Empirical Evidence on Inflation Expectations in the New Keynesian Phillips Curve

Citation

Mavroeidis, Sophocles, Mikkel Plagborg-Møller, and James H. Stock. 2014. Empirical Evidence on Inflation Expectations in the New Keynesian Phillips Curve. *Journal of Economic Literature* 52, no. 1: 124–188.

Published Version

doi:10.1257/jel.52.1.124

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:22795845>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Open Access Policy Articles, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#OAP>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

Empirical evidence on inflation expectations in the new Keynesian Phillips curve*

Sophocles Mavroeidis
University of Oxford

Mikkel Plagborg-Møller
Harvard University

James H. Stock
Harvard University

First version: February 23, 2013

This version: November 28, 2013

Abstract

We review the main identification strategies and empirical evidence on the role of expectations in the new Keynesian Phillips curve, paying particular attention to the issue of weak identification. Our goal is to provide a clear understanding of the role of expectations that integrates across the different papers and specifications in the literature. We discuss the properties of the various limited-information econometric methods used in the literature and provide explanations of why they produce conflicting results. Using a common data set and a flexible empirical approach, we find that researchers are faced with substantial specification uncertainty, as different combinations of various a priori reasonable specification choices give rise to a vast set of point estimates. Moreover, given a specification, estimation is subject to considerable sampling uncertainty due to weak identification. We highlight the assumptions that seem to matter most for identification and the configuration of point estimates. We conclude that the literature has reached a limit on how much can be learned about the new Keynesian Phillips curve from aggregate macroeconomic time series. New identification approaches and new data sets are needed to reach an empirical consensus.

1 Introduction

The idea that there is a trade-off between the rates of inflation and unemployment (or related measures of real economic activity), at least in the short run, is widely accepted in the economics profession and guides monetary policy making by major central banks. Phillips (1958) provided the first formal statistical evidence on this trade-off using data on wage inflation in the United

---

*Mavroeidis: University of Oxford and INET at Oxford. Plagborg-Møller: Harvard University. Stock: Harvard University. We benefited from comments by Guido Ascari, Gary Chamberlain, Janet Currie, Roger Farmer, Jeff Fuhrer, Max Kasy, Greg Mankiw, Kristoffer Nimark, Adrian Pagan, Barbara Rossi, Francisco Ruge-Murcia, Mark Watson, Mike Woodford, Francesco Zanetti, five anonymous referees, and seminar participants at Bath, Birkbeck, the ECB, Harvard, Oxford, UPF, the 2013 AEA Meetings, the 2013 DAEiNA meeting and the 2013 BMRC-QASS conference at Brunel University. Mavroeidis acknowledges financial support from the Leverhulme trust through a Philip Leverhulme Prize, and the European Commission through a Marie Curie Fellowship CIG 293675.

Kingdom, though the idea had existed well before that.¹ Samuelson and Solow (1960) extended what they called “Phillips’ curve” to U.S. data and to price inflation. The subsequent study of the nature of the Phillips trade-off and its implications for monetary policy and business cycle fluctuations has been one of the most active research areas in economics over the last fifty years.

The Phillips curve has a fascinating history, marked by landmark contributions and heated policy debates; see Gordon (2011) for an insightful recent survey. The literature is so large that it is impossible to address all major contributions to it in any single survey article. Instead, we focus on what is currently the most widely used model of this kind, the new Keynesian Phillips curve (NKPC), which has gained its popularity from its appealing theoretical microfoundations and what appeared to be early empirical success. The theory of the NKPC was laid out mostly in the 1980s and 1990s, and it is, by now, a standard feature of modern macroeconomic textbooks, see, e.g., Woodford (2003) and Galí (2008). The key property of the NKPC is that inflation is primarily a forward-looking process, driven by expectations of future real economic activity, rather than past shocks. From a policy perspective, this severely limits the scope for actively exploiting the Phillips trade-off. Instead, the forward-looking behavior provides a central role for monetary policy rules and opens the door to expectations management and communications as tools of monetary policy.² Indeed, the early empirical success of the NKPC, along with its rigorous microfoundations, has led to its widespread adoption as the key price determination equation in policy models used at central banks around the world.

This survey focuses on the empirics of the NKPC, and more specifically on the recent literature on the estimation of the NKPC using limited-information methods such as the generalized method of moments (GMM). The literature, which dates back to the seminal papers by Roberts (1995), Fuhrer and Moore (1995), Galí and Gertler (1999) and Sbordone (2002), is vast. By focusing on limited-information methods, we exclude estimation of the NKPC using full-information (system) methods, in which the NKPC is one of multiple structural equations within a simultaneous system, typically a dynamic stochastic general equilibrium (DSGE) model. By imposing a theoretical model on all the variables in the system, full-information methods have the potential to improve estimator precision, but they also introduce the risk of misspecification in other equations, inducing bias or inconsistency of the NKPC parameters of interest (Lindé, 2005; Beyer et al., 2008; Fukač and Pagan, 2010). In contrast, because they do not impose economic structure elsewhere in the model, the limited-information methods we consider here are robust to extraneous model misspecification. Full-information methods have been the subject of recent reviews by An and Schorfheide (2007) and Schorfheide (2011). Another promising strand of the literature uses detailed micro data to study price dynamics, see the review by Nakamura and Steinsson (2013). These papers are outside the scope of our review, which focuses only on studies that use macro data.

