

Adjustment Costs and the Identification of Cobb Douglas Production Functions

Stephen Bond

Institute for Fiscal Studies and Nuffield College, Oxford

Måns Söderbom

Centre for the Study of African Economies,
Department of Economics, University of Oxford,
and Institute for Fiscal Studies

January 31, 2005

Abstract

Cobb Douglas production function parameters are not identified from cross-section variation when inputs are perfectly flexible and chosen optimally, and input prices are common to all firms. We consider the role of adjustment costs for inputs in identifying these parameters in this context. The presence of adjustment costs for all inputs allows production function parameters to be identified, even in the absence of variation in input prices. This source of identification appears to be quite fragile when adjustment costs are deterministic, but more useful in the case of stochastic adjustment costs. We illustrate these issues using simulated production data.

JEL Classification: D20, D24, C23.

Key words: Production functions, adjustment costs, identification.

Acknowledgement: This paper is dedicated to the memory of Tor Jakob Klette, a fine scholar and an inspiration to those who knew him. We thank Victor Aguirregabiria, Bronwyn Hall, participants in a conference at the University of Oslo and a seminar at the University of Göteborg for helpful comments. Financial support from the ESRC Centre for Public Policy at the Institute for Fiscal Studies is gratefully acknowledged. The CSAE is funded by the ESRC.

1 Introduction

This paper considers the role of adjustment costs for inputs in the identification and estimation of parameters for a correctly specified Cobb Douglas production function. We focus on the case where input prices are common to all firms at the same point in time, and where any observed time series variation in these common input prices cannot be used to identify production function parameters due to the inclusion of time-specific intercepts. In this case Akerberg and Caves (2003) have recently shown that the control function estimation procedures proposed by Olley and Pakes (1996) and Levinsohn and Petrin (2003) do not identify parameters on inputs that are perfectly flexible. We note that this identification problem is more general. If all inputs are chosen optimally and are perfectly flexible in the sense that they can be varied immediately without incurring any costs, then all inputs are perfectly collinear with the productivity shocks observed by firms. If some inputs are predetermined so that they cannot be adjusted in response to the current productivity shock, the remaining variable inputs are linearly dependent on the productivity shock and the value of these predetermined inputs. In both cases the parameters on the perfectly flexible inputs are not identified, regardless of the estimation technique considered.

We note that this identification problem becomes less acute if all inputs are costly to adjust. The presence of adjustment costs and productivity shocks that vary across firms implies that the shadow prices of inputs vary across firms, even if all firms face common purchase or rental prices. This breaks the collinearity between productivity shocks and optimally chosen input levels. More importantly, this breaks the collinearity between the levels of different inputs, at least in cases where these inputs are subject to different levels of adjustment costs, or where adjustment costs are themselves stochastic. As a result, parameters on inputs that are subject to adjustment costs can be estimated consistently using instrumental variables methods, such as those considered by Blundell and Bond (2000). This

is also convenient as it does not require the precise form of the adjustment cost functions to be specified. In contrast, the presence of adjustment costs for all inputs implies that the control functions proposed by Olley and Pakes (1996) and Levinsohn and Petrin (2003) are mis-specified, and these methods do not yield consistent estimates of production function parameters.

We illustrate these issues using simulated data on optimal inputs and outputs for firms with a two-factor Cobb Douglas technology and (possibly stochastic) quadratic adjustment costs for both capital and labour. The results of a small Monte Carlo experiment suggest that consistent estimates of both production function parameters can be obtained using predetermined lagged inputs as instruments for the endogenously determined current inputs. With deterministic adjustment costs, identification is shown to become weak in three cases: i) where adjustment costs for both inputs are too low, so that both inputs become highly collinear; ii) where adjustment costs for both inputs are too high, so that there is insufficient variation in the input levels; iii) where adjustment costs for the two inputs are too similar, which again results in a high level of collinearity. With stochastic adjustment costs, the performance of these instrumental variables estimators improves considerably, and we are also able to estimate dynamic specifications with serially correlated productivity shocks.

The paper is organised as follows. Section 2 reviews how identification fails when one or more of the inputs is perfectly flexible. Section 3 considers the behaviour of inputs in the presence of adjustment costs. Section 4 considers the behaviour of simple instrumental variables estimators using simulated production data. Section 5 concludes.

2 Identification Issues in the Absence of Adjustment Costs

We first review the identification issues that arise when some or all of the inputs are perfectly flexible. Consider a two-factor Cobb Douglas technology for firm i

$$Y_i = A_i K_i^\alpha L_i^\beta \quad (1)$$

where Y_i, K_i, L_i are observed measures of output, capital and labour respectively, and A_i is a productivity characteristic that is observed by the firm but not by the econometrician.¹ The firm chooses inputs and output to maximise net revenue

$$R_i = PY_i - UK_i - WL_i \quad (2)$$

where input and output prices are common to all firms and, for simplicity, assumed to be exogenously given. The firm hires labour at the wage W per unit and, for simplicity, is assumed to rent capital at the rental price U per unit. We assume $\alpha + \beta < 1$ to ensure this problem has a solution.

The first order conditions for optimal input choices are the standard marginal productivity conditions

$$P \left(\frac{\partial Y_i}{\partial K_i} \right) = U \quad (3)$$

$$P \left(\frac{\partial Y_i}{\partial L_i} \right) = W \quad (4)$$

Solving these, we can express the optimal capital and labour inputs as log linear functions of real input prices and the productivity characteristic

$$k_i = \left(\frac{1 - \beta}{1 - \alpha - \beta} \right) \ln \alpha + \left(\frac{\beta}{1 - \alpha - \beta} \right) \ln \beta - \left(\frac{1 - \beta}{1 - \alpha - \beta} \right) (u - p) - \left(\frac{\beta}{1 - \alpha - \beta} \right) (w - p) + \left(\frac{1}{1 - \alpha - \beta} \right) a_i \quad (5)$$

$$l_i = \left(\frac{\alpha}{1 - \alpha - \beta} \right) \ln \alpha + \left(\frac{1 - \alpha}{1 - \alpha - \beta} \right) \ln \beta - \left(\frac{\alpha}{1 - \alpha - \beta} \right) (u - p) - \left(\frac{1 - \alpha}{1 - \alpha - \beta} \right) (w - p) + \left(\frac{1}{1 - \alpha - \beta} \right) a_i \quad (6)$$

¹We omit time subscripts in this section to simplify the notation.

