The changing response of US industry to monetary policy*

Michael Ehrmann[†] Martin Ellison[‡]

August 2002

Abstract

In this paper we investigate whether recent developments in production technology have changed the way US industry responds to changes in expected demand. We construct a simple dynamic model which predicts that improvements in the ability of industry to adjust production quickly may actually lead to a slower response than before. In the past, firms responded quickly to expected changes in demand. Nowadays, it is much easier to change production levels so firms can simply wait until future demand conditions actually materialise before adjusting production. To test the predictions of the model, we analyse the response of US industry to monetary policy assuming that monetary policy changes give a signal about future demand conditions. We do so using a variety of empirical approaches and find evidence of a slower response to monetary policy.

*We thank seminar participants at the Bank of Finland, the European Central Bank and the University of Warwick for helpful comments and suggestions. The views expressed are our own and not necessarily those of the European Central Bank.

[†]European Central Bank, *e-mail address* michael.ehrmann@ecb.int

[‡]University of Warwick, *e-mail address* martin.ellison@warwick.ac.uk

1 Introduction

Recent progress in production technology has prompted many commentators to suggest that US industry has become increasingly flexible over the past years. They argue that advances such as just-in-time production, lean manufacturing and improved inventory management enable firms to adjust their production levels more quickly, easily and at lower cost. As Chairman Alan Greenspan (2001) recognised,

> "New technologies for supply-chain management and flexible manufacturing imply that businesses can perceive imbalances in inventories at an early age – virtually in real time – and can cut production promptly in response to the developing signs of unintended inventory building."

This paper investigates the extent to which these developments may have changed the response of US industry to expected changes in demand. We begin by analysing a simple dynamic model which makes explicit the link between production technology and the response to such expected demand changes. The model predicts that if firms are able to adjust production *more quickly,* then they will choose to respond *more slowly* to future demand conditions. In the past when production was less flexible, firms reacted to signals that anticipated changes in future demand. If demand was expected to pick up, firms would already adjust their production ready to meet this higher demand in the future. Nowadays, since it is much easier to adjust production levels, there is less need for firms to take such preemptive actions. They can just wait until future demand actually materialises before increasing production directly to meet sales.[1]

To test the predictions of the model, we analyse the response of US industry to monetary policy shocks. Monetary policy shocks act with a lag on the

[1] Ahmed *et al.* (2002) refer to this as "good practices" and suggest it may explain some of the reduction in US output volatility observed in recent years. The hypothesis is also supported by Kahn *et al.* (2000).

economy, affecting future demand conditions. For the empirical test, we employ a multivariate model with capacity utilisation, consumer price inflation, the federal funds rate and commodity prices. Capacity utilisation is preferred over industrial production since it more clearly reflects the adjustments that firms make *within* their production processes. Since our measures of capacity utilisation are derived directly from survey data we also avoid detrending issues.

To capture potential changes in the response we use three different approaches. In the first approach, we estimate the response pre and post 1984. The break point of 1984 was chosen because of clear evidence of structural change in the US economy at that time. McConnell and Perez Quiros (2000) report a marked decline in output volatility and Sensier and van Dijk (2001) extend this analysis to several US macroeconomic time series, confirming that the volatility of several real variables has dropped considerably in the early 1980's. In the second approach we search for more gradual evolutionary changes rather than a one-off structural break. We employ rolling-regression techniques to identify any smooth changes in the nature of the response. In contrast, the third approach is based on a Markov-switching model, which allows for potentially repeated structural changes.

The remainder of this paper is organised as follows. Section 2 describes the simple dynamic model and its key predictions. In Section 3 we present some descriptive evidence using disaggregated industry data to show that the response of industry to actual changes in demand has been changing. Section 4 puts the dynamic model to empirical testing. It describes how the model relates to our empirical specification and identification scheme, and provides the results with the three different approaches. A final section concludes.

2 A simple dynamic model

To illustrate the link between production technology and the response of US industry to expected changes in demand we construct a simple two-period model. Industry firms operate in a perfectly competitive market, choosing the amount of labour N_t to employ each period and producing according to the production function $f(N_t)$. Labour is supplied perfectly elastically at wage w in each period and we abstract from the role of inventories by assuming that the finished good is non-storable.[2]

Production technology is introduced by assuming that firms face a cost of changing the level of production between the two periods, which is modelled as $\frac{(N_{t+1}-N_t)^{\gamma+1}}{\gamma+1}$. The convexity of this adjustment cost depends on the production technology parameter $\gamma \geq 0$. A decrease in γ implies less convexity and so more flexible production. $\gamma = 0$ and $\gamma = 1$ describe the cases of linear and quadratic adjustment costs respectively.

An expected change to demand is assumed to have no impact on the market price p_t in the first period but does affect the market price p_{t+1} in the second period. In the case of increasing demand, this implies $p_{t+1} > p_t$. There is no further change to expected demand in period two (i.e., no further change to demand in period three). We assume that the prices p_t and p_{t+1} are known with certainty by the firm before any decisions are made. The profit maximisation problem of the representative firm is given in (1).[3]

$$\max_{N_t, N_{t+1}} \left\{ p_t f(N_t) - wN_t + p_{t+1} f(N_{t+1}) - wN_{t+1} - \frac{(N_{t+1} - N_t)^{\gamma+1}}{\gamma + 1} \right\} \tag{1}$$

We solve the maximisation problem of the representative firm by backward induction. When the firm reaches the second period it treats the quan-

[2] Allowing for storable goods would merely strengthen the predictions of the model.

