Modelling manufacturing inventories

Working Paper no. 284

John D Tsoukalas

Changes in the stock of firms' inventories are an important component of the business cycle. In fact, discussion about the timing of a recovery following economic recessions often focuses on inventories. But there is no consistent explanation for their behaviour. Most modelling work focuses on the so-called production smoothing model — where the firm maintains a smooth production plan and uses inventories to satisfy unforeseen changes in demand. Moreover, this model has generally only been applied to manufacturers' finished goods inventories.

This paper offers an extended stage-of-fabrication inventory model that considers not only finished goods inventories, but also input inventories — the sum of raw materials and work-in-process inventories. Stylised facts for UK manufacturing reveal that input inventories are empirically more important than finished goods inventories. This is true not only in terms of their size but also in terms of their volatility.

One of the key facts of the UK manufacturing sector is the significant interaction between finished goods and input inventories. The covariance between input and finished goods (or output) inventories can explain over one quarter of the variance in manufacturing inventory investment. This is an important finding because it points to linkages between different aspects of production. More importantly, it implies that finished goods inventories cannot be considered in isolation from input inventories. Intuitively, an optimising firm that decides to draw down finished goods inventories (as often happens following an unexpected demand shock) will typically increase production in the future to correct this imbalance. This correction will affect input inventories as well because the firm has to draw down input inventories in order to increase production.

The paper demonstrates that ignoring input inventories yields misleading results. In particular, the precision and plausibility (relative magnitudes) of the estimated parameters in the joint model differs from those when input inventories are ignored. To estimate the model, a maximum likelihood approach is used that is shown to be superior to the often-used generalised method of moments estimators (GMM). The sizable interaction between input and finished goods inventories yields very precise estimates. One of the key findings of the model is the familiar production-smoothing result. The estimation results suggest that firms satisfy unexpectedly strong demand from finished goods inventories, resulting in the latter falling below companies' desired levels. Given that estimated costs of changing production are large relative to stockout costs (of deviations of inventories from their desired level), this imbalance corrects rather slowly, implying that inventories deviate from target for long periods. On the cost side of the model, when materials become more expensive companies prefer to cut production temporarily: cutting production implies that companies save on the expensive materials. With sales unchanged, this shortfall is satisfied out of finished goods inventories causing them to fall from their desired level. Moreover, despite the presence of input inventories — where fixed costs of ordering may be substantial — the estimated aggregate marginal cost function is a rising function of output, thus implying decreasing returns to scale in manufacturing.

The New Keynesian Phillips Curve in the United States and the euro area: aggregation bias, stability and robustness

Working Paper no. 285

Bergljot Barkbu, Vincenzo Cassino, Aileen Gosselin-Lotz and Laura Piscitelli

The traditional Phillips curve relates current inflation to lagged inflation and a cyclical indicator, such as the output gap or the unemployment rate. This specification has efficiently characterised the pattern of inflation over most of the post-war period in most industrialised economies. Two concerns have however been raised. First, the traditional Phillips curve is subject to the Lucas critique — its coefficients may not be invariant to changes in policy regimes. Second, the traditional Phillips curve explains recent data for the United States and the euro area less well, where inflation has been low despite positive output gaps. In an attempt to deal with the shortcomings of the traditional approach, the New Keynesian Phillips Curve (NKPC) literature uses microfoundations to derive a relationship between inflation, expectations of future inflation and the current value of the cyclical indicator. However, the pure NKPC lacks sufficient inertia to adequately explain the path of actual inflation. As a result, many authors attempt to improve the degree of fit of the NKPC by inserting a lagged inflation term. Such a curve is often referred to as a 'hybrid' NKPC. We estimate this using generalised method of moments techniques for the United States, the euro area and the three largest euro-area economies.

This relationship between output and inflation is of key importance to monetary policy authorities concerned with price stabilisation. In particular, understanding whether the relationship is stable and how it evolves over time is a primary concern. Therefore, it is somewhat surprising that while numerous papers have estimated the NKPCs for these countries there has been relatively little emphasis on testing the stability of their parameters over time. In other words, the robustness of the NKPC to the Lucas critique has not yet been subject to proper statistical testing. We aim to fill this gap by conducting comprehensive stability and structural break analysis, performing rolling and recursive estimation and applying standard tests for structural breaks.

