

CENTRE FOR APPLIED MACROECONOMIC ANALYSIS

The Australian National University



CAMA Working Paper Series

April, 2010

**INVENTORIES, INFLATION DYNAMICS AND THE NEW KEYNESIAN
PHILLIPS CURVE**

Thomas A. Lubik

Federal Reserve Bank of Richmond

Wing Leong Teo

National Taiwan University

CAMA Working Paper 13/2010

<http://cama.anu.edu.au>

Inventories, Inflation Dynamics and the New Keynesian Phillips Curve*

Thomas A. Lubik
Federal Reserve Bank of Richmond[†]

Wing Leong Teo
National Taiwan University[‡]

January 2010

Abstract

We introduce inventories into an otherwise standard New Keynesian model and study the implications for inflation dynamics. Inventory holdings are motivated as a means to generate sales for demand-constrained firms. We derive various representations of the New Keynesian Phillips curve with inventories and show that one of these specifications is observationally equivalent to the standard model with respect to the behavior of inflation when the model's cross-equation restrictions are imposed. However, the driving variable in the New Keynesian Phillips curve - real marginal cost - is unobservable and has to be proxied by, for instance, unit labor costs. An alternative approach is to impute marginal cost by using the model's optimality conditions. We show that the stock-sales ratio is linked to marginal cost. We also estimate these various specifications of the New Keynesian Phillips curve using GMM. We find that predictive power of the inventory-specification at best approaches that of the standard model, but does not improve upon it. We conclude that inventories do not play a role in explaining inflation dynamics within our New Keynesian Phillips curve framework.

JEL CLASSIFICATION: E24; E32; J64.

KEY WORDS: Phillips curve; GMM; marginal costs; inventories.

*The authors wish to thank seminar participants at the University of Sydney for helpful comments. Most of this research was conducted while Lubik was the Geoffrey Harcourt Visiting Professor of Economics at the University of Adelaide. The views expressed in this paper are those of the authors and should not be interpreted as those of the Federal Reserve Bank of Richmond or the Federal Reserve System.

[†]Research Department, P.O. Box 27622, Richmond, VA 23261. Email: thomas.lubik@rich.frb.org.

[‡]Department of Economics, National Taiwan University, 21 Hsu Chow Road, Taipei 100, Taiwan. Phone: (886) 2351-9641-ext 449, Fax: (886) 2351-1826, Email: wlteo@ntu.edu.tw.

1 Introduction

The behavior of inflation and the sources of its fluctuations have been the subject of intense research in the last couple of years. The development of the New Keynesian monetary model has provided empirical macroeconomists with an internally consistent framework to study the linkages between inflation dynamics and economic activity variables. These are captured by the New Keynesian Phillips curve (NKPC) which postulates a relationship between inflation, expected inflation and real marginal cost. The NKPC is derived from optimal price-setting behavior of monopolistically competitive firms, and thereby provides more structure for interpreting the data than earlier efforts based on accelerationist Phillips curves.

In their seminal empirical treatment of the NKPC, Galí and Gertler (1999) demonstrate that it describes inflation dynamics to a reasonable degree. They also highlight two problematic issues. First, marginal cost is unobservable and has to be proxied or related via economic theory to observable variables. The second issue is that, in general, marginal cost is less volatile and persistent than inflation. Galí and Gertler (1999) resolve these issues by using the labor share as a proxy instead of unit labor cost (both of which can be constructed from the production function), and by adding indexation in price setting, which introduces a lagged inflation term in the NKPC. While the thus modified NKPC describes inflation dynamics reasonably well, follow-up research has uncovered various problems.¹

The literature has addressed these shortcomings by branching out in two directions. Many papers approach inflation dynamics from a system perspective,² in which marginal cost is implicitly constructed through the restrictions imposed by the rest of the model. The second direction introduces additional features into the underlying model to modify the behavior of marginal cost. This delivers a theoretical rationale for adding additional driving forces for inflation to empirical specifications of the NKPC. A recent example of this approach is the New Keynesian model with search and matching frictions in the labor market developed by Krause et al. (2008a, 2008b). The driving forces of inflation in the thus modified NKPC are a host of labor market variables besides marginal cost. Our paper is in the spirit of this approach.

Specifically, we modify the standard New Keynesian model by introducing inventory behavior on part of the firms. We motivate inventory holdings as a way for firms to generate sales, as in Bils and Kahn (2000). When potential buyers approach a firm, it need not

¹Nason and Smith (2008) provide a concise summary and exposition.

²Schorfheide (2008) discusses this approach in some detail.

change current production, but can instead satisfy demand by drawing from its inventory. This decouples current production from sales, and introduces an intertemporal aspect in which firms jointly decide on the level of profit-maximizing prices and the desired level of inventory holdings. Production towards inventory thus provides opportunities for increased sales in the future. Since monopolistically competitive firms have price-setting power they thus face an intertemporal production and an intertemporal pricing trade-off which feeds into aggregate inflation dynamics.

We derive a version of the NKPC in this framework with inventories, and show that the modified equation contains additional activity variables besides marginal cost and that this also affects the coefficients. However, the structural relationship derived from the model allows several representations of the NKPC, one of which is observationally equivalent to the standard NKPC. By means of a data-driven calibration analysis we construct implied series for marginal cost and the driving process in the NKPC and contrast these with typical proxies used in the literature. We show that *a priori* the inclusion of inventories can improve the predictive power of the marginal cost series, but only if the cross-equation restrictions from the rest of the model are ignored.

In the next step, we use the firm's optimality condition for inventories to construct an implied marginal cost series from observable variables, as mandated by theory. We then use the constructed marginal cost series as an explanatory variable in the NKPC and estimate it using a generalized methods of moments approach. We find that introducing inventories does not affect inflation dynamics as seen through the NKPC.

Finally, we jointly estimate the NKPC and the optimal inventory condition with GMM. The structural parameters of the NKPC are within the bounds of previously established results, while the inventory parameters either take on implausible values or are not identified. These findings are robust for various specification changes. We consequently argue in this paper that inventory holdings of the kind we discuss do not hold promise for explaining inflation dynamics. However, we also discuss the role that a limited information approach plays in generating these findings.

There has been a recent surge of papers studying the behavior of inventories in monetary business cycle models. The papers closest to ours are Hornstein (2005), Jung and Yun (2006) and Boileau and Letendre (2008). The former two combine Calvo-type price setting in a monopolistically competitive environment with the approach to inventories as introduced by Bilal and Kahn (2000). The use of the Calvo-approach to modeling nominal rigidity allows them to discuss the importance of strategic complementarities in price

setting at the cost of a less transparent reduced-form specification of the NKPC. Boileau and Letendre (2008) compare various approaches to introducing inventories in a sticky-price model. Their empirical analysis is purely calibration-based and they do not focus on the specific implications for the NKPC. Moreover, they do not allow for backward indexation in price setting, and the potential for depreciation of the inventory.

The rest of the paper is organized as follows. In the next section, we present the New Keynesian model from which we derive the NKPC, and discuss and motivate how we introduce inventories. In section 3, we show how the NKPC can be derived from the model's optimality conditions. Section 4 discusses the data we use in the empirical application and presents some stylized facts for the relationship between inflation, marginal cost and inventories. We also calibrate the model and use proxies for marginal cost to compute implied series for the marginal cost term and the overall driving process in the NKPC. The core part of the paper is section 5, where we take a more structural approach. We first back out the marginal cost series from the model's optimality conditions, and then estimate the inventory optimality condition jointly with the NKPC. Section 6 contains a discussion of the limits of a reduced-form and partial equilibrium approach in the context of our New Keynesian inventory model. Section 7 concludes.

2 The Model

We introduce inventories in the manner suggested in Bils and Kahn (2000). Inventories are assumed to help facilitate sales as firms can rely on the stock of previously produced goods when demand rises. This can be motivated by a firm's desire to avoid stock-outs, in which case the firm would face marginal production cost or the loss of marginal revenue. Moreover, a larger stock can facilitate matching with potential buyers and thus increase sales. To motivate the existence of sticky price, we assume that firms are monopolistically competitive and set their optimal price along a downward-sloping demand curve. We capture these elements by modeling aggregate sales s_t as a Dixit-Stiglitz aggregator of firm-specific sales s_{it} and the stock of goods available for sales a_{it} :

$$s_t = \left(\int_0^1 \left(\frac{a_{it}}{a_t} \right)^{\frac{\theta}{\epsilon}} (s_{it})^{\frac{\epsilon-1}{\epsilon}} di \right)^{\frac{\epsilon}{\epsilon-1}}. \quad (1)$$

$\epsilon > 1$ is the elasticity of substitution between differentiated goods, while θ is the sales demand elasticity with respect to the stock available for sales. The aggregator function

implies the following individual demand function for sales of good i at price P_{it} :

$$s_{it} = \left(\frac{a_{it}}{a_t} \right)^\theta \left(\frac{P_{it}}{P_t} \right)^{-\epsilon} s_t. \quad (2)$$

Note that we assume that sales of firm i depends on the ratio of its stock of goods available to the aggregate stock of goods available instead of on its stock alone. This assumption allows us to simplify the algebras and it can be motivated by assuming that a larger stock can facilitate sales for a firm only if its stock is larger relative to its competitors.