[3] We abstract from including a discount factor in the final three terms in (1). If we were to include it, it would not appear in the log-linearised first order conditions (4) and (6). Although the steady-state around which we log-linearise would be affected, the behaviour of deviations from steady state would not change. The qualitative conclusions of the model are therefore not affected by this simplification.

tity of labour employed in the previous period, N_t, as predetermined and the profit maximisation problem reduces to (2).

$$\max_{N_{t+1}} \left\{ p_{t+1} f(N_{t+1}) - wN_{t+1} - \frac{(N_{t+1} - N_t)^{\gamma+1}}{\gamma + 1} \right\} \tag{2}$$

Differentiating with respect to N_{t+1} gives the second-period first order condition (3).

$$p_{t+1} f'(N_{t+1}) - w - (N_{t+1} - N_t)^{\gamma} = 0 \tag{3}$$

We log-linearise the first order condition (3) around the symmetric solution of the model in which prices are the same in both periods. In the symmetric solution, $p^* = p_t = p_{t+1}$, and the firm chooses the same amount of labour to employ in each period, $N^* = N_t = N_{t+1}$. The log-linearised first order condition (4) describes the optimal choice of labour employed in the second period. The hat notation indicates percentage deviation from the symmetric solution and the parameter θ is defined by $\theta = -f''(N^*)/f'(N^*) > 0$.

$$\hat{N}_{t+1} = \frac{1}{\theta + \gamma}(\hat{p}_{t+1} + \gamma \hat{N}_t) \tag{4}$$

Equation (4) shows that a higher price in the second period increases the quantity of labour employed in the second period. Similarly, the presence of adjustment costs means that if a high quantity of labour was employed in the first period then labour employed in the second period is also higher. We now proceed backwards to solve for the optimal choice of labour employed in the first period. Differentiating the initial maximisation problem (1) with respect to N_t gives the first-period first order condition(5).

$$p_t f'(N_t) - w + p_{t+1} f'(N_{t+1}) \frac{\partial N_{t+1}}{\partial N_t} - w \frac{\partial N_{t+1}}{\partial N_t} - (N_{t+1} - N_t)^{\gamma} \left(\frac{\partial N_{t+1}}{\partial N_t} - 1\right) = 0 \tag{5}$$

To obtain an explicit solution to the first order condition we use the optimal choice of labour employed in the second period (4) to internalise the effects of actions in period one on outcomes in period two. After substituting out for N_t and $\partial N_{t+1}/\partial N_t$ and log-linearising, the optimal choice of labour employed in the first period is given by equation (6).

$$\hat{N}_t = \frac{\theta+\gamma}{\theta(\theta+3\gamma)}\hat{p}_t + \frac{2\gamma}{\theta(\theta+3\gamma)}\hat{p}_{t+1} \tag{6}$$

We are now in a position to analyse the response of US industry to a change in expected demand. Increased future demand has no effect on prices in period one but raises prices in period two. The expected future price increase already has an effect on production and labour employed in the first period, as shown by equation (7).

$$\frac{\partial \hat{N}_t}{\partial \hat{p}_{t+1}} = \frac{2\gamma}{\theta(\theta+3\gamma)} > 0 \tag{7}$$

This preemptive behaviour is driven by the desire of the firm to avoid making large changes in the level of its production between the two periods. The firm increases production in period one in anticipation of high demand and production in period two. The extent to which this occurs depends on the current state of production technology, as measured by the parameter γ. As production technology becomes more flexible, γ falls and the degree of preemptive actions is reduced, as shown by equation (8). For linear adjustment costs, $\gamma = 0$ and the effect disappears completely.

$$\frac{\partial \left[\frac{\partial \hat{N}_t}{\partial \hat{p}_{t+1}}\right]}{\partial \gamma} = \frac{2}{(\theta+3\gamma)^2} > 0 \tag{8}$$

The model therefore predicts that improvements in production technology allow firms to delay acting until the expected change in demand realises. There is less need to anticipate future demand conditions so we would expect to see a slower response from US industry in recent times.

3 Descriptive evidence from disaggregated data

Before putting the model to empirical testing by analysing the response of US industry to monetary policy, this section provides some preliminary, descriptive evidence using disaggregated data on the industry level. The US Census

Bureau's Manufacturer's Shipments, Inventories and Orders (M3) survey provides monthly data on the US manufacturing sector. These are available for several industry categories as defined in the 1997 North American Industry Classification System (NAICS). With data on new orders and unfilled orders, a first, preliminary, test of potential changes in the production structure of US manufacturing can be performed.

If the production process has become more flexible, we would expect that a shock to new orders can be processed more quickly. Hence, less of these orders should be moved to the stock of unfilled orders in the subsequent months, and for a shorter period of time. This means that we test whether technological process allows for a quicker response of industry to changes in *actual* demand. This test is somewhat different from the core question of this paper, namely how industry responds to a change in *expected* demand. However, it can provide a first impression whether technological changes might indeed be at work. Eventually, the test on changes in *expected* demand is preferable, though, because it is likely to involve longer lags (to a change in *actual* demand, we would always expect a rapid adjustment of production, which should be less sensitive to changes in adjustment costs), and thus to be more informative.

The measures of production lags are obtained by running separate bivariate VARs for each industry category available, with new orders and unfilled orders (both deflated by the PPI). From these, we calculate impulse responses to a standardised shock to new orders (obtained from a Choleski decomposition with new orders odered first, i.e. shocks to new orders affect unfilled orders instantaneously, whereas shocks to unfilled orders do not have a contemporaneous effect on new orders).

