.e.* where each industry has equal weight). Comparing these two methods gives an indication whether variation in these data reflect intrinsic variation between industries, and whether the distribution of firms across industries plays a role.

5.2 Aggregation

The data capture a large share of business investment, despite the small number of firms in the dataset. The 142 firms in the large-T balanced panel accounted for around 11 percent of nominal business investment in 1999 while the 261 firms in the large-N balanced panel accounted for 16 percent.[37] The average growth rate of aggregate investment in our sample over the period 1979-2000 is 6 percent, compared to 8 percent for the ONS measure, but the standard deviation of annual growth rates for our series are about twice as large as the mean growth of the series and of the standard deviation of the ONS measure. There are many corresponding movements in the two series and the correlation of large-n balanced panel and the national accounts measure of nominal business investment is around 0.40. Our estimates indicate a larger fall in output and investment in the early 1990s than the ONS measures. Of course, we would not expect the two series to be identical due to differences in sample, sectoral coverage, definition, and other methodological issues. Apart from any other factors, there is the fundamental difference that the firm-level data record investment done by firms resident in the UK, as

[37] This further emphasises the skewness in the size distribution of firms and the greater likelihood of continuing in the sample of large firms.

Figure 3: Annual investment growth

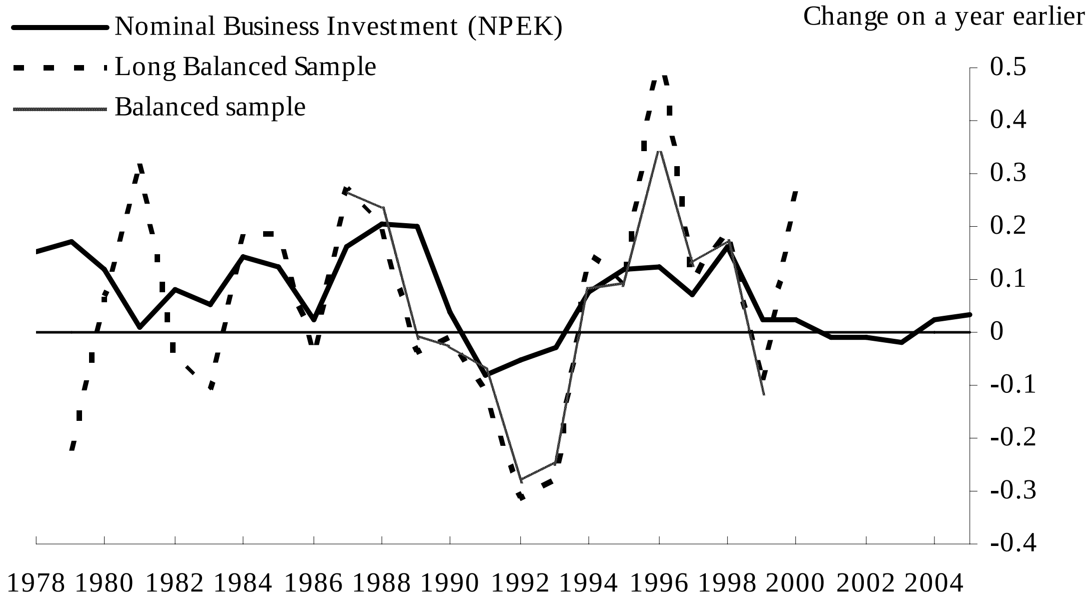

opposed to investment in the UK performed by both foreign and domestic firms resident in the UK, which is what is captured in the ONS data.

A balanced panel would not be expected to capture the variation in aggregate investment due to new entrants. In a similar way, large firms are over-represented in our sample. Such firms are likely to account for a large share of investment, but investment growth is likely to be stronger for small (expanding) firms (see Rossi-Hansberg and Wright (2005) for US evidence), not all of which will survive. This could explain the lower average growth rate of our sample of companies than for the ONS aggregate measure. In our sample, there is a bias towards firms with a large number of continuous observations, which may be larger on average than the sample of all firms as these are the firms that were either larger to begin with or have successfully expanded.

There is some correspondence between 'investment bursts', defined as observations for which the investment/capital ratio is greater than 0.5, and growth of aggregate investment (similar results were found for the United States by Doms and Dunne (1994)). This may also provide some insight into mergers and acquisitions (M&A) related measurement problems. In particular, ONS estimates of the number of mergers and acquisitions appear to

Figure 4: Incidence of 'investment bursts'

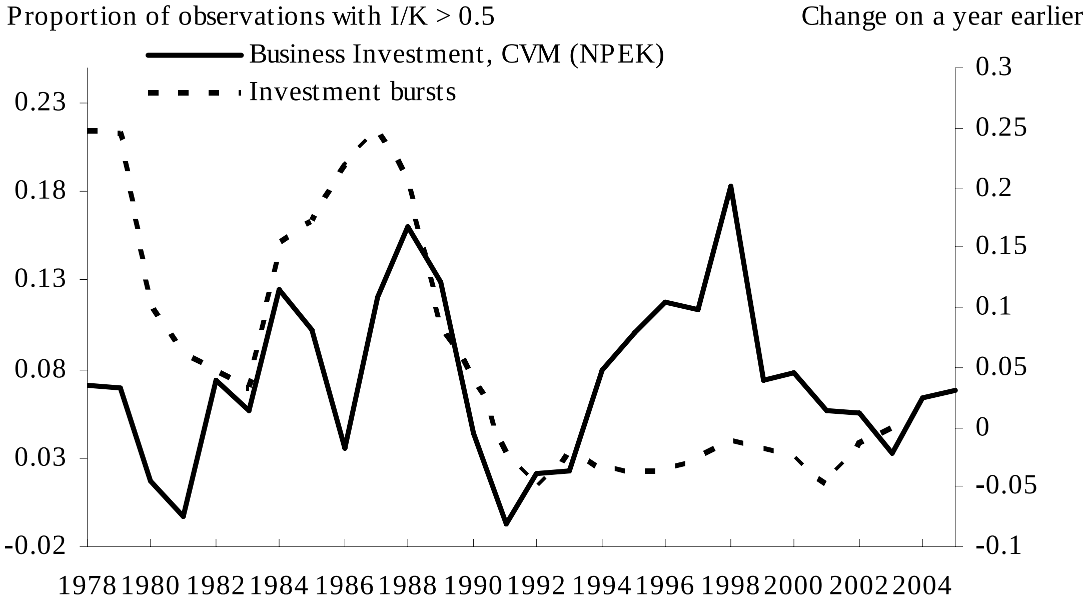

have followed business investment growth closely in some periods such as the late 1980s but not others, such as the mid-1990s (Figure 5). The series for 'investment bursts' in our sample appears as if it is more closely related to aggregate M&A activity than investment growth. This suggests a possibility of mismeasurement.