We can also derive the aggregate price index P_t consistent with aggregate sales as:

$$P_t = \left(\int_0^1 \left(\frac{a_{it}}{a_t} \right)^\theta (P_{it})^{1-\epsilon} di \right)^{\frac{1}{1-\epsilon}}. \quad (3)$$

We define the law of motion for the stock of goods available for sale a_{it} as follows:

$$a_{it} = y_{it} + (1 - \delta)(a_{it-1} - s_{it-1}), \quad (4)$$

where y_{it} is the output produced by firm i , and $0 < \delta < 1$ is the rate of depreciation of the inventory stock. The inventory stock at the end of period $t - 1$ may be defined as $x_{it-1} = a_{it-1} - s_{it-1}$.³ Furthermore, we assume that production uses labor h_{it} as its only input (or, alternatively, that capital is firm-specific and fixed over the relevant decision period) and is subject to shifts in aggregate productivity z_t :

$$y_{it} = z_t h_{it}^{1-\alpha}. \quad (5)$$

$0 < \alpha < 1$ is the labor elasticity.

We assume that each monopolistically competitive firm is subject to nominal rigidity in the form of a quadratic cost of adjusting its optimal price relative to a geometric index composed of steady state inflation π and lagged inflation π_{t-1} with weight $0 < \eta < 1$. That is, price adjustment cost is $\frac{\varphi}{2} \left(\frac{P_{it}}{\pi_{t-1}^\eta \pi^{1-\eta} P_{it-1}} - 1 \right)^2 s_t$, with $\varphi > 0$. Note that we scale the cost function by aggregate sales instead of output since the former is the relevant activity variable in a model with inventories.

A firm's current profits are then given by:

$$\frac{P_{it}}{P_t} s_{it} - w_t h_{it} - \frac{\varphi}{2} \left(\frac{P_{it}}{\pi_{t-1}^\eta \pi^{1-\eta} P_{it-1}} - 1 \right)^2 s_t, \quad (6)$$

³The law of motion for inventories is therefore:

$$x_{it} = y_{it} - s_{it} + (1 - \delta)x_{it-1},$$

where the net addition to inventories is unsold output.

where w_t is the competitive wage. Firms maximize the present value of (6) which they evaluate at the discount factor $\beta^t \lambda_t$, where $0 < \beta < 1$, and λ_t is the marginal utility of household consumption. They choose their optimal price P_{it} , the desired level of goods for sale a_{it} , and labor input h_{it} , subject to the demand function (2), the law of motion (4), and the production function (5).

The first-order conditions of the firm's profit maximization problem are:

$$\begin{aligned} \varphi \left(\frac{P_{it}}{\pi_{t-1}^\eta \pi^{1-\eta} P_{it-1}} - 1 \right) \frac{1}{\pi_{t-1}^\eta \pi^{1-\eta} P_{it-1}} s_t &= E_t \beta \frac{\lambda_{t+1}}{\lambda_t} \varphi \left(\frac{P_{it+1}}{\pi_t^\eta \pi^{1-\eta} P_{it}} - 1 \right) \frac{P_{it+1}}{\pi_t^\eta \pi^{1-\eta} P_{it}^2} s_{t+1} \\ &\quad + E_t \beta \frac{\lambda_{t+1}}{\lambda_t} mc_{it+1} (1 - \delta) \epsilon \frac{s_{it}}{P_{it}} + (1 - \epsilon) \frac{s_{it}}{P_{it}} \\ \theta \frac{P_{it}}{P_t} \frac{s_{it}}{a_{it}} - mc_{it} &= E_t \beta \frac{\lambda_{t+1}}{\lambda_t} mc_{it+1} \left[(1 - \delta) \left(\theta \frac{s_{it}}{a_{it}} - 1 \right) \right], \end{aligned} \quad (8)$$

and

$$w_t = (1 - \alpha) z_t h_{it}^{-\alpha} mc_{it}. \quad (9)$$

mc_{it} is the multiplier on the firm's consolidated budget constraint, which is obtained by substituting the production function and the demand function into the inventory accumulation equation. It is also the firm's marginal cost. Eq. (7) resembles a typical optimal price setting condition in a New Keynesian model with convex price adjustment costs (e.g. Krause and Lubik, 2007). The main difference is that marginal cost now enters the pricing relationship in expectations due to the presence of an inventory of unsold goods. The third condition simply equates a firm's marginal product with the real wage. We use this relationship later on to construct a time series for the unobservable marginal cost.

The optimality condition for optimal stocks (8) relates the sales-stock ratio $\gamma_{it} = s_{it}/a_{it}$ to the time path of marginal cost. Imposing symmetry, we can rewrite this condition as:

$$\frac{\theta \gamma_t - mc_t}{\theta \gamma_t - 1} = \beta (1 - \delta) E_t \frac{\lambda_{t+1}}{\lambda_t} mc_{t+1}. \quad (10)$$

An expected increase in marginal cost, which is a positive inflationary shock per Eq. (7), is matched by a fall in current marginal cost, other things being equal. The inventory model does therefore not necessarily deliver positive contemporaneous comovement of marginal cost with inflation. Alternatively, the sale-stock ratio might increase. The relative movements of γ_t and mc_t crucially depend on their relative sizes. We will delve further into this in the linearized version.

3 Deriving the NKPC with Inventories

We now use the first-order conditions of the firm's optimization problem to derive a reduced-form specification of the NKPC, which we then use as a data-generating process for our empirical analysis. We first compute the steady state. From optimal price setting we find that marginal cost:

$$mc = \frac{(\epsilon - 1)/\epsilon}{\beta(1 - \delta)}. \quad (11)$$

In the standard New Keynesian framework without inventories marginal cost is equal to the inverse of the (gross) markup $\frac{\epsilon}{\epsilon - 1}$. This is adjusted here by the factor $\beta(1 - \delta)$, which takes into account that current production has an intertemporal effect on sales due to inventory holdings. If we impose that $1/mc > 1$, this requires $\frac{\epsilon - 1}{\epsilon} < \beta(1 - \delta)$. Without inventory depreciation, this restriction holds for typical parameterizations. For $\delta > 0$, however, the restriction becomes more binding and requires a lower ϵ , that is, a less competitive product market. We can next derive the steady state sales-stock ratio γ from Eq. (10):

$$\gamma = \frac{\epsilon - 1}{\theta} \frac{1 - \beta(1 - \delta)}{\beta(1 - \delta)}. \quad (12)$$

In the next step we (log-)linearize the first-order conditions around the respective steady states of the endogenous variables. Denote $\tilde{x}_t = \log x_t - \log x$ as the log deviation of a variable x_t from its steady state. The linearized price-setting equation is then given by:

$$(1 + \beta\eta)\tilde{\pi}_t = \beta E_t \tilde{\pi}_{t+1} + \eta \tilde{\pi}_{t-1} + \frac{\epsilon - 1}{\varphi} E_t (\tilde{\lambda}_{t+1} - \tilde{\lambda}_t + \tilde{mc}_{t+1}). \quad (13)$$

This specification of the NKPC is close to the standard version, except for the driving process of inflation, which now involves *expected* marginal cost and marginal utility $\tilde{\lambda}$. We can derive a more typical version of the NKPC by using the linearized equation (10):

$$[1 - \beta(1 - \delta)]\tilde{\gamma}_t = \tilde{mc}_t - \beta(1 - \delta)(1 - \theta\gamma)E_t (\tilde{\lambda}_{t+1} - \tilde{\lambda}_t + \tilde{mc}_{t+1}). \quad (14)$$

Substituting Eq. (14) into Eq. (13) results in our benchmark version of the NKPC with inventories:

$$(1 + \beta\eta)\tilde{\pi}_t = \beta E_t \tilde{\pi}_{t+1} + \eta \tilde{\pi}_{t-1} + \varphi^{-1} \frac{1}{1 - \beta(1 - \delta)} \frac{\theta\gamma}{1 - \theta\gamma} \tilde{mc}_t - \varphi^{-1} \frac{\theta\gamma}{1 - \theta\gamma} \tilde{\gamma}_t. \quad (15)$$

There are two components in the driving process relevant for explaining inflation: marginal cost and the sales-stock ratio. The (conditional) response of inflation to movements in these variables depends on the sign of $1 - \theta\gamma$. Using the steady state expression from above, we find that this coefficient is positive if $\frac{\epsilon - 1}{\epsilon} < \beta(1 - \delta)$, which is the same restriction

for a strictly positive mark-up. In our empirical analysis we will therefore only consider parameterizations that are consistent with this.

