

INVESTMENT DECISIONS IN MANUFACTURING: ASSESSING THE EFFECTS OF REAL OIL PRICES AND THEIR UNCERTAINTY

KONSTANTINOS DRAKOS^a AND PANAGIOTIS TH. KONSTANTINOU^{b,c*}

^a *Department of Accounting and Finance, and Department of Economics, Athens University of Economics and Business, Greece*

^b *Department of Economics and Finance, Brunel University, Uxbridge, UK*

^c *Department of Economics, University of Macedonia, Thessaloniki, Greece*

SUMMARY

We investigate the effects of real oil prices and their uncertainty on investment decisions. Making use of plant-level data, we estimate dynamic, discrete-choice models that allow modeling investment inaction, under different assumptions related to initial conditions and unobserved heterogeneity. We find that increases in real oil price changes and in real oil price uncertainty significantly reduce the likelihood of investment action, in line with the predictions of irreversible investment theory. We also document that investment decisions exhibit strong, pure state dependence and are also significantly affected by initial conditions. Copyright © 2012 John Wiley & Sons, Ltd.

Received 10 August 2008; Revised 23 April 2012



Supporting information may be found in the online version of this article.

1. INTRODUCTION

Disaggregate investment decisions are largely characterized by persistent and non-smooth behavior; i.e. prolonged periods during which investment is actually zero are followed by prolonged periods in which investment is positive (Caballero *et al.*, 1995; Doms and Dunne, 1998; Nilsen and Schiantarelli, 2003).¹ Recent advances in the theoretical analysis of investment behavior have made substantial progress in the last three decades, deviating from the frictionless neoclassical benchmark, towards models that allow investment irreversibility, or, in general, non-convex adjustment costs.² In such models, the investment decision becomes a discrete choice between investing and staying put. For instance, in the presence of non-convex adjustment costs in uncertain environments, delaying the implementation of an investment project might emerge as an optimal choice, when the decision maker prefers to wait until part of the uncertainty is resolved.

Conventional wisdom suggests that adjustments in investment spending by firms are expected to be affected by energy price shocks. For instance, energy price shocks cause reductions in consumer expenditure (Edelstein and Kilian, 2009; Hamilton, 2009; Kilian, 2009b), subsequently lowering demand for firms' output and hence lower investment spending. In addition, such shocks are thought

* Correspondence to: Panagiotis Th. Konstantinou, Department of Economics and Finance, Brunel University, Uxbridge UB8 3PH, UK. E-mail: Panagiotis.Konstantinou@brunel.ac.uk

¹ The findings that (i) investment inaction is not rare, (ii) there is a substantial degree of irreversibility, and (iii) that investment decisions show substantial persistence have been documented, for instance, in Barnett and Sakellaris (1999), Bontempi *et al.* (2004), Cooper and Haltiwanger (2006), Gelos and Isgut (2001), and Sakellaris (2004).

² One class of such models explicitly introduces fixed investment costs and (partial) irreversibility (e.g. Abel and Eberly, 1994, 1996; Caballero and Engel, 1999). Another class of models, the so-called real options theory, suggests that a decision maker with an opportunity to invest possesses an option similar to a financial call option. If the decision maker proceeds with the irreversible investment, the lost option value is an opportunity cost that must be reflected in the cost of investment (e.g. McDonald and Siegel, 1986; Pindyck, 1988; Dixit and Pindyck, 1994).

to lead to increases in the marginal cost of production, also resulting in lower investment spending, although, with a few exceptions, this is not empirically verified (see Edelman and Kilian, 2007, for a discussion). On the other hand, changes in oil prices are thought to create uncertainty about future oil prices, causing firms to postpone irreversible investment decisions (Bernanke, 1983; Pindyck and Rotemberg, 1983; Pindyck, 1991).

Do higher real oil prices significantly reduce the likelihood of investment? Does an increase in real oil price uncertainty lead to postponing investment? In this paper, we focus on these two questions by analyzing the dynamic behavior of investment decisions at a disaggregate level, utilizing plant-level data for the Greek manufacturing sector. Drawing on prior empirical work on investment, we separate investment decisions between activity (either positive or negative investment) and inactivity (zero investment episodes).

Our analysis has a number of novel and distinct features. First, we make use of plant-level data to analyze the effects of real oil prices on investment, whereas existing studies operate at a higher level of aggregation.³ Second, we explicitly evaluate the existence and direction of the effects of real oil prices and real oil price uncertainty on the dynamics of the investment decision process, which—to the best of our knowledge—have not been explored before at such a disaggregate level.⁴ Third, we do so by using dynamic binary-choice models of investment behavior that disentangle the effect of real oil prices and their uncertainty from persistence due to unobserved heterogeneity or state dependence.

Our findings show that increases in real oil price changes and real oil price uncertainty adversely affect investment decisions of manufacturing plants. In more detail, we find that rising real oil prices significantly reduce the probability of investment action. This finding is robust across different estimators employed. Additionally, we find that increases in real oil price uncertainty raise significantly the probability of investment inaction. This finding is robust not only across estimators, but also when employing different measures of uncertainty such as the one-sided ‘risk’ measures suggested by Kilian and Manganello (2007).

Moreover, in one set of robustness experiments we allow for the effect of the unexpected real oil price change (a ‘shock’) and find that it reduces significantly the probability of investment. That is, we find that there are significantly negative *level* effects from unexpected changes in real oil prices, without making the adverse effects of increases in real oil price uncertainty less important. This piece of evidence can be considered as complementary to those in Edelman and Kilian (2007), who show, however, that there is no empirical support for theoretical models of the effects of uncertainty on business fixed investment expenditures.⁵ Our results show that there are indeed strong uncertainty effects on investment dynamics, but they are discernible at a more disaggregate level of analysis.

Finally, as a by-product of our analysis, we document the existence of strong state dependence in investment. We find that estimates of state dependence in investment are affected, as expected, by the assumptions made regarding initial conditions and the treatment of unobserved heterogeneity. Despite the sensitivity to these assumptions, we find that the likelihood of investment action is significantly positively correlated with investment action in the last period, across all estimators examined.

³ The use of a micro-level panel dataset is essential to avoid the problem of aggregation over production units, which results when investment decisions are observed at a higher aggregation level that masks investment discontinuity. The use of such a dataset makes it more likely that zeros (investment inaction) will be observed.

⁴ This relates our work—at least in spirit—to studies that examine the effects of aggregate uncertainty on disaggregate investment decisions (e.g. Campa, 1993, 1994; Campa and Goldberg, 1995). The importance of aggregate uncertainty on investment decisions is also studied in Pindyck (1993), who shows that, under identical technology and market conditions, industry-wide uncertainty induces investment inaction.

⁵ Of course, our results are not directly comparable, for two reasons. The first relates to the different time period analyzed, and more importantly to our focus on the Greek manufacturing sector, rather than the US manufacturing sector. The second relates to our use of crude oil prices rather than retail/firm energy prices.