Overall, our estimates of the structural and reduced-form coefficients on the lagged and expected future inflation terms are broadly in line with previous studies for both the United States and the euro area. One notable exception is the discount factor obtained for the euro area, which is lower than that found in most other studies. On the question of stability, rolling and recursive estimations produce stable and plausible estimates for the United States, but unstable parameters for the euro area. The breakpoint test analysis does not reveal any significant shift in any of the coefficients associated with past and expected future inflation and real marginal cost for the United States. For the euro area, on the other hand, there is some tentative evidence of a structural break affecting the coefficients on past and expected future inflation in the late 1980s, possibly related to the German re-unification. In the disaggregated euro-area analysis, rolling estimation produces unstable estimates for Germany, and, albeit to a lower degree, for Italy. The estimates for France appear to be considerably more stable over the period considered. Consistent with these results, the breakpoint test analysis points to instability in the late 1970s for Italy and in the early 1980s for Germany. There is no evidence of structural breaks affecting inflation dynamics in France. These conflicting country-level results could indicate the presence of an aggregation bias in the results obtained with euro-area data, which could explain the implausibly low estimate of the discount factor obtained for the euro area.

There are several implications for monetary policy makers. Overall, our results suggest that policymakers should treat the forecasts generated by Phillips curves with some caution, as the structural parameters underlying the estimated relationships may be unstable over time. For the euro area, in particular, it may be useful to look at the results of individual countries, in addition to the aggregate results. Moreover, policymakers should examine the results of a broad range of estimation methodologies to assess whether the forecasts generated by a Phillips curve model agree with other evidence. This is consistent with the approach currently taken in most major central banks.

Modelling the cross-border use of collateral in payment systems

Working Paper no. 286

Mark J Manning and Matthew Willison

Over the past decade, there has been a decisive shift towards real-time gross settlement (RTGS) of high-value payments, typically across accounts at a central bank. Settlements in such a system can only be completed if the paying bank has sufficient funds in its account; ie if it has adequate liquidity. Hence, the focus has shifted away from credit risk and towards liquidity risk.

Central banks typically address this risk by making intraday liquidity available to settlement banks on favourable terms. A commonly adopted policy is to provide intraday credit on an unlimited, free, but fully collateralised basis. But given the potential for high collateral costs to encourage payment delays and risk gridlock in the system, it will be optimal for central banks to accept a wide variety of assets as collateral, thereby allowing the efficient management of commercial banks' collateral portfolios.

Restricted eligible collateral lists are a particular issue for banks operating in multiple countries and hence facing settlement obligations in a number of payment systems. If, in each country, the central bank accepts domestic securities only, a bank must hold sufficient (costly) collateral assets to meet its expected liquidity needs wherever it is active. Once liquidity demands have been realised, a bank may find itself with a shortage of collateral in one market and abundant collateral in another. This collateral will then lay idle while the bank obtains additional assets in the country in which it experiences a shortfall. Such a mismatch is inefficient. It also contributes to liquidity risk in the payment system because there could be some disruption to the bank's payments activity while it enters the market to acquire the necessary eligible assets.

One potential policy response is then to broaden the eligible collateral list to include foreign assets; ie to allow cross-border use of collateral. But any change in policy on collateral eligibility may result in a change in market participants' incentives, so the implications of cross-border use of collateral for liquidity risk in payment systems are best analysed in the context of a model that captures full optimising behaviour on the part of commercial banks. In this regard, we develop a stylised two-country, two-bank model in which risk-neutral banks minimise expected costs with respect to their collateral choice in each country. Banks are active in both countries' payment systems and make collateral choices under uncertainty as to both the size and the location of their liquidity needs. In our baseline model, we assume that each bank only realises a liquidity demand in one country at any one time; in other words, banks' liquidity demands are negatively correlated across countries. Using our baseline model, we compare outcomes for liquidity risk for cases in which: (i) there is no cross-border use of collateral; and (ii) both central banks permit cross-border use of collateral. A number of key results emerge.

First we show that, when both countries permit cross-border use of collateral (which we refer to as the symmetric cross-border use of collateral), banks will concentrate their holdings in the country with the lowest collateral costs and may, with sufficiently high costs of experiencing a shortfall (relative to start-of-day collateral costs), reduce collateral holdings in each country. Importantly, even with a decline in total collateral holdings, we find that liquidity risk, as measured by expected collateral shortfalls, will fall in both countries. This reflects the fact that it will always be optimal for a bank to hold a larger amount of collateral across two connected countries than in a single unconnected location. Hence, there will always be a larger pool from which to draw to meet a liquidity need in a single country.