The introduction of inventories has two effects on the NKPC. First, it changes the responsiveness of inflation to marginal cost. In the standard version of the NKPC this coefficient is equal to $\frac{\epsilon-1}{\varphi}$, which in Eq. (13) captures the impact of expected marginal cost. We will analyze the quantitative differences between these two specifications below. Second, inflation dynamics now also directly depend on the sales-stock ratio γ_t . In the standard NKPC model, an increase in a firm's marginal cost leads to higher prices as the firm passes on the cost of inputs to consumers. In the inventory model, however, there is an additional channel as firms can draw from their inventory stock to meet sales. Holding constant marginal cost, higher sales-stock ratio impacts current inflation negatively since the higher sales can be met out of inventories, which constitute additional goods supply. Finally, note that inflation dynamics is purely forward-looking when $\eta = 0$.

4 The Cyclical Behavior of Inventories and Marginal Cost

We now take a closer look at the joint behavior of inflation, marginal cost and inventories over the business cycle. We use the NKPC with inventories (15) as organizing principle of our discussion. We first present some stylized facts about the variables involved. We then calibrate the structural parameters in the NKPC and construct a time series of marginal cost, which is essentially unobservable, in various ways. In the next step, we contrast the driving processes of inflation in the standard NKPC with that in our benchmark version with inventories. The former consists only of real marginal costs, while the latter also includes the sales-stock ratio and different coefficients.

4.1 Data

Our full sample period ranges from 1947:1 - 2008:4. We also consider a sub-sample from 1984:1 onwards, which covers the Great Moderation during which the behavior of many macroeconomic time series changed. The data are from the U.S. National Income and Product Accounts (NIPA) and are extracted from the Haver Analytics database. The key series in our framework is the sales-stock ratio γ_t , which is not directly available. We construct it instead by using the definition of inventories $x_t = a_t - s_t$, so that $\gamma_t = \left(1 + \frac{x_t}{s_t}\right)^{-1}$, where the inventory-sales ratio $\frac{x_t}{s_t}$ is taken from NIPA. We use observations for both total private inventories and for the non-farm business sector.

Similarly, we have to construct a series for unobservable marginal cost. We can do this

in various ways. First, marginal cost can be derived from the labor demand condition (9): $mc_t = \frac{w_t h_t}{(1-\alpha)y_t}$, where the right-hand side is real unit labor cost, i.e. the wage divided by the marginal product. Alternatively, we can use the definition of the labor share $\frac{w_t h_t}{y_t} = (1-\alpha)mc_t$. Both proxies for marginal costs have been used in the literature (e.g. Galí and Gertler, 1999). Finally, we can also construct marginal cost from (10) by solving this equation forward and using observations on γ_t and the marginal utility of consumption λ_t . This approach, however, requires specifying a stochastic process for these variables. We discuss this issue further below.

Our output measure is real per capita GDP, which we compute by dividing real GDP by the civilian non-institutional population aged 16 and over. Our measure of inflation is the (log) change in the consumer price index. We consider an alternative price level measure, the GDP deflator, in the robustness section. We also utilize a consumption series, which is private sector consumption of non-durables. We pass all quantity variables through an HP-filter with smoothing parameter $\lambda = 1600$ for quarterly data. The second moments of the data series are reported in Table 1. The moments are computed for the non-farm business sector; the results for the full private sector (not reported) are virtually identical.

Over the full sample, inflation is half as volatile as output and they comove positively, albeit not strongly. The sales-stock ratio γ is somewhat more volatile than GDP , but completely acyclical ($corr(y, \gamma) = -0.01$).⁴ The marginal cost proxies, unit labor cost ULC and the labor share LS , both exhibit a small negative contemporaneous correlation with output, but a small positive one with inflation. Both are less volatile than output, albeit not by much. The fact that the sales-stock ratio not only comoves negatively with the proxies, but also with inflation, would favor the role of inventories in explaining inflation dynamics as the coefficient on γ in (15) is negative. Note, however, that the correlation between ULC and LS is only 0.74. The choice of the marginal cost proxy therefore clearly matters for the empirical analysis of the NKPC.

The results for the sub-sample from 1984 on are broadly consistent in terms of the comovement pattern. The standard deviations of all series decline by half, however, which reflects the Great Moderation. The cyclical patterns of the series do not change either except for the extent of comovement. For instance, the sale-stock ratio now shows a higher

⁴These numbers differ somewhat from the results reported in other papers. For instance, Bils and Kahn (2000) emphasize that the sales-stock ratio is procyclical while it is acyclical or mildly countercyclical in our analysis. This difference is due to differences in the sample period, data frequency, and different aggregation levels of the data. Bils and Kahn (2000) use monthly data of aggregate manufacturing over the period 1959-1997, and obtain a correlation of 0.68. Using our data for the same sample period we obtain a correlation between the sales-stock ratio and GDP of 0.27.

degree of negative correlation with GDP and inflation, but much less so than with the marginal cost proxies. Again, this possibly reflects the effect of the Great Moderation on macroeconomic series.⁵

4.2 Calibration

We now assign numerical values to the structural parameters. We use this calibration to construct a series for the driving process in the NKPC (15), i.e. the weighted sum of marginal cost and the sales-stock ratio, but we also keep some parameters fixed in the estimation exercise. Each period corresponds to a quarter. The calibrated parameter values are reported in Table 2.

We set the discount factor $\beta = 0.99$ which implies a 4 percent annual real interest rate. We choose $\epsilon = 11$ as our benchmark value for the demand elasticity parameter. In the standard model, this would imply a steady state gross markup $mc^{-1} = \frac{\epsilon}{\epsilon-1} = 1.10$, that is, a mark-up of 10%. In the presence of inventories, however, this has to be adjusted by the factor $\beta(1-\delta)$. We consider two cases for the depreciation rate, $\delta = 0$ and $\delta = 0.05$ as in Jung and Yun (2006). The latter case implies a mark-up of 3.5%, which is just at the lower bound of the range suggested by Basu and Fernald (1997). We set the steady state sale-stock ratio γ to 0.25, which is the mean of the data series over the full sample. This allows us to compute the sales elasticity with respect to stock of goods available parameter θ from the steady state relationship (12), which in the benchmark case with no depreciation implies $\theta = 0.40$. Increasing the depreciation rate leads to a higher implied sales elasticity, e.g. with $\delta = 0.05$, we have $\theta = 2.53$.⁶ Intuitively, when the inventory stock is less persistent and declines over time, a higher responsiveness of sales is required to maintain a given sales-stock ratio.

We now use these calibrated parameter values to construct a series for the driving process in the modified NKPC (15) that includes inventories. We compare this imputed series with the typically used proxies in empirical studies. The driving process in the standard NKPC (derived from an optimal price-setting specification with quadratic price adjustment costs) is simply $\frac{\epsilon-1}{\varphi} \widetilde{mc}_t$. In the specification with inventories, the modified driving process in (15) is:

$$\frac{\epsilon-1}{\varphi} \left[\frac{(\epsilon-1)^{-1}}{1-\beta(1-\delta)} \frac{\theta\gamma}{1-\theta\gamma} \widetilde{mc}_t - \frac{1}{\epsilon-1} \frac{\theta\gamma}{1-\theta\gamma} \widetilde{\gamma}_t \right]. \quad (16)$$

⁵Incidentally, McConnell and Perez-Quiros (2000) suggest that one of the causes of the Great Moderation is changes in firm's inventory management.

⁶An alternative calibration assumes a higher sales-stock ratio of $\gamma = 0.30$, which is the mean for the sub-sample after 1984. In this case $\theta = 0.34$ for $\delta = 0$ and $\theta = 2.11$ for $\delta = 0.05$.