where lower case letters denote natural logarithms of the terms denoted by the corresponding upper case letters. Thus if input prices are common to all firms and capital and labour are both chosen optimally with no adjustment costs or frictions, the levels of capital and labour will vary across firms only with the factor neutral productivity characteristic. For the log linear production function

$$y_i = \alpha k_i + \beta l_i + a_i \quad (7)$$

this implies that the two inputs are perfectly collinear with each other, and with the error term.²

Somewhat different issues arise when the level of one of the inputs, say capital, is chosen before the productivity characteristic a_i is observed by the firm. In this case the optimal level of the perfectly flexible labour input can be expressed as a log linear function of the real wage, the productivity characteristic and the predetermined level of capital

$$l_i = \left(\frac{1}{1-\beta} \right) \ln \beta - \left(\frac{1}{1-\beta} \right) (w - p) + \left(\frac{\alpha}{1-\beta} \right) k_i + \left(\frac{1}{1-\beta} \right) a_i \quad (8)$$

If the real wage is common to all firms, this implies that there are no valid instruments for the endogenous labour input in (7) that are informative after conditioning on the predetermined level of k_i . The only variation in l_i after conditioning on k_i is due to the productivity characteristic a_i , but any valid instruments in (7) must be orthogonal to a_i . Consequently the production function parameters α and β are not identified in this case also.

The failure of the control function approaches proposed by Olley and Pakes (1996) and Levinsohn and Petrin (2003) to identify parameters on perfectly flexible inputs in a Cobb Douglas production function is thus an example of this more

²The problem of perfect collinearity with the error term in the production function could be avoided by assuming an additive measurement error in log output. Identification still fails as a result of perfect collinearity between the two inputs. The early literature on cross-section production functions recognised this problem, and ‘solved’ it somewhat arbitrarily by invoking optimisation errors. See, for example, Marschak and Andrews (1944) and Mundlak and Hoch (1965).

general identification problem. The essence of these approaches is to augment the basic production function specification (7) with an additional function of observed variables that controls for the unobserved variation in productivity. Akerberg and Caves (2003) have noted that any correctly specified control function must be perfectly collinear with the flexible inputs, conditional on the levels of any predetermined inputs. To illustrate this, suppose we observe a variable z_i that is proportional to a_i with some unknown factor of proportionality (i.e. $z_i = a_i/\phi$), and we include this term as an additional control. We also introduce an additive measurement error in the natural logarithm of observed output, denoted e_i , so that the augmented specification remains stochastic

$$y_i = \alpha k_i + \beta l_i + \phi z_i + e_i \quad (9)$$

Clearly if both inputs are perfectly flexible, they are perfectly collinear with each other and with z_i . If capital is predetermined and labour is perfectly flexible, l_i is linearly dependent on k_i and z_i , and we again have perfect collinearity between the included terms in (9). The production function parameters α and β are not identified in either case.

We emphasise here that this problem is not unique to the control function approach. If the Cobb Douglas production function is correctly specified, then variation in prices is required to identify parameters on inputs that are perfectly flexible and optimally chosen. However the presence of unobserved variation across firms in prices basically rules out the strategy of controlling for the productivity shocks using any function of observed inputs that is common to all firms.³ Optimisation errors would need to be of a very special form for the control function to remain correctly specified, with optimisation errors for any of the inputs used in the control function being ruled out, and optimisation errors for the remaining inputs being required. Otherwise the identification of Cobb Douglas production

³Specifically there can be no unobserved variation in the price of inputs that are used in the control function, and no persistent variation in the price of the remaining inputs.

function parameters using the control function approach requires that none of the inputs are perfectly flexible and chosen in response to the current productivity shock. These issues are discussed further in Akerberg and Caves (2003).

In the remainder of this paper we consider the identification of Cobb Douglas production function parameters in the presence of explicit adjustment costs for all inputs. Since the control functions proposed by Olley and Pakes (1996) and Levinsohn and Petrin (2003) do not control fully for the unobserved productivity shocks when all inputs are costly to adjust, we focus on the scope for obtaining consistent estimates using instrumental variables methods.

3 Identification Issues in the Presence of Adjustment Costs

Adjustment costs weaken the dependence of current input levels on current productivity shocks, and introduce dependence on the history of past shocks. Provided these productivity shocks are not completely common to all firms, this introduces variation across firms in the shadow prices of each input, even if all firms face the same purchase or rental prices. This suggests that the collinearity problems discussed in the previous section will be less extreme if all inputs are subject to adjustment costs. We illustrate this in this section by considering a simple dynamic problem with a two-factor Cobb Douglas production function and strictly convex adjustment costs for each of the inputs.

The firm now chooses inputs and output to maximise the present value of current and expected future net revenues, which we write recursively as

$$\begin{aligned}
 V_t(K_{t-1}, L_{t-1}) = & P_t F_t(K_t, L_t) - P_t^K I_t - P_t^K G_t(I_t, K_t) \\
 & - W_t L_t - W_t C_t(H_t, L_t) + \psi_t E_t[V_{t+1}(K_t, L_t)]
 \end{aligned} \tag{10}$$

where output in period t (Y_t) is given by the two-factor stochastic production function $F_t(K_t, L_t)$, which will be assumed to have the Cobb Douglas functional

form as in (1).⁴ $E_t[\cdot]$ denotes an expected value given information in period t , and ψ_t is the discount factor giving the value in period t of revenue in period $t + 1$. Capital and labour inputs evolve according to the equations of motion

$$K_t = (1 - \delta)K_{t-1} + I_t \quad (11)$$

$$L_t = (1 - q)L_{t-1} + H_t \quad (12)$$

where I_t is gross investment in period t , H_t is gross hiring, and δ and q are exogenously given depreciation and quit rates respectively. As the firm's problem is now inherently dynamic, there is little simplification from assuming capital to be rented, and we assume instead that capital is purchased at the price P_t^K per unit in period t . Labour is again hired at the wage W_t per unit in period t . Investment and hiring incur additional costs of adjustment given by the (possibly stochastic) functions $G_t(I_t, K_t)$ and $C_t(H_t, L_t)$ respectively. Notice that these take the form of additional financial costs, rather than the form of foregone production. Any costs of adjustment in the form of lost output would clearly affect the relationship between inputs and output, leaving the simple static specification in (7) fundamentally mis-specified. Note also that the current productivity shock (A_t) is known by the firm in period t , when making its investment and hiring decisions. All prices are again assumed to be exogenously given and common to all firms.