The rest of the paper is organized as follows. Section 2 describes our econometric methodology for modeling the investment process and for measuring real oil price uncertainty. Section 3 gives a brief overview of the data employed, discusses our core empirical findings as well as various extensions and robustness experiments. Section 4 concludes.

2. EMPIRICAL METHODOLOGY

In our work we make use of dynamic random-effects models to model the probability of investment action, which include the previous state, to allow for state dependence. Special attention is paid to the treatment of unobserved heterogeneity and initial conditions. The former relates to whether the observed persistence of investment is the outcome of ‘pure’ or ‘spurious’ state dependence.⁶ The initial conditions are important, as in short panels, like ours, they have an impact on the entire path of outcomes.

The empirical specification for modeling the investment decision takes the form of a dynamic binary-choice model:

$$y_{it} = 1 \left\{ \mathbf{x}'_{it} \boldsymbol{\beta} + \gamma y_{it-1} + c_i + u_{it} > 0 \right\}, \quad i = 1, \dots, N; t = 1, \dots, T_i \quad (1)$$

where y_{it} is a binary indicator variable for investment action by plant $i = 1, \dots, N$ in year t , the vector \mathbf{x}_{it} contains explanatory variables affecting the propensity to trigger investment, while c_i denotes a time-invariant component capturing plant-specific heterogeneity and u_{it} is a well-behaved random term.

The random-effects (RE) specifications we employ require that the distributional properties of c_i and u_{it} , as well as their relationship to the explanatory variables, be specified, along with the initial conditions of the dynamic process. In what follows we assume that \mathbf{x}_{it} is strictly exogenous for u_{it} (conditional on c_i), and more specifically that $u_{it} | \mathbf{X}, \mathbf{c} \sim \text{n.i.i.d.}(0, \sigma_u^2)$.⁷ The standard random-effects model assumes that $c_i | \mathbf{x}_i \sim \text{n.i.i.d.}(0, \sigma_c^2)$. An alternative following Mundlak (1978) and Chamberlain (1984) is to allow for correlation between c_i and the observed characteristics, assuming a relationship of the form $c_i = \bar{\mathbf{x}}'_i \boldsymbol{\zeta} + \alpha_i$, with $\alpha_i \sim \text{n.i.i.d.}(0, \sigma_\alpha^2)$ being independent of \mathbf{x}_{it} and u_{it} for all i and t and $\bar{\mathbf{x}}_i \equiv T_i^{-1} \sum_{t=1}^T \mathbf{x}_{it}$ —the correlated random-effects (CRE) model. In this instance, model (1) may be written as

$$y_{it} = 1 \left\{ \mathbf{x}'_{it} \boldsymbol{\beta} + \gamma y_{it-1} + \bar{\mathbf{x}}'_i \boldsymbol{\zeta} + \alpha_i + u_{it} > 0 \right\}, \quad i = 1, \dots, N; t = 1, \dots, T_i \quad (2)$$

The random-effects specification (2) implies that the correlation between the composite error $v_{it} = \alpha_i + u_{it}$ in any two periods will be the same, namely $\rho = \text{corr}(v_{it}, v_{is}) = \sigma_\alpha^2 / (\sigma_\alpha^2 + \sigma_u^2)$ for $t, s = 1, \dots, T_i$ and $t \neq s$. Moreover, since y is binary, a convenient normalization is $\sigma_u^2 = 1$. If $\gamma = 0$, model (2) involves only a single integral, by conditioning on the individual effect and integrating it out, so parameters can be estimated by maximum likelihood (ML) using Gaussian–Hermite quadrature (Butler and Moffitt, 1982).

⁶ Pure state dependence would imply that the probability of investment in year t depends on the outcome in year $t - 1$, after controlling for unobserved heterogeneity.

⁷ Here $\mathbf{c} = (c_1, \dots, c_N)'$, and $\mathbf{X} = (\mathbf{x}'_1, \dots, \mathbf{x}'_N)'$ with $\mathbf{x}_i = (\mathbf{x}_{i1}, \dots, \mathbf{x}_{iT_i})'$.

In order to estimate the model when $\gamma \neq 0$, it is necessary to make an assumption about the relationship between the initial observation, y_{i0} , and the individual-specific effect. One possibility is to assume that y_{i0} is exogenous, i.e. a non-random starting position for each i . In this case, likelihood can be decomposed into two independent factors and the joint probability for $t=1, \dots, T_i$, and can be maximized without reference to that for $t=0$. However, if the initial conditions are correlated with α_i , this method of estimation overstates state dependence (Chay and Hyslop, 2000).

In our work we explore two alternative approaches that treat the initial observations as endogenous following Heckman (1981) and Wooldridge (2005), respectively.⁸ Heckman (1981) suggests specifying a linearized reduced-form equation for the initial value:

$$y_{i0} = 1 \left\{ \mathbf{z}'_{i0} \lambda + \theta \alpha_i + u_{i0} > 0 \right\} \tag{3}$$

where $\mathbf{z}_{i0} = (\mathbf{x}'_{i0}, \bar{\mathbf{x}}'_i)'$ and u_{i0} is assumed to be independent of α_i , with the former satisfying the same distributional assumptions as u_{it} for $t \geq 1$. A test of $\theta=0$ provides a test of exogeneity of the initial condition in this model.

Equations (2) and (3) together specify a complete model for a random sample (y_0, y_1, \dots, y_T) . One can then marginalize the likelihood with respect to α_i , obtaining the appropriate likelihood function for the maximization. For instance, the contribution to the likelihood for plant i in the model is given by

$$L_i = \int \left\{ \Phi \left[(\mathbf{z}'_{i0} \lambda + \theta \alpha_i) (2y_{i0} - 1) \right] \prod_{t=1}^{T_i} \Phi \left[(\mathbf{x}'_{it} \beta + \gamma y_{it-1} + \bar{\mathbf{x}}'_i \zeta + \alpha_i) (2y_{it} - 1) \right] \right\} d\Phi(\alpha_i) \tag{4}$$

where Φ is the standard normal cumulative distribution function. As α_i is normally distributed, the above integral can be evaluated using Gaussian–Hermite quadrature (Butler and Moffitt, 1982).