The imputed series for our benchmark calibration with $\delta = 0$ is depicted in Figure 1 in the top panel.⁷ The marginal cost series is unit labor cost from the non-farm business sector. Notably, the two series overlay each other almost perfectly. Their correlation is 0.99, while the standard deviation of the imputed series at 1.81% is slightly higher than that of unit labor cost (1.60%). Adding inventories to the benchmark model seemingly has no effect on the driving process of inflation compared to the standard NKPC.

The picture changes when we set the inventory depreciation rate to $\delta = 0.05$, but keep the other parameter values the same.⁸ This case is depicted in the bottom panel of Figure 1. The imputed series is now much more volatile than the marginal cost proxy, with a standard deviation of 4.85%, albeit still with a correlation coefficient of 0.99. Since we are keeping the mean sales-stock ratio γ fixed across the two specifications, this implies a sales elasticity of 2.53, much higher than in the benchmark. Note, however, that the quantitative differences between the two specifications are purely driven by the differences in the depreciation rate. A higher depreciation rate increases both the weight on the sales-stock ratio and the weight on the marginal cost in the imputed series, see Eqs. (12) and (16). However, while it raises the volatility of the driving process, it does not increase its comovement with marginal cost and thereby inflation.

This statement comes with the caveat that it strictly applies only to the experiment where we impute the series for the driving process and use the imputed series as an exogenous regressor. Both $\widetilde{mc}_t$ and $\widetilde{\gamma}_t$ are, however, endogenously determined with inflation in the larger system, which may impose additional cross-equation restrictions on their joint behavior. We will return to this issue in the estimation section. Any empirical treatment of the NKPC with inventories will therefore confront similar challenges as the standard NKPC in that dimension, but might offer improvement in terms of matching up the volatilities of inflation with its driving process. The key parameter is, of course, the inventory depreciation rate δ , which we attempt to estimate more formally in the next section.

⁷In the figures, we factored out the scale factor $\frac{\varepsilon-1}{\varphi}$ since it is the same for both specifications, that is, with and without inventories. This allows for a more direct comparison of the marginal cost proxies and the imputed driving process.

⁸This parameterization does not violate the restriction for a strictly positive mark-up:

$$0.91 = \frac{\varepsilon-1}{\varepsilon} < \beta(1-\delta) = 0.94.$$

5 Inflation and Inventories: A Limited Information Approach

We now proceed to conduct a more formal analysis of the NKPC. We pursue a limited information approach in that we do not use all the information available in the full general equilibrium model. Instead, we concentrate on the NKPC (15) and treat it as an equation describing the dynamics of inflation driven by marginal cost and inventories. To be more precise, we do not impose the cross-equation and cross-coefficient restrictions on the comovement of the endogenous variables that the full model would prescribe. Instead, we treat the driving variables as exogenously determined.

We pursue several approaches in taking the NKPC to the data. First, we note that Eq. (15) is an expectational difference equation that can be solved forward. Given stochastic processes for the driving variables, we can then describe inflation dynamics as a general autoregressive model, the reduced-form coefficients of which can be estimated by least squares. Given convenient parameterizations of some structural parameters, it is then possible to identify key parameters. Our second approach treats the NKPC as a moment condition which we estimate with an instrumental variable approach such as GMM.

We are interested in two questions. First, how well does the modified NKPC explain inflation dynamics when compared with the standard specification. The second question is related to the first and deals with the degree of indexation in price-setting. This is captured in the theoretical model by the presence of lagged inflation in the NKPC and the value of the parameter η . Inflation exhibits a fairly high degree of persistence, which the standard NKPC has difficulty explaining due to the lack of persistence in marginal cost. This motivated Galí and Gertler (1999) to introduce indexation into the basic framework.

5.1 Constructing a Marginal Cost Measure

Marginal cost is unobservable to the econometrician. In the previous section, we therefore used the labor share and unit labor cost as proxies. Both have support by theory since in the log-linearized framework detailed above they are, in fact, exactly equal to marginal cost. In this section, we use an alternative theory-based method of constructing a marginal cost series. We note that Eq. (14) is an expectational difference equation in $\widetilde{mc}_t$, with driving processes $\widetilde{\gamma}_t$ and $(\widetilde{\lambda}_{t+1} - \widetilde{\lambda}_t)$. Given observations on the latter, we can project them on the former to construct a time-series for marginal cost. This idea is similar to the present-value literature on the empirical evaluation of asset-pricing, government debt and current-account stability, and has also been used by Jung and Yun (2006).

We first solve (14) forward:

$$\widetilde{mc}_t = \sum_{j=0}^{\infty} [\beta(1-\delta)(1-\theta\gamma)]^j E_t \left\{ [1-\beta(1-\delta)] \widetilde{\gamma}_{t+j} + \beta(1-\delta)(1-\theta\gamma) (\widetilde{\lambda}_{t+1+j} - \widetilde{\lambda}_{t+j}) \right\}. \quad (17)$$

Given time-series processes for the variables on the right-hand side, we can solve out the conditional expectation and compute the discounted infinite sum for calibrated parameter values.⁹ The expression on the right-hand side is an optimal predictor for the unobservable marginal cost.

To illustrate this point, we make two simplifying assumptions. First, assume that the marginal utility of wealth $\widetilde{\lambda}_t = -\sigma \widetilde{c}_t$, where c_t is consumption and $\sigma > 0$ is the inverse of the intertemporal substitution elasticity. This specification can be derived from a simple household optimization problem with CRRA-preferences. We then have $E_t (\widetilde{\lambda}_{t+1+j} - \widetilde{\lambda}_{t+j}) = -\sigma E_t (\widetilde{c}_{t+1+j} - \widetilde{c}_{t+j})$. If consumption is a martingale, then this term is exactly equal to zero as there are no predictable components in future consumption growth. Nevertheless, because of consumption smoothing, movements in this term will be small and we therefore disregard it in this simple calculation. The second assumption specifies an (exogenous) law of motion for the sales-stock ratio $\widetilde{\gamma}_t$, which we assume to be AR(1) with a zero mean innovation: $\widetilde{\gamma}_t = \rho \widetilde{\gamma}_{t-1} + \epsilon_t$. It then follows that $E_t \widetilde{\gamma}_{t+j} = \rho^j \widetilde{\gamma}_t$.

We can substitute this into the present value condition (17) and find:

$$\widetilde{mc}_t = \frac{1 - \beta(1-\delta)}{1 - \beta(1-\delta)(1-\theta\gamma)\rho} \widetilde{\gamma}_t. \quad (18)$$

The imputed marginal cost series is simply a scaled version of the sales-stock ratio with the same autoregressive coefficient ρ , but a smaller innovation variance on account of the coefficient on $\widetilde{\gamma}_t$. This expression can then be used in Eq. (16) to find the reduced-form driving process:

$$\frac{1}{\epsilon - 1} \frac{\theta\gamma}{1 - \theta\gamma} \frac{\beta(1-\delta)(1-\theta\gamma)\rho}{1 - \beta(1-\delta)(1-\theta\gamma)\rho} \widetilde{\gamma}_t. \quad (19)$$

Note that both the constructed marginal cost series (18) and the driving process (19) keep the cyclical behavior of $\widetilde{\gamma}_t$, which is countercyclical and comoves negatively with CPI inflation.¹⁰ This is clearly a strike against the inventory model since the typical marginal cost proxies, unit labor cost and the labor share, comove positively with CPI inflation. This reduced-form representation of the marginal cost equation also resolves the issue discussed

⁹ Given the restrictions on the parameter space discussed in the calibration section, it is straightforward to show that the infinite sum exists and is bounded.

¹⁰ It is straightforward to show that the coefficient on $\widetilde{\gamma}_t$ is positive. A sufficient condition for this is that the steady state mark-up is positive, which holds for the restriction $\frac{\epsilon-1}{\epsilon} < \beta(1-\delta)$.

above, namely to what extent the inclusion of the sales-stock ratio reinforces the cyclical nature of the marginal proxies. The logic of the model suggests that when the optimal inventory condition is taken into account, this renders marginal cost at best acyclical with respect to inflation. The obvious caveat is that this interpretation rests on the simplifying assumptions we made above.