Despite the apparent early empirical success of the NKPC, the literature which we survey is full of puzzles. What should be relatively innocuous changes in instruments used, in vintages of data and in model specification all seem to matter significantly for the results. For example, simply reestimating the benchmark NKPC in Galí and Gertler (1999), using the same data series, method,

---

¹Mankiw (2001, p. C46) offers a prescient quote from the work of Hume (1752). Irving Fisher demonstrated a strong correlation between U.S. inflation and unemployment already in 1926, cf. Fisher (1973), but he interpreted the causality as running from prices to real economic activity.

²The policy relevance of this research agenda is highlighted in two speeches by Federal Reserve chairman Bernanke (2007, 2008).

and time period, but with revised data, reduces the estimate on the activity variable (real marginal cost) by half and makes the coefficient no longer statistically significant (Rudd and Whelan, 2007). This is but a single example of a high degree of sensitivity in this literature to minor econometric changes. Our goals in this review are to understand the reasons for this sensitivity and, more specifically, to provide a clear understanding of the role of expectations that integrates across the different papers and specifications. We do so first by reviewing the papers in the literature and the econometric theory underlying their approaches, then by estimating multiple specifications using a common data set. Since the first empirical work on the NKPC, there have been significant methodological developments in the area of estimation with weak instruments, and our analysis draws heavily on these methods to help explain the puzzles in the literature. Earlier surveys on the NKPC include Henry and Pagan (2004), Ólafsson (2006), Rudd and Whelan (2007), Nason and Smith (2008*b*) and Tsoukis et al. (2011). We extend these surveys by emphasizing the many econometric issues raised by estimation of the NKPC, and our empirical analysis spans a much wider range of estimation approaches and specifications than what previous individual papers have considered (in fact, we suspect we estimated more NKPC specifications than the entire preceding literature combined).

The outline of the paper is as follows. Section 2 briefly reviews the derivation of the NKPC under the Calvo (1983) assumption on price setting, followed by a description of the main extensions and empirical specifications. We emphasize that uncertainty about the NKPC parameters translates into significant uncertainty about the new Keynesian model's policy implications. Departing from the rational expectations assumption has non-trivial consequences for the model. Due to space constraints, we do not discuss recent movements away from the NKPC, such as state dependent pricing (Dotsey et al., 1999) or imperfect information (Mankiw and Reis, 2011).

Section 3 reviews the various limited-information econometric methods that have been used in the study of the NKPC, including instrumental variables, minimum distance, maximum likelihood and the use of survey data on inflation expectations. Furthermore, we propose a new idea for identification using data revisions as external instruments that obviates the need to impose ad hoc exclusion restrictions on the dynamics. We compare the different methods under both strong and weak instruments. In the case of strong instruments, we provide results, not previously explicitly available in the literature, that permit comparison of the various estimators, and we highlight the trade-off between efficiency and robustness. We pay particular attention to methods that are robust to weak instruments. The expectation of future inflation is the key endogenous covariate in the NKPC. Because inflation is notoriously hard to forecast, it is difficult to find exogenous (i.e., lagged) economic variables that correlate strongly with expected future inflation; in other words, potential instruments that satisfy the exclusion restriction will likely be weak. We show that when this is the case, even estimators that do not explicitly rely on instrumental variable techniques can be severely biased. Hence, weak instrument issues provide a unifying explanation of the sensitivity of NKPC estimates and of the puzzling disagreement between analyses based on standard inference procedures. We also discuss complications arising from misspecification of the NKPC.

Section 4 surveys the vast empirical literature on the NKPC, covering over 100 papers we have found on this topic. The initial success of the rational expectations specification (with the labor share as the proxy for firm marginal cost) around the turn of the millennium was quickly followed by doubts about robustness to data choices and estimation methods. A plethora of extensions of the basic model have been pursued with no clear universal consequences. Approaches that

exploit the non-rationality of expectations or positive trend inflation have recently been gaining traction, while a parallel strand of the literature has emphasized the weak identification issues inherent in estimating the NKPC. Due to differences across papers in data sets, instrument choices, specifications, estimators and attention to the weak identification problem, no consensus has been reached on parameter values or the reasons for the variability in estimates, and policy implications are entirely unclear. To date, few papers have sought to compare or integrate more than a couple of the empirical strategies, leaving the research program in considerable disarray.

Section 5 provides a new set of empirical results based on a common data set and a flexible empirical strategy that spans multiple popular approaches in the literature. Like most papers, we focus on quarterly post-war U.S. data. Apart from the standard data series, we have assembled a unique real-time data set on the labor share. By computing point estimates of the NKPC parameters from a comprehensive set of specifications that combine data and model choices from the literature, we find that the specification uncertainty is vast: Almost any parameter combination that is even remotely close to the range considered in the literature can be generated by some a priori unobjectionable specification. Furthermore, given a particular specification, the sampling uncertainty is large. We show this by computing weak identification robust confidence sets for several benchmark specifications. One type of specification that appears to be typically better identified uses survey forecasts as proxies for inflation expectations; however, such specifications are only microfounded if survey forecasts are rational, which does not seem to be the case empirically. Survey specifications are also less suitable for counterfactual policy analysis and forecasting.