This problem has two control variables (I_t and H_t) and two state variables (K_{t-1} and L_{t-1}). Using the equations of motion (11) and (12) to eliminate the current values of the state variables from (10) and then differentiating with respect to I_t , K_{t-1} , H_t and L_{t-1} gives the first order conditions

$$0 = P_t \frac{\partial F_t}{\partial K_t} - P_t^K - P_t^K \frac{\partial G_t}{\partial I_t} - P_t^K \frac{\partial G_t}{\partial K_t} + \psi_t E_t \left[\frac{\partial V_{t+1}}{\partial K_t} \right] \quad (13)$$

$$\frac{\partial V_t}{\partial K_{t-1}} = (1 - \delta) P_t \frac{\partial F_t}{\partial K_t} - (1 - \delta) P_t^K \frac{\partial G_t}{\partial K_t} + (1 - \delta) \psi_t E_t \left[\frac{\partial V_{t+1}}{\partial K_t} \right] \quad (14)$$

$$0 = P_t \frac{\partial F_t}{\partial L_t} - W_t - W_t \frac{\partial C_t}{\partial H_t} - W_t \frac{\partial C_t}{\partial L_t} + \psi_t E_t \left[\frac{\partial V_{t+1}}{\partial L_t} \right] \quad (15)$$

⁴We omit firm subscripts in this section to simplify the notation.

$$\frac{\partial V_t}{\partial L_{t-1}} = (1-q)P_t \frac{\partial F_t}{\partial L_t} - (1-q)W_t - (1-q)W_t \frac{\partial C_t}{\partial L_t} + (1-q)\psi_t E_t \left[\frac{\partial V_{t+1}}{\partial L_t} \right] \quad (16)$$

Letting $\lambda_t^K = \frac{1}{1-\delta} \frac{\partial V_t}{\partial K_{t-1}}$ denote the shadow value in period t of inheriting one additional unit of capital (after depreciation) from the previous period,⁵ we can combine (13) and (14) to obtain

$$\begin{aligned} \lambda_t^K &= P_t \frac{\partial F_t}{\partial K_t} - P_t^K \frac{\partial G_t}{\partial K_t} + (1-\delta)\psi_t E_t [\lambda_{t+1}^K] \\ &= P_t^K + P_t^K \frac{\partial G_t}{\partial I_t} \end{aligned} \quad (17)$$

Similarly letting $\lambda_t^L = \frac{1}{1-q} \frac{\partial V_t}{\partial L_{t-1}}$, the analogous expressions for labour are

$$\begin{aligned} \lambda_t^L &= P_t \frac{\partial F_t}{\partial L_t} - W_t - W_t \frac{\partial C_t}{\partial L_t} + (1-q)\psi_t E_t [\lambda_{t+1}^L] \\ &= W_t \frac{\partial C_t}{\partial H_t} \end{aligned} \quad (18)$$

We can rewrite the first line of (18) as

$$P_t \frac{\partial F_t}{\partial L_t} - W_t \frac{\partial C_t}{\partial L_t} = W_t + \lambda_t^L \left(1 - (1-q)\psi_t E_t \left[\frac{\lambda_{t+1}^L}{\lambda_t^L} \right] \right) \quad (19)$$

Notice that if there are no adjustment costs for labour, we have $\frac{\partial C_t}{\partial H_t} = \frac{\partial C_t}{\partial L_t} = 0$. This implies, using the second line of (18), that $\lambda_t^L = 0$. In this case, (19) reduces to the standard marginal productivity condition, as in (4). More generally, with adjustment costs, (19) indicates that the marginal revenue product of labour is equated with the wage plus an additional term which depends on the current and expected future levels of the shadow value of inheriting a higher labour force from the previous period. The first line of (18) further indicates that this shadow value depends on the current productivity shock (A_t), which affects the marginal physical product of labour ($\frac{\partial F_t}{\partial L_t}$). Thus even if all firms face a common wage rate (W_t), idiosyncratic productivity shocks will generate variation across firms in the relevant shadow price of labour.

⁵Equivalently λ_t^K is the shadow value associated with the constraint (11).

Similar observations apply to the optimal choice of the capital stock in the presence of adjustment costs. Here we can rewrite the first line of (17) as

$$P_t \frac{\partial F_t}{\partial K_t} - P_t^K \frac{\partial G_t}{\partial K_t} = \lambda_t^K \left(1 - (1 - \delta)\psi_t E_t \left[\frac{\lambda_{t+1}^K}{\lambda_t^K} \right] \right) \quad (20)$$

With no adjustment costs for capital, we have $\lambda_t^K = P_t^K$ and (20) reduces to the standard marginal productivity condition, in which $P_t^K \left(1 - (1 - \delta)\psi_t E_t \left[\frac{P_{t+1}^K}{P_t^K} \right] \right)$ is the user cost of capital, with which the marginal revenue product is equated.⁶ With adjustment costs, the marginal revenue product of capital is equated with a term that depends on current and expected future levels of the shadow value of inheriting a higher capital stock. Again, using the first line of (17), this shadow value depends on the current productivity shock, and thus varies across firms provided the productivity shocks are not completely common.

These first order conditions also indicate that the optimal levels of the inputs in the presence of adjustment costs will depend on past as well as current productivity shocks, and will therefore be persistent, even if the productivity shocks themselves are serially uncorrelated. To illustrate this, we focus on the intertemporal optimality condition for labour (18), and simplify by assuming adjustment costs to have the form $C_t(H_t, L_t) = cH_t^2/2$, so that $(\partial C_t/\partial H_t) = cH_t$ and $(\partial C_t/\partial L_t) = 0$. Simplifying further by assuming all prices and the discount factor to be constant, we can write the first line of (18) as

$$H_t = \frac{1}{c} \left(\frac{P}{W} \frac{\partial F_t}{\partial L_t} - 1 \right) + (1 - q)\psi E_t[H_{t+1}] \quad (21)$$

Letting $H_{t+1} = E_t[H_{t+1}] + \varepsilon_{t+1}$, where ε_{t+1} is the error made in forecasting H_{t+1} using information available in period t , and solving for H_{t+1} gives

$$H_{t+1} = \left(\frac{1}{(1 - q)\psi} \right) H_t - \left(\frac{1}{c(1 - q)\psi} \right) \left(\frac{P}{W} \frac{\partial F_t}{\partial L_t} - 1 \right) + \varepsilon_{t+1} \quad (22)$$

⁶Notice our timing assumption that this period's investment contributes immediately to production. Differences between (19) and (20) reflect the assumptions that labour is hired and capital is owned by the firm.