A different approach to the initial conditions problem is proposed by Wooldridge (2005), who suggests a conditional maximum likelihood (CML) estimator, considering the distribution conditional on the initial period value and exogenous covariates. More specifically, instead of specifying a model for y_{i0} given \mathbf{x}_i and α_i , a model is specified for α_i given \mathbf{x}_i and y_{i0} . In particular, it is assumed that

$$\alpha_i = \delta_0 + \delta_1 y_{i0} + a_i \tag{5}$$

as the Mundlak specification above already includes \mathbf{x}_i . Substituting into (2) gives

$$y_{it} = 1 \left\{ \mathbf{x}'_{it} \beta + \gamma y_{it-1} + \delta_0 + \delta_1 y_{i0} + \bar{\mathbf{x}}'_i \zeta + a_i + u_{it} > 0 \right\}, \quad i = 1, \dots, N; \quad t = 1, \dots, T_i \tag{6}$$

In this model, the contribution to the likelihood function for individual i is given by

$$L_i = \int \left\{ \prod_{t=1}^{T_i} \Phi \left[(\mathbf{x}'_{it} \beta + \gamma y_{it-1} + \delta_0 + \delta_1 y_{i0} + \bar{\mathbf{x}}'_i \zeta + a_i) (2y_{it} - 1) \right] \right\} d\Phi^*(a_i) \tag{7}$$

⁸ There is yet another approach, due to Orme (2001). A comparison of these three approaches is discussed in Arulampalam and Stewart (2009). See also Stewart (2007) for a discussion of the approaches of Heckman (1981) and Wooldridge (2005). We discuss results from all three approaches in an online supplement (supporting information).

where Φ^* is the normal distribution function of the new unobservable individual-specific heterogeneity a_i given in (5). Thus (6) is again a one-factor probit model that can be easily estimated by ML using Gaussian quadrature procedures. In Wooldridge's method, the exogeneity of the initial condition is tested by the significance of the coefficient on y_{i0} .

In all specifications above, u_{it} is assumed i. i. d. In a recent contribution Hsiao *et al.* (2012), following Pesaran (2004), propose a simple test statistic to assess the null of cross-sectional independence. In particular, they suggest using

$$CD = \sqrt{\frac{2}{N(N-1)}} \left(\sum_{i=1}^N \sum_{j=i+1}^{N-1} \sqrt{T_{ij}} \hat{r}_{ij} \right) \quad (8)$$

where T_{ij} is the number of common time series observations available for any pair of plants i and j , and \hat{r}_{ij} is the correlation coefficient computed using the generalized residuals estimated under the null hypothesis. They show that under the null hypothesis of cross-sectional independence $CD \xrightarrow{d} N(0, 1)$ for $N, T \rightarrow \infty$ and that the CD statistic has exactly mean at zero for fixed values of N and T , under a wide range of panel data models, including heterogeneous models, non-stationary and dynamic panels.

3. DATA AND EMPIRICAL FINDINGS

3.1. Data and Benchmark Measure of Real Oil Price Uncertainty

The data used in this paper come from the Annual Industrial Survey (AIS) for Greece, which surveys plants belonging to all firms with more than 10 employees across 21 manufacturing industries. The sample constitutes an unbalanced panel of plants, built from data collected in the 12 AISs for the period 1994–2005 (51,881 plant–year observations). The investment indicator (y_{it}) is constructed on the basis of the difference between (gross) values for capital acquisitions and disposals by plant, reported by the AIS. The vector of explanatory variables (\mathbf{x}_{it}) we use includes plant-specific characteristics such as sales, cash flow, equity and loans as ratios to value-added and (log) employment, all lagged one period to avoid simultaneity. In addition, it includes the percentage change in real oil prices and a measure of real oil price uncertainty. The former is obtained by using annual data on Brent, quoted in US dollars, converted into euros and then deflated by the producer price index of manufacturing goods.⁹

To obtain our uncertainty metric, we estimate a GARCH(1,1) model using the same data on the percentage change of real oil prices, on a monthly frequency. In particular the conditional mean is chosen to be a restricted AR(10), to ensure that no autocorrelation is present in the residuals. The model is estimated recursively, in each case utilizing monthly observations up to December of year $t - 1$. Using these estimates, we forecast the conditional standard deviation for the 12 months in year t and then use the average predicted volatility in year t as our benchmark measure of real oil price uncertainty. Note that this uncertainty measure, albeit backward looking, reflects that economic agents, upon deciding, have to make forecasts about the uncertainty they will be facing, and in addition it is a measure known at the beginning of year t .

3.2. Empirical Results

In all specifications the set of explanatory variables includes industry and year dummies, while the CRE estimators also include time averages of each plant-specific (time-varying) characteristic included

⁹ As discussed in Edelstein and Kilian (2007), there is a subtle difference between firm energy prices and the crude oil prices we employ here. As the former are unavailable, we employ the latter in our analysis as a proxy.

Table I. Dynamic models of investment activity

Covariate	Dynamic RE (exogenous initial conditions)	Dynamic CRE (exogenous initial conditions)	Dynamic CRE (endogenous initial conditions)	
			Heckman	Wooldridge
$y_{i,t-1}$	0.330*** [35.611]	0.314*** [33.796]	0.194*** [23.509]	0.275*** [30.902]
$\hat{\sigma}_{i t-1}^{oil}$	-0.749*** [-5.591]	-0.500*** [-3.843]	-0.454*** [-3.160]	-0.460*** [-3.537]
π_t^{oil}	-0.106*** [-4.916]	-0.072*** [-3.425]	-0.071*** [-3.087]	-0.070*** [-3.353]
$SL_{i,t-1}$	0.037*** [7.012]	0.015*** [2.230]	0.013 [1.492]	0.015*** [2.293]
$CF_{i,t-1}$	0.019*** [5.180]	0.008* [1.702]	0.007 [1.263]	0.008* [1.713]
$EMP_{i,t-1}$	0.114*** [33.831]	0.090*** [12.866]	0.083*** [9.993]	0.092*** [13.194]
$EQ_{i,t-1}$	0.035*** [13.521]	0.007*** [2.717]	0.005 [1.522]	0.007** [2.456]
$LO_{i,t-1}$	0.037*** [3.001]	-0.013 [-0.981]	-0.001 [-0.051]	-0.012 [-0.936]
<i>Time-averaged plant characteristics (observed heterogeneity)</i>				
\overline{SL}_i		0.027*** [2.725]	0.036*** [2.600]	0.020** [2.017]
\overline{CF}_i		0.031*** [4.097]	0.029** [2.545]	0.033*** [4.320]
\overline{EMP}_i		0.008 [1.079]	0.027*** [2.873]	-0.007 [-0.937]
\overline{EQ}_i		0.101*** [19.209]	0.123*** [16.625]	0.095*** [18.133]
\overline{LO}_i		0.278*** [8.736]	0.314*** [7.033]	0.267*** [8.545]
<i>Initial conditions</i>				
θ			1.019*** [13.210]	
$y_{i,0}$				0.136*** [17.841]
<i>Diagnostics</i>				
ρ	0.282*** [20.589]	0.256*** [19.102]	0.318*** [19.892]	0.259*** [20.505]
log L	-15812.196	-15478.078	-12717.619	-15275.266
No. obs.	42,794	42,794	47,997	42,794
CD-test	612.352***	766.229***	797.184***	762.235***
\bar{r}	0.342	0.358	0.363	0.359
<i>Predicted probabilities</i>				
Pred. prob. \hat{p}_0	0.412	0.386	0.419	0.366
Pred. prob. \hat{p}_1	0.925	0.934	0.947	0.938
APE = $\hat{p}_1 - \hat{p}_0$	0.512	0.547	0.528	0.572
PPR = \hat{p}_1 / \hat{p}_0	2.242	2.416	2.259	2.562