Section 6 concludes by summarizing the main lessons from the literature and our empirical exercise. We recommend that future research pursue substantially new types of data sets, as well as estimation approaches that are tailored to handle the identification problem.

2 Economic Theory, Specifications and Policy Implications

The origins of the NKPC can be traced back to the late seventies, in the work of Fischer (1977) and Taylor (1979). The NKPC is a forward-looking model of inflation, according to which current inflation is determined by expected future inflation and marginal costs. It implies that monetary policy can affect inflation through the management of inflation expectations. This contrasts sharply with the traditional ‘old’ Phillips curve, which yields a strongly path dependent inflation process, so that disinflation can be slow and costly (Mankiw, 2001). The importance of expectations was highlighted early on by Phelps (1967) and Friedman (1968). Their so-called ‘expectations-augmented’ Phillips curves emphasized that the inflation/unemployment trade-off shifts with expected inflation, a property shared by the Phillips curves of the New Classical literature of the 1970s (e.g., Sargent and Wallace, 1975). The key difference of these Phillips curves from the NKPC is that past (and thus predetermined) expectations about current inflation matter, not expectations about the future.

2.1 Economic foundations of the model

A simple derivation of the NKPC can be obtained as follows.³ The basic ingredients for the derivation of the NKPC is a microeconomic environment with identical monopolistically competitive firms facing constraints on price adjustment. We consider here only time-contingent pricing constraints. The details of the constraints do not matter much for the final form of the NKPC (Roberts, 1995), so we focus on the assumption of Calvo (1983), which is analytically attractive. Prices are expressed in logs and inflation in percentage points. All variables, except prices and inflation, are expressed in percent deviations from a zero-inflation steady state. We discuss the assumption of zero steady-state inflation below. The Calvo framework assumes that each firm in the economy has a constant probability $1 - \theta$ of optimally adjusting its price in any given period. Because the economy consists of a continuum of identical firms, by the law of large numbers it follows that a fraction θ of firms cannot change their prices in any given period, and that prices remain fixed on average for $1/(1 - \theta)$ periods. Therefore, the parameter $\theta \in (0, 1)$ is an index of price rigidity. Assume that each firm produces a differentiated product and faces a constant price elasticity of demand $\varepsilon > 1$ for its product. Let $p_{i,t}^*$ denote the optimal price chosen by a firm $i \in [0, 1]$ if it gets to reoptimize in period t . By the law of large numbers, the aggregate price level p_t evolves as a convex combination of last period's price and the cross-sectional average of the current reset prices:

$$p_t = \theta p_{t-1} + (1 - \theta) \int_0^1 p_{i,t}^* di. \quad (1)$$

In the absence of price rigidity, monopolistically competitive firms set prices as a markup over their nominal marginal costs. With price rigidity, maximization over all expected discounted future profits induces firms to take into account the probability that they will not be able to reset their prices optimally in the future. Let β denote the common subjective discount factor. The optimal reset price can, in a first-order log-linearization, be expressed as a mark-up over a weighted average of current and expected future marginal costs:

$$p_{i,t}^* = (1 - \theta\beta) \sum_{j=0}^{\infty} (\theta\beta)^j E_{i,t} (mc_{i,t+j,t}^n), \quad (2)$$

where $mc_{i,t+j,t}^n$ is the nominal marginal cost faced at time $t + j$ for a firm i that was last able to reset its price optimally at time t , and $E_{i,t}$ denotes the expectation with respect to the beliefs of firm i . Relating this to the aggregate marginal cost mc_t^n requires a specification of the production function. Suppose the production function is Cobb-Douglas with labor elasticity $1 - \alpha$. Then it can be shown that, to a first-order approximation,

$$mc_{i,t+j,t}^n = mc_{t+j}^n + \frac{\alpha\varepsilon}{1 - \alpha} (p_{i,t}^* - p_{t+j}). \quad (3)$$

Substituting (3) into (2) and rearranging yields

$$p_{i,t}^* - p_{t-1} = (1 - \theta\beta) \sum_{j=0}^{\infty} (\theta\beta)^j [\kappa E_{i,t} (mc_{t+j}^n) + E_{i,t} (\pi_{t+j})], \quad (4)$$

---

³For detailed derivations, see Woodford (2003, ch. 3) or Galí (2008, ch. 3).

where $\kappa = \frac{1-\alpha}{1-\alpha+\alpha\epsilon} \leq 1$, inflation is given by $\pi_t = p_t - p_{t-1}$, and $mc_t = mc_t^n - p_t$ denotes aggregate real marginal costs. Inserting (4) into (1), we find

$$\pi_t = (1-\theta)(1-\theta\beta) \sum_{j=0}^{\infty} (\theta\beta)^j \left[\kappa \hat{E}_t(mc_{t+j}) + \hat{E}_t(\pi_{t+j}) \right], \quad (5)$$

where $\hat{E}_t = \int_0^1 E_{i,t} di$ is the *cross-sectional average expectation operator*.