6 Results

We now turn to the results. We estimate using two alternative conditioning variables. We begin by following Chirinko et al. (2004) and using sales as a proxy for output, but we also examine results using value-added.

6.1 Conditioning on real sales

We examine three permutations of the data, described above. The first is the unbalanced panel containing 403 firms; the second a balanced panel comprising 261 firms but with a relatively short period (14 years); the third a smaller panel with a longer period (26 years: 1975 to 2000.)

Figure 5: Investment and M&A activity

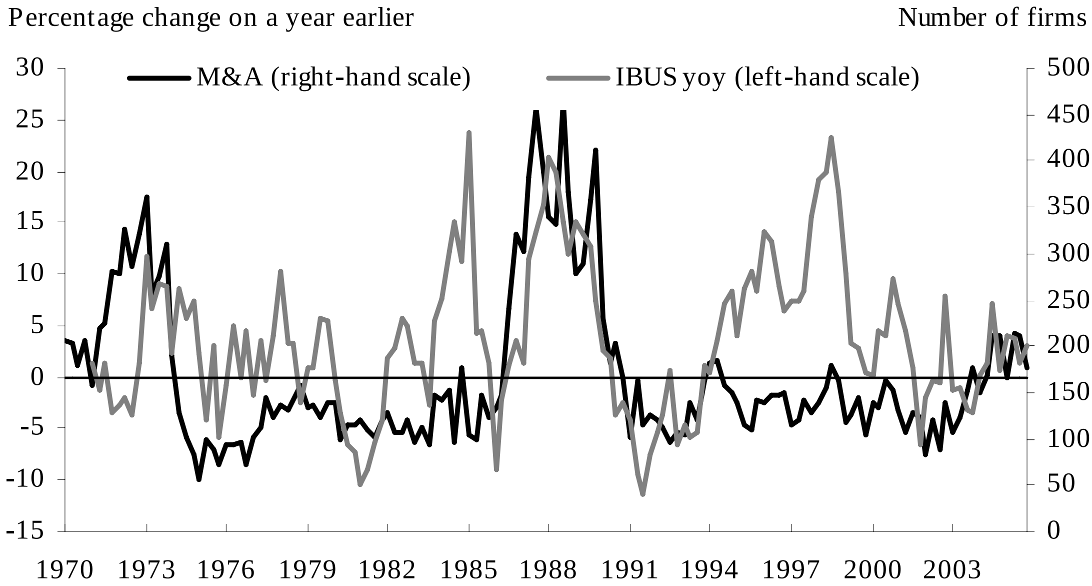

6.1.1 Time-averaged data: sales

Beginning with the time-averaged results, Panel A of Table 3 reports our baseline results. The unconstrained results without sectoral dummies provide a well-determined point estimate for σ of 0.32. When sectoral dummies are added, the point estimate falls to 0.13, but the standard error increases. This value is not significantly different from the estimate without dummies, or zero.[38] Given the construction of the user cost, the variation is by industry, and the sectoral dummies are effectively capturing the variation.

The intriguing result, however, is the coefficient on sales. This is below unity, which implies increasing returns to scale. Returns are calculated as $\eta = (1 - \sigma)/(\beta - \sigma)$. Constant returns to scale (CRS) requires $\beta = 1$ for all values of σ, so a sufficient test for CRS is $\beta = 1$, which is clearly rejected. The point estimate of η is reported in the table, together with the standard error calculated by the Delta method, as in Chirinko et al. (2004). In the baseline case the point-estimate is 2.90, numerically distant from unity. Moreover, the difference appears to be highly significant. However, inference may be suspect here. The Delta method is a first order approximation, but η is

[38]This result is similar to that in Chirinko et al. (2004) (point 1 in Section 3.1).

non-linear in σ and β. To compound the problem, there is a high degree of curvature in the relation at the parameter values we observe, increasing as $\beta - \sigma$ approaches zero. Intuitively, when $\beta - \sigma$ is small changes in the parameters are highly levered.[39] This means that although we reject CRS, it is hard to be confident about the degree to which η exceeds unity. It is also worth pointing out that the estimate of σ remains well below unity when CRS is imposed.

This result is of course of some independent interest. Chirinko et al. (2004) find increasing returns, but our estimates are larger. Are these results plausible? Hall (1990) reports estimates showing widespread increasing returns at one and two digit level, although Caballero and Lyons (1990) in a similar exercise, looking at European data including the UK, argue that the evidence points to external (industry level) returns, and Basu and Fernald (1997) suggest that constant or decreasing returns are more common in US industries. But strictly comparable estimates of returns are in fact hard to come by. Many investigators using firm-level data condition on Tobin's Q rather than the user cost. And those who do use the user cost frequently impose CRS, examining only capital (or investment) to output ratios.[40] In the final column we impose this restriction in our data set. The estimate of σ rises, but is still below 0.5.[41] We return to the issue of returns to scale in the next section.

The time-average method does not require a balanced panel, but as a robustness test and to help compare the results with those from the panel, we report two balanced samples. The results are not much changed, although in Panel B the point estimate of η is extremely large. But this is very badly determined (because, as discussed above, $\beta - \sigma$ is close to zero, at 0.03). As a further robustness test, we split the sample into large and small firms (using employment as the indicator). Table 4 shows the key results. For the unconstrained and CRS cases without sectoral dummies the estimates lie within roughly one standard error of the complete sample estimate. For the case with sectoral dummies, the results are much more dispersed for the elasticity, but not for the scale estimate. In this case the correlation between industry dummies and user cost is more acute of course, due the restricted samples.