We make a number of extensions to the baseline model to relax some of its simplifying assumptions. First, in the absence of co-ordinated policy, it may be that only one central bank permits cross-border use of collateral. In this case of asymmetric cross-border use of collateral, we show that banks' collateral choices will be driven by two potentially offsetting factors. On the one hand, banks will shift collateral holdings towards the collateral that is eligible in both countries. On the other, banks will still be inclined to accumulate larger holdings of the cheaper collateral. When the cheaper collateral can be used across borders, these two factors are mutually reinforcing and the outcome will be the same as if there were symmetric cross-border use of collateral. When the collateral eligible in both countries is only slightly more expensive, banks will still hold only this collateral, but slightly less will be held overall than in the symmetric case. Again, liquidity risk will decline in both countries. Finally, when the collateral eligible in both countries is significantly more expensive, collateral will be held in both countries and the expected shortfall in the country accepting foreign collateral will be the same as in the case with no cross-border use.

Another extension allows some probability that a bank experiences a liquidity need in both countries simultaneously. In our model, banks adjust collateral holdings to take account of this possibility. But as there remains a chance that banks could experience a liquidity need in just one country, it may, under certain conditions, still be optimal to reduce total collateral holdings relative to the case with no cross-border use of collateral. Such a reduction would imply higher expected shortfalls in at least one country when a bank faces simultaneous liquidity needs, compared to the case with no cross-border use of collateral. The size of the respective shortfalls experienced in each country will depend on how the available collateral is ultimately allocated between countries.

Finally, we also consider an extension in which central banks have the option of accepting collateral in stressed situations only. Under such a regime, and with a sufficiently low probability that the emergency facility will be triggered, banks' reductions in collateral holdings may be more muted than if cross-border use of collateral were allowed routinely. As a result, should a stressed situation arise in one country, banks may have a larger pool of collateral to draw upon than they would have in the case of routine cross-border use of collateral. Expected shortfalls would, in such a case, be lower. If central banks place a higher weight on liquidity risk mitigation in times of stress, and recognise that it may be more difficult to access additional collateral during a crisis, such a policy may be attractive.

Assessing central counterparty margin coverage on futures contracts using GARCH models

Working Paper no. 287

Raymond Knott and Marco Polenghi

Financial markets benefit from a sound and stable environment. For this reason, central banks follow developments in the financial markets and their associated infrastructure closely. Central counterparties (CCPs), which help to protect market participants against counterparty default have become an increasingly important part of this infrastructure, as they have expanded into new markets, and undergone both vertical and horizontal consolidation. Many regulators and central banks concerned with systemic stability have, as a result, increased their focus on how CCPs manage their risks. The present paper examines some key aspects of risk control in these institutions.

In order to reduce counterparty credit risk for market participants, exchange-traded and some non exchange traded derivatives contracts are guaranteed against counterparty failure by a CCP. In providing this service, however, CCPs themselves become exposed to the risk of counterparty default. To protect themselves, they have developed a system of collateralisation, or margining, by which members of the CCP are required to place a sum of initial margin in a CCP account when they register positions. The initial margin is designed to provide protection against potential changes in the market value of a member's positions over a time horizon of one or more days.

CCPs typically select the appropriate level of initial margin by inspecting the historical distribution of price movements, focusing particularly on recent price changes. After being set, however, initial margin levels will often remain unchanged for some time. During this period, the coverage they provide can change substantially, varying according to market conditions.

This paper describes a model that can be used to assess the coverage provided by initial margins. Previous studies have largely concentrated on assessing long-run average coverage levels. The present study shows how coverage can be assessed on a day-by-day basis. In order to measure variations in coverage, we use a model of returns which assumes a time-varying volatility (a so-called GARCH process). This is used to model the returns of two heavily traded derivatives contracts, the Brent oil and the FTSE 100 futures contracts. Different variants of the GARCH process are estimated, which assume that the changes in volatility are distributed according to either Student t, extreme value or historical distributions. To select the best-performing variant, a backtesting procedure is applied in which the models' forecasts of returns are compared against actual outcomes. Overall, across all coverage levels, we find that the Student t and historical distribution variants offer the best fit to actual returns.

The modelling approach described allows us to estimate the probability that initial margin will be used up, as that probability changes from day to day. Although we find that the average probability of exhaustion for initial margins is low, we note that the probability can increase in volatile markets. That suggests a need to reset the initial margin more often in such circumstances, as most CCPs do in practice.