46 series of the M3 survey data are available for both variables. In order to gain an assessment of the evolution within the cross-section, we compare three indicators derived from these impulse responses: the initial response of unfilled orders, its maximum response, and the time to maximum. The comparison across time will be performed by estimating these measures for

two subsamples: 1960-1983, and 1984-2000.[4] Doing so, we follow the evidence of significant changes in the structure of the US economy in 1984 documented by McConnell and Perez Quiros (2000) and Sensier and van Dijk (2001). At this time there was a marked reduction in the volatility of many real variables in the US economy.

Table 1 reports the results. The initial response to a standardised shock to new orders changes across industries, although only little and against what we would expect: there seems to be a higher fraction of new orders that get moved into unfilled orders: the average initial response increases from 0.22 to 0.29. We can find a decrease in the initial response, as we would have expected under more flexible technology, only for 12 of the 46 series.

On the contrary, the two other measures support the notion of a more flexible production: the average maximum response has decreased substantially, from 0.63 to 0.41. Furthermore, the lag to the maximum response has decreased on average by 3.5 months. A smaller maximum response is observed in 39 of the 46 series, a shorter lag in 29. Even stronger evidence in favour of more flexible production is obtained when looking at the number of series for which there is either a decrease in the maximum response or a shorter lag. This is the case for 42 of the 46 series. There is thus some cross-sectional evidence that new orders are passed through to shipments more quickly, and that therefore unfilled orders react less to shocks to new orders, which is consistent with the occurence of technological changes.

	mean 1960-1983	mean 1984-2000	no. of decreases
initial response	0.22	0.29	12
maximum response	0.63	0.41	39
time to maximum (months)	13	9.5	29

Table 1: Average response of unfilled orders to a shock to new orders across 46 industry series

[4] 8 of the 46 series are available from 1970 only.

4 The response of US industry to monetary policy

4.1 Baseline empirical specification

To put the dynamic model of section 2 to empirical testing, we need to analyse the response of industry to a change in *expected* demand. We will do so by analysing the response of industry to monetary policy. There is ample evidence that monetary policy shocks affect demand conditions, and do so with a lag.[5] Although the individual firms are unlikely to directly observe monetary policy shocks and thus cannot be expected to react to the shocks as such, we assume that in the course of the transmission of monetary policy to demand, the firm will receive signals that future demand is changing, e.g., through changes in leading indicators, business sentiments, consumer confidence, etc.

The model predicts a slower response to monetary policy in more recent times. To perform this test, we need to assume that the lag in transmission to demand has either remained constant, or has accelerated. As long as the transmission to demand has not slowed down recently, findings of a slower response of US industry support our model predictions. To our knowledge, there is no evidence on recent changes in the transmission lags to demand. However, we would expect that if there have been changes in the speed of transmission, that they have led to a faster response, e.g., because of a potentially increased speed in the pass-through of policy rates to retail lending rates in a deregulated banking sector, or because of the more direct exposure of consumers through relatively large stock holdings.[6]

[5] See, e.g., Leeper et al. (1996), Christiano et al. (1999).

[6] There is some evidence about recent changes in the overall transmission mechanism in the US, without decomposing supply and demand sides. However, this is somewhat conflicting. Boivin and Giannoni (2002) report a reduction in the strength of the effects. Evidence consistent with the hypothesis in our model is provided in Batini and Nelson (2002), who show that the lag from monetary policy actions to inflation in the US has

We estimate a series of vector autoregression models in capacity utilisation, consumer price inflation, the federal funds rate and a world commodity price index.[7] We model capacity utilisation rather than alternative measures such as total industrial production since it better captures the adjustments firms make *within* their production processes. Our focus is on how firms utilise their capacity rather than how capacity itself is determined. By using capacity utilisation we also avoid problems related to the detrending of total industrial production.[8]

Monetary policy shocks are identified in our model by a Choleski decomposition of the system. We order the variables as capacity utilisation, consumer price inflation and the federal funds rate, with commodity prices assumed to be exogenous. This identification scheme implies that neither output nor prices respond immediately to monetary policy shocks. They only react with a lag. In addition, the central bank takes contemporaneous output and inflation into account when setting monetary policy. Our choice of identification scheme follows many previous studies.

Two series for capacity utilisation are available, for total industry and for manufacturing. Total industry capacity utilisation (TCU) is our preferred indicator, since it covers a somewhat broader set of industry than manufacturing capacity utilisation (MCU). Since the series for MCU are available further back in time, we will also use those data to gain a longer-term perspective. Our sample starts in 1960 for MCU and 1970 for TCU. We measure consumer price inflation by log-differences of the consumer price index. All data are seasonally adjusted.

The baseline empirical specification of our model is shown in equation (9). $\mathbf{X}_t$ denotes the vector of endogenous variables (capacity utilisation,

increased for three of the four monetary policy measures they consider.

[7] It has become standard in the literature to include commodity prices in small dimensional VAR models to ameliorate price puzzles.