Until this point we have not imposed any restrictions on the nature of firms' beliefs about future economic conditions. Assume now that firms have identical, rational expectations (RE), i.e., $E_{i,t} \equiv E_t$. Then the cross-sectional expectation equals the rational expectation, $\hat{E}_t = E_t$. If we shift equation (5) forward by one period and take time- t expectations on both sides, we get

$$E_t(\pi_{t+1}) = (1-\theta)(1-\theta\beta) \sum_{j=0}^{\infty} (\theta\beta)^j [\kappa E_t(mc_{t+j+1}) + E_t(\pi_{t+j+1})], \quad (6)$$

where we have, importantly, used the law of iterated expectations, $E_t[E_{t+1}(\cdot)] = E_t(\cdot)$. The infinite sum on the right-hand side of (6) is closely related to the infinite sum on the right-hand side of (5) (with $\hat{E}_t = E_t$). Combining these two equations gives rise to an expectational difference equation for inflation:

$$\pi_t = \beta E_t(\pi_{t+1}) + \frac{(1-\theta)(1-\theta\beta)}{\theta} \kappa mc_t. \quad (7)$$

Empirical analyses of the NKPC use proxies for the real marginal cost measure mc_t . Under the already exploited assumption of Cobb-Douglas production technology, mc_t is proportional both to the labor share of income (nominal labor compensation divided by nominal output) and the output gap (the deviation of real output from the level that would obtain if prices were fully flexible). Letting x_t denote a candidate proxy for real marginal cost and adding an unrestricted unobserved disturbance term u_t , we can rewrite the model as

$$\pi_t = \beta E_t(\pi_{t+1}) + \lambda x_t + u_t. \quad (8)$$

This is the baseline purely forward-looking NKPC that we will refer to in subsequent sections. The disturbance term u_t can be interpreted as measurement error or any other combination of unobserved cost-push shocks, such as shocks to the mark-up or to input (e.g., oil) prices.

Non-rational expectations While we did not need to impose assumptions on firms' individual beliefs to arrive at the expression (5) for inflation, the derivation of equation (6) – and thus the difference equation (7) – crucially relied on the fact that the rational expectation operator E_t satisfies the law of iterated expectations. However, this law does not hold in general for the cross-sectional average expectation operator $\hat{E}_t$, even if individual firm expectations do satisfy the law. Hence, under general non-rational expectation formation, the difference equation specification of the NKPC in equation (7) is not consistent with the above microeconomic foundations that constitute the standard new Keynesian modeling framework (a similar argument is given by Preston, 2005).⁴

---

⁴It is possible for non-RE models to result in a difference equation NKPC. Adam and Padula (2011) derive a difference equation under a set of assumptions on the possibly non-rational cross-sectional expectation operator; their “Condition 1” essentially invokes the law of iterated expectations at the aggregate level.

This will be the case, for instance, if firms' expectations are not based on the same information set or if they are not perfectly model-consistent. As we discuss in section 3.1, this has implications for empirical tests of the NKPC that use survey forecasts to proxy for the expectation term. Preston (2005), Angeletos and La'O (2009) and Kurz (2011) have derived microfounded inflation equations in certain models with non-rational or heterogeneous expectations.

2.2 Extensions

It was recognized early on that the purely forward-looking NKPC (8) has difficulty fitting aggregate US inflation dynamics, see Fuhrer and Moore (1995) and Galí and Gertler (1999). This led to specifications that included lagged inflation terms in the model. This is often called 'intrinsic inflation persistence'. Galí and Gertler (1999) introduce lagged terms by assuming that a fraction of firms update their prices using some backward-looking rule of thumb, while Fuhrer and Moore (1995) generate persistence through staggered relative wage contracts. Another popular device is to assume that the fraction θ of firms that are unable to re-optimize their prices in the Calvo model instead index prices to past inflation, see Christiano et al. (2005) and Sbordone (2005, 2006).⁵

Because it could be thought of as a combination of new and old Phillips curves, such a specification is referred to as a 'hybrid NKPC'. In principle, if the objective is to nest traditional Phillips curves, one could allow for any number of lagged inflation terms in the model. An appropriate baseline hybrid specification that nests most traditional Phillips curve specifications would take the form

$$\gamma(L)\pi_t = \gamma_f E_t(\pi_{t+1}) + \lambda x_t + \eta' w_t + u_t, \quad (9)$$

where $\gamma(L) = 1 - \gamma_1 L - \gamma_2 L^2 - \dots - \gamma_l L^l$ is a lag polynomial, x_t is the main forcing variable, w_t denotes additional controls, and u_t is an unobserved shock. If the lag polynomial only features one lag, the coefficient γ_1 is often denoted γ_b . Equation (9) nests the pure NKPC (8) with $\gamma(L) = 1$ and $\eta = 0$. With $\gamma_f = 0$, it also nests the backwards-looking "old" Phillips curve, and in particular, Gordon's (1990) "triangle" model. It is more general than the typical hybrid NKPC specifications that only include one lag of inflation, such as Galí and Gertler (1999), Sbordone (2005, 2006) and Christiano et al. (2005). The latter are based on indexation to last quarter's inflation, but this is clearly arbitrary and can be easily generalized to include more general indexation schemes, such as a weighted average of inflation over the previous four quarters (this includes, as a special case, indexation to last *year's* inflation, which nests the Atkeson and Ohanian, 2001, specification), or richer rule-of-thumb behavior by backward-looking firms, see Zhang and Clovis (2010). Restrictions on $\gamma(L)$ are exclusion restrictions on the dynamics of inflation, which are typically used to provide instruments for identification. Therefore, such exclusion restrictions are not innocuous. A popular restriction, which is seldom rejected by the data, is that the inflation coefficients sum to 1, i.e., $\gamma_f = 1 - (\gamma_1 + \gamma_2 + \dots + \gamma_l) = \gamma(1)$.