Current hiring thus depends on past productivity shocks, both through past hiring and through the previous period's marginal product of labour, as well as on the innovation to the productivity process in the current period, which is reflected here in the forecast error. We can also note that the persistence in hiring, and hence in the level of employment, will depend *inter alia* on the adjustment cost parameter c . This suggests that, at least provided capital and labour are subject to different levels of adjustment costs, there will be variation in the capital-labour ratio.

To see this more intuitively, consider the firm's response to a large, permanent increase in productivity, assuming that adjustment costs for capital are higher than those for labour. Eventually the firm will want to have higher levels of both inputs, as in the case where inputs are flexible. However the capital stock will adjust more slowly than employment, since adjusting the capital stock is relatively expensive. This will generate a lower capital-labour ratio during the period after the shock when significant adjustments are occurring. Similarly, in response to a large, temporary increase in productivity, both inputs will be temporarily higher, but the adjustments (upward initially and downward subsequently) in capital will be more muted. Again the capital-labour ratio will tend to be lower in the period immediately after a positive productivity shock. Conversely, capital-labour ratios will tend to be higher following (temporary or permanent) negative productivity shocks, as in this case capital adjusts downwards more slowly. As a result, the combination of productivity shocks with different levels of adjustment costs for different inputs will generate variation both across firms and over time in, for example, the capital-labour ratio.

This variation will of course be endogenous, in that both capital and labour and the capital-labour ratio will be correlated with the current productivity shock in the error term of simple log linear production function specifications like (7). However the presence of adjustment costs also makes the variation in these input levels persistent, so that lagged levels of the inputs provide informative instru-

ments. Following a temporary positive productivity shock, for example, the subsequent downward adjustment of the capital stock occurs gradually to minimise the present value of adjustment costs. The presence of adjustment costs thus generates not only variation but also predictable variation in capital, labour and the capital-labour ratio, even if the productivity shocks themselves are serially uncorrelated. As a result, lagged levels of the inputs can be used to obtain valid and informative instruments, and the production function parameters are identified using standard instrumental variables estimators.

This source of identification is likely to be strengthened if the costs of adjustment are themselves subject to stochastic variation. Stochastic adjustment costs are a standard assumption in the Q model of investment (see Hayashi, 1982), and have been considered more recently in the context of fixed adjustment cost levels by, for example, Caballero, Engel and Haltiwanger (1995). With the combination of (convex) adjustment costs and stochastic productivity shocks, input levels are generally in the process of adjusting. Shocks to the level of adjustment costs that, for example, make it relatively inexpensive for the firm to adjust its capital stock in the current period, will thus introduce an additional and relatively exogenous source of variation into the capital-labour ratio. These exogenous adjustment cost shocks are unlikely to be observed, and so cannot be used directly as instruments. However their presence is expected to improve the properties of standard instrumental variables estimators of the production function parameters, using lagged internal instruments, to the extent that they generate more (predictable) variation in the relative levels of different inputs.

4 Simulation Evidence

The previous section suggests that, even without variation in input prices, the parameters of a Cobb Douglas production function are identified provided there are adjustment costs for each of the inputs considered, and provided there is

variation across firms in the productivity shocks. In this section we illustrate this using simulated data for a two-factor Cobb Douglas production function where capital and labour inputs are each subject to quadratic costs of adjustment.

To generate the simulated data, each firm chooses investment and hiring to maximise its value in (10) subject to the constraints in (11) and (12). The production function is

$$Y_{it} = F(K_{it}, L_{it}) = e^{a_{it}} K_{it}^{\alpha} L_{it}^{\beta}$$

The prices of inputs and output are common to all firms, constant over time, and normalised to one. The productivity shocks a_{it} are initially drawn from an iid normal distribution with mean zero and variance σ_a^2 , so that the static log linear production function that we estimate

$$y_{it} = \alpha k_{it} + \beta l_{it} + a_{it} \tag{23}$$

is correctly specified.

4.1 Deterministic adjustment costs

The adjustment cost functions are initially specified to be deterministic, with net changes to the input levels subject to increasing marginal costs

$$G(I_{it}, K_{it}) = \frac{g}{2} \left(\frac{I_{it}}{K_{i,t-1}} - \delta \right)^2 K_{i,t-1}$$

$$C(H_{it}, L_{it}) = \frac{c}{2} \left(\frac{H_{it}}{L_{i,t-1}} - q \right)^2 L_{i,t-1}$$

The properties of the simulated input and output series depend on the parameters $(\alpha, \beta, g, c, \delta, q, \psi, \sigma_a^2)$. We fix $\delta = q = 0.1$ and $\psi = 1/1.1$. For the simulations reported in Table 1, we also fix $\alpha = 0.4$, $\beta = 0.5$ and $\sigma_a^2 = 0.017$. Our main interest is in the properties of instrumental variables estimators of α and β in (23), and how these change as we vary the adjustment cost parameters g and c .

Table 1 reports Monte Carlo results for four experiments using different values of these adjustment cost parameters. In each case we generate data for a panel

of 500 firms and 10 time periods.⁷ The numerical solution method used to obtain this simulated data is described in the appendix. This procedure is repeated 100 times using different draws of the productivity shocks. Table 1 reports the mean and the standard deviation over these 100 replications for each of the estimated coefficients.

The reported two-stage least squares estimates use the first lags of both inputs and output ($k_{i,t-1}$, $l_{i,t-1}$ and $y_{i,t-1}$) as instruments.⁸ Very similar estimates were obtained using further lags of the series as additional instruments. As well as these instrumental variables estimates of α and β , we also report OLS estimates of the autoregressive coefficients in simple AR(2) specifications for k_{it} and l_{it}

$$k_{it} = \rho_1^K k_{i,t-1} + \rho_2^K k_{i,t-2} + v_{it}^K$$

$$l_{it} = \rho_1^L l_{i,t-1} + \rho_2^L l_{i,t-2} + v_{it}^L$$

and OLS estimates of the coefficients in simple specifications that relate the natural logarithms of the input levels chosen in the presence of adjustment costs to the corresponding levels that would have been chosen in the absence of adjustment costs

$$k_{it} = a^K + b^K k_{it}^* + \varepsilon_{it}^K$$

$$l_{it} = a^L + b^L l_{it}^* + \varepsilon_{it}^L$$

where k_{it}^* and l_{it}^* are calculated using the true parameter values and productivity shocks using (5) and (6) respectively. The former illustrate the serial correlation in the optimally chosen inputs that (in Table 1) results solely from the presence of adjustment costs. The latter indicate how closely the actual inputs track the levels that would be chosen in the absence of adjustment costs, which provides a more intuitive indication of the size of the adjustment cost parameters that

⁷In fact we generate data for 110 periods and discard the first 100 observations, to minimise any impact of the starting values. Almost identical results are obtained using a start up period of 200 observations.