[39] Simulations suggest that in the parameter regions we observe the distribution of η is highly skewed with fat tails.

[40] *E.g.*, Caballero et al. (1995).

[41] This result is similar to that obtained by Chirinko et al. (2004) (point 2 in Section 3.1).

Our data are likely to be poorly measured. Is this a source of bias for returns to scale? Chirinko et al. (2004) show that in their sample the Rao (1973) test indicates little bias as a consequence. Repeating those tests on our data also suggests there is little bias. The estimated degree of bias in β in Table 4 suggests the absolute bias never exceeds 0.011.

6.1.2 Dynamic panel estimates: sales

The time averaging method ignores all information about short-run dynamics. As discussed in Section 3.2 above, an alternative method of estimating long-run parameters is the Pooled Mean Group, which may be more efficient, and which we therefore prefer.

Given the non-stationarity of the data, this method requires the existence of a cointegrating relationship. Consequently, we begin by reporting Pedroni (1999, 2004) panel tests based on the two balanced panels, where this is straightforward. As explained above in Section 3.2.3, there are seven tests. Table 5 shows that the majority of tests cannot reject the null of no cointegration in all the series. But the small sample performance of these tests is variable and, as Pedroni (2004) shows, the size and power of the test is very sensitive to sample size. In particular, the time series is very short for the large N case and this may lead to failure to reject the null (Type I error), as (speaking somewhat loosely) the tests are based on mean-reversion.

The 'group' tests in the last three rows of Table 5 based on the 'between-dimension' may be considered the most apprioriate as they allow for heterogeneity in the autoregressive process. Even among these tests, only the parametric ADF type test (group t) rejects the null of no cointegration for both samples, although it does so quite strongly. The group non-parametric t test also rejects the null, albeit at only the 10% level. In a recent paper, Westerlund and Basher (2006) suggest that even these test statistics are far from robust. Evidence presented in Pedroni (2004) indicates that the group t statistics may be undersized for small T and the other statistics oversized. Taken together, the evidence in Table 5 cannot firmly reject the hypothesis of no cointegration. However, the evidence from these tests is ambiguous and weak. The most relevant large T group evidence does clearly reject the null. Given this result and the fact that the dynamic model reported below identifies significant negative error correction terms (that is, adjusting towards a long-run cointegrating relationship), we proceed to estimation on the assumption that there is cointegration and that the PMG method is therefore valid.

Table 3: Time-averaged results (sales)

Panel A. Unbalanced panel: N = 403			
	Unconstrained		CRS
	Benchmark	Sectoral dummies	
σ	0.32	0.14	0.46
	0.05	*0.13*	*0.06*
β	0.55	0.54	1
	0.03	*0.03*	
η	2.90	2.15	1
	0.48	*0.37*	
$\bar{R}^2$	0.46	0.53	
Root MSE	0.33	0.34	0.43
Panel B. Balanced panel large N: N = 261, T = 14			
	Unconstrained		CRS
	Benchmark		
σ	0.42		0.55
	0.07		*0.09*
β	0.45		1
	0.04		
η	23.8		1
	68.9		
$\bar{R}^2$	0.36		
Root MSE	0.33		0.43
Panel C. Balanced panel large T: N = 142, T =26			
	Unconstrained		CRS
	Benchmark		
σ	0.31		0.49
	0.08		*0.10*
β	0.49		1
	0.05		
η	3.80		1
	1.37		
$\bar{R}^2$	0.40		
Root MSE	0.38		0.49

Standard errors italicised

Table 4: Time-averaged results (sales): split sample

Unbalanced panel of large firms: N = 190		
	Unconstrained	CRS
	Benchmark	
σ	0.40	0.52
	0.06	*0.08*
β	0.59	1
	0.04	
Unbalanced panel of small firms: N = 213		
	Unconstrained	CRS
	Benchmark	
σ	0.24	0.39
	0.07	*0.08*
β	0.52	1
	0.05	

Standard errors italicised

Table 5: Pedroni panel cointegration tests

	N = 142, T = 26	N = 261, T = 14
panel v	-7.11	-8.19
panel ρ	1.12	6.37
panel PP t	-1.59*	0.65
panel t	-0.95	-0.05
group ρ	3.63	13.19
group PP t	-1.48*	4.85
group t	-2.98***	-4.97***

No trends
* indicates significant at 10% level
*** indicates significant at 1% level

The PMG results are presented in Panel A of Table 6 for the longer balanced panel. We report results for a range of lag structures.[42] In each case the estimates of σ are similar, and similar to those from the time averaged method reported above. Note that the average estimate of the ECM term is within the appropriate range, and well determined: as just observed, we take this as evidence for cointegration. The Hausman misspecification test rejects the PMG specification in the case where lags are selected by the Akaike information criteria, which might be due to over-parametrisation, but in the other three cases is not rejected. As in the previous results, the point estimates for β are significantly below unity, but in this case are closer to CRS than in the time-averaged case, and similar to those reported in Chirinko et al. (2004). Once again, the Delta method produces large standard errors for η so inference about the precise value is difficult, but we are able to reject constant returns (from the estimate for β).

Given the Hausman test results, the mean group results are redundant, but we nevertheless report them in Panel B. The inefficiency is apparent in the dispersed estimates, large standard errors and economically inadmissable values. In practice, the Mean Group estimates may be very sensitive to outliers (Pesaran et al. (1999)), and this seems to be the case here.[43]

We also report conventional panel estimates: for a purely static model; and a dynamic model where all parameters are pooled. The results continue to suggest that σ lies well below unity.

Does the estimated degree of returns to scale reduce the plausibility our results? As discussed above, it is difficult to make comparisons with the literature because most studies maintain constant returns. As mentioned

[42]The AIC is the Akaike Information Criteria, and the SBC the Schwarz Bayesian Criteria. Both impose a parameter penalty and tend to deliver similar lag structures, although the AIC tends to have longer lags. In both cases the maximum lag is restricted to 3. The median lag on the dependent variable is 2 for the AIC and 1 for the SBC: the median lag on both of the explanatory variables is 1 and 0 for the AIC and SBC respectively.