The parameters of equation (9) are often referred to as 'reduced-form' or 'semi-structural' because they are functions of the deeper structural parameters of the microfounded model. For example, when the discount factor is one, the Galí and Gertler (1999) specification has $\gamma_f = \theta / (\theta + \omega)$

---

⁵Sheedy (2010) and Yao (2011) observe that intrinsic persistence arises when the hazard rate of price resetting (the Calvo parameter) is not constant. See also Kozicki and Tinsley (2002) for further discussion on the sources of lagged inflation dynamics in the NKPC.

and $\gamma_1 = 1 - \gamma_f$, where ω is the fraction of price setters who are backward-looking. Restrictions on the admissible range of the deep parameters can affect the range of the semi-structural ones and thus have nontrivial implications for inference.

Trend inflation The derivation of equation (8) follows from log-linearizing firms' optimizing conditions around a zero-inflation steady state. Allowing for non-zero steady state inflation, often referred to as "trend inflation", has important implications for the specification of the NKPC, as established by Kozicki and Tinsley (2002), Ascari (2004) and Cogley and Sbordone (2008). Trend inflation $\bar{\pi}_t$ corresponds to long-run inflation expectations, i.e., $\bar{\pi}_t = \lim_{T \rightarrow \infty} E_t(\pi_{t+T})$. With non-zero trend inflation, the NKPC cannot in general be written in the difference equation form (9), as extra forward-looking terms enter on the right-hand side and the semi-structural parameters are functions of trend inflation. However, Cogley and Sbordone (2008, p. 2105) show that if non-resetting firms' prices are indexed to a mixture of past inflation π_{t-1} (weight ρ) and current trend inflation $\bar{\pi}_t$ (weight $1 - \rho$), then

$$\hat{\pi}_t = \gamma_f E_t \hat{\pi}_{t+1} + \gamma_b \hat{\pi}_{t-1} + \lambda x_t + \gamma_b (\beta E_t \Delta \bar{\pi}_{t+1} - \Delta \bar{\pi}_t), \quad (10)$$

where $\gamma_f = \beta/(1 + \beta\rho)$, $\gamma_b = \rho/(1 + \beta\rho)$, and we define the *inflation gap* $\hat{\pi}_t = \pi_t - \bar{\pi}_t$. If trend inflation is constant, $\bar{\pi}_t \equiv \bar{\pi}$, the last term above drops out and we are left with a standard NKPC in which the inflation gap replaces raw inflation.⁶ If furthermore $\beta = 1$, then $\gamma_f + \gamma_b = 1$, so the NKPC can be expressed in terms of the change in inflation $\Delta\pi_t = \Delta\hat{\pi}_t$, causing the constant trend inflation to drop out of the relation altogether. Suppose instead that trend inflation is time-varying. If $\beta = 1$ and $E_t \Delta \bar{\pi}_{t+1} = \Delta \bar{\pi}_t$ (the change in trend inflation is unforecastable), the last term on the right-hand side of (10) vanishes and we are left with an NKPC relation in terms of the inflation gap. Such a specification also obtains if non-reset prices are fully indexed to trend inflation ($\rho = 0$).

In the rest of this paper we will focus on inflation gap specifications of the NKPC, i.e., relations of the form (10) without the last term on the right-hand side. We do this to keep the exposition simple but acknowledge that we do not give the trend inflation issue as much attention as it deserves. Interested readers are referred to the review by Ascari and Sbordone (2013).

2.3 Policy implications

Ideally, estimation uncertainty in the NKPC parameters would only translate into limited ambiguity about our understanding of the effects of shocks and policy interventions on the broader economy. Unfortunately, this is not the case for the range of NKPC parameter estimates reported in the literature. Figure 1 displays impulse responses of inflation and the output gap to a 25 basis point monetary policy shock in the canonical three-equation new Keynesian model (Galí, 2008). The model and calibration are described in section A.1 in the Appendix, and they are based on a hybrid NKPC with one lag of inflation whose coefficient γ_1 is equal to $1 - \gamma_f$. By a monetary policy shock we mean a shock to the innovation in the AR(1) process for the Taylor rule disturbance. We treat the semi-structural parameters γ_f and λ as being variation free from the remaining structural parameters (which are calibrated as in Galí, 2008, ch. 3.4) and vary them over a set of values that

---

⁶This is equivalent to an NKPC written in terms of raw inflation π_t with an intercept that depends on $\bar{\pi}$, as established by Yun (1996) for the special case $\rho = 0$.

is consistent with the spread of estimates reported in the literature, cf. section 4 below, namely $\gamma_f = 0.3, 0.4, \dots, 0.8$ and $\lambda = 0.01, 0.03, 0.05$. As Figure 1 shows, this leads to a wide range of possible impulse responses, with substantially different short-run dynamics and steady state return times. For given λ , lower values of γ_f imply more sluggish adjustment and more hump-shaped dynamics in inflation. The disparity between the high- γ_f and low- γ_f dynamics increases the lower is λ . For $\lambda = 0.03$ (the thick curves in the figure), the effect of the monetary policy shock is felt for about twice as long for $\gamma_f = 0.3$ than for $\gamma_f = 0.8$. The most negative *cumulative* 15-quarter inflation impulse response in Figure 1 is 5 times larger in magnitude than the least negative cumulative inflation response; the ratio between the most and least negative cumulative output gap responses is 5.7. This sensitivity of key economic measures to the NKPC parameters extends beyond the simple model considered here, as demonstrated by Fuhrer (1997), Mankiw (2001) and Estrella and Fuhrer (2002).⁷ If indeed the NKPC is a good approximation to actual price setting, it is therefore highly desirable from a policy perspective to obtain precise estimates of the NKPC coefficients.⁸