Note: The oil price uncertainty metric, $\hat{\sigma}_{i|t-1}^{oil}$, is constructed as a 12-month average of the predicted one-year-ahead monthly real oil price volatility. π_t^{oil} denotes the percentage change of the real oil price in year t relative to year $t-1$. The set of controls also includes industry and time dummies. In the first two specifications the initial condition is taken to be exogenous, while $y_{i,0}$ denotes the initial condition, as in Wooldridge (2005). In the Heckman (1981) estimator, the initial period is modeled as a function of $SL_{i,t-1}$, $CF_{i,t-1}$, $EMP_{i,t-1}$ and all time-averaged plant-specific characteristics. log L in the Heckman (1981) estimator is for the joint model for all periods (1994–2005), whereas in all other models it corresponds to the period 1995–2005, which also explains the difference in the number of observations. CD-test denotes the test of cross-sectional independence proposed by Hsiao *et al.* (2012), and \bar{r} indicates the average pair-wise correlation coefficients of the generalized residuals. \hat{p}_0 denotes the average predicted probability for investment action in year t , given inaction in the previous year; \hat{p}_1 denotes the average predicted probability for investment action in year t , given action in the previous year; APR stands for average partial effect and PPR stands for predicted probability ratio. The numbers in square brackets denote z -scores, while *, **, and *** indicate statistical significance at the 10%, 5% and 1% level.

in \mathbf{x}_{it} . Table I reports the estimated marginal effect of each covariate for all dynamic random-effects probit models, holding all other covariates at their respective sample means—with industry and time effects also evaluated at their mean values. The second column reports estimates from the standard

random-effects model; the third column estimates from the CRE model and the last two columns report estimates from the CRE estimators of Heckman and Wooldridge.¹⁰

The signs of the control variables suggest that larger plants (higher employment) show higher probability of investment. Similar results hold for higher level of equity financing, higher sales and higher operating profits (cash flow), which are also associated with higher probability of investment action—although the estimates from the Heckman estimator indicate that their effects are insignificant. A higher loan to value-added ratio, on the other hand, is found to increase the probability of investment only when employing the simple random-effects probit estimators.

Furthermore, time-averaged variables—representing fixed underlying differences between plants in the CRE specifications—play a key role in the model, accounting for the potential correlation between the unobserved individual-specific heterogeneity and observable characteristics. Most of their marginal effects are individually statistically significant, suggesting that the CRE specifications are more appropriate. In addition, these estimates carry the same sign with the marginal effects of the corresponding year-specific variables. Two notable differences are average employment—being significant only in the Heckman specification—and average loan to value-added—an increase in which results in a significantly higher probability of investment action in all three CRE estimators.

Moreover, we find strong evidence of endogeneity of the initial conditions, which turn out to be strong determinants of the subsequent investment decision process. In particular, looking at the results from both the Heckman and Wooldridge estimators, we reject the null hypothesis of exogeneity of initial conditions—as θ for the former and the marginal effect of $y_{i,0}$ for the latter are strongly significant.

Examining the issue of state dependence, across all four specifications the lagged investment activity variable is highly significant, reflecting strong persistence. We find that the size of the relevant estimated marginal effect decreases somewhat when we take into account heterogeneity and especially when initial conditions are treated as endogenous. In addition, there are a number of ways in which the partial effect of y_{it-1} on $Pr(y_{it}=1)$ may be assessed in models like the ones considered here. The approach we take is based on estimates of counterfactual outcome probabilities taking y_{it-1} as fixed at 0 and at 1, and evaluated at $\mathbf{x}_{it} = \bar{\mathbf{x}}$ (with industry and time effects also evaluated at their averages). That is, we calculate \tilde{p}_0 and \tilde{p}_1 , which stand for the predicted probabilities of investment action in year t , given investment inaction or action in $t-1$, respectively. Then, the magnitude of the effect of past investment activity can be assessed using the concepts of the average partial effect (APE), defined as $(\tilde{p}_1 - \tilde{p}_0)$, and the predicted probability ratio (PPR), defined as $(\tilde{p}_1/\tilde{p}_0)$.

The estimated probabilities are reported in the bottom panel of Table I, along with the APEs and the PPRs, for each model. The predicted probability of being active in investment at year t , conditional on being active in year $t-1$, is estimated to be in the range between 92% (RE probit) and 95% (Heckman estimator), while the predicted probability of being active in investment in year t , conditional on being inactive in $t-1$, ranges between 37% (Wooldridge estimator) and 42% (Heckman estimator). Hence the APE is between 51 and 57 percentage points, while the PPR is between 2.24 and 2.56. Thus, on average, and controlling for heterogeneity, past investment action is associated with a difference in the probability of current investment action by more than 50 percentage points or, put differently, the probability of investment action is at least some 2.2 times higher if there has been some investment action during last period.

We next turn to the two oil-related variables of interest in our specifications: the percentage change in real oil prices and real oil price uncertainty. First, we find that—in our models where the percentage change of oil prices enters linearly—an increase in real oil prices reduces significantly the probability of investment action, a finding which holds across all estimators employed. For instance, focusing on

¹⁰ The initial period in the Heckman estimator is modeled as a function of sales, cash flow and employment and time averages of all covariates included in the model. The rest of the covariates as well as industry dummies had to be dropped for identification purposes.

the last column of Table I, we find that an increase in real oil prices by one percentage point reduces the probability of investment action by 0.07%. Second, we also find that an increase in real oil price uncertainty reduces significantly the probability of investment action, irrespective of the estimator employed. For example, focusing again on the last column of Table I, we find that an increase in our measure of real oil price uncertainty by 0.01 (roughly 11% relative to its average value) reduces the probability of investment action by 0.46%.¹¹ Moreover, even when allowing for unobserved heterogeneity to be correlated with observable characteristics, as well as explicitly modeling initial conditions, increases in real oil prices and real oil price uncertainty retain their negative effect on the probability of investment activity. More importantly, though, we see that the estimated magnitude of these effects is robust across all estimators employed.

Finally, we evaluate the extent to which the assumption of cross-sectional independence of the error term is valid, by means of the CD-test. For all four estimators, we find that the null is strongly rejected. In addition, the estimated average cross-sectional correlation of the generalized residuals is above 0.34. However, there is no well-established technique that allows us to correct for this deviation from the i.i.d. assumption.¹² To this end, our results should be interpreted keeping this caveat in mind.

3.3. Extensions and Sensitivity Analysis

In this subsection we examine various extensions, such as using different measures of real oil price uncertainty, assessing the existence of asymmetry of oil-related effects, and expanding the set of controls to include plant-specific uncertainty, the business cycle and industry-wide uncertainty.¹³

3.3.1. Alternative Measures of Real Oil Price Uncertainty

Thus far, we have employed a measure of real oil price uncertainty that is derived from a GARCH model of conditional volatility, which, despite being based on out-of-sample forecasts over a one-year horizon, might not fully capture the ‘risk’ a decision maker is facing. On the one hand, this measure converges quickly to the unconditional volatility of real oil prices (Kilian and Vigfusson, 2011) and, on the other hand, in the context of investment decisions, the risk of real oil price increases rather than a simple increase in variance of real oil prices is probably more relevant.