⁸For a given set of instruments, two-stage least squares provides an asymptotically efficient Generalised Method of Moments estimator in this context with iid errors.

we consider.⁹ We also report the simple correlation coefficient between k_{it} and l_{it} ; recall that the identification problem when both inputs are perfectly flexible results from this correlation being one. Finally, for comparison, we also report estimates of the coefficient β on the labour input obtained using a simple control function estimator. Specifically we use OLS to estimate the augmented specification

$$y_{it} = \beta l_{it} + \theta_1 i_{it} + \theta_2 i_{it}^2 + \theta_3 k_{it} + \theta_4 k_{it}^2 + \theta_5 i_{it} k_{it} + \xi_{it} \quad (24)$$

where $i_{it} = I_{it}/K_{i,t-1}$ is the rate of investment. Here we follow the suggestion of Olley and Pakes (1996) to proxy for the unobserved productivity shock using a flexible function of investment and capital, exploiting the fact that current investment decisions are also influenced by these shocks. Given our timing assumptions and the presence of adjustment costs, these additional terms are not expected to control fully for the correlation between l_{it} and a_{it} , but it is nevertheless interesting to consider how this approach fares in a case where the controls are mis-specified.

Column (i) of Table 1 reports the results for a baseline case in which we set $g = 1$ and $c = 0.2$. Adjustment costs for capital are thus five times greater than adjustment costs for labour. The resulting capital and labour series are both serially correlated, with the capital series displaying more persistence. Conversely the level of employment is more responsive to the current productivity shock.¹⁰ The two inputs are highly but not perfectly correlated with each other, indicating that there is variation in the capital-labour ratio. Given these properties, it is not surprising that lagged inputs provide valid and informative instruments for the current input levels included in (23), and the reported two-stage least squares coefficients estimate the true production function parameters quite well using this simulated dataset.

⁹Note that, as g and c approach zero, so a^K and a^L approach zero and b^K and b^L approach one.

¹⁰Although neither series is highly responsive to the (serially uncorrelated) productivity shocks. A shock that would produce 20% increases in both labour and capital in the absence of adjustment costs here produces an increase in employment of just over 1% and an increase in the capital stock of less than $\frac{1}{2}\%$.

Columns (ii), (iii) and (iv) illustrate three cases where identification of the production function parameters becomes much weaker. In column (ii), both the adjustment cost parameters are twenty times smaller than they are in column (i). Thus although adjustment costs are higher for capital than for labour, they are now very low for both. The resulting input series are less persistent and more responsive to the current productivity shocks. As a result, they are also much more highly correlated with each other. This property is inherited by the predicted values after projecting on the lagged instruments, and the two-stage least squares estimates of the production function parameters are therefore much more imprecise in this case. The basic problem is that when adjustment costs for both inputs are too low, we are too close to the outcome with perfectly flexible inputs, in which capital and labour are perfectly collinear and identification fails.

In column (iii), both the adjustment cost parameters are ten times higher than they are in column (i). The input series are much more persistent, with the capital series in particular being close to having a unit root. They are also much less responsive to the current productivity shocks. The basic problem here is that there is now very little variation at all in the optimal input levels; adjusting the capital stock, in particular, is now so expensive that the firms undertake very limited adjustments. Again the result is to make the instrumental variables estimates of the production function parameters, particularly that on capital, much less precise.

In column (iv) the levels of the two adjustment cost parameters are neither too low nor too high, but now they are more similar to each other, with adjustment costs for capital only 50% higher than adjustment costs for labour. This makes the time series properties of the two inputs, and their responses to current productivity shocks, much more similar. In this case the two inputs are again found to be highly collinear, and the instrumental variables estimates of the production function parameters are very imprecise.

We note that, in all four cases, the inclusion of additional investment and capital terms in the augmented specification (24) does not control fully for the

correlation between current employment and the current productivity shock. This is not surprising in the presence of adjustment costs for both inputs, in which case there is no function of investment and capital that can be inverted to control for the productivity shock.

Table 2 considers a more general design in which the productivity shocks are serially correlated. Specifically we now have

$$a_{it} = \rho a_{i,t-1} + u_{it}$$

where the innovations u_{it} are iid normal with mean zero and variance 0.017. For $\rho \neq 0$, the serially correlated error term in (23) implies that lagged levels of the inputs (and output) are no longer valid instruments. For example, $l_{i,t-1}$ is influenced by the realisation of $u_{i,t-1}$, which is clearly correlated with a_{it} . In this case, two-stage least squares estimates of α and β obtained using lagged inputs as instruments in (23) are seriously biased.¹¹

Consistent estimates can in principle be obtained by estimating the serial correlation coefficient ρ jointly with the production function parameters α and β . We can quasi-difference (23) to obtain

$$(y_{it} - \rho y_{i,t-1}) = \alpha(k_{it} - \rho k_{i,t-1}) + \beta(l_{it} - \rho l_{i,t-1}) + u_{it} \quad (25)$$

in which the error term is again serially uncorrelated and orthogonal to lagged levels of both the inputs and output. The instrumental variables estimates reported in Table 2 are non-linear two-stage least squares estimates, obtained using $k_{i,t-1}$, $l_{i,t-1}$ and $y_{i,t-1}$ as instruments in (25).

Column (i) reports results using the same adjustment cost parameters as in column (i) of Table 1, with ρ set at 0.3. Although the instrumental variables procedure yields reasonable estimates of ρ , the estimates of the production function

¹¹For example, in the experiment reported in column (i) of Table 2, the linear 2SLS estimates of α in the static production function specification (23) have a mean (standard deviation) of -0.42 (.046), and the corresponding estimates of β are 1.33 (.035).

parameters are now very imprecise. Columns (ii)-(iv) report similarly disappointing results for the other configurations of adjustment cost parameters considered in Table 1.