[43]As mentioned in Section 3.2, estimation by OLS is feasible, although inefficient. Using only one lag of the dependent variable and imposing CRS, $\hat{\sigma} = 0.20$, with a standard error of 0.03. Allowing the individual returns to scale to be free (but with no lags other than the error correction, as estimation is not otherwise feasible) there is a similar result: 0.17, with standard error 0.03. In this case the average value of $\hat{\beta}$ is 0.91, but this mean group estimate is very poorly determined: the standard error is 8.45. If we allow both long-run parameters to be free the point estimates of β and σ are 0.53 and 0.59 respectively: these are outside the admissable set but come with extremely large mean-group standard errors. Given the inefficiency, the results for σ are broadly consistent with those from the ML method.

above, independent evidence (Caballero and Lyons (1990)) tends to suggest that, in industries where there are increasing returns, it is due to external effects and should not identified within a firm, although it may be that this effect is in practice relegated to a trend term spuriously captured by the capital stock. But there are also reasons that might explain why we find increasing returns. One characteristic of our sample is that firms are unrepresentatively large. In the light of evidence that Gibrat's Law (that the size and rate of growth of firms are independent) does not hold,[44] it it possible that firm size is related to returns to scale and that our sample of larger firms is characterised by increasing returns to some degree. Another possibility is that the dynamic panel is misspecified inasmuch as it assumes labour-augmenting technological progress.[45] To the extent that technological developments increase output through other channels, these effects may be captured as increasing returns to scale in our specification.

We also apply the mean-group method of Pesaran (2006). As the method is valid only for static models, we are implicitly estimating a cointegrating relationship by an augmented static regression. The augmentation does not appear to improve the precision of the estimates relative to those reported in Table 6. In Table 7 we report mean estimates for all the results ('unrestricted'), and for cases where σ is restricted to lie in a meaningful range. The mean estimates are poorly determined, but there is no evidence of a larger estimate for σ or β. The sample with the shorter T must be even more subject to caveats than those reported in Table 6. In Table 8 we report the pooled estimates. The estimates are much more precise. Those for β are virtually unchanged, while those for sigma differ rather more. When CRS is imposed, the estimate of σ rises.

6.2 Conditioning on value added

There are fewer data available for the value-added case. Otherwise the permutations are as before: an unbalanced panel containing 273 firms; a balanced panel comprising 201 firms but with a relatively short period; the third a smaller panel (106 firms) with a longer period.

[44] See Sutton (1997) for a survey, although empirical studies have tended to focus on the manufacturing sector.

[45] Although the time-averaged results allow very general processes for technological progress.

Table 6: Panel results (sales): T = 26, N = 142

Panel A. PMG				
	1 lag	2 lags	AIC*	SBC*
σ	0.42	0.32	0.32	0.36
	0.02	*0.02*	*0.02*	*0.02*
β	0.87	0.84	0.91	0.88
	0.02	*0.02*	*0.01*	*0.02*
η	1.29	1.30	1.15	1.24
	6.10	*5.33*	*4.10*	*5.03*
ECM	-0.13	-0.13	-0.16	-0.15
	0.01	*0.00*	*0.025*	*0.014*
Hausman test	2.2	0.3	5.7	4.0
p-values	0.34	0.87	0.05	0.14
Panel B. MGE				
	1 lag	2 lags	AIC*	SBC*
σ	-0.68	0.32	0.76	0.91
	1.12	*0.19*	*0.20*	*0.29*
β	6.87	1.11	0.34	0.68
	4.21	*0.55*	*0.30*	*0.00*
η	0.22	0.86	-0.58	0.33
	0.11	*2.26*	*4.92*	*0.00*
ECM	-0.13	-0.21	-0.22	-0.20
	0.01	*0.02*	*0.02*	*0.02*
Panel C.	Static: Fixed Effects		Dynamic: Fixed Effects	
			1 lag	2 lags
σ	0.64		0.44	0.38
	0.01		*0.03*	*0.04*
β	0.73		0.65	0.64
	0.01		*0.37*	*0.04*
η	3.79		2.61	2.42
	89.5		*27.2*	*21.1*
ECM			-0.11	-0.11
			0.01	*0.01*

Standard errors italicised
* Lags change across groups and variables
ECM estimates are mean group averages

Table 7: Heterogeneous panel with cross-sectional correlation: mean group estimates

T = 26, N =142		
	Unrestricted	Restricted
σ	0.09	0.25
	0.25	*0.22*
β	0.46	0.44
	0.57	*0.28*

T = 14, N =261		
	Unrestricted	Restricted
σ	0.18	0.29
	0.50	*0.24*
β	0.26	0.35
	0.44	*0.23*

Standard errors italicised
Constrained excludes values of σ or β outside the range $[0, 1]$.

Table 8: Heterogeneous panel with cross-sectional correlation: pooled estimates

T = 26, N =142		
	Unconstrained	CRS
σ	0.02	0.14
	0.01	*0.02*
β	0.50	1
	0.01	-

T = 14, N =261		
	Unconstrained	CRS
σ	0.13	0.25
	0.01	*0.02*
β	0.27	1
	0.01	-

Standard errors italicised
Constrained excludes values of σ or β outside the range $[0, 1]$.

6.2.1 Time-averaged data: value added

Panel A of Table 9 again reports our baseline results. The unconstrained results without sectoral dummies provide a well-determined point estimate for σ of 0.35, very close to that with sales. The coefficient on value added is again below unity. The point-estimate of η is 2.76, also near the previous results. Maintaining CRS, the estimate of σ again rises, but remains below 0.5. The balanced panel results, using a smaller sample, are rather more variable than for sales. But the overall message of $\sigma < 1$ and increasing returns is unchanged.

6.2.2 Dynamic panel estimates: value added

The dynamic panel results are presented in Panel A of Table 10 for the longer balanced panel. N is reduced to 106 and, perhaps more critically, T to only 20, which may make the results somewhat less robust - the time series spans the period 1983, an unusual point near the trough of the 1980's recession, to 2002. The point estimates for σ are generally lower and tend to vary more than for the sales case, as for the estimates for η. The Hausman tests for poolability are comfortably accepted in all cases. As with sales, the results for the mean group estimates are therefore redundant: although we report the results for completeness, in some cases the parameters are outside the admissable region. Again, we report conventional panel estimates. The results are consistent with the time averaging cases, with generally slightly lower estimates of σ and β.