In deriving such one-sided ‘risk’ measures we have two options. The first is to follow Kilian and Manganelli (2007) and define the ‘risk’ of excessive real oil price increase h periods from date τ , above a specific threshold value, $\bar{\pi}$, as

$$\text{EIR}_{\zeta\tau}(\bar{\pi}) = \int_{\bar{\pi}}^{+\infty} (\pi_{\tau+h} - \bar{\pi})^{\zeta} dF(\pi_{\tau+h}) \quad (9)$$

where $F(\cdot)$ is the probability distribution function of future real oil price change outcomes ($\pi_{\tau+h}$), estimated by the empirical distribution of real oil price change forecasts. Note that this class of risk measures is defined in terms of percentage increases in real oil prices, which squares well with standard financial planning models and practice (Ross *et al.*, 2005). In such models, one usually employs forecasts of the growth rates of all the relevant variables (such as sales, cost, etc.) as inputs, so risk measures like (9) seem more appropriate. The second, which is more conventional in the economics

¹¹ To understand better the magnitude of these effects, note that an increase of sales by 1% of value-added increases the probability of investment action by 0.15%!

¹² We have already included time effects as the least possible remedy for the existence of cross-sectional dependence.

¹³ We briefly discuss results when using alternative measures of uncertainty/risk. The rest of our estimation results are available as supporting information in the online supplement.

literature, is to define the risk measures in terms of the real oil price (the relative price of oil), as this would show up in many standard profit maximization problems. In this instance, we may define the ‘upside risk’ that real oil prices h periods from date τ , $R_{\tau+h}$, will be above a threshold value, \bar{R} , as

$$\text{UR}_{\zeta\tau}(\bar{R}) = \int_{\bar{R}}^{+\infty} (R_{\tau+h} - \bar{R})^{\zeta} dF^*(R_{\tau+h}) \quad (10)$$

where $F^*(\cdot)$ is the predictive distribution of real oil prices. As both these classes of risk measures are useful in different contexts, we report results for both.

Before proceeding, note that for $\zeta=1$ both (9) and (10) reduce to tail conditional expectations, multiplied by the corresponding tail probabilities: $\text{EIR}_{1\tau}(\bar{\pi}) = \text{E}(\pi_{\tau+h} - \bar{\pi} | \pi_{\tau+h} > \bar{\pi}) \text{Pr}(\pi_{\tau+h} > \bar{\pi})$ and $\text{UR}_{1\tau}(\bar{R}) = \text{E}(R_{\tau+h} - \bar{R} | R_{\tau+h} > \bar{R}) \text{Pr}(R_{\tau+h} > \bar{R})$; while for $\zeta=2$ these reduce to the (one-sided) variance about the target again multiplied by the corresponding tail probability: $\text{EIR}_{2\tau}(\bar{\pi}) = \text{E}[(\pi_{\tau+h} - \bar{\pi})^2 | \pi_{\tau+h} > \bar{\pi}] \text{Pr}(\pi_{\tau+h} > \bar{\pi})$ and $\text{UR}_{2\tau}(\bar{R}) = \text{E}[(R_{\tau+h} - \bar{R})^2 | R_{\tau+h} > \bar{R}] \text{Pr}(R_{\tau+h} > \bar{R})$. The excessive increase risk measures (EIR) can be computed as in Kilian and Manganeli (2007), and the upside risk measures (UR) can be calculated as discussed in Alquist *et al.* (2011), for different values of ζ . In calculating such risk measures we have chosen $\bar{\pi}$ to be 20% and \bar{R} to be 50 euros (in constant 2005 prices)—our results not being sensitive to this choice—and focus at a four-year-ahead horizon. Albeit limited in nature, as many business fixed investment projects tend to have lifetimes well beyond four years, this choice is intended to capture—to the extent that is possible—that the relevant measure of risk should reflect the lifetime of the investment project.

Table II summarizes the estimated marginal effects from employing the dynamic CRE estimators of Heckman and Wooldridge and these one-sided risk measures, leaving the rest of the controls the same. When examining real oil prices, we see that their estimated marginal effects are closely in line with those reported in Table I and significant. One exception is when we employ the Heckman estimator and the UR measures of one-sided risk: in this case increases in real oil prices do not influence significantly the probability of investment.¹⁴

On the other hand, when we employ the EIR_1 or the EIR_2 measures, we find that any increase in these translates to a significantly lower probability of investment action. Instead, when we employ the UR measures, we again find that the probability of investment is lowered, but not in a significant manner. Note, however, that the estimated marginal effects of EIR and UR are closely aligned. Focusing on the estimator of Wooldridge, an increase of EIR_1 by 0.01 (about 8.19% above its mean) reduces the probability of investment action by 0.22%, while an equiproportional increase in UR_1 (relative to its mean) leads to a reduction of the probability of investment by 0.15%. Similarly, an increase in EIR_2 by 0.01 (about 17% above its mean) reduces the probability of investment action by 0.74%, while an equiproportional increase of UR_2 (relative to its mean) reduces the probability by 0.5%.

3.3.2. Other Robustness Experiments

Our findings thus far are robust to a number of different extensions. First, in line with the empirical macroeconomics literature which aims at identifying the unpredictable component in real oil prices (‘shocks’) as the relevant measure that affects spending decisions (Edelstein and Kilian, 2007, 2009), we isolate the unpredictable component of the four-year percentage change in real oil prices,

¹⁴ The marginal effects of the other covariates are similar in terms of magnitude and statistical significance to those discussed in Table I.