3 Econometric Methods

In this section we describe the main estimators that have been used in the literature and discuss their properties under strong and weak identification. We focus on the “semi-structural” parametrization of the NKPC, as opposed to the underlying structural parameters, to facilitate comparison across different specifications.⁹ For ease of exposition, we focus in this section on estimation of the pure NKPC (8), but all our points generalize to the hybrid specification (9).

3.1 Estimators

A glance at the NKPC (8) reveals two immediate estimation issues. First, as noted by Roberts (1995), the forcing variable x_t may be correlated with the structural error term u_t (e.g., they may both be driven in part by cost-push shocks). Second, the inflation expectation term $E_t(\pi_{t+1})$ is certainly endogenous, and – even worse – it is unobservable. Different empirical approaches in the literature differ mainly in the way they deal with inflation expectations. They can be usefully categorized as follows:

1. Replace expectations by realizations and use appropriate instruments (GIV).
2. Derive expectations from a particular reduced-form model (VAR).
3. Use direct measures of expectations (Survey).

---

⁷Adding estimation uncertainty in the non-NKPC equations to the mix will of course generate even greater uncertainty about appropriate impulse responses.

⁸The degree of inflation indexation also influences the relative optimality of policies that flexibly target the price level or inflation rate (Woodford, 2003, ch. 8.2.1).

⁹Moreover, estimation of the structural parameters raises some additional issues if the mapping from the semi-structural to structural parameters is not injective, so that the latter are not globally identified even when the former are (Ma, 2002).

We also propose a new strategy based on the use of data revisions as instruments. This can be thought of as a variant of the first approach. All of these approaches can be implemented using the Generalized Method of Moments (Hansen, 1982). We use GMM as a common unifying framework in our discussion and empirical work, which is convenient because weak identification robust methods are readily available for GMM. GMM estimation is briefly described in section A.2.1 in the Appendix.

GIV This approach was originally proposed for the estimation of rational expectation (RE) models by McCallum (1976). Hansen and Singleton (1982), who studied estimation of Euler equation models, called it *Generalized Instrumental Variable* estimation (GIV). It has been popularized in the estimation of the NKPC by the seminal contributions of Roberts (1995) and Galí and Gertler (1999). It is the most common approach in the literature because it is simple to implement and more robust than the alternatives. Identification is obtained via exclusion restrictions, i.e., excluding lags of variables from the model and using them as instruments.

The simplest and most common implementation is to replace the rational expectation $E_t(\pi_{t+1})$ in the difference equation (8) by the realization π_{t+1} . This yields

$$\pi_t = \beta\pi_{t+1} + \lambda x_t + \underbrace{u_t - \beta[\pi_{t+1} - E_t(\pi_{t+1})]}_{\tilde{u}_t}, \quad (11)$$

where the residual $\tilde{u}_t$ differs from u_t because it includes the future one-step-ahead inflation forecast error. Let $\vartheta = (\beta, \lambda)$ and define the ‘residual’ function

$$h_t(\vartheta) = \pi_t - \beta\pi_{t+1} - \lambda x_t. \quad (12)$$

Suppose Z_t is a vector of valid instruments, such that

$$E[Z_t h_t(\vartheta)] = 0 \quad (13)$$

holds at $\vartheta = \vartheta_0$, the true parameter value. Efficient GMM estimation (see section A.2.1 in the Appendix) is based on the sample moments $f_T(\vartheta) = T^{-1} \sum_{t=1}^T Z_t h_t(\vartheta)$ and a heteroskedasticity and autocorrelation consistent (HAC) estimator of their variance, because $h_t(\vartheta_0)$ is generally autocorrelated due to the presence of the inflation forecast error.

The most common identifying assumption in the literature, which is the one used by Galí and Gertler (1999), is that the cost-push shock u_t satisfies $E_{t-1}(u_t) = 0$. Under the RE assumption, this implies $E_{t-1}(\tilde{u}_t) = 0$ by the law of iterated expectations. This yields unconditional moment restrictions of the form (13) with $Z_t = Y_{t-1}$, for any vector of predetermined (i.e., known at time $t-1$) variables Y_{t-1} . Any predetermined variables can be used as instruments, and different implementations of GIV estimation differ mainly in the choice of instruments. For example, Rudd and Whelan (2005) obtain alternative GMM moment conditions by iterating the NKPC forward. As we show in section A.2.2 in the Appendix, this is equivalent to estimating the original NKPC difference equation (8) with transformed instruments.

VAR Almost all of the papers that use the second approach rely on the assumption that the reduced-form dynamics of the variables can be represented by a finite-order vector autoregression (VAR). This is why we refer to it as the VAR approach.