For illustration purposes we now estimate ρ with least squares on the full sample and compute the marginal cost series for the benchmark and alternative calibrations. We find $\hat{\rho} = 0.81$ with a 95% confidence interval $[0.73, 0.88]$. The imputed series are depicted in Figures 2 and 3. The former compares the imputed marginal cost series under the benchmark and alternative calibrations with unit labor costs. The differences to the previous calibrations are striking. The marginal cost series are excessively smooth, and an order of magnitude less volatile than unit labor cost. Increasing the depreciation rate has virtually no impact on the behavior of the imputed series in striking contrast to the experiment depicted in Figure 1. This pattern is confirmed in Figure 3 which reports the driving processes for both calibrations and for both imputation methods. The excess smoothness of imputed marginal cost series using the present-value relationship with the sales-stock ratio is reminiscent of results in the asset pricing and intertemporal current account literature. It sheds doubt on the validity and correct specification of the optimality condition (14), which we used to back out the marginal cost series. However, this conclusion is subject to the caveat that we disregarded the contribution of consumption growth and that we projected marginal cost based on a univariate model for $\tilde{\gamma}_t$ alone. We therefore check the robustness of our conclusions by adding observations on consumptions and by using a multivariate forecasting model for the sales-stock ratio.¹¹

Consider the generic data vector x_t , which contains in our case the stock-sales ratio $\tilde{\gamma}_t$, the growth rate of consumption $\Delta\tilde{c}_t$, other variables we judge as useful for forecasting marginal cost, such as GDP, and lags thereof. We assume that the process for x_t is described by the VAR: $x_t = Ax_{t-1} + \epsilon_t$, which is either a true first-order process or the companion-form representation of a higher order process. Conditional expectations are therefore given by: $E_t x_{t+j} = A^j x_t$. Denote the extraction vector for some variable z_t as ι_z , so that $E_t \tilde{\gamma}_{t+j} = \iota_\gamma A^j x_t$ and $E_t (\tilde{\lambda}_{t+1+j} - \tilde{\lambda}_{t+j}) = -\sigma E_t \Delta\tilde{c}_{t+1+j} = -\sigma \iota_c A^j x_t$. The implied marginal cost series can then be computed as follows:

$$\widetilde{mc}_t = [1 - \beta(1 - \delta)\iota_\gamma - \beta(1 - \delta)(1 - \theta\gamma)\sigma\iota_c][I - \beta(1 - \delta)(1 - \theta\gamma)A]^{-1}x_t. \quad (20)$$

We contrast the imputed series with the marginal cost proxies in Figure 4. Parameter values

¹¹This follows the approach in Jung and Yun (2006).

for the benchmark and alternative calibrations are the same as before. The differences to the simple exercise above are striking. The imputed marginal cost series is now as volatile as unit labor cost. The standard deviation of the imputed series is 1.23, which is slightly less than that of unit labor cost at 1.60, but more than that of the labor share. The correlation coefficient of the two series is 0.40, although the correlation coefficient with inflation is still negative (-0.12). This stands in contrast to the imputed marginal cost series for the simple example above, where we assumed that the stock-sales ratio $\tilde{\gamma}_t$ is an exogenous AR(1) process. The method of imputing the present-value relationship thus clearly matters. We return to this point in Section 6.

Given the thus imputed marginal cost we can next compute the implied driving process as in Eq. (16), which are depicted in Figure 5. There are two notable observations. First, a higher depreciation rate makes the imputed series more volatile as we already noticed in the exercises above. Estimating this parameter will therefore be of prime importance. Second, the correlation of the driving processes in the specification with and without inventories is now only 0.21, while the contemporaneous correlation with inflation is -0.12 . This raises doubts as to whether the inventory specification can match inflation dynamics.

Before we move on to the estimation of the NKPC, we briefly summarize what we have done so far. We have constructed time series for marginal cost and for the driving process in the NKPC. The former was constructed using the optimality condition for inventories in the theoretical model. When we impute the marginal cost and the driving process based on an exogenous AR(1) process for stock-sales ratio, the resulting imputed series are excessively smooth. However, when we use a VAR to impute the marginal cost, the resulting series exhibits volatility of the same order of magnitude as unit labor cost and the labor share. Nevertheless, the correlation between the inventory-based marginal cost series and inflation is mildly negative. We also showed that the driving process can be written in terms of marginal cost and the sales-stock ratio or in terms of either of these variables if additional information from the model's optimality conditions is brought to bear. The time-series properties of the various marginal cost series we use, however, carry over to the driving process.

5.2 GMM Estimation of the NKPC with Inventories

We now estimate the NKPC using a generalized method of moments (GMM) approach. To provide a benchmark for our inventory specification, we first estimate both restricted and unrestricted version of the standard NKPC with various proxies for marginal cost as in

the original model of Galí and Gertler (1999). Our focus will be on the importance of the backward-looking term in the NKPC, and thus on the value of the indexation parameter η . This will give an indication as to whether inflation is more explained by its own intrinsic price dynamics (i.e., a large value of η) or if extrinsic sources of persistence are more dominant. Of secondary interest is the coefficient on marginal cost in the NKPC, which captures the strength of the transmission mechanism between the real and nominal side and indicates the degree of price stickiness.

We estimate the following NKPC for the specification without inventories:

$$\tilde{\pi}_t = \gamma_f E_t \tilde{\pi}_{t+1} + \gamma_b \tilde{\pi}_{t-1} + \kappa \tilde{mc}_t. \quad (21)$$

Inflation depends on past and expected future inflation, while the forcing variable is real marginal cost. This is the hybrid NKPC from the benchmark model in Galí and Gertler (1999). The coefficient on expected inflation $\gamma_f = \beta/(1 + \beta\eta)$, while the coefficient on past inflation $\gamma_b = \eta/(1 + \beta\eta)$; finally, the slope coefficient $\kappa = (\epsilon - 1)/[\varphi(1 + \beta\eta)]$. Note that when $\eta = 0$ the specification reduces to the purely forward looking NKPC. Moreover, when $\beta = 1$, then $\gamma_f + \gamma_b = 1$.

In our empirical exercises, we report results from three specifications. First, we estimate a highly restricted version of the NKPC (21) in which the coefficients γ_f and γ_b is restricted to sum to one. We relax this restriction in the second exercise, where we estimate reduced-form coefficients, γ_f, γ_b and κ . Finally, we also estimate the structural parameters in the structural version of NKPC (21) with $\gamma_f = \beta/(1 + \beta\eta)$, $\gamma_b = \eta/(1 + \beta\eta)$ and $\kappa = (\epsilon - 1)/[\varphi(1 + \beta\eta)]$. We impose $\beta = 0.99$ in this case. We note, however, that the parameters in the coefficient $(\epsilon - 1)/\varphi \equiv \kappa'$ are not separately identifiable in the baseline specification as the coefficient simply scales the marginal cost term and appears nowhere else.

Our econometric approach is relatively straightforward. Let $\mathbf{z}_t$ denote a vector of variables observed at time t . The NKPC (21) then defines a set of orthogonality conditions:

$$E_t [\tilde{\pi}_t - \gamma_f \tilde{\pi}_{t+1} - \gamma_b \tilde{\pi}_{t-1} - \kappa \tilde{mc}_t] \mathbf{z}_t = 0. \quad (22)$$

Given these conditions, we can estimate the model using the generalized method of moments (GMM). We choose our instruments from the set $\mathbf{z}_t$ which includes lags of the inflation rate $\tilde{\pi}$, proxies for marginal cost $\tilde{mc}$, and real per capita GDP. In the inventory specification we also use the sales-stock ratio $\tilde{\gamma}$ as an instrument. The actual choice of the instrument in each regression is informed by parsimony, high p-values in overidentification tests, and a high correlation with the endogenous variables in the first-step regression. The weighting

matrix is computed from the estimated heteroskedasticity- and autocorrelation-adjusted covariance matrix. The sample period for the empirical analysis is 1960:1 - 2008:4.

The GMM-estimation results for the standard NKPC are reported in Table 3. We use both real unit labor cost and the labor share as proxies for real marginal costs. The estimates are very similar to those found in the literature. For the labor share, the coefficient γ_f on expected inflation is 0.75 across specifications. This is consistent with a structural estimate of $\eta = 0.34$, i.e., a fraction of backward-looking price setters of about 1/3. While the forward-looking coefficient is very precisely estimated, the indexation parameter and the marginal cost coefficient κ are less precisely estimated. Furthermore, the J-test statistics for overidentifying restrictions suggest that the model is well-specified given the parsimonious set of instruments. When we use unit labor costs in the estimation, the results for the inflation coefficients are virtually identical, while the NKPC coefficient κ is generally smaller. However, the J-test statistics show less, albeit still convincing, support for this specification. The differences in the estimates reflect the different time-series properties of the respective proxies.