Table 9: Time-averaged results (value added)

Panel A. Unbalanced panel: N = 273		
	Unconstrained	CRS
	Benchmark	
σ	0.35 *0.07*	0.40 *0.09*
β	0.59 *0.11*	1
η	2.76 *0.70*	1
$\bar{R}^2$ Root MSE	0.48 0.39	 0.46
Panel B. Balanced panel large N: N = 201		
	Unconstrained	CRS
	Benchmark	
σ	0.26 *0.07*	0.45 *0.09*
β	0.47 *0.05*	1
η	3.65 *1.11*	1
$\bar{R}^2$ Root MSE	0.34 0.30	 0.39
Panel C. Balanced panel large T: N = 106		
	Unconstrained	CRS
	Benchmark	
σ	0.50 *0.12*	0.67 *0.14*
β	0.58 *0.07*	1
η	5.97 *7.78*	1
$\bar{R}^2$ Root MSE	0.41 0.38	 0.44

Standard errors italicised

Table 10: Panel results (value added): : T = 20, N = 106

Panel A. PMG				
	1 lag	2 lags	AIC*	SBC*
σ	0.19	0.29	0.15	0.09
	0.02	*0.03*	*0.01*	*0.02*
β	0.81	0.64	0.64	0.50
	0.03	*0.03*	*0.01*	*0.02*
η	1.32	2.02	1.72	2.22
ECM	-0.16	-0.16	-0.19	-0.22
	0.01	*0.01*	*0.025*	*0.03*
Hausman test	2.15	0.96	3.32	0.79
p-values	0.34	0.62	0.19	0.71
Panel B. MGE				
	1 lag	2 lags	AIC*	SBC*
σ	3.87	0.68	2.60	0.39
	3.48	*0.49*	*1.61*	*0.36*
β	-16.01	0.46	-3.90	0.78
	16.79	*0.43*	*5.44*	*0.54*
η	0.14	-1.39	0.25	0.33
	0.11	*2.26*	*4.92*	*0.00*
ECM	-0.20	-0.25	-0.25	-0.35
	0.02	*0.02*	*0.03*	*0.03*

Panel C.	Static: Fixed Effects	Dynamic: Fixed Effects	
		1 lag	2 lags
σ	0.47	0.35	0.39
	0.02	*0.04*	*0.05*
β	0.61	0.66	0.59
	0.02	*0.04*	*0.04*
η	3.72	2.11	3.03
ECM		-0.17	-0.15
		0.01	*0.08*

Standard errors italicised
* Lags change across groups and variables
ECM estimates are mean group averages

7 Conclusions

The elasticity of the capital stock with respect to the user cost is determined by the elasticity of substitution between capital and other factors. There has been a debate about the value of this parameter, with much of the focus on how best to estimate the long-run parameter. Using a data set spanning over thirty years and with 403 firms, we estimate the long-run relationship determining the capital stock. We find that there is robust evidence for a user cost elasticity that is substantially below unity. Our results suggest the parameter lies in the region of 0.4, which is consistent with the aggregate time-series evidence for the United Kingdom, and with some firm-level estimates for the United States.

Another clear result is that there is increasing returns to scale, although the non-linear relationship between the returns and the estimated parameters make inference difficult. In our preferred results, the point estimate is similar to that found in Chirinko et al. (2004), although generally our estimates are higher - arguably, implausibly high, given alternative evidence on returns. However, whether this is particular to our data set is hard to say, as the hypothesis of constant returns is generally maintained in comparable research. When we impose that restriction, we continue to obtain estimates for the user-cost elasticity that continue to lie well below unity. While that elasticity seems to be robustly determined,the degree of increasing returns suggests further investigation is warranted.

A Definitions of Variables and Industries

This appendix describes in detail the construction of each of the variables discussed in Subsection 4.1 and details the industry classification. The sample for company accounts from DataStream/Worldscope runs from 1970 to 2005.[46]

Company accounts dated up to the end of May in any given year are reassigned to the previous year, as the reported data are likely to relate primarily to economic activity during the previous calendar year.[47] National Accounts-based data in the BEID are from the 2005 Blue Book.

A.1 Construction of variables

Investment: Gross investment is defined according to the discussion in 4.1 up to 1991 as:

$$INV = ds435(TotalNewFixedAssets)$$

Following tax/accounting changes in 1992, INV is defined as:

$$INV = ds1024(Payments : FixedAssets)$$

This excludes $ds479(FixedAssets(Subs)$, for which data are no longer available. An alternative definition allowing for disposals would be defined as $INV2 = ds435 - ds423(SalesOfFixedAssets)$ up to 1991 and thereafter $INV2 = ds1026(Netpayment : FixedAssets)$.[48] Construction of the capital stock is discussed separately below.

Real sales: These are used to proxy for real value added and are constructed as:

$$s_{t,f} = ds104(TotalSales)_{t,f}/ppi_{t,i}$$

[46] It was downloaded in April 2006 and augmented by earlier versions of the same dataset to incorporate either firms that are no longer active or some series for which Worldscope no longer provides data.

[47] Most accounts in the dataset are dated in March or December (this is the most common month). Where a firm changes accounting date so that there are two data observations for a given year, the earlier observations are dropped.

[48] Chirinko and Schaller (2004b) make further adjustments where there are large acquisitions/divestitures (see also Chirinko et al. (1999)).

where $ppi_{t,i}$ is the industry-level output price deflator from the BEID.

Real value added: There is some limited information on employment costs and operating profits that we use to construct a measure of value added:

$$va_{t,f} = (ds117(EmploymentCosts)_{t,f} + ds137(Operatingprofits)_{t,f})/ppi_{t,i}$$

where $ppi_{t,i}$ is the industry-level output price deflator from the BEID.