[8] Boivin and Giannoni (2002) use detrended GDP as a measure of output and so results are sensitive to whether recent developments are assumed to have changed the trend growth in GDP.

consumer price inflation and the federal funds rate) and $\mathbf{Z}_t$ contains the commodity price index and a linear trend to allow for any potential trends in consumer price inflation or the federal funds rate. $\mathbf{u}_t$ is a 3-dimensional vector of fundamental disturbances that are normally distributed and uncorrelated at all leads and lags. Allowing 4 lags in the VAR is sufficient to avoid any autocorrelation in the residuals.

$$\mathbf{X}_t = \mathbf{v} + \mathbf{B}_1\mathbf{X}_{t-1} + \ldots + \mathbf{B}_4\mathbf{X}_{t-4} + \sum_{i=0}^{4}\mathbf{C}_i\mathbf{Z}_{t-i} + \mathbf{A}_0\mathbf{u}_t \tag{9}$$

$$\mathbf{u}_t \sim \mathbf{N}\left(\mathbf{0}; \mathbf{I}_3\right)$$

We will use several variants of this model to test for changes in the response of US industry to monetary policy.

4.2 Evidence with a sample split

In this section, we split the available sample into two subsamples. Again, we follow the evidence of a marked reduction in the volatility of many real variables in the US economy in 1984 documented by McConnell and Perez Quiros (2000) and Sensier and van Dijk (2001). Figure 1 shows that this is also reflected in our measures of capacity utilisation. We therefore divide the sample into 1970:1-1983:12 (1960:1-1983:12 for MCU) and 1984:1-1999:10. The VAR model is estimated independently for each period and structure is imposed through separate Choleski decompositions.

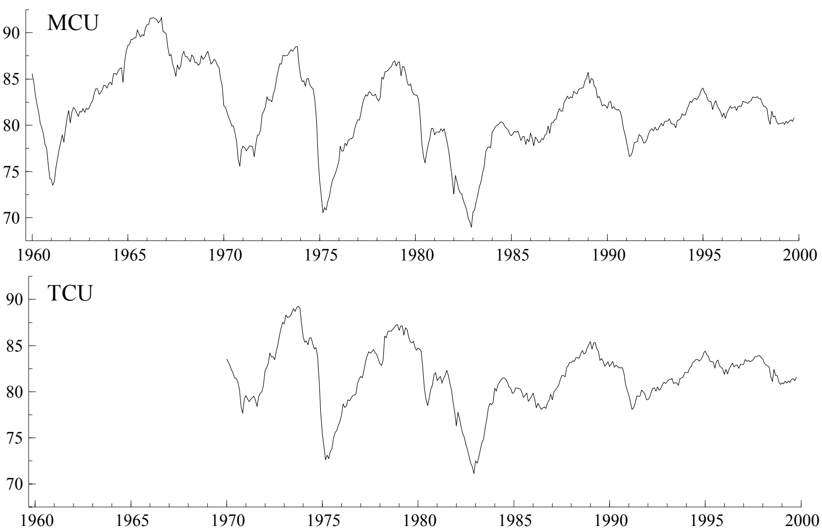

Figure 1: Measures of capacity utilisation.

Figures 2 and 3 show the estimated response to a monetary policy shock using total capacity utilisation (TCU) and manufacturing capacity utilisation (MCU) respectively. The left column of each figure reports the responses in the first sample, whereas those for the second sample are shown in the right column. In all cases the shock to the federal funds rate is normalised to one percentage point to make the magnitude of the responses comparable. We derive 90% error bands for all figures in this paper using standard bootstrapping techniques.

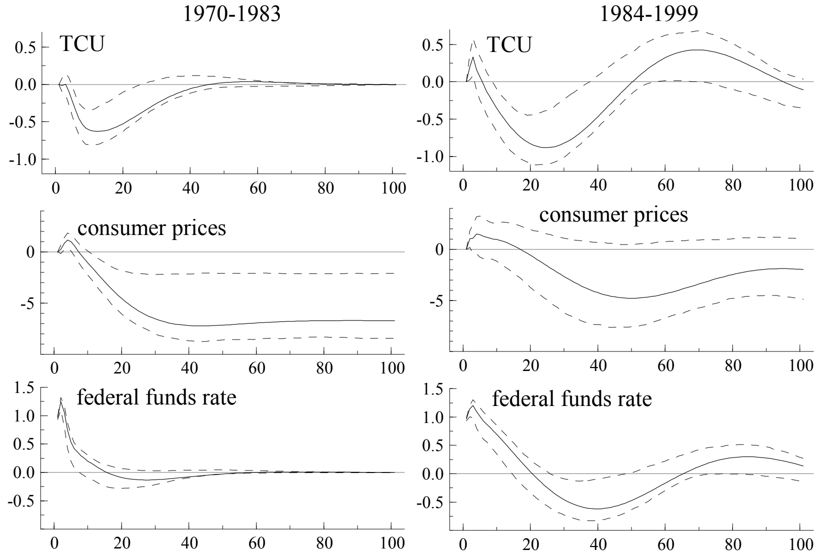

Figure 2: Responses to a monetary policy shock: split sample model with TCU.

The signalling effect of monetary policy for future demand conditions can be seen in the response of consumer prices to a monetary policy shock. Consumer prices decline significantly after a tightening of monetary policy in the first subsample, but only with a lag. The response of consumer prices after 1984 is much less precisely estimated, but seems to follow a very similar pattern. Due to the lack of statistically significant differences in the price response, we conclude that there is no evidence that the change in future demand conditions signalled by a monetary policy shock has changed. We therefore only need to discuss the response of capacity utilisation in the remainder of this paper.