Regulators and central banks are also interested in understanding what the impact on the market, or the CCP's post-margin resources, would be of more extreme price moves. In particular, they would like to know what the additional liquidity demands might be if a margin-exhausting price change were to occur. To illustrate how the model could be applied to this question, we calculate the conditional expected loss for the FTSE 100 and Brent futures contracts, ie the expected loss the CCP would suffer when the initial margin is used up completely. This is then used to generate an estimate of the additional liquidity demands that each market would experience. We find that, if the initial margin for the Brent contract were to be exceeded, it would require a greater percentage increase in margin, compared to the FTSE 100, largely due to the higher tail-thickness of the Brent return distribution. We note that, for the sample window chosen, the additional liquidity demands are relatively modest compared to typical intraday margin calling mechanisms.

The price puzzle: fact or artefact?

Working Paper no. 288

Efrem Castelnuovo and Paolo Surico

The initial positive response of prices to a contractionary monetary policy shock is a stylised fact of most empirical studies measuring the effects of monetary policy on the aggregate economy. This behaviour is often referred to as 'puzzling' because macroeconomic models either cannot explain it theoretically (eg a standard sticky-price model) or, even when capable of explaining it in principle, they do not produce a positive price response empirically (eg models of the cost channel transmission of monetary policy).

The presence of a price puzzle is important because it casts serious doubts on the possibility of correctly identifying a monetary policy shock. If the central bank monitors and responds to a larger information set than that of the econometrician, what may be referred to as a policy shock by the latter is actually a combination of a genuine policy shock and some endogenous policy reactions. The result of this omission is that a policy tightening in anticipation of future inflation could be wrongly interpreted by the econometrician as a policy shock, delivering a spurious correlation between a tightening of policy and a rise in inflation: the price puzzle.

In a speech as Fed Governor, Bernanke offered a new interpretation of the mis-identification of structural shocks. He noted that '*[...] changes in inflation expectations, which are ultimately the product of the monetary policy regime, can also be confused with truly exogenous shocks in conventional econometric analysis. [...] insufficiently anchored inflation expectations have led to periodic 'inflation scares', in which inflation expectations have risen in an apparently autonomous manner. Increases in inflation expectations have the flavor of adverse aggregate supply shocks in that they tend to increase the volatility of both inflation and output, in a combination that depends on how strongly the monetary policy makers act to offset these changes in expectations*'.

This paper offers a theoretically and empirically consistent explanation for the price puzzle using a micro-founded New Keynesian model and structural vector autoregressions (VARs). A major contribution is to show that the price puzzle has been historically a feature of specific monetary policy regimes. These regimes are the period prior to the appointment of Paul Volcker as Fed Chairman in August 1979 for the United States and the period prior to the introduction of the inflation-targeting regime in 1992 for the United Kingdom. This result is robust to using two different identification strategies, as well as augmenting the VARs with unit labour costs and a commodity price index. Moreover, the subsample evidence on the price puzzle is found to be independent from using real GDP, detrended output, the output gap or output growth as a measure of real activity.

A wide number of contributions to the empirical literature on monetary policy rules finds that a shift in the conduct of monetary policy occurred at the end of 1979 in the United States and at the end of 1992 in the United Kingdom. We therefore investigate the link between these shifts in the conduct of policy and our results about the price puzzle, modelling monetary policy using simple mechanical rules. It should be emphasised that this paper does not suggest that monetary policy in the United States and United Kingdom was in fact conducted using these mechanic policy rules: rather, they are simply a useful empirical representation of monetary policy.

A sticky price model is calibrated to the magnitude of the historical shift in the conduct of US monetary policy. This model is used to simulate artificial data and then the structural VARs are estimated on the artificial data. A main finding is that only when the central bank does not raise the interest rate sufficiently in response to inflation (and thus inflation expectations are not well anchored in the theoretical model) do the structural VARs estimated on the artificial data generate a sizable price puzzle. In contrast, the theoretical model is not capable of generating an initial positive response of the price level to a monetary policy shock, even when the nominal interest rate responds less than fully to inflation.

Our results suggest that the price puzzle is in fact an artificial result that arises from expected inflation being omitted from the VAR. Expected inflation is remarkably more persistent when expectations are not fully stabilised by the monetary policy. Furthermore, such omitted variable bias is found to account for the apparently puzzling response of inflation to a policy shock observed on actual data, consistently with Bernanke's argument.

Our results suggest that when the policy framework does not mandate sufficient response to inflation, the behaviour of private sector expectations gives rise to perverse dynamics, like inflation persistence and the price puzzle, that are not necessarily intrinsic characteristics of the economy.