Weighted-Average Cost of Capital (WACC): the cost of finance is the weighted-average cost of capital defined as an annual average of end-month calculations for:

$$r_t = \varpi \left(dy_t \left(1 + g \right) \right) + \left(1 - \varpi \right) \left(r_t^{LT} + s_t \right)$$

where ϖ is the share of equities in total liabilities since 1987 Q2 (based on National Accounts data) and assumed to be constant at the 1987 Q2 level of 0.14 since 1982. dy is the dividend yield for all UK non-financial firms and g is the assumed real rate growth rate of profits (set at $g = 0.025$ to be consistent with standard practice.[49] r^{LT} is the 10-year real interest rate on government debt, derived from index-linked government bonds using the VRP methodology since 1985 and the Svensson methodology previously. s is the spread of corporate borrowing rates over the risk-free rate from the Merrill Lynch UK investment-grade corporate option-adjusted spread (OAS) from 1997, but using an estimated series for earlier dates based on subtracting 10-year government yields from various corporate bond data found in the Global Financial Data dataset. Prior to 1982, the WACC series is the backcast data generated in earlier work in the Bank (see Ellis and Price (2004)).[50] This backcasting method has the obvious disadvantage of not being based directly on data but has the advantage of not requiring explicit assumptions to be made about inflation expectations in the 1970s.

User cost: The industry-level user cost of capital is defined as:

$$c_{t,i} = \sum_{w=1}^{w=7} \left(w_{t,i,n} c_{t,i,n} \right) \text{ where } c_{t,i,n} = p_{t,i,n} \left[r_t + d_{t,i} - \Delta p_{t,i,n} \right] T_{t,n}$$

where $c_{t,i,n}$ is the industry-specific user cost of asset n, $w_{t,i,n}$ is the share of each asset in the capital stock of that industry, $p_{t,i,n}$ is the price of a given type of capital asset in a particular industry, r_t is the cost of finance (WACC), $d_{t,i}$ is the asset-specific rate of deprecation and $T_{t,n}$ is a tax factor

[49] For example, Chirinko et al. (2004) use a value of 0.024.

[50] This was used to backcast an earlier series. The end points are very close so we do not splice the data.

(see Basu et al. (2003)). $\Delta p_{t,i,n}$ is calculated on a forward-looking basis (*i.e.*, $\Delta p_{t,i,n} = p_{t+1,i,n} - p_{t,i,n}$).[51]

A.2 Other variables

We also constructed some standard variables not used in this analysis.

Tobin's Q: This is approximated in the standard way using the average level of Q, which implicitly includes short-term debt (consistent with Bond et al. (2004)):[52]

$$Q_t = \frac{dsMV(MarketValue) + ds321(TotalLoanCapital) - ds390(NetCurrentAssets)}{K}$$

Cash flow: This is defined as:

$$cf_t = ds175(AfterTaxProfit.Adjusted) + ds136(Depreciation)$$

Note that Bond et al. (2004) use an alternative definition of after tax profits ($ds182$).

Cash stock is total cash and cash equivalent (*ds375*).

Dividends are defined as ratio of dividends (*ds187*) to cash flow. Negative payout ratios (due to negative profits) are recorded so that they are considered as extremely high payout ratios.

B Construction of capital stock

This appendix discusses construction of the real capital stock ('economic' or 'economist's' capital stock) from book-value measures along the lines proposed by Chirinko and Schaller (2004a). For a given initialisation, the real capital stock is constructed using the PIM method:

$$k_{t,f} = k_{t-1,f}\,(1 - d_{t,i}) + inv_{t,f} \tag{B1}$$

where $d_{t,i}$ is the industry-depreciation rate from the BEID. The key issue is how to initialise the firm-level capital stock ($k_{0,f}$). The procedure aims to

[51] Strictly, this should be $E_t(\Delta p_{t,i,n})$. This series is fairly persistent so the practical relevance of distinction may be limited.

[52] This methodology derives from Blundell et al. (1992).

find the best estimate of this in terms of the economic real capital $k^E_{0,f}$ at replacement cost, given $k^A_{0,f}$ (accountant's capital). Estimates of the capital stock derived using the PIM under either measure will converge at some point, but Chirinko and Schaller (2004a) find that satisfactory convergence in the *growth rate* of capital is only achieved by $t = 15$. By applying suitable adjustments to accountant's capital, all observations of the derived capital stock prior to $t = 15$ can in principle be used as estimates of economic capital.

We correct for (i) the price-level, (ii) the 'price distortion' between book-value and historic cost due to inflation, and (iii) the 'depreciation distortion' due to the different rates of depreciation applied by accountants and economists. The overall approach is summarised by equation A-1 from the paper by Chirinko and Schaller (2004a), where the initial capital stock $\widehat{k}^E_{f,0}$ is set using historic cost estimates from company accounts ($K^H_{f,0}$):

$$\widehat{k}^E_{f,0} = \frac{K^H_{f,0}}{p^I_{i,0}} \Gamma_{i,0} \left[\omega[\delta], \pi\right] IVC_{i,0} \tag{B2}$$

The first term on the right-hand side $\left(\frac{K^H_{f,0}}{p^I_{i,0}}\right)$ is an adjustment to historic cost estimates to give accountant's capital in real terms (k^A, which is adjusted for the price-level in period 0), the second term (Γ) is an adjustment for the wedge between historic and replacement cost due to inflation, and the third term is a correction for the basic Initial Value Problem ('depreciation distortion': IVC). i refers to industry variables and f refers to firm variables. The hat denotes estimated rather than measured variables.