The results suggest a faster response of US industry to monetary policy in the earlier subsample. In the 1970's, capacity utilisation reacted very rapidly to a monetary policy shock, reaching a maximum within 13 months. This is consistent with firms responding in anticipation of the changes in future demand signalled by monetary policy shocks. Following a tightening of monetary policy, firms would quickly cut production ready for the lower demand that is expected in the future. By the 1990's, the response of capacity utilisation to monetary policy shocks was much slower, with a peak only after

26 months. This supports the view that improved production technology now allows firms to wait until the lower demand actually materialises before cutting their production directly in line with demand. The magnitude of the response of capacity utilisation appears smaller in the first subsample, although this is not a very robust finding. Whereas the differences in the timing of the response can be replicated using MCU (see figure 3), the differences in magnitude disappear.

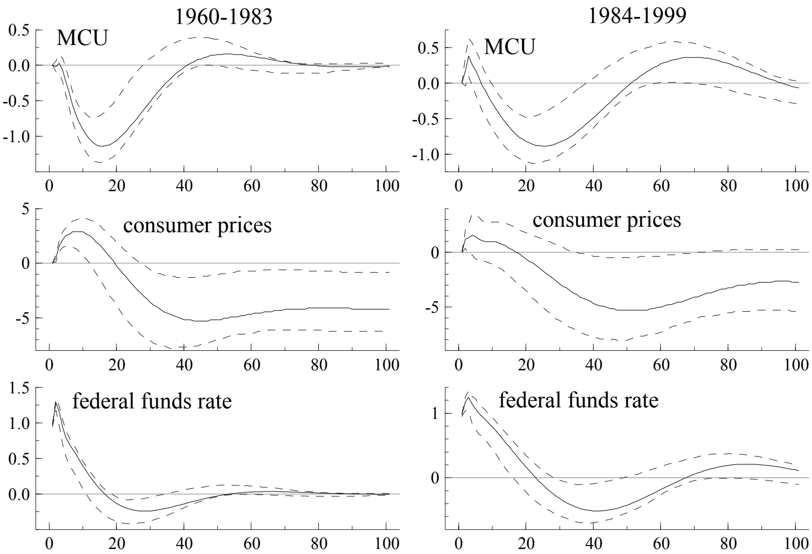

Figure 3: Responses to a monetary policy shock: split sample model with MCU.

One problem remains with this analysis. In figures 2 and 3 it is clear that the federal funds rate is much more persistent in the second subsample than the first. The fact that the federal funds rate is above baseline for longer is potentially an explanation for the slower response of capacity utilisation. To test this hypothesis we ask how capacity utilisation would have reacted in the first subsample if the federal funds rate had been more persistent. In table 2 we report summary statistics for the response of capacity utilisation in each subsample and for a normalisation of the first subsample in which the persistence of the federal funds rate is matched to that in the second subsample.

	Time to maximum	Maximum response
TCU		
1970:1-1983:12	13	-0.6
1984:1-1999:10	26	-0.9
normalised 1970:1-1983:12 (persistence of federal funds rate matched to 1984:1-1999:10)	25	-4.3
MCU		
1960:1-1983:12	15	-1.1
1984:1-1999:10	26	-0.9
normalised 1960:1-1983:12 (persistence of federal funds rate matched to 1984:1-1999:10)	31	-7.6

Table 2: Summary statistics for the response of capacity utilisation to a monetary policy shock

After correcting for the increased persistence of the federal funds rate in the second subsample, the timing of the response of capacity utilisation is more similar over the subsamples. In table 2 the maximum response of total capacity utilisation in the normalised first subsample is observed at 25 months, close to that in the second subsample. However, the normalisation creates a different problem. In the normalised subsamples the maximum response of capacity utilisation is way in excess of that observed in the data. In table 2, the normalised response is at least five times greater than that estimated for either subsample. We conclude that the most likely reason for the slower reaction in the second subsample is not the increased persistence of the federal funds rate. If it was then there must have been a very dramatic change in the magnitude of the response across the two subsamples.

4.3 Evidence from an evolutionary perspective

The previous section assumed a structural break in the US economy in 1984. Although several studies have identified this date as a time of statistically significant change in the US economy, it is likely that developments in production technology occur gradually rather than suddenly. Progress tends to be characterised by regular, small innovations which only have incremental effects, in which case a structural break model may not be the most appropriate. In this section we therefore take an evolutionary perspective in assessing the evidence for changes in the behaviour of US industry.

We begin by estimating the model for a wider range of subsamples. Figure 4 shows the response of manufacturing capacity utilisation (MCU) to monetary policy shocks in six different periods.

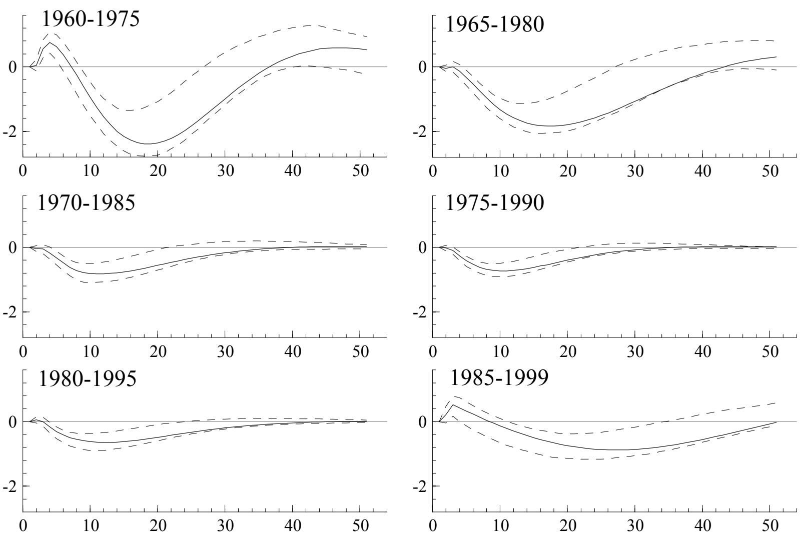

Figure 4: Responses of manufacturing capacity utilisation (MCU) to a monetary policy shock, various subsamples

Figure 4 reveals considerable variation in the timing and magnitude of the response over different subsamples. During the subsample 1960-1975 it took 19 months for the response of capacity utilisation to reach a maximum. This quickened to 10 months over the next three subsamples, before slowing again to 29 months in the last subsample (1985-1999).