The results for the modified NKPC with inventories are reported in Table 4. We first estimate a reduced form version of (15) where we treat the coefficients on the inflation terms, marginal cost (κ_{mc}), and the sales-stock ratio (κ_γ) as reduced form coefficients:

$$\tilde{\pi}_t = \gamma_f E_t \tilde{\pi}_{t+1} + \gamma_b \tilde{\pi}_{t-1} + \kappa_{mc} \tilde{mc}_t - \kappa_\gamma \tilde{\gamma}_t. \quad (23)$$

We use observed data on $\tilde{\gamma}_t$ and the same proxies for $\tilde{mc}_t$ as before; that is, we do not impose the cross-equation restrictions imposed by the optimality condition (14). Expected inflation now carries an estimated coefficient of $\gamma_f = 0.80$, which implies a lower degree of intrinsic inflation persistence. However, this value is not statistically different from the lower estimate in the standard model. The marginal-cost coefficient $\kappa_{mc} = 0.019$, while $\kappa_\gamma = 0.028$. Both coefficients are imprecisely estimated, although the J-statistic indicates that the specification would not be rejected at conventional significance levels. Note, however, that the sales-stock ratio is negatively correlated with both marginal cost proxies. Hence, movements in the two observables tend to go in the same direction and reinforce each other in driving inflation. By adding a persistent variable to the right-hand side this mechanically reduces the importance of the lagged inflation term.¹² The results for

¹²This finding is reminiscent of the point made by Krause et al. (2008a,b), who add search and matching frictions in the labor market to the standard New Keynesian framework and study the impact on the NKPC. This modifies the concept of the marginal cost term as in the present paper and adds additional terms to the right-hand side of the NKPC. However, the impact on inflation dynamics is negligible, although it reduces the importance of the lagged-inflation term slightly.

unit labor costs are again broadly similar. κ_{mc} is close to zero and statistically insignificant, while the coefficient on $\tilde{\gamma}_t$ is the same as for the labor-share specification.

We will now delve deeper into whether this specific structural model is consistent with the data. Since the inventory-specification is more richly parameterized, we proceed in several steps. We first estimate the optimality condition for the sales-stock ratio (14) with GMM, again using the two proxies for marginal cost. We impose the steady state restriction $\theta\gamma = (\epsilon - 1)\frac{1-\beta(1-\delta)}{\beta(1-\delta)}$ and the consistency condition $\frac{\epsilon-1}{\epsilon} < \beta(1-\delta)$. As in the calibration exercise above, we assume that $\tilde{\lambda}_t = -\sigma\tilde{c}_t$ and use data on non-durable consumption in the estimation. We also set $\beta = 0.99$ and $\sigma = 1$. This allows us to identify the depreciation rate δ and the demand elasticity ϵ . The results are reported in Table 4.

The depreciation rate δ is estimated to be zero for both proxies. If we had not imposed the non-negativity condition, the estimate would be -0.02 with a very small standard deviation. This suggests that the estimation algorithm settles on a corner solution. The demand elasticity ϵ is tightly estimated at 34 (23) for labor share (unit labor cost) data, which implies a mark-up of 3% (4.5%). This is much smaller than typically found in either the calibration or estimation literature. The overidentification test barely finds in favor of either specification with unit labor cost or the labor share as a proxy for marginal cost. We experimented widely with the instrument set, but could not find p-values above 30%. We would conclude at this point that the optimal inventory condition is only weakly consistent with the data, which suggests that an approach of backing out the unobserved marginal cost from the sales-stock series may be ill-advised. We will investigate this point further when we use the imputed marginal cost series in the estimation below.

We now estimate the inventory condition (14) jointly with the NKPC (15). As before, we fix $\beta = 0.99$ and $\sigma = 1$. The two moment conditions are estimated on the inflation rate, the respective marginal cost proxies, the sales-stock ratio and non-durable consumption. The instrument set includes lags of these variables and real GDP. Overall, this specification clearly fails the overidentification test for any instrument set we tried. The parameter estimates, however, are in line with the results from the previous two specifications. δ and ϵ are virtually identical to those found before, while the estimate of the indexation parameter $\eta = 0.29$ is consistent with less intrinsic inflation persistence in the inventory specification. The implied marginal cost coefficient $(\epsilon - 1)/\varphi = 0.050$ is twice as large as in the standard NKPC. In the case of unit labor costs, $\eta = 0.24$ while the implied marginal cost coefficient $(\epsilon - 1)/\varphi = 0.039$, both lower than the case of labor share. However, the overidentification test clearly rejects the two-equation specification for both proxies.

In the next step, we avoid using proxies for the marginal cost term and instead use the imputed marginal cost series from the optimal inventory-condition, computed from Eq.(20). We perform three different exercises. First, we estimate the standard NKPC (21) where we use the imputed marginal cost series as the driving process. Second, we estimate the NKPC specification with inventories (23) using the imputed series and the observed sales-stock ratio. Finally, we estimate the NKPC using the imputed driving process (16). The latter two specifications use, strictly speaking, redundant information as the stocks-sale ratio is already incorporated in the imputed marginal cost series. The results are reported in Table 5.

We find that across the board the performance of the model with imputed marginal cost is worse than for the standard proxies, as captured by the overidentification statistic. However, the estimates of the NKPC parameters are broadly in line with all previous specifications. The forward-looking coefficient in the standard NKPC specification is 0.73 and thus statistically identical to the estimates using the marginal cost proxies. Similar findings are obtained for the other specifications and for the marginal cost coefficients. Specifically, the coefficient on the sales-stock ratio is larger than the coefficient on marginal cost, but both are fairly imprecisely estimated. This leads us to observe that the NKPC is a robust description of inflation dynamics as captured by forward- and backward-looking behavior, but this appears almost independent of the specific driving process. We already noted that the imputed marginal cost series does not exhibit the same time-series behavior as the marginal cost proxies which explains the comparatively worse performance of the former specification.

To turn this argument around, the GMM estimates suggest that the introduction of inventories into an otherwise standard NKPC framework does not markedly alter the implications for inflation dynamics. When compared to the standard framework, we find that the degree of backward-looking price setting, and hence intrinsic inflation persistence, decreases by a small amount. However, this is dependent on the marginal cost proxy used and on whether we estimate a reduced-form or structural representation of the NKPC. The same conclusion applies to the effect on the marginal cost coefficient. The main caveat is that we did not impose the full set of cross-equation restrictions. Interestingly, when we include the sales-stock ratio as an additional regressor besides the imputed marginal cost, the coefficient on the latter declines substantially. This suggests that inventories do have *some* role to play for explaining movements in inflation.

We also estimate the optimal inventory condition, both as a single equation and com-

bined with the NKPC. We find that both specifications do not satisfy the overidentification test, which leads us to conclude that inventory dynamics are not well-captured by this optimality condition. The parameter estimates for the inventory depreciation rate are robustly zero across all specifications, while the demand elasticity parameter implies a low mark-up. These estimates go in the opposite direction of the calibration above, which produced more volatility in the driving process.

5.3 Robustness

We assess the robustness of our GMM estimates in two directions. First, we consider an alternative inflation concept in the NKPC, namely the GDP deflator. The correlation between the CPI-based inflation rate and the change in the GDP-deflator is 0.89, which suggests that the choice of price series does matter to some extent for the analysis of inflation dynamics. These impressions are confirmed by the results from our estimation exercise.¹³ Across all specifications, the degree of backward-looking behavior is larger when we use the GDP-deflator as price variable. Moreover, the Phillips-curve coefficients are smaller in each specification, too. We find, for example, that the coefficient on expected inflation in the unrestricted reduced-form specification of the standard NKPC is $\gamma_f = 0.65$, while it is 0.76 with CPI data. The coefficient on marginal cost at $\kappa = 0.16$ is less than half that under the benchmark. This pattern is also reflected in the estimates of the parameters, with $\eta = 0.56$ and $\kappa' = 0.01$.

Turning to the NKPC with inventories, the estimate of the forward-looking coefficient drops to 0.64, while the coefficients on the two components of the driving process, marginal cost and the sales-stock ratio fall by half. However, we found in the benchmark specification that the inclusion of inventories increases the weight on the forward-looking component. In this robustness check, adding inventories reduces the forward-looking coefficient, albeit not in a statistically significant manner. Finally, we estimate the 2-equation system composed of the optimal inventory condition and the structural version of the NKPC. As in the benchmark case, the overidentification test clearly rejects this specification. We find, however, that the estimate of the indexation parameter $\eta = 0.19$ is much smaller than in the benchmark version (at $\eta = 0.29$) and smaller than the single-equation estimates using the GDP-deflator.