The correction for the depreciation distortion ($IVC_{i,0}$) is calculated as the ratio of 'true' economist's capital to accountant's capital, where 'true' is defined as observations for $t = 15$, after which the capital stock should have achieved satisfactory convergence. We calculate $\widehat{k}^E_{f,15}$ using the PIM initialised with $\widehat{k}^A_{f,0}$ = actual book value (B2 with $\Gamma_{i,0} = 1$). We then calculate industry averages of the ratio of $\frac{\widehat{k}^E_{f,15}}{\widehat{k}^A_{f,15}}$ to $t > 15$ using a a balanced panel:[53]

$$IVC_{i,t} = \frac{\widehat{k}^E_{i,t}}{\left(K^H_{i,t}/p^I_{i,t}\right)} \tag{B3}$$

where K^N is the nominal counterpart of $\widehat{k}^E$. We set $t = 15$, as $\widehat{k}^E_{15}$ is assumed to be a good approximation to k^E_{15}. It is further assumed that this ratio is

[53] The use of industry averages is less compelling in our dataset than for Chirinko and Schaller (2004a) as there are considerably fewer observations for each industry in this case.

a constant so that $IVC_{i,0} = IVC_{i,15}$. This allows us to estimate a suitable correction for period $t = 0$ to the initial value problem.

In principle we also need to adjust for the 'price distortion', $\Gamma[\omega[\delta], \pi]$. Chirinko and Schaller (2004a)[54] use national accounts industry-level data:

$$\Gamma[\omega[\delta], \pi] = K^N_{i,0}/K^H_{i,0} \tag{B4}$$

where K^N_0 is capital at nominal replacement cost. Unfortunately, there are no longer data in the UK national accounts for $K^H_{i,t}$ (capital at historical cost) so this method is not available to us. Γ is a function of π, which varies markedly between $t = 0$ and $t = 15$ (*e.g.* the mid-1970s and the late 1980s). But we assume this is swept up into the $t = 15$ adjustment.

The period when $t = 15$ is typically around 1988/89. This coincides with and follows a period of intense investment activity, which might tend to bring historic cost and economic capital relatively close to each other (as the role of depreciation of existing assets becomes relatively less important). Simulations, however, suggest that this would not make a large amount of difference to the ratio between the two.

We initialise capital stocks for the first observation of each firm in our sample (rather than the first year of a particular variant of the dataset). This both maximises the number of observations for each firm and partly helps to counter bias in our estimation of the capital stock by introducing some random variation in the starting dates.

C Definition of industries by three-digit SIC code and BEID industry

Firms in the DSW database are allocated a three-digit SIC (Standard Industrial Classification, 1992) on the basis of the sector with the largest share of its sales.[55] The BEID uses a 32-industry classification. This table shows how two-digit SIC industries correspond to those in the BEID.

[54] See Equation 21 of their paper.

[55] SIC92 is the 1992 version of the U.K.'s Standard Industrial Classification. It is identical to the European NACE system. Details on SIC92 industry codes can be found at http://www.statistics.gov.uk/methods_quality/sic/contents.asp

Table C1: Industrial classification

Number	Name	SIC	Number of obs in 1998
1	Agriculture	01,02,05	
2	Oil and gas	11,12	
3	Coal & other mining	10,13,14	7
4	Manufactured fuel	23	
5	Chemicals & pharmaceuticals	24	11
6	Non-metallic mineral products	26	
7	Basic metals & metal goods	27,28	
8	Mechanical engineering	29	
9	Electrical engineering & electronics	30,31,32,33	52
10	Vehicles	34,35	45
11	Food, drink & tobacco	15,16	
12	Textiles, clothing & leather	17,18,19	
13	Paper, printing and publishing	21,22	5
14	Other manufacturing	20,25,36,37	25
15	Electricity supply	40.1	
16	Gas supply	40.2,40.3	
17	Water Supply	41	5
18	Construction	45	10
19	Wholesale, vehicle sales & repairs	50,51	37
20	Retailing	52	
21	Hotels & catering	55	
22	Rail transport	60.1	
23	Road transport	60.2,60.3	
24	Water transport	61	37
25	Air transport	62	
26	Other transport services	63	
27	Communications	64	19
28	Finance	65,66	35
29	Business Services	67, 70,71,72,73,74	8
30	Public administration and defence	75	1
31	Education	80	
32	Health and social work	85	
33	Waste treatment	90	
34	Miscellaneous services	91-99	31

References

Antràs, P. (2004). Is the U.S. aggregate production function Cobb-Douglas? New estimates of the elasticity of substitution. *Contributions to Macroeconomics 4*.

Arellano, M. and S. Bond (1991). Some tests of specification for panel data: Monte Carlo evidence and an application to employment equations. *Review of Economic Studies 58*, 277–320.

Baltagi, B. H. and J. M. Griffin (1997). Pooled estimators vs. their heterogeneous counterparts in the context of dynamic demand for gasoline. *Journal of Econometrics 77*, 303–27.

Barnett, W. and T. Sakelleris (1998). Non linear responses of firm investment to Q: testing a model of convex and non-convex adjustments costs. *Journal of Monetary Economics 42*, 261–88.

Barro, R. J. and X. Sala-i-Martin (1995). *Economic Growth*. McGraw-Hill.

Basu, S. and J. G. Fernald (1997). Returns to scale in U.S. production: estimates and implications. *The Journal of Political Economy 105*, 249–83.

Basu, S., J. G. Fernald, N. Oulton, and S. Srinivasan (2003). The case of the missing productivity growth: or, does information technology explain why productivity accelerated in the US but not the UK? *NBER Working Paper*.

Blundell, R., S. Bond, M. Devereux, and F. Schiantarelli (1992). Investment and Tobins Q: evidence from company panel data. *Journal of Econometrics 51*, 233–57.

Bond, S., D. Harhoff, and J. Van Reenen (1999). Investment, R and D, and financial constraints in Britain and Germany. *IFS Working Paper 55.*

Bond, S., A. Klemm, R. Newton-Smith, M. Syed, and J. Vlieghe (2004). The roles of expected profitability, Tobin's Q and cash flow in econometric models of company investment. *Bank of England Working Paper*.

Breitung, J. and M. Pesaran (2005). Unit roots and cointegration in panels. In L. Matyas and P. Sevestre (Eds.), *The Econometrics of Panel Data* (Third ed.). Kluwer Academic Publishers.

Caballero, R. (1999). Aggregate investment. In J. Taylor and M. Woodford (Eds.), *Handbook of Macroeconomics*. Amsterdam: North Holland. Chapter 12.

Caballero, R. J. (1994). Small sample bias and adjustment costs. *The Review of Economics and Statistics 76*, 52–8.