To describe the evolutionary nature of the changes in more detail we proceed by applying rolling regression techniques. We estimate the model over a 12 year rolling window and report the results in Figure 5.[9] We show estimates of the magnitude and timing of the response of manufacturing capacity utilisation (MCU) to a monetary policy shock. We also include a measure of the persistence of monetary policy shocks to reflect our concerns in the previous section that increased persistence may cause a slowing in the response of capacity utilisation. We report the half-life of shocks to the federal funds rate, which measures the time that elapses after a shock before the federal funds rate has returned halfway back towards its original level. The time axis in Figure 5 denotes the mid-point of the rolling window. For example, the results of the first rolling window regression, 1960:1-1972:12, are presented at the date 1966:7.

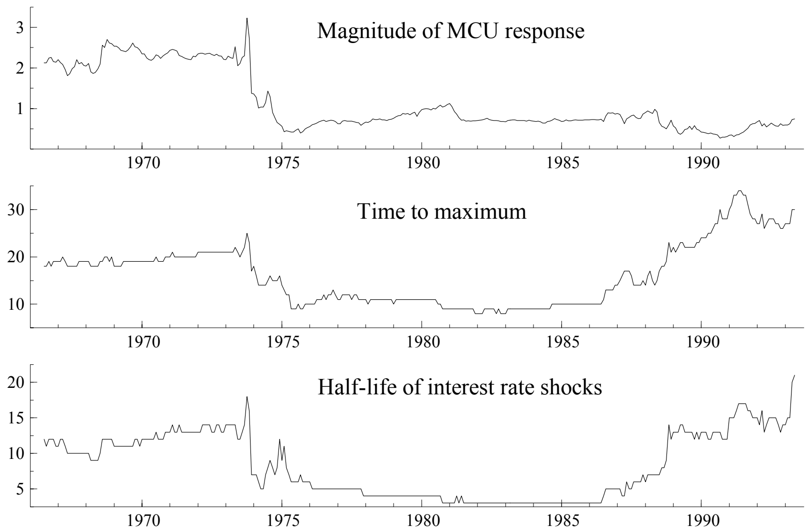

Figure 5: Evolution of the responses to a monetary policy shock

The rolling window regression results confirm considerable evolution in the response of US industry to monetary policy shocks over the sample period. The magnitude of the response jumped in the early 1970's and has

[9]Note that the top panel of Figure 5 shows the magnitude of the maximum (negative) response.

remained remarkably constant since. The time to maximum has gradually increased towards the end of the sample period to reach unprecedented levels in the 1990's. The persistence of monetary policy shocks has also evolved. It dropped in the early 1970's and did not return to the earlier high levels of persistence until the late 1980's.

The jump in the characteristics of the response in the early 1970's can be attributed to changes in monetary policy as the Fed attempted to cope with the highly volatile oil prices of that time. Monetary policy shocks became less persistent, reducing their impact and compressing the period in which they affected the economy. Firms reacted rationally to the new policy by reacting less but more quickly. All of this is observed in the data. In this case, it appears that changes in the response of US industry were caused by a shift in the nature of monetary policy.

The strong evolutionary trends we observe in the response in the 1980's cannot easily be explained by a shift in monetary policy. The persistence of monetary policy does increase again but this time there is no corresponding change in the magnitude of the response. Similarly, the timing of the response is considerably slower than that seen in the 1960's. It is much easier to reconcile this evidence with improvements in production technology than a shift in the nature of monetary policy. Developments in production technology are not expected to have an effect on the magnitude of the response but, as production becomes more flexible, firms are increasingly able to wait for anticipated demand to actually materialise before reacting.

The results in this section point to evolutionary changes in both the 1970's and 1980's. However, the evolution observed in the 1980's is not simply a reversal of that experienced in the 1970's. Mirroring the results of the previous section, we find that the changes in the response of industry cannot sufficiently be explained by corresponding changes in the persistence of monetary policy in the 1980's. More likely, the changes reflect improvements in production technology. The evidence suggests that US industry has not simply reverted to 1960's-style behaviour. This furthermore suggests that the changes we observe are indeed caused by different response patterns of industry, and not simply due to the particularities of the 1970s and early 1980s, where the oil crises were dramatically disrupting the US economy.[10]

4.4 Evidence with a Markov-switching model

This section develops an alternative approach and allows for potentially repeated structural changes in the response of US industry to monetary policy. We cast our baseline model in a Markov-switching framework, in which the system is characterised by several different regimes. The behaviour of such a model is a combination of within-regime dynamics and shifts between the regimes.[11]

To investigate the evidence for Markov-switching effects we use the technique developed by Ehrmann *et al.* (2002). Their procedure enables us to simultaneously estimate the within-regime dynamics and date the regime shifts. The procedure involves first estimating an unrestricted Markov-switching vector autoregression and then identifying monetary policy shocks and impulse responses separately for each regime.