Table II. Alternative uncertainty/one-sided risk measures

Covariate	Risk measure employed							
	EIR _{1t} ^{0,20}		EIR _{2t} ^{0,20}		UR _{1t} (50)		UR _{2t} (50)	
	Heckman	Wooldridge	Heckman	Wooldridge	Heckman	Wooldridge	Heckman	Wooldridge
$y_{i,t-1}$	0.194*** [23.510]	0.275*** [30.902]	0.194*** [23.836]	0.275*** [30.902]	0.194*** [23.511]	0.275*** [30.902]	0.194*** [23.508]	0.275*** [30.902]
$\hat{\sigma}_t^{\text{oil}}$	-0.213*** [-3.160]	-0.216*** [-3.539]	-0.730*** [-3.191]	-0.739*** [-3.695]	-0.256 × 10 ⁻² [-0.335]	-0.956 × 10 ⁻² [-1.385]	-0.012 × 10 ⁻² [-0.335]	-0.045 × 10 ⁻² [-1.365]
π_t^{oil}	-0.042** [-2.212]	-0.041** [-2.376]	-0.048** [-2.458]	-0.047*** [-2.648]	-0.038 [-1.309]	-0.066*** [-2.597]	-0.038 [-1.308]	-0.067*** [-2.605]
$SL_{i,t-1}$	0.013 [1.492]	0.015** [2.293]	0.013 [1.492]	0.015** [2.293]	0.013 [1.492]	0.015** [2.293]	0.013 [1.492]	0.015** [2.293]
$CF_{i,t-1}$	0.007 [1.263]	0.008* [1.713]	0.007 [1.263]	0.008* [1.713]	0.007 [1.263]	0.008* [1.713]	0.007 [1.263]	0.008* [1.713]
$EMP_{i,t-1}$	0.083*** [9.993]	0.092*** [13.194]	0.083*** [10.018]	0.092*** [13.194]	0.083*** [9.993]	0.092*** [13.194]	0.083*** [9.993]	0.092*** [13.194]
$EQ_{i,t-1}$	0.005 [1.522]	0.007** [2.456]	0.005 [1.522]	0.007** [2.456]	0.005 [1.522]	0.007** [2.456]	0.005 [1.522]	0.007** [2.456]
$LO_{i,t-1}$	-0.001 [-0.051]	-0.012 [-0.936]	-0.001 [-0.051]	-0.012 [-0.936]	-0.001 [-0.051]	-0.012 [-0.936]	-0.001 [-0.051]	-0.012 [-0.936]
<i>Time-averaged plant characteristics (observed heterogeneity)</i>								
SL_i	0.036*** [2.600]	0.020** [2.017]	0.036*** [2.601]	0.020** [2.017]	0.036*** [2.600]	0.020** [2.017]	0.036*** [2.600]	0.020** [2.017]
\overline{CF}_i	0.029** [2.545]	0.033*** [4.320]	0.029** [2.545]	0.033*** [4.320]	0.029** [2.545]	0.033*** [4.320]	0.029** [2.545]	0.033*** [4.320]
\overline{EMP}_i	0.027*** [2.873]	-0.007 [-0.937]	0.027*** [2.874]	-0.007 [-0.937]	0.027*** [2.873]	-0.007 [-0.937]	0.027*** [2.873]	-0.007 [-0.937]
\overline{EQ}_i	0.123*** [16.626]	0.095*** [18.133]	0.123*** [16.742]	0.095*** [18.133]	0.123*** [16.626]	0.095*** [18.133]	0.123*** [16.625]	0.095*** [18.133]
\overline{LO}_i	0.314*** [7.033]	0.267*** [8.545]	0.314*** [7.041]	0.267*** [8.545]	0.314*** [7.033]	0.267*** [8.545]	0.314*** [7.033]	0.267*** [8.545]

<i>Initial conditions</i>					
θ	1.019*** [13.210]	1.019*** [13.210]	1.019*** [13.210]	1.019*** [13.210]	1.019*** [13.210]
$\gamma_{i,0}$	0.136*** [17.841]	0.136*** [17.841]	0.136*** [17.841]	0.136*** [17.841]	0.136*** [17.841]
<i>Diagnostics</i>					
ρ	0.318*** [19.892]	0.318*** [19.892]	0.318*** [19.892]	0.318*** [19.892]	0.318*** [19.892]
$\log L$	-12717.619	-12717.619	-12717.619	-12717.619	-12717.619
No. obs.	47,997	47,997	47,997	47,997	47,997
CD-test	797.184***	797.184***	797.184***	797.184***	797.184***
\bar{r}	0.363	0.363	0.363	0.363	0.363
<i>Predicted probabilities</i>					
Pred. prob. \bar{p}_0	0.419	0.419	0.419	0.419	0.419
Pred. prob. \bar{p}_1	0.947	0.947	0.947	0.947	0.947
APE = $\bar{p}_1 - \bar{p}_0$	0.528	0.528	0.528	0.528	0.528
PPR = \bar{p}_1/\bar{p}_0	2.259	2.259	2.259	2.259	2.259

Note: $EIR_{\zeta\tau}^{0,xx}$ denotes the excessive increase risk of real oil price changes increasing above $x\tau\%$, whereas $UR_{\zeta\tau}(\bar{R})$ denotes the upside risk of real oil prices being above \bar{R} . In particular, let the excessive increase risk be $EIR_{\zeta\tau}(\bar{\pi}) = \int_{\bar{\pi}}^{+\infty} (\pi_{\tau+h} - \bar{\pi})^{\zeta} dF(\pi_{\tau+h})$, where $\pi_{\tau+h}$ denotes the four-year change in real oil prices (with monthly observations), $F(\cdot)$ is the distribution function of future real oil price change outcomes, $\pi_{\tau+h}$ and $\bar{\pi}$ measures the particular threshold relative to which 'excessiveness' is defined. Similarly let $UR_{\zeta\tau}(\bar{R}) = \int_{\bar{R}}^{+\infty} (R_{\tau+h} - \bar{R})^{\zeta} dF^*(R_{\tau+h})$, where $R_{\tau+h}$ denotes real oil prices 48 months ahead (with monthly observations), $F^*(\cdot)$ is the distribution function of future real oil price outcomes, $R_{\tau+h}$ and \bar{R} measures the particular threshold, relative to which upside risk is defined. Letting $\zeta = 1$, we obtain $EIR_{1\tau}(\bar{\pi}) = E[\pi_{\tau+h} - \bar{\pi} | \pi_{\tau+h} > \bar{\pi}] Pr(\pi_{\tau+h} > \bar{\pi})$ and $UR_{1\tau}(\bar{R}) = E[(R_{\tau+h} - \bar{R}) | R_{\tau+h} > \bar{R}] Pr(R_{\tau+h} > \bar{R})$; letting $\zeta = 2$ we obtain $EIR_{2\tau}(\bar{\pi}) = E[(\pi_{\tau+h} - \bar{\pi})^2 | \pi_{\tau+h} > \bar{\pi}] Pr(\pi_{\tau+h} > \bar{\pi})$ and $UR_{2\tau}(\bar{R}) = E[(R_{\tau+h} - \bar{R})^2 | R_{\tau+h} > \bar{R}] Pr(R_{\tau+h} > \bar{R})$. The EIR quantities are estimated along the lines of Kilian and Manganello (2007), and they measure the risk a decision maker is facing that, at month τ , real oil prices after 48 months will increase by more than the chosen threshold, $\bar{\pi}$. The UR quantities are estimated as described in Alquist *et al.* (2011), and they measure the risk a decision maker is facing that, at month τ , real oil prices after 48 months will be above the chosen threshold, \bar{R} . The measures we employ are defined as the 12-month averages of $EIR_{\zeta\tau}(\bar{\pi})$ and $UR_{\zeta\tau}(\bar{R})$ respectively. That is, $EIR_{\zeta\tau}^{\bar{\pi}} \equiv (1/12) \sum_{\tau=1}^{12} EIR_{\zeta\tau}(\bar{\pi})$ and $UR_{\zeta\tau}^{\bar{R}} \equiv (1/12) \sum_{\tau=1}^{12} UR_{\zeta\tau}(\bar{R})$. See also note to Table 1.

as the residual from a first-order autoregression.¹⁵ Our findings are almost identical to those reported in Table II, showing that there is a strong negative effect of oil price ‘surprises’ to the probability of investment. These findings can be considered to be complementary to those in Edelstein and Kilian (2007), as we provide evidence that oil price ‘shocks’ are important determinants of investment decisions at the plant level.