Suppose the information set consists of current and lagged values of some n -dimensional vector z_t (this includes at least π_t, x_t and any other variables that are to be used as instruments), and that z_t admits a finite-order VAR representation of order l , which can be written in companion form as

$$Y_t = AY_{t-1} + V_t, \quad (14)$$

where Y_t, V_t are $nl \times 1$ and A is $nl \times nl$. If $E_t(V_{t+1}) = 0$, equation (14) implies that $E_t(\pi_{t+1}) = \zeta'Y_t$, where ζ is the row of A that corresponds to the inflation equation in the VAR. Substituting this in the NKPC (8), yields the moment conditions $E[(\pi_t - \beta Y_t' \zeta - \lambda x_t) Y_{t-1}] = 0$, which determine the structural parameters $\vartheta = (\beta, \lambda)$ given ζ , and ζ is identified by the reduced-form equation $E[(\pi_t - Y_{t-1}' \zeta) Y_{t-1}] = 0$. The VAR assumption suggests using Y_{t-1} as instruments, so the model can be estimated by GMM with the $2nl$ moment conditions

$$E[\tilde{h}_t(\vartheta_0, \zeta_0) \otimes Y_{t-1}] = 0, \quad \text{where} \quad (15)$$

$$\tilde{h}_t(\vartheta, \zeta) = \left(\pi_t - \beta Y_t' \zeta - \lambda x_t : \pi_t - Y_{t-1}' \zeta \right)'. \quad (16)$$

Here ζ_0 is the true value of ζ . The seminal papers in this strand of the literature are Fuhrer and Moore (1995) and Sbordone (2002), and they use two different econometric implementations: maximum likelihood (VAR-ML) and minimum distance (VAR-MD), respectively. We describe these methods in section A.2.3 in the Appendix. The realization that the VAR assumption to identification can be easily imposed in GMM, which we refer to as VAR-GMM, appears to be new in the literature.¹⁰ VAR-GMM, VAR-ML and VAR-MD are numerically identical if the model is just-identified, i.e., if the dimension of Y_t is the same as the dimension of the parameter vector ϑ .

Bayesian inference is mostly used for full-information analysis of DSGE models (e.g., Lubik and Schorfheide, 2004), which is beyond the scope of the present survey. A few notable limited-information Bayesian studies of the NKPC are reviewed in section 4. These all rely on the VAR assumption.

Surveys This approach was introduced by Roberts (1995) and, after a lag, has seen increasing popularity. Under the survey approach, direct measures of expectations from surveys, e.g., the Survey of Professional Forecasters or the Federal Reserve's "Greenbooks", are used as proxies for inflation expectations in the NKPC. Let $\pi_{t+j|t}^s$ denote the j -step-ahead survey forecast of inflation at time t . The most common implementation substitutes the one-quarter-ahead forecast $\pi_{t+1|t}^s$ for $E_t(\pi_{t+1})$ in the NKPC (8) to get

$$\pi_t = \beta \pi_{t+1|t}^s + \lambda x_t + u_t + \beta \varepsilon_t, \quad \text{where} \quad (17)$$

$$\varepsilon_t \equiv E_t(\pi_{t+1}) - \pi_{t+1|t}^s. \quad (18)$$

The survey error ε_t can be a combination of measurement error and news shocks, the latter arising when survey responses are based on a smaller information set than the one the agents in the model

---

¹⁰Guerrieri et al. (2010) use a VAR-GMM procedure to estimate an open-economy NKPC, but they do not compare their estimates to non-VAR specifications. See also Sbordone (2005) for a discussion of the relation between VAR-MD and GIV.

use. Identification depends on the properties of this survey error ε_t as well as on the correlation between $\pi_{t+1|t}^s$ and the cost-push shock u_t . Some authors treat $\pi_{t+1|t}^s$ as exogenous, but we argue below that the assumptions underlying this are too strong. Alternatively, one may use $\pi_{t+1|t-1}^s$, which is certainly predetermined, instead of $\pi_{t+1|t}^s$, which is typically measured within the quarter. Another possibility is to treat survey forecasts as endogenous and use predetermined variables as instruments, as in the GIV approach.

Because *estimation* of the NKPC using survey forecast data obviates the need for modeling inflation expectations, some authors interpret the procedure as allowing for non-rational price setting. When the NKPC is to be used for policy purposes, such as in forecasting, the lack of a dynamic model for expectations becomes a disadvantage. To date, few papers have attempted to model non-rational survey expectations formation (see Fuhrer, 2012). As argued in section 2.1, the difference equation NKPC (8) is in general inconsistent with the standard new Keynesian framework if firm expectations are not model-consistent or if they are based on dispersed information. Consequently, survey specifications of the form (17) are only microfounded if the inflation forecasts firms rely on are rational and identical across firms. However, common empirical findings are that survey forecasts of inflation violate testable implications of rationality and the dispersion of expectations across individual forecasters is large, see for example Thomas (1999) and Mankiw et al. (2004). This point does not seem to have been taken to heart by the empirical NKPC literature. While the proper microfoundations for price setting under non-rational expectation formation are lacking, the survey forecast specification may still be taken as a primitive, which nicely summarizes our intuition about price setting being partially forward-looking, partially backward-looking as well as being responsive to aggregate demand conditions.

Nunes (2010), Fuhrer and Olivei (2010) and Fuhrer (2012) estimate versions of the NKPC where inflation expectations are specified as a combination of rational and survey expectations, replacing $E_t(\pi_{t+1})$ with $\phi E_t(\pi_{t+1}) + (1 - \phi) \pi_{t+1}^s$ in the NKPC (8).¹¹ The parameter ϕ is not identified if survey expectations are rational.