To begin we estimate the unrestricted vector autoregression (10). We assume the existence of two distinct regimes s_t and allow intercepts, autoregressive parameters, variances and covariances to all switch between regimes. No *a priori* restrictions are imposed on the characteristics of the regimes. Information criteria support including two lags in the model.[12]

[10] During 1979-1982, the federal funds rate was not used as monetary policy instrument by the Fed. In a robustness check, we have therefore performed the rolling windows regression with dummies for this period. The results remain qualitatively unchanged.

[11] The Markov-switching model encompasses a model with a single structural break. Estimation of a true structural break model would reveal a two-regime model in which the final regime is absorbing.

[12] To allow for better identification of the monetary policy shock we detrend the federal funds rate in the Markov-switching estimations. This removes the low-frequency inverse U-shaped trend in interest rates, which is likely to have been factored into expectations. We detrend using a linear-quadratic filter, although similar results are obtained with alternative high pass filters. To minimise the effect of this we base our regime inference

$$\mathbf{X}_t = \begin{cases} \mathbf{v}_1 + \mathbf{B}_{11}\mathbf{X}_{t-1} + \mathbf{B}_{21}\mathbf{X}_{t-2} + \sum_{i=0}^{2}\mathbf{C}_{i1}\mathbf{Z}_{t-i} + \mathbf{A}_{01}u_t & \text{if } s_t = 1 \\ \mathbf{v}_2 + \mathbf{B}_{12}\mathbf{X}_{t-1} + \mathbf{B}_{22}\mathbf{X}_{t-2} + \sum_{i=0}^{2}\mathbf{C}_{i2}\mathbf{Z}_{t-i} + \mathbf{A}_{02}u_t & \text{if } s_t = 2 \end{cases}$$

$$\mathbf{u}_t \sim \mathbf{N}\left(\mathbf{0}; \mathbf{I}_3\right) \tag{10}$$

The regime s_t in Markov-switching models is assumed to follow a hidden Markov process with a finite number of states. The probability of being in regime j next period conditional on the current regime i is assumed constant and exogenous. Equation (11) defines the 2×2 conditional transition probabilities for our 2-state model and collects them in an exogenous transition matrix $\mathbf{P}$.

$$\Pr(s_{t+1} = j \,|s_t = i) = \rho_{ij} \quad \text{s.t.} \quad \sum_{j=1}^{2}\rho_{ij} = 1 \text{ for } i = 1, 2$$

$$\mathbf{P} = \begin{bmatrix} \rho_{11} & \rho_{12} \\ \rho_{21} & \rho_{22} \end{bmatrix} \tag{11}$$

The estimate[13] (12) of the transition probability matrix shows that the two estimated regimes are highly persistent. Regimes 1 and 2 have expected durations of 43 and 51 months respectively.

$$\hat{\mathbf{P}} = \begin{bmatrix} 0.98 & 0.02 \\ 0.02 & 0.98 \end{bmatrix} \tag{12}$$

Figure 6 plots total capacity utilisation (TCU) and estimated smooth probabilities of being in regime 1. The smooth probability represents the optimal inference of when any switches in regime occurred. The pattern of

solely on the behaviour of capacity utilisation and inflation.

[13]Estimations were performed using the MSVAR 0.99 package for Ox 2.10.

regime switches closely resembles the sample split chosen for Section 4. A very clear regime switch is precisely dated to early 1984, at which time the volatility of total capacity utilisation diminished dramatically. However, the resemblance to a structural break is not perfect because there have already been two switches back into the "old" regime. The latter coincided with the recession of 1990/1991, suggesting that the "new" regime may not be perfectly absorbing.

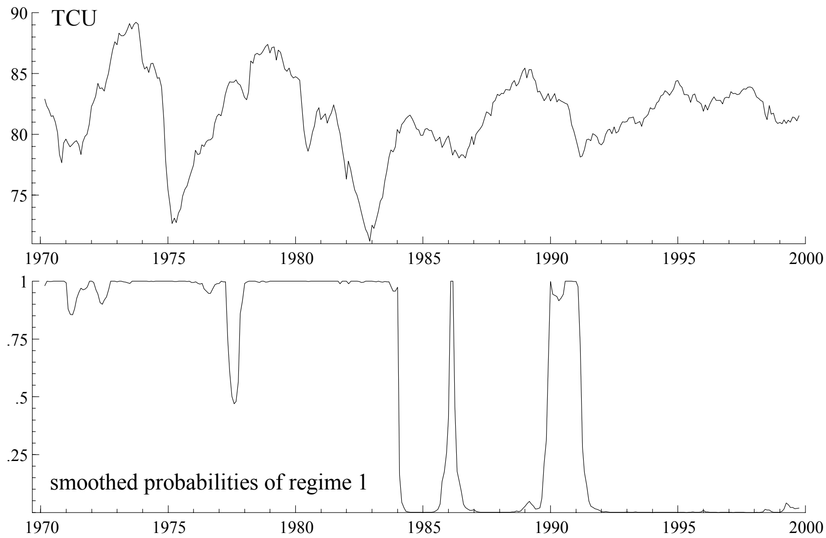

Figure 6: Results from the Markov-switching model.

To see if US industry responds differently to monetary policy within the two identified regimes, we follow the Ehrmann *et al.* (2001) procedure and calculate regime-dependent impulse response functions. These show how capacity utilisation reacts to monetary policy shocks conditional on the regime. They are obtained by imposing separate restrictions on the two regime-dependent variance-covariance matrices in Equation (10). We use a separate Choleski decomposition for each regime.