Second, we explore the possibility of asymmetric effects of real oil prices and real oil price uncertainty on plant-level investment, along two dimensions, namely size (Campa, 1994; Ghosal and Loungani, 2000) and oil intensity in production. Here we proxy size by the level of employment, classifying a plant as small (large) if its number of employees is below (above) the median. In assessing the importance of oil intensity one should focus on the indirect energy share (Lee and Ni, 2002), which, however, has been shown not to be a key factor by Kilian and Park (2009). As data on the indirect energy share are unavailable, we proxy (direct) oil intensity by the share of plant petrol expenditures to total energy expenditures. As far as real oil prices are concerned, we find no differential effect on either small or highly oil-dependent plants, in line with the evidence in Kilian and Park (2009). Similarly, we find that rising real oil price uncertainty does not have a differential effect on plants that are highly dependent on oil. We do find, however, that smaller plants are indeed influenced more severely by rising real oil price uncertainty, documenting the differential effect of uncertainty on the investment decisions of smaller production units.

Third, in order to assess whether plant-specific uncertainty makes any difference to our results, we obtain measures of plant-specific uncertainty, by estimating a (pooled) AR(1) process for profits, allowing for time-varying conditional volatility by means of a pooled-panel GARCH (PP-GARCH) in the spirit of Cermeno and Grier (2006).¹⁶ Based on this, we produce one-year-ahead predictions of conditional standard deviation of profits, which we use as an extra control in our analysis. Moreover, to account for the fact that both the real oil price volatility and the profit volatility may vary with the business cycle, we also include the economy-wide output gap as a control in our analysis.¹⁷ We find that the inclusion of plant-specific uncertainty and the output gap does not affect our previous results in any substantial manner. In addition, increases in plant-specific uncertainty do not reduce significantly the probability of investment action. Moreover, when output is above trend, the probability of investment action increases significantly; in other words, the probability of investment action increases when output rises above trend output (the economy is in a boom), whereas it is less likely that investment will take place when output is cyclically below trend (in a recession).

Finally, we evaluate whether the inclusion of industry-wide uncertainty in our controls makes a significant difference to our results. To measure industry uncertainty we follow Bloom *et al.* (2007) and use the unconditional standard deviation of daily stock returns from the Industrials Price Index, in year t , which is a forward-looking indicator incorporating the impact of different sources of uncertainty on the whole manufacturing sector. As this measure may partly reflect noise unrelated to fundamentals, we also consider a measure that normalizes the daily industry returns by the return on all shares to eliminate the effect of any aggregate stock market bubbles. Utilizing both these measures, we find that increasing industry-wide uncertainty reduces significantly the probability of investment action. In both cases, however, our conclusions regarding real oil price changes and real oil price uncertainty remain unaffected.

¹⁵ The alternative is to obtain such ‘shocks’ from a vector autoregression model, as in Kilian (2008, 2009a), but this is beyond the scope of our analysis.

¹⁶ As the data we use are confidential, we are unable to match plants with specific firms, and hence the use of uncertainty measures based on stock return volatility (see, for example, Leahy and Whited, 1996; Bloom *et al.*, 2007) is not possible. See also Ghosal and Loungani (2000) for an application using industry profits. These authors obtain their uncertainty variable by modeling the profit rate as a (panel) AR(2) process.

¹⁷ Output gap is defined as the ratio of actual to potential real gross domestic product for the whole economy.

4. CONCLUSIONS

The goal of this paper has been to investigate the effects of real oil prices and real oil price uncertainty on the dynamic behavior of investment decisions at a disaggregate level, using plant-level data for the Greek manufacturing sector. We find that increases in both these variables reduce significantly the probability of investment. The finding that rising real oil price uncertainty increases the likelihood of investment inaction is in line with the predictions of irreversible investment theory. We have assessed the sensitivity of these effects to a number of modeling assumptions, such as the modeling of unobserved heterogeneity and the endogeneity of initial conditions, and found that these are robustly estimated as being negative and significant in all cases.

Moreover, the negative effects of increasing real oil prices and real oil price uncertainty are also robust across different measures of real oil price uncertainty, including one-sided risk measures, which capture the risk of real oil price increases a decision maker is facing. These findings are also robust across a number of extensions, including taking into account industry-wide uncertainty, controlling for the business cycle and for plant-specific uncertainty. In addition, our analysis provides evidence that the effect of rising real oil price uncertainty is non-uniformly distributed across decisions makers, since the resulting reduction of the probability of investment action is amplified for smaller plants, whereas no such evidence is found for real oil prices.

Finally, we document the existence of strong state dependence in investment. We find that estimates of state dependence in investment are affected by the assumption made regarding initial conditions and the treatment of unobserved heterogeneity. Despite the sensitivity to these assumptions, we find that the likelihood of triggering investment is significantly positively correlated with investment action in the last period, suggesting that its presence significantly affects the time trajectory of investment decisions.

ACKNOWLEDGEMENTS

We have benefited from discussions with Fragiskos Archontakis, Rafael De Hoyos, Stelios Fountas, Nicholas Gianakopoulos, Giannis Karagiannis, Theodore Panagiotidis, Andreas Pick, Costas Roumanias, Vasilis Sarafidis, Mark Stewart and Athanasios Tagkalakis, as well as from comments from participants of the 2010 ASSET Meeting in Alicante. We are also indebted to M. Hashem Pesaran and two anonymous referees of this journal for comments on earlier versions, which helped considerably in improving the content of this paper. Previous versions of this paper were circulated under the titles 'Oil Price Uncertainty and Business Fixed Investment: A Real Options Approach' and 'Investment Decisions of Manufacturing Plants: Investigating State Dependence and Oil Uncertainty Effects'. Any remaining errors and ambiguities remain our own responsibility.

REFERENCES

- Abel A, Eberly J. 1994. A unified model of investment under uncertainty. *American Economic Review* **84**: 1369–1384.
- Abel A, Eberly J. 1996. Optimal investment with costly reversibility. *Review of Economic Studies* **63**: 581–593.
- Alquist R, Kilian L, Vigfusson R. 2011. Forecasting the price of oil. CEPR Discussion Paper No. 8388, Centre for Economic Policy Research, London.
- Arulampalam W, Stewart M. 2009. Simplified implementation of the Heckman estimator of the dynamic probit model and a comparison with alternative estimators. *Oxford Bulletin of Economics and Statistics* **71**: 659–681.
- Barnett S, Sakellaris P. 1999. A new look at firm market value, investment, and adjustment costs. *The Review of Economics and Statistics* **81**: 250–260.
- Bernanke B. 1983. Irreversibility, uncertainty, and cyclical investment. *Quarterly Journal of Economics* **98**: 85–106.