In the second robustness exercise, we estimate the NKPC over a sub-sample, specifically the period from 1985:1 onwards. The starting date coincides with a commonly chosen

¹³These results are not reported in a table, but are available from the authors upon request.

break date in the behavior of the Federal Reserve (see Lubik and Schorfheide, 2004) and the accompanying decline in macroeconomic volatility. The results are reported in Table 6. We only report estimates for the inventory model with labor share proxy as the results with unit labor costs follow the same pattern as in the benchmark specification. In the specification where we only estimate reduced-form coefficients, the weight on expected inflation increases to $\gamma_f = 0.88$, while $\gamma_b = 0.12$. This is consistent with the findings in the previous literature that inflation dynamics became more extrinsic and less persistent during the Great Moderation. Both the estimates for κ_{mc} and κ_γ are not significantly different from zero, which is similar to our results for the full sample. The sub-sample estimates differ, however, for the structural parameters. In the case of the optimal inventory condition we find that the depreciation rate is estimated at $\delta = 0.04$, while the demand elasticity parameter $\varepsilon = 4.95$. Both estimates are statistically significant. Compared to the baseline, the depreciation is higher, and the demand elasticity is lower. Furthermore, the J-test statistic does not reject this specification. This suggests that the behavior of inventory variables has, in fact, changed for this sample period, which is consistent with the statistical evidence presented in Table 1. This does not, however, improve the performance of the two-equation specification (the optimal inventory condition together with the NKPC) as captured by the J-test. However, the point estimates of the structural parameters generally fall in line with the general pattern established above. Noticeably, the indexation parameter η is tightly estimated at 0.10, which reflects the decline in inflation persistence over the sub-sample.

6 Pitfalls of a Limited Information Approach

In our empirical analysis, we have pursued a limited information approach in that we did not incorporate all information potentially available from the underlying theoretical model. Specifically, we only focused on the NKPC (15), which is derived from the firm's price setting problem alone and does not incorporate information from the households' optimization problem or resource constraints in the economy. In other words, we did not fully impose cross-equation and cross-coefficient restrictions from the full model. This approach is preferable if there are doubts about the overall validity of the underlying model, or if the model or parts of it are likely to be mis-specified. However, this approach has some pitfalls, which we now illustrate by means of a simple example.

Assume that production is linear in labor input. We also abstract for illustration purposes from movements in exogenous productivity, which we normalize to unity. Then, the first-order condition (9) implies that $w_t = mc_t$. Moreover, standard CRRA-utility on

household consumption implies the relationship $\lambda_t = c_t^{-\sigma}$, that is, the marginal utility of consumption is equal to the multiplier on the budget constraint. Both optimality conditions can be connected by the household's labor supply condition. Assuming that labor supply is perfectly elastic and separable from consumption, this implies $w_t = \lambda_t^{-1}$. We therefore find that $\widetilde{mc}_t = -\widetilde{\lambda}_t = \sigma \widetilde{c}_t$. We can substitute these expressions - which, we want to emphasize, have been derived from the rest of the model and have not been used in the derivations above - into the (linearized) NKPC in Eq. (13) to obtain:¹⁴

$$\widetilde{\pi}_t = \beta E_t \widetilde{\pi}_{t+1} + \frac{\epsilon - 1}{\varphi} \widetilde{mc}_t. \quad (24)$$

This is, of course, the same representation of the NKPC that has been used numerous times in the literature. Since marginal cost is unobservable, it is typically proxied by unit labor cost (as we did above), or linked to an aggregate activity variable such as output via the production function. Seemingly, one representation of the NKPC derived from the inventory-model is observationally equivalent to the standard NKPC. However, the pitfall in this line of reasoning is that the full inventory model imposes additional restrictions that are obscured by focusing on this single-equation representation alone. For instance, the expression for the multiplier $\widetilde{\lambda}_t$ can be used in the optimality condition (14) to derive the model-consistent reduced-form representation for marginal cost:

$$\widetilde{mc}_t = \frac{1 - \beta(1 - \delta)}{1 - \beta(1 - \delta)(1 - \theta\gamma)} \widetilde{\gamma}_t. \quad (25)$$

This reveals the direction of misspecification we committed in the example above, where we imputed marginal cost by assuming an AR(1) process for the sales-stock ratio and by disregarding movements in the stochastic discount factor. Interestingly, the coefficients on $\widetilde{\gamma}_t$ differ by the AR(1)-parameter ρ . Since it is less than one, the imputed version implies a less volatile marginal cost series than the model-consistent series. This observation turned out to be correct as we saw when we imputed the marginal cost series from the present-value relationship in Figure 2. In other words, a simple AR(1) specification is not enough to capture the reduced-form dynamics in $\widetilde{\gamma}_t$. It is, of course, well known that capturing the time-series properties of the driving variables in present-value computations are crucial (see Nason and Smith, 2008). This potential pitfall can be avoided by backing out the implied, and internally consistent, marginal cost series from a fully-specified general equilibrium model.

However, this approach is not innocuous either since the full model is likely to be misspecified, and may impose conflicting restrictions on the series to be imputed. We illustrate

¹⁴We abstract from indexation in price-setting for simplicity: $\eta = 0$.

this now in a simple example. The model-consistent marginal cost-series is derived in Eq. (25). Alternatively, marginal cost can also be written in terms of an aggregate activity variable since $\widetilde{mc}_t = \widetilde{c}_t = \widetilde{s}_t$.¹⁵ We can use the law of motion for the stock of goods (4) to then express marginal cost in terms of output and inventories:

$$\widetilde{mc}_t = \frac{y}{a} \widetilde{y}_t + \widetilde{\gamma}_t - (1 - \delta) \widetilde{\gamma}_{t-1} + (1 - \delta)(1 - \gamma) \widetilde{s}_{t-1}. \quad (26)$$

Both marginal cost series are model-consistent in the sense that they are derived from equilibrium relationships. However, they are derived from different relationships. A priori, there is nothing in the model to guarantee that both methods result in an identical series for the unobserved marginal cost. We plot both implied series in Figure 6. The volatility of the second marginal cost series (26), labeled MC (GE2) in the graph, is of the same order of magnitude as the VAR based series MC (VAR), while the series based on Eq. (25) MC (GE1) is much less volatile. Moreover, the correlation between the two series is -0.12 . Similarly, the correlation between the VAR based series and (26) is -0.27 , which suggests that a full-information estimation approach would have some difficulty matching the inventory and inflation data. Finally, the correlation between unit labor cost and the imputed series (26) is 0.02 . We conclude at this point that our specification of the New Keynesian model with inventories is not able to capture inflation dynamics.

7 Conclusion

We introduce inventories into a New Keynesian monetary model and show how this implies a NKPC Phillips curve that is driven by marginal cost. The key theoretical point of our paper is that the presence of inventories changes the notion of marginal cost and the driving process of inflation in the NKPC. We show that NKPC can be written in a variety of representations, some of which are observationally equivalent to the standard version. However, the inventory model provides a way of backing out a marginal cost series from a firm's optimality conditions. We find, however, that the imputed series are of not much help in capturing inflation dynamics via the NKPC. This leads us to a discussion of the shortcomings of a partial equilibrium approach to modeling inflation dynamics.

Although the conclusion in our paper is mainly of a negative part, it is also by its nature model-specific. We would therefore consider research in the following directions as useful. First, inventories can be introduced into the model in alternative ways. A chief candidate

¹⁵The last equality is derived from the resource constraint $c_t + \frac{\varphi}{2} \left(\frac{\pi_t}{\pi} - 1 \right)^2 s_t = s_t$. We also assume $\sigma = 1$ for simplicity.

would be inventories in a production (as opposed to in final goods as in our setup). A second approach would add more structure to the production side of the economy, such as capital and variable capacity utilization. Alternatively, the production side could be motivated by introducing rigidities in the labor market, such as search and matching frictions, which by themselves affect marginal cost and the driving process in the NKPC.

From an empirical perspective, it would also be useful to study the implications of the model using the information in the entire equation system. This would allow the researcher to implicitly construct the correct marginal cost series, assuming that the theoretical model is not mis-specified, without having to rely on choices for a convenient semi-structural or reduced-form specification of the NKPC. Lubik and Teo (2009) pursue such an approach to estimate a New Keynesian model with inventory behavior using Bayesian methods.