The regime-dependent impulse response functions are shown in Figure 7. Total capacity utilisation (TCU) shows the familiar slower response in the second regime, which roughly corresponds to the second sample in Section 4. The main advantage of estimating the model in this form is the possibility of testing for differences across the regimes. A t-test of whether the response in regime 1 is as slow as in regime 2 is rejected with a p-value of 0.04.[14]

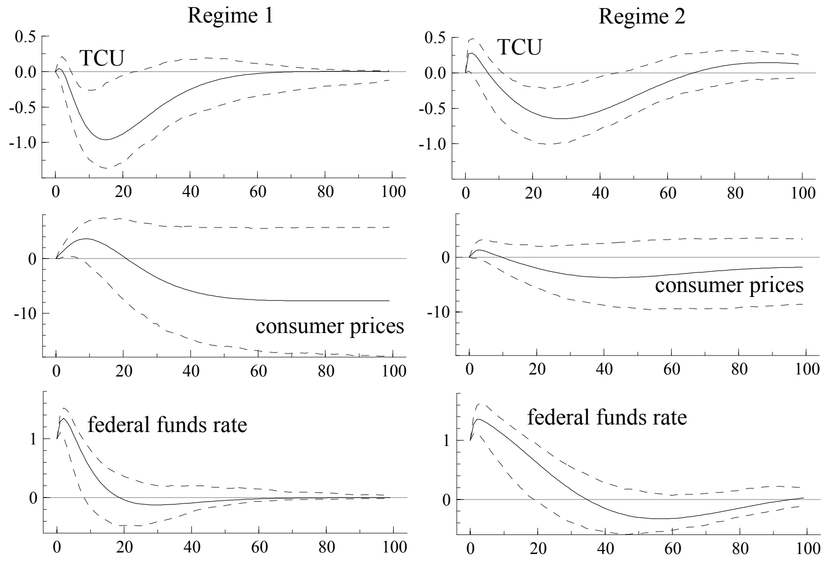

Figure 7: Regime-dependent impulse responses

The results in this section complement those from Sections 4 and 5. The Markov-switching model dates a regime switch very exactly to early 1984 and within-regime dynamics are very similar to those observed in the sample split of Section 4. However, the estimates imply that the system has not necessarily settled in a "new" absorbing regime yet. It has already switched back into the "old" regime twice since 1984.

[14]More precisely, we test whether the magnitude of the response is significantly larger in regime 1 than regime 2. The reported p-value is obtained for period 9, which is the maximum level of significance. The low level of significance arises from the relatively wide error bands around the impulse response functions, a problem encountered in most VAR analyses with macroeconomic data.

5 Conclusions

The results of the empirical investigation support the prediction of the simple dynamic model that recent improvements in production technology imply a slower response of US industry to monetary policy. Using a variety of approaches, we find evidence that firms do respond more slowly to monetary policy shocks now than before. This is consistent with a reduced need for firms to react to the signal of future demand conditions contained in monetary policy. Nowadays, flexible production technology allows firms to simply wait until future demand actually materialises before increasing production directly in line with demand. In other words, the fact that firms can adjust production *more quickly* enables them to choose to respond *more slowly* to monetary policy.

Our results also show that the slower response cannot be attributed to a change in the nature of monetary policy. Whilst changes in the 1970's appear to have been caused by policy, the 1980's changes clearly show a different pattern. There is no support for the argument that the US economy has simply reverted to 1960's-style behaviour.

References

[1] Ahmed, S., Levin, A. and B.A. Wilson (2002). Recent US Macroeconomic Stability: Good Policies, Good Practices, or Good Luck?, Board of Governors of the Federal Reserve System International Finance Discussion Paper No. 2002-730

[2] Batini, N. and E. Nelson (2001). The Lag from Monetary Policy Actions to Inflation: Friedman Revisited. *International Finance* 4(3), 381-400

[3] Boivin, J. and M. Giannoni (2002). Has Monetary Policy Become Less Powerful?, mimeo, Columbia University and Federal Reserve Bank of New York

[4] Christiano, L.J., Eichenbaum, M. and C.L. Evans (1999). Monetary Policy Shocks: What Have We Learned and to What End?, in: Taylor, J.B. and M. Woodford (eds.), *Handbook of Macroeconomics Volume 1A*, Amsterdam: Elsevier Science North Holland, 65-148

[5] Ehrmann, M., Ellison, M. and N. Valla (2002). Regime-Dependent Impulse Response Functions in a Markov-Switching Vector Autoregression Model, *Economics Letters*, forthcoming

[6] Greenspan, A. (2001) Testimony to the Congress, 28th February

[7] Kahn, J., McConnell, M.M. and G. Perez Quiros (2000). The Reduced Volatility of the US Economy: Policy or Progress?, mimeo, Federal Reserve Bank of New York and European Central Bank

[8] Leeper, E.M., Sims, C.A. and T. Zha (1996). What Does Monetary Policy Do?, *Brookings Papers on Economic Activity* 0(2), 1-63

[9] McConnell, M.M. and G. Perez Quiros (2000). Output Fluctuations in the United States: What Has Changed Since the Early 1980s?, *American Economic Review*, 90, 1464-1476

[10] Sensier, M. and D. van Dijk (2001). Short-term Volatility versus Long-term Growth: Evidence in US Macroeconomic Time Series, *Centre for Growth and Business Cycle Research Discussion Paper No. 8,* University of Manchester.