The main problem here appears to be that there is insufficient information to estimate the three parameters jointly, rather than the effect of the serially correlated productivity shocks on the properties of the inputs. To illustrate this, we generated the production data with $g = 1$, $c = 0.2$ and $\rho = 0$, as in column (i) of Table 1, but we attempted to estimate ρ jointly with the production function parameters α and β , as we do in Table 2. In this case we obtained means (standard deviations) of the non-linear instrumental variables estimates of α and β of 0.413 (0.250) and 0.480 (0.228) respectively.¹² Comparison with the results in column (i) of Table 1 shows that the standard deviations of the estimated production function parameters are approximately doubled when we estimate ρ as in Table 2, rather than imposing the correct value of zero as in Table 1.¹³

4.2 Stochastic adjustment costs

In Tables 3 and 4 we introduce stochastic shocks to the adjustment costs for capital, and we allow these adjustment cost shocks to themselves be serially correlated.¹⁴ The adjustment cost function for capital used here is

$$G(I_{it}, K_{it}) = \frac{g}{2} \left(\frac{I_{it}}{K_{i,t-1}} - \delta - \omega_{it} \right)^2 K_{i,t-1} \quad (26)$$

$$\omega_{it} = \tau \omega_{i,t-1} + v_{it}$$

¹²In this case, the mean estimate of ρ was zero to three decimal places, with a standard deviation of 0.038.

¹³Conversely we can estimate the production function parameters quite well with serially correlated productivity shocks, if we impose the correct value of ρ . For example, imposing $\rho = 0.3$ and estimating α and β using (linear) 2SLS in (25) yields estimates of 0.39 (.11) and 0.50 (.09) with the adjustment cost parameters used in column (i) of Table 2.

¹⁴Our specification is common in the literature on the Q model of investment. See, for example, Blundell et al. (1992).

Thus, if the firm is adjusting its capital stock upwards, a positive value of ω_{it} reduces the cost of additional investment in the current period.¹⁵ The innovations v_{it} are iid normal with mean zero and variance 0.02, independent of the productivity shocks, and we set the serial correlation parameter for these adjustment cost shocks to $\tau = 0.3$.

Table 3 reports results for static production functions (23), in which the productivity shocks (a_{it}) are serially uncorrelated and this is correctly imposed on the econometric specifications (as in Table 1). Column (i) of Table 3 uses the same adjustment cost parameters (g and c) as in column (i) of Table 1, so that the only difference between these two cases is the presence of stochastic adjustment costs for capital in the design of Table 3. As expected, this results in a considerable improvement in the precision of the linear two-stage least squares estimates of the production function parameters. Interestingly, the standard deviation of the estimated coefficient on labour is halved relative to that in Table 1, although we have only introduced stochastic adjustment costs for capital. The improvement in the precision with which we estimate the coefficient on capital is even greater.¹⁶

Column (ii) of Table 3 shows that identification again becomes much weaker in the case where the adjustment cost parameters (g and c) are very low. The intuition for this result is quite clear: if the level of adjustment costs is sufficiently low, the presence of stochastic shocks to these adjustment costs is relatively unimportant, and we remain very close to the case with frictionless adjustment in which these parameters are not identified. Columns (iii) and (iv) however show that, with stochastic adjustment cost shocks, the properties of the instrumental variables estimates of the production function parameters are much less sensitive to the presence of adjustment cost levels that are either ‘too high’ or ‘too similar’.

¹⁵Alternatively one could think of introducing stochastic variation in the depreciation rate (or the quit rate), although this would also have to be accounted for in the equation(s) of motion for the input(s).

¹⁶Broadly similar results for this static production function specification were obtained in an experiment in which we introduced serially uncorrelated adjustment cost shocks (i.e. setting $\tau = 0$).

The intuition for the latter result is that stochastic shocks to the adjustment costs for capital are generating variation in the capital-labour ratio that is not present in the corresponding design in Table 1.¹⁷ The intuition for the former result comes from observing that variation in the capital-labour ratio due to the exogenous adjustment cost shocks becomes more important, relative to the endogenous productivity shocks, as the level of adjustment costs increases.¹⁸ Indeed the precision of the IV estimates is better here than in column (i). We can also note that the biases found for simple OLS estimates of the production function parameters are much smaller here than in our other designs, and that, as a result, the bias in the control function estimate of the coefficient on labour is also much smaller.

Table 4 reports results for the dynamic production function specification (25), in which the productivity shocks (a_{it}) are serially correlated and the autoregressive parameter (ρ) is estimated jointly with the parameters of the Cobb Douglas production function (as in Table 2). Column (i) of Table 4 again uses our baseline adjustment cost parameters, as in column (i) of Table 2, so that these cases differ only due to the presence of stochastic adjustment costs for capital in Table 4. For this dynamic specification, the improvement attributable to (serially correlated) adjustment cost shocks is much greater than was the case for the static specification. The non-linear two-stage least squares estimates of all three parameters now have negligible biases, and their standard deviations are dramatically reduced. In contrast to the findings of Table 2, the introduction of these stochastic adjustment costs results in sufficient information to estimate these parameters jointly, with reasonable precision.¹⁹ As expected, this does not hold when the levels of the

¹⁷Hence we would not expect the same finding if there were adjustment cost shocks for both inputs that were highly correlated with each other.

¹⁸Note from (26) that, in the limit as $g \rightarrow \infty$, *all* the variation in investment rates would be driven by the exogenous adjustment cost shocks. In this experiment, the R^2 for the regression of k_{it} on k_{it}^* is almost zero, indicating that the capital input here responds very little to the productivity shocks.

¹⁹We note that this was not found to be the case in experiments with serially uncorrelated adjustment cost shocks, in which the improvements relative to the deterministic case were found to be much more modest.

adjustment costs are ‘too low’, as illustrated in column (ii) of Table 4. However, as for the static specification reported in Table 3, we find that with stochastic adjustment cost shocks, the properties of these instrumental variables estimates of the production function parameters are robust to adjustment cost levels that are similar for the two inputs considered (column (iv)), and actually improve as we consider higher levels of adjustment costs (column (iii)).

5 Conclusions

The parameters of Cobb Douglas production functions are not identified when inputs are perfectly flexible and chosen optimally, and input prices are common to all firms. This paper has shown that the presence of adjustment costs for all inputs considered allows Cobb Douglas production function parameters to be identified, even in the absence of variation in input prices. When adjustment costs are deterministic, this source of identification is shown to be quite fragile, and can become weak if the levels of adjustment costs for different inputs are too low, too high or too similar. The properties of instrumental variables estimators of the production function parameters are shown to improve considerably when there is exogenous stochastic variation across firms in the level of adjustment costs, particularly when these adjustment cost shocks are serially correlated. In this case we show that quite complex stochastic specifications of the production function may be identified, even when input prices are common to all firms.