References

- [1] Basu, Susanto and John G. Fernald (1997): “Returns to Scale in U.S. Production: Estimates and Implications”. *Journal of Political Economy*, 105, 249-283.
- [2] Bilal, Mark, and James A. Kahn (2000): “What Inventory Behavior Tells Us About Business Cycles”. *American Economic Review*, 90(3), 458-481.
- [3] Boileau, Martin, and Marc-André Letendre (2008): “Inventories, Sticky Prices, and the Persistence of Output and Inflation”. Manuscript.
- [4] Galí, Jordi and Mark Gertler (1999): “Inflation Dynamics: A Structural Econometric Analysis”. *Journal of Monetary Economics*, 44(2), 195-222.
- [5] Hornstein, Andreas (2005): “Sticky Prices and Real Rigidities/Flexibilities”. Manuscript.
- [6] Jung, YongSeung, and Tack Yun (2006): “An Inventory-Theoretic Approach to the New Keynesian Phillips Curve”. Manuscript.
- [7] Krause, Michael U., and Thomas A. Lubik (2007): “The (Ir)relevance of Real Wage Rigidity in the New Keynesian Model with Search Frictions”. *Journal of Monetary Economics*, 54(3), 706-727.
- [8] Krause, Michael U., David J. Lopez-Salido, and Thomas A. Lubik (2008a): “Inflation Dynamics with Search Frictions: A Structural Econometric Analysis”. *Journal of Monetary Economics*, 55, pp. 892-916.

- [9] Krause, Michael U., David J. Lopez-Salido, and Thomas A. Lubik (2008b): “Do Search Frictions Matter for Inflation Dynamics?” *European Economic Review*, 52, 1464-1479.
- [10] Lubik, Thomas A., and Frank Schorfheide (2004): “Testing for Indeterminacy: An Application to U.S. Monetary Policy”. *American Economic Review*, 94, 190-217.
- [11] Lubik, Thomas A., and Wing Leong Teo (2009): “Inventory Dynamics and Monetary Policy in an Estimated New Keynesian Model”. Manuscript.
- [12] McConnell, Margaret M. and Gabriel Perez-Quiros (2000): “Output Fluctuations in the U.S.: What Has Changed Since the Early 1980s?” *American Economic Review*, 90, 1464-1476.
- [13] Nason, James M., and Gregor W. Smith (2008): “The New Keynesian Phillips Curve: Lessons from Single-Equation Econometric Estimation”. *Federal Reserve Bank of Richmond Economic Quarterly*, 94(4), 361-395.
- [14] Schorfheide, Frank (2009): “DSGE Model-Based Estimation of the New Keynesian Phillips Curve”. *Federal Reserve Bank of Richmond Economic Quarterly*, 94(4), 397-433.

Table 1. Business Cycle Statistics

<i>Sample Period : 1947 – 2008</i>						
<i>Variable</i>	<i>s.d.(%)</i>	<i>Cross-Correlation</i>				
		<i>y</i>	<i>π</i>	<i>ulc</i>	<i>ls</i>	<i>γ</i>
<i>GDP</i>	1.68	1	0.27	-0.39	-0.26	-0.01
<i>CPI</i>	0.82		1	0.07	0.11	-0.29
<i>ULC</i>	1.60			1	0.74	-0.61
<i>LS</i>	1.06				1	-0.30
<i>Sales/Stock</i>	2.04					1

<i>Sample Period : 1984 – 2008</i>						
<i>Variable</i>	<i>s.d.(%)</i>	<i>Cross-Correlation</i>				
		<i>y</i>	<i>π</i>	<i>ulc</i>	<i>ls</i>	<i>γ</i>
<i>GDP</i>	0.94	1	0.32	-0.23	-0.17	-0.30
<i>CPI</i>	0.48		1	0.06	0.03	-0.37
<i>ULC</i>	1.02			1	0.84	-0.20
<i>LS</i>	0.88				1	-0.01
<i>Sales/Stock</i>	1.42					1

Table 2. Parameter Values and Steady State

Parameter	Definition	Benchmark	Alternative
β	Discount Factor	0.99	0.99
ϵ	Elasticity of Demand	11	11
δ	Inventory Depreciation	0	0.05
γ	Steady State Sale-Stock Ratio	0.25	0.25

Table 3. GMM Estimates: Standard NKPC

Specification	Marginal Cost Proxy			Unit Labor Cost	
	Labor Share				
Restricted NKPC $\gamma_f + \gamma_b = 1$	γ_f	$J(1)$	γ_f	κ	$J(8)$
	0.749 (0.077)	0.004 (0.951)	0.755 (0.069)	0.024 (0.010)	6.691 (0.570)
Unrestricted NKPC	γ_b	$J(1)$	γ_b	κ	$J(7)$
	0.764 (0.196)	0.004 (0.951)	0.741 (0.076)	0.023 (0.010)	6.707 (0.460)
Structural NKPC $\beta = 0.99$	η	$J(2)$	η	κ'	$J(8)$
	0.343 (0.138)	0.043 (0.978)	0.331 (0.120)	0.017 (0.008)	6.751 (0.564)

Note: The numbers in parenthesis are standard errors. For J-statistics, the numbers in parentheses are p-values. For structural NKPC,

$$\kappa' \equiv (\epsilon - 1)/\varphi.$$

Table 4. GMM Estimates: NKPC with Inventories

Specification	Labor Share					Marginal Cost Proxy					Unit Labor Cost				
	γ_f	γ_b	κ_{mc}	κ_γ	$J(1)$	γ_f	γ_b	κ_{mc}	κ_γ	$J(6)$	γ_f	γ_b	κ_{mc}	κ_γ	$J(6)$
Reduced form NKPC	0.800 (0.162)	0.192 (0.128)	0.019 (0.024)	0.028 (0.015)	0.001 (0.971)	0.721 (0.063)	0.247 (0.064)	0.002 (0.012)	0.029 (0.011)	6.363 (0.384)					
Optimal Inventory		δ		ε	$J(8)$	δ				$J(10)$					
$\beta = 0.99$		0.000 (0.000)		34.06 (2.085)	12.89 (0.116)	0.000 (0.000)				22.91 (1.300)					11.96 (0.288)
Structural NKPC	η	δ	ε	φ	$J(14)$	η	δ	ε	φ	$J(12)$					
$\beta = 0.99$	0.295 (0.001)	0.000 (0.000)	33.60 (0.046)	648.4 (1.263)	46.14 (0.001)	0.242 (0.001)	0.000 (0.000)	21.79 (0.087)	532.6 (2.785)	58.37 (0.000)					

Note: The numbers in parenthesis are standard errors. For J-statistics, the numbers in parentheses are p-values.

Table 5. GMM Estimates: Imputed Marginal Cost

Specification					
Standard NKPC	γ_f	γ_b	κ_{mc}		$J(7)$
	0.731	0.265	0.039		8.778
	(0.036)	(0.045)	(0.012)		(0.269)
NKPC with Inventories	γ_f	γ_b	κ_{mc}	κ_γ	$J(6)$
	0.695	0.271	0.017	0.027	8.727
	(0.017)	(0.012)	(0.019)	(0.005)	(0.189)
NKPC with Inventories (solved out DP)	γ_f	γ_b	κ_{mc}		$J(7)$
	0.748	0.282	0.028		8.625
	(0.029)	(0.061)	(0.049)		(0.283)

Note: The numbers in parenthesis are standard errors. For J-statistics, the numbers in parentheses are p-values. The imputed marginal cost is computed using Eq. (20).

Table 6. Robustness: Sub-Sample 1985:1-2008:4

Specification	NKPC with Inventories				
Restricted NKPC	γ_f	γ_b	κ_{mc}	κ_γ	$J(1)$
	0.878	0.120	0.000	0.022	0.636
	(0.209)	(0.093)	(0.034)	(0.019)	(0.425)
Optimal Inventory	δ		ε		$J(8)$
	0.044		4.947		7.014
	(0.017)		(1.732)		(0.535)
Structural NKPC	η	δ	ε	φ	$J(16)$
	0.108	0.030	6.699	1573	28.31
	(0.004)	(0.001)	(0.333)	(107.7)	(0.013)

Note: The numbers in parenthesis are standard errors. For J-statistics, the numbers in parentheses are p-values.




Figure 1: Imputed Driving Process in the NKPC with Inventories

Note: The imputed driving process is computed using Eq. (16) with unit labor cost acts as a proxy for marginal cost.




Figure 2: Imputed Marginal Cost

Note: The imputed marginal cost is computed using Eq. (18).




Figure 3: Imputed Driving Process in the NKPC with Imputed MC

Note: The imputed driving process using imputed MC is computed using Eq. (19). The imputed driving process using ULC refers to the imputed driving process in Figure 1.




Figure 4: VAR Based Imputed Marginal Cost

Note: The imputed marginal cost is computed using Eq. (20).




Figure 5: VAR Based Imputed Driving Process

Note: The VAR based imputed driving process is computed using Eqs. (16) and (20).




Figure 6: Comparison of Imputed Marginal Costs