Caballero, R. J., E. Engel, and J. C. Haltiwanger (1995). Plant-level adjustment and aggregate investment dynamics. *Brookings Papers on Economic Activity*, 1–54.

Caballero, R. J. and R. K. Lyons (1990). Internal versus external economies in european industry. *European Economic Review 34*, 805–30.

Carpenter, R. and A. Guariglia (2003). Cash flow, investment opportunities: New tests using UK panel data. *University of Nottingham mimeo*.

Chirinko, R. S. (1993). Business fixed investment spending: modeling strategies, empirical results and policy implications. *Journal of Economic Literature 31*, 1875–911.

Chirinko, R. S., S. Fazzari, and A. Meyer (2004). That elusive elasticity: a long-panel approach to estimating the price sensitivity of business capital. *Unpublished*.

Chirinko, R. S., S. M. Fazzari, and A. P. Meyer (1999). How responsive is business capital formation to its usercosts? an exploration with micro data. *Journal of Public Economics 74*, 53–80.

Chirinko, R. S. and D. Mallick (2006). The substitution elasticity, growth theory, and the low-pass filter panel model. *Unpublished*.

Chirinko, R. S. and H. Schaller (2004a). The initial value problem. *Unpublished*.

Chirinko, R. S. and H. Schaller (2004b). The irreversibility premium. *Mimeo*.

Coen, R. M. (1969). Tax policy and investment behavior: comment. *American Economic Review 59*, 370–79.

Criscuolo, C., J. Haskel, and R. Martin (1999). Building the evidence base for productivity policy using business data linking. *American Economic Review 89*.

Doms, M. and T. Dunne (1994). Capital adjustment patterns in manufacturing plants. *Center for Economic Studies Discussion Paper 94.*

Ellis, C. and S. Price (2004). UK business investment and the user cost of capital. *Manchester School 72*, 72–93.

Goolsbee, A. (1998). Investment tax incentives, prices, and the supply of capital goods. *Quarterly Journal of Economics 113*, 121–48.

Hall, R. (1990). Invariance properties of Solow's productivity residual. In P. Diamond (Ed.), *Growth/Productiviuty/Unemployment: Essays to Celebrate Bob Solow's Birthday.* Cambridge, Mass.: Harvard University Press.

Hall, R. E. (1995). Comment on 'plant-level adjustment and aggregate investment dynamics'. *Brookings Papers on Economic Activity*, 47–51.

Hamermesh, D. S. (1993). *Labor Demand.* Princeton University Press, Princeton.

Harrison, R., K. Nikolov, M. Quinn, G. Ramsay, A. Scott, and R. Thomas (2005). *The Bank of England Quarterly Model.* Bank of England.

Hayashi, F. (1982). Tobin's marginal q and average q: a neoclassical interpretation. *Econometrica 50*, 213–44.

Holly, A. (1982). A remark on Hausman's specification test. *50*, 749–60.

Holly, S., M. H. Pesaran, and T. Yamagata (2006). A spatio-temporal model of house prices in the US. *Unpublished*.

Jones, C. I. (2003). Growth, captal shares and a new perspective on the production function. *Unpublished, UC Berkeley*.

Jorgensen, D. W. (1963). Capital theory and investment behaviour. *American Economic Review, Papers and Proceeedings 247-59.*

Kapetanios, G., M. H. Pesaran, and T. Yamagata (2006). Panels with nonstationary multifactor error structures. *Unpublished*.

Kiyotaki, N. and K. West (1996). Business fixed investment and the recent business cycle in japan. *NBER Macroeconomics Annual*, 277–323.

Klump, R., P. M. and A. Willman (2007). The long-term sucCESs of the neoclassical growth model. *Oxford Review of Economic Policy,* forthcoming.

Krusell, P., L. E. Ohanian, J.-V. Rios-Rull, and G. L. Violante (2000). Capital-skill complementarity and inequality: a macroeconomic analysis. *68*, 1029–53.

Nadiri, M. I. (1970). Some approaches to theory and measurement of total factor productivity: a survey. *Journal of Economic Literature 8*, 1117–77.

Nerlove, M. (1967). Recent empirical studies of the CES and related production functions. In M. Brown (Ed.), *The Theory and Empirical Analysis of Production*. Columbian University Press, New York.

Pedroni, P. (1999). Critical values for cointegration tests in heterogeneous panels with multiple regressors. *Oxford Bulletin of Economics and Statistics 61*, 653–70.

Pedroni, P. (2004). Panel cointegration; asymptotic and finite sample properties of pooled time series tests, with an application to the PPP hypothesis. *Econometric Theory 20*, 597–625.

Pesaran, M. H. (2006). Estimation and inference in large heterogeneous panels with a multifactor error structure. *Econometrica 74*, 967–1012.

Pesaran, M. H., Y. Shin, and R. P. Smith (1999). Pooled estimation of dynamic heterogeneous panels. *Journal of the American Statistical Association 94*, 621–34.

Pesaran, M. H. and R. P. Smith (1995). Estimating long-run relationships from dynamic heterogeneous panels. *Journal of Econometrics 68*, 79–113.

Rao, P. (1973). Some notes on the errors-in-variables model. *The American Statistician 27*, 217–18.

Rossi-Hansberg, E. and M. L. J. Wright (2005). Firm size dynamics in the aggregate economy. *NBER Working Paper No. 11261*.

Sargent, T. (1979). *Macroeconomic theory*. Academic Press, London.

Schaller, H. (2006). Estimating the long-run user cost elasticity. *Journal of Monetary Economics 53*, 725–36.

Shapiro, M. D. (1986). Investment, output, and the cost of capital. *Brookings Papers on Economic Activity*, 111–52.

Sutton, J. (1997). Gibrat's legacy. *Journal of Economic Literature 35*, 40–59.

Westerlund, J. and S. A. Basher (2006). Mixed signals among tests for panel cointegration. *Unpublished*.

Whelan, K. (2003). A guide to the use of chain aggregated NIPA data. *Federal Reserve Board Finance and Economics Discussion Series*.

Zellner, A. (1962). An efficient method of estimating seemingly unrelated regressions and tests of aggregation bias. *Journal of the American Statistical Association 57*, 348–68.