References

- [1] Akerberg, D. and Caves, K. (2003) “Structural identification of production functions”, mimeo, UCLA.
- [2] Blundell, R. and Bond, S. (2000) “GMM estimation with persistent panel data: an application to production functions”, *Econometric Reviews*, 19, 321-340.
- [3] Blundell, R., Bond, S., Devereux, M.P. and Schiantarelli, F. (1992) “Investment and Tobin’s Q: evidence from company panel data”, *Journal of Econometrics*, 51, 233-257.
- [4] Caballero, R., Engel, E. and Haltiwanger, J. (1995) “Plant-level adjustment and aggregate investment dynamics”, *Brookings Papers on Economic Activity*, 1995(2), 1-54.
- [5] Fafchamps, M. and Pender, J. (1997) “Precautionary saving, credit constraints, and irreversible investment: theory and evidence from semi-arid India”, *Journal of Business and Economics Statistics*, 15, 180-94.
- [6] Hayashi, F. (1982) “Tobin’s average q and marginal q: a neoclassical interpretation”, *Econometrica*, 50, 213-224.
- [7] Judd, K.L. (1998) *Numerical Methods in Economics*, Cambridge, Massachusetts; London, England: The MIT Press.
- [8] Levinsohn, J. and Petrin, A. (2003) “Estimating production functions using inputs to control for unobservables”, *Review of Economic Studies*, 70, 317-342.
- [9] Marschak, J. and Andrews, W. (1944) “Random simultaneous equations and the theory of production”, *Econometrica*, 12, 143-205.
- [10] Mundlak, Y. and Hoch, I. (1965) “Consequences of alternative specifications of Cobb-Douglas production functions”, *Econometrica*, 33, 814-828.
- [11] Olley, S. and Pakes, A. (1996) “The dynamics of productivity in the telecommunications equipment industry”, *Econometrica*, 64, 1263-1297.

Table 1. Static Production Functions; Deterministic Adjustment Costs

		(i)	(ii)	(iii)	(iv)
$N = 500$	$\alpha = 0.4$	$g = 1$	$g = 0.05$	$g = 10$	$g = 0.6$
$T = 10$	$\beta = 0.5$	$c = 0.2$	$c = 0.01$	$c = 2$	$c = 0.4$
Instrumental Variables	α	0.390 (.126)	0.411 (.482)	0.363 (.343)	0.401 (.420)
	β	0.506 (.098)	0.488 (.526)	0.505 (.240)	0.494 (.394)
Control Function	β	-0.353 (.001)	0.114 (.004)	-1.358 (.002)	-0.404 (.001)
AR(2) Capital	ρ_1^K	1.001 (.015)	0.505 (.016)	1.069 (.017)	0.903 (.016)
	ρ_2^K	-0.098 (.016)	-0.012 (.014)	-0.084 (.018)	-0.028 (.017)
AR(2) Labour	ρ_1^L	0.671 (.014)	0.365 (.016)	0.880 (.015)	0.817 (.016)
	ρ_2^L	0.067 (.016)	0.035 (.014)	0.031 (.015)	0.021 (.017)
Target Capital	a^K	2.634 (.002)	1.934 (.007)	2.683 (.001)	2.614 (.002)
	b^K	0.021 (.001)	0.282 (.003)	0.002 (2e-04)	0.028 (.001)
Target Labour	a^L	2.750 (.003)	1.859 (.007)	2.889 (.001)	2.807 (.002)
	b^L	0.056 (.001)	0.362 (.002)	0.008 (3e-04)	0.036 (.001)
Corr(k_{it}, l_{it})		0.918	0.991	0.878	0.995

Means (standard deviations) of estimated coefficients in 100 replications

Table 2. Dynamic Production Functions; Deterministic Adjustment Costs

		(i)	(ii)	(iii)	(iv)
$N = 500$	$\alpha = 0.4$	$g = 1$	$g = 0.05$	$g = 10$	$g = 0.6$
$T = 10$	$\beta = 0.5$	$c = 0.2$	$c = 0.01$	$c = 2$	$c = 0.4$
	$\rho = 0.3$				
Instrumental Variables	α	0.636 (.586)	-0.114 (.200)	0.413 (.293)	0.979 (1.85)
	β	0.218 (.657)	1.045 (.249)	0.496 (.251)	-0.099 (1.89)
	ρ	0.332 (.096)	0.202 (.214)	0.302 (.022)	0.319 (.081)
Control Function	β	-0.077 (.001)	0.163 (.003)	-0.804 (.001)	-0.171 (.001)
AR(2) Capital	ρ_1^K	1.272 (.016)	0.802 (.016)	1.373 (.014)	1.194 (.015)
	ρ_2^K	-0.348 (.015)	-0.157 (.016)	-0.384 (.014)	-0.287 (.015)
AR(2) Labour	ρ_1^L	0.965 (.016)	0.675 (.016)	1.176 (.013)	1.113 (.015)
	ρ_2^L	-0.139 (.016)	-0.083 (.015)	-0.235 (.014)	-0.227 (.016)
Target Capital	a^K	2.584 (.005)	1.648 (.010)	2.675 (.001)	2.541 (.005)
	b^K	0.039 (.002)	0.386 (.004)	0.004 (5e-04)	0.051 (.002)
Target Labour	a^L	2.663 (.006)	1.574 (.010)	2.867 (.002)	2.718 (.005)
	b^L	0.086 (.002)	0.458 (.004)	0.014 (.001)	0.062 (.002)
Corr(k_{it}, l_{it})		0.942	0.994	0.895	0.996

Means (standard deviations) of estimated coefficients in 100 replications

Table 3. Static Production Functions; Stochastic Adjustment Costs

		(i)	(ii)	(iii)	(iv)
$N = 500$	$\alpha = 0.4$	$g = 1$	$g = 0.05$	$g = 10$	$g = 0.6$
$T = 10$	$\beta = 0.5$	$c = 0.2$	$c = 0.01$	$c = 2$	$c = 0.4$
	$\tau = 0.3$				
Instrumental Variables	α	0.400 (.032)	0.367 (.268)	0.400 (.009)	0.400 (.030)
	β	0.500 (.042)	0.537 (.296)	0.500 (.013)	0.499 (.045)
Control Function	β	1.873 (.015)	1.084 (.002)	.568 (.013)	1.995 (.025)
AR(2) Capital	ρ_1^K	1.130 (.013)	0.515 (.016)	1.252 (.017)	1.039 (.016)
	ρ_2^K	-0.229 (.014)	-0.018 (.016)	-0.269 (.017)	-0.178 (.016)
AR(2) Labour	ρ_1^L	0.996 (.015)	0.372 (.016)	1.759 (.011)	1.061 (.017)
	ρ_2^L	-0.090 (.016)	0.031 (.016)	-0.763 (.011)	-0.153 (.016)
Target Capital	a^K	2.632 (.010)	1.924 (.007)	2.676 (.030)	2.571 (.007)
	b^K	0.021 (.004)	0.282 (.003)	0.002 (.011)	0.028 (.003)
Target Labour	a^L	2.748 (.007)	1.848 (.007)	2.881 (.022)	2.761 (.005)
	b^L	0.056 (.002)	0.363 (.002)	0.008 (.007)	0.036 (.002)
Corr(k_{it}, l_{it})		0.935	0.989	0.957	0.906