- Bloom N, Bond S, van Reenen J. 2007. Uncertainty and investment dynamics. *Review of Economic Studies* **74**: 391–415.
- Bontempi E, Del Boca A, Franzosi A, Galeotti M, Rota P. 2004. Capital heterogeneity: does it matter? Fundamental Q and investment on a panel of Italian firms. *The RAND Journal of Economics* **35**: 674–690.
- Butler JS, Moffitt R. 1982. A computationally efficient quadrature procedure for the one-factor multinomial probit model. *Econometrica* **50**: 761–764.
- Caballero R, Engel E. 1999. Explaining investment dynamics in US manufacturing: a generalised (S,s) approach. *Econometrica* **67**: 783–826.
- Caballero R, Engel E, Haltiwanger J. 1995. Plant-level adjustment and aggregate investment dynamics. *Brookings Papers on Economic Activity* **2**: 1–54.
- Campa J. 1993. Entry by foreign firms in the United States under exchange rate uncertainty. *The Review of Economics and Statistics* **75**: 614–622.
- Campa J. 1994. Multinational investment under uncertainty in the chemical processing industries. *Journal of International Business Studies* **25**: 557–578.
- Campa J, Goldberg L. 1995. Investment in manufacturing, exchange rates and external exposure. *Journal of International Economics* **38**: 297–320.
- Cermeno R, Grier K. 2006. Conditional heteroscedasticity and cross-sectional dependence in panel data: an empirical study of inflation uncertainty in the G7 countries. In *Contributions in Economic Analysis*, Baltagi BH (ed.). North-Holland: Amsterdam; 259–278.
- Chamberlain G. 1984. Panel data. In *Handbook of Econometrics*, Griliches S, Intriligator M (eds). North-Holland: Amsterdam; 1247–1318.
- Chay K, Hyslop D. 2000. Identification and estimation of dynamic binary response panel data models: empirical evidence using alternative approaches. Working paper, University of California at Berkeley.
- Cooper R, Haltiwanger J. 2006. On the nature of capital adjustment costs. *Review of Economic Studies* **73**: 611–633.
- Dixit A, Pindyck R. 1994. *Investment under Uncertainty*. Princeton University Press: Princeton, NJ.
- Doms M, Dunne T. 1998. Capital adjustment patterns in manufacturing plants. *Review of Economic Dynamics* **1**: 409–429.
- Edelstein P, Kilian L. 2007. The response of business fixed investment to changes in energy prices: a test of some hypotheses about the transmission of energy price shocks. *BE Journal of Macroeconomics* **7**.
- Edelstein P, Kilian L. 2009. How sensitive are consumer expenditures to retail energy prices? *Journal of Monetary Economics* **56**: 766–779.
- Gelos G, Isgut A. 2001. Fixed capital adjustment: is Latin America different? *The Review of Economics and Statistics* **83**: 717–726.
- Ghosal V, Loungani P. 2000. The differential impact of uncertainty on investment in small and large businesses. *The Review of Economics and Statistics* **82**: 338–349.
- Hamilton JD. 2009. Causes and consequences of the oil shock of 2007–08. *Brookings Papers on Economic Activity* **1** Spring: 215–261.
- Heckman J. 1981. The incidental parameters problem and the problem of initial conditions in estimating a discrete time-discrete data stochastic process. In *Structural Analysis of Discrete Data with Econometric Applications*, Manski C, McFadden D (eds). MIT: Cambridge, MA; 179–195.
- Hsiao C, Pesaran MH, Pick A. 2012. Diagnostic tests of cross section independence for limited dependent variable panel data models. *Oxford Bulletin of Economics and Statistics* **74**: 253–277.
- Kilian L. 2008. Exogenous oil supply shocks: how big are they and how much do they matter for the U.S. economy? *The Review of Economics and Statistics* **90**: 216–240.
- Kilian L. 2009a. Not all oil price shocks are alike: disentangling demand and supply shocks in the crude oil market. *American Economic Review* **99**: 1053–1069.
- Kilian L. 2009b. Comment on ‘Causes and consequences of the oil shock of 2007-08’ by J.D. Hamilton. *Brookings Papers on Economic Activity* **1** Spring: 267–278.
- Kilian L, Manganelli S. 2007. Quantifying the risk of deflation. *Journal of Money, Credit, and Banking* **39**: 561–590.
- Kilian L, Park C. 2009. The impact of oil price shocks on the U.S. stock market. *International Economic Review* **50**: 1267–1287.
- Kilian L, Vigfusson R. 2011. Nonlinearities in the oil price–output relationship. *Macroeconomic Dynamics* **15**: 337–363.
- Leahy J, Whited T. 1996. The effect of uncertainty on investment: some stylized facts. *Journal of Money, Credit, and Banking* **28**: 64–83.
- Lee K, Ni S. 2002. On the dynamic effects of oil price shocks: a study using industry level data. *Journal of Monetary Economics* **49**: 823–852.

- McDonald R, Siegel D. 1986. The value of waiting to invest. *Quarterly Journal of Economics* **101**: 707–727.
- Mundlak Y. 1978. On the pooling of time series and cross-section data. *Econometrica* **46**: 69–85.
- Nilsen A, Schiantarelli F. 2003. Zeros and lumps in investment: empirical evidence on irreversibilities and nonconvexities. *The Review of Economics and Statistics* **85**: 1021–1037.
- Orme C. 2001. Two-Step inference in dynamic non-linear panel data models. Mimeo, University of Manchester.
- Pesaran MH. 2004. General diagnostic tests for cross section dependence in panels. Cambridge working paper in economics No. 0435, Faculty of Economics, University of Cambridge.
- Pindyck R. 1988. Irreversible investment, capacity choice, and the value of the firm. *American Economic Review* **78**: 969–985.
- Pindyck R. 1991. Irreversibility, uncertainty, and investment. *Journal of Economic Literature* **29**: 1110–1148.
- Pindyck R. 1993. A note on competitive investment under uncertainty. *American Economic Review* **83**: 273–277.
- Pindyck R, Rotemberg J. 1983. Dynamic factor demands and the effects of energy price shocks. *American Economic Review* **73**: 1066–1079.
- Ross SA, Westerfield RW, Jaffe J. 2005. *Corporate Finance* (7th edn). McGraw-Hill: New York.
- Sakellaris P. 2004. Patterns of plant adjustment. *Journal of Monetary Economics* **51**: 425–450.
- Stewart M. 2007. Inter-related dynamics of unemployment and low-wage employment. *Journal of Applied Econometrics* **22**: 511–531.
- Wooldridge JM. 2005. Simple solutions to the initial conditions problem in dynamic, nonlinear panel data models with unobserved heterogeneity. *Journal of Applied Econometrics* **20**: 39–54.