Means (standard deviations) of estimated coefficients in 100 replications

Table 4. Dynamic Production Functions; Stochastic Adjustment Costs

		(i)	(ii)	(iii)	(iv)
$N = 500$	$\alpha = 0.4$	$g = 1$	$g = 0.05$	$g = 10$	$g = 0.6$
$T = 10$	$\beta = 0.5$	$c = 0.2$	$c = 0.01$	$c = 2$	$c = 0.4$
	$\rho = 0.3$				
	$\tau = 0.3$				
Instrumental Variables	α	0.410 (.072)	0.021 (.553)	0.399 (.013)	0.404 (.043)
	β	0.488 (.093)	0.900 (.645)	0.501 (.021)	0.499 (.058)
	ρ	0.304 (.032)	0.281 (.172)	0.301 (.016)	0.299 (.018)
Control Function	β	1.519 (.008)	1.040 (.002)	0.630 (.019)	1.639 (.017)
AR(2) Capital	ρ_1^K	1.150 (.016)	0.801 (.015)	1.254 (.014)	1.085 (.015)
	ρ_2^K	-0.247 (.015)	-0.159 (.017)	-0.271 (.014)	-0.211 (.016)
AR(2) Labour	ρ_1^L	1.122 (.016)	0.677 (.015)	1.741 (.010)	1.202 (.014)
	ρ_2^L	-0.226 (.017)	-0.087 (.018)	-0.744 (.010)	-0.295 (.014)
Target Capital	a^K	2.574 (.013)	1.620 (.011)	2.667 (.035)	2.503 (.011)
	b^K	0.042 (.005)	0.396 (.004)	0.005 (.013)	0.053 (.004)
Target Labour	a^L	2.649 (.010)	1.547 (.011)	2.858 (.026)	2.677 (.008)
	b^L	0.090 (.003)	0.467 (.004)	0.016 (.009)	0.065 (.003)
Corr(k_{it}, l_{it})		0.916	0.994	0.957	0.907

Means (standard deviations) of estimated coefficients in 100 replications

Appendix: A Sketch of the Solution Method Used to Generate the Simulated Production Data

The firm chooses inputs and output to maximise the present value of current and expected future net revenues:

$$V(K_{t-1}, L_{t-1}; a_{t-1}) = \max_{K_t, L_t} P_t F_t(K_t, L_t) - P_t^K I_t - P_t^K G_t(I_t, K_t) - W_t L_t - W_t C_t(H_t, L_t) + \psi_t E_t[V(K_t, L_t; a_t)].$$

This is the Bellman equation. Our task is to find the solution $\{K_t^*, L_t^*\}$, which will depend on the parameters of the model and the state variables $\{K_{t-1}, L_{t-1}; a_{t-1}\}$. To do this we use numerical dynamic programming techniques. We describe this approach briefly here.

Begin by rewriting the Bellman equation in more compact form as

$$V(K_{t-1}, L_{t-1}; a_{t-1}) = \max_{K_t, L_t} \pi(K_t, L_t; a_t) + \psi_t E_t[V(K_t, L_t; a_t)]. \quad (27)$$

While the profit function $\pi(K_t, L_t; a_t)$ is known from the parameterisation of the model, we have to obtain $E_t[V(K_t, L_t; a_t)]$ using numerical methods. To do this we use value function iteration, which is a slow but robust method. See Chapter 12 of Judd (1998) for a rigorous discussion of numerical dynamic programming, including details on value iteration. Very informally, the principles of value function iteration are as follows:

1. Start with a guess for the true value function $V(x)$, where $x = \{K, L, a\}$. Call this guess $V^1(x)$. Use it on the right-hand side of the Bellman equation (27), and find the optimal choice rule $u = \{K^*, L^*\}$. We take as our first guess $V(x) = 0$, and so the solution at this stage will be equivalent to that of the static profit maximisation problem $\max_{K_t, L_t} \pi(K_t, L_t; a_t)$.
2. Update the guess for the true value function using the solution obtained in the previous step, $u(x)$. Call this updated guess $V^2(x)$. Check if $V^2(x) =$

$V^1(x)$. If true, you will have converged to the true function and so iteration can stop; if not go to step 3.

3. For $j = 3, 4, \dots$, use $V^{j-1}(x)$ on the right-hand side of the Bellman equation and calculate the optimal choice rule $u = \{K^*, L^*\}$. Update the guess of the value function, $V^j(x)$. Check if $V^j(x) = V^{j-1}(x)$. If true, you will have converged to the true function and so iteration can stop; if not, $j = j + 1$ and repeat step 3.

While straightforward in principle, the mechanics of this method are complicated for two reasons: first, while discretisation of the state space is a necessity in numerical dynamic programming, we need to allow for the fact that we are dealing with continuously distributed variables (capital and labour); second, we need to calculate the expected value of the firm in the next period.

We deal with the first problem by following the approach adopted by Fafchamps and Pender (1997). Specifically, we discretise the state space in such a way as to set the chosen values equal to the optimal nodes of a Chebyshev polynomial. We then interpolate between nodes using the Chebyshev iterative formula. We do this in R_+^q space where q is equal to the number of state variables. We use 6^q nodes in all simulations reported in the paper. For details on the Chebyshev approximation approach, see Chapter 6 of Judd (1998).

We deal with the second problem by numerical integration, using a Gauss-Hermite quadrature. This involves evaluating $V(\cdot; a_t)$ at a finite number of values of a_t and summing the results using a set of weights. The weights and the positions of the nodes are determined by the Gauss-Hermite quadrature. We use three nodes throughout. For details on numerical integration, see Chapter 7 of Judd (1998).