External instruments Many of the time series typically used to estimate the NKPC, such as GDP deflator inflation, the labor share and the output gap, undergo large revisions over time. Because firms' expectations of current and future economic conditions feature prominently in the new Keynesian model, it is crucial to keep in mind the availability of data at different points in time when estimating the NKPC. With access to real-time data, i.e., data sets of different vintages, one possible identification strategy is to use past data vintages of inflation and the forcing variable (and perhaps other time series) as instruments.

We appear to be the first to consider estimation of the NKPC using such external instruments. An advantage of the approach is that it avoids placing exclusion restrictions on the NKPC. This is discussed more formally in section A.2.4 of the Appendix. Regardless of the number and types of right-hand side variables in the hybrid specification (9), past-vintage instruments should not affect current inflation if we control for the latest-vintage data, and thus are plausibly uncorrelated with the GIV error $\tilde{u}_t$ in (11). Hence, real-time instruments are plausibly exogenous, although they

¹¹Nunes (2010) argues that this nesting can be motivated by a variant of the argument of Galí and Gertler (1999), where a fraction of rule-of-thumb firms set their prices using published professional inflation forecasts as opposed to lagged inflation. This justification cannot apply to the Federal Reserve's Greenbook forecasts, as they are only released to the public with lags of several years.

could potentially be very weak. In section 5 we empirically evaluate the success of the external instruments approach.

3.2 Comparison of estimators under conventional asymptotics

To compare the properties of the sampling distributions of the various estimators, we start out by outlining the trade-offs between efficiency and robustness under the conventional asymptotic approximations. The conventional asymptotic theory, which is the main analytical tool in graduate econometrics textbooks, implies that GMM estimators of the parameters ϑ are consistent and asymptotically normal under certain regularity conditions, see Newey and McFadden (1994). In linear instrumental variable (IV) models this asymptotic theory has the first-stage F statistic, which measures the explanatory power of the instruments for the endogenous variable, tending to infinity at the same rate as the sample size, so we refer to the theory as *strong-instrument asymptotics*. In section 3.3 below we will argue that the strong instrument approximation is not empirically relevant in the present context, as it does not capture the kind of near-irrelevance of the instruments that is prevalent in the estimation of the NKPC. Still, it is useful to first establish the properties of the estimators in the most familiar analytical framework.

Most of the estimators are obtained from conditional moment restrictions of the form $E_{t-1}[h_t(\vartheta_0)] = 0$, for which the theory of optimal instruments of Chamberlain (1987) provides an efficiency bound. $E_{t-1}[h_t(\vartheta_0)] = 0$ implies that all predetermined variables Y_{t-j} , $j \geq 1$, are admissible instruments. The optimal choice of instruments is derived in section A.2.1 of the Appendix, and we summarize the findings below.

GIV For GIV estimators, the residual function $h_t(\vartheta_0) = u_t - \beta[\pi_{t+1} - E_t(\pi_{t+1})]$ is generally autocorrelated, so optimal instruments are given by an infinite-order moving average of Y_{t-1} . Because their derivation requires modeling the conditional mean and variance of the data, none of the papers in this literature has attempted to use optimal instruments.¹² Therefore, we cannot rank GIV estimators reviewed in this paper in terms of efficiency. Indeed, claims about their relative efficiency are not formally justified.

VAR In contrast, the VAR-GMM estimator is based on the moment conditions (15). Let $v_{\pi t}$ denote the VAR error term in the reduced-form inflation equation in (14). When the VAR assumption holds, the VAR residuals $\tilde{h}_t(\vartheta_0, \zeta_0) = (u_t, v_{\pi t})'$ are serially uncorrelated, unlike the GIV residuals $h_t(\vartheta_0)$ in (12). If the VAR residuals are also conditionally homoskedastic, then the VAR-GMM estimator does indeed use optimal instruments, and is therefore asymptotically more efficient than the corresponding GIV estimators that do not impose the VAR assumption. But note that the key conditions for this, strong instruments and conditional homoskedasticity, are arguably too strong in this application.

Fuhrer and Moore (1995), and several later papers, estimate the NKPC by ML. To evaluate the likelihood, they combine the NKPC with reduced-form equations for all variables other than

---

¹²Fuhrer and Olivei (2005) propose a GMM implementation of a restricted VAR approach, which they refer to as producing optimal instruments. Because these instruments are linear combinations of Y_{t-1} , they are not optimal in Chamberlain's sense, cf. section A.2.1.

inflation to form a complete ‘limited-information’ system of equations, and use an algorithm by Anderson and Moore (1985) that finds a RE solution for any given value of the parameters. For certain parameter combinations there may be multiple stable RE solutions, a situation known as indeterminacy (see, e.g., Lubik and Schorfheide, 2004), and so the likelihood is not uniquely determined by the NKPC and remaining reduced-form parameters. When the solution is unique (determinacy), the reduced-form is a (restricted) finite-order VAR. The Fuhrer-Moore approach restricts the parameter space to the region in which indeterminacy does not occur. The determinacy assumption is standard in full-information estimation of DSGE models – for example, it is imposed by Dynare (Adjemian et al., 2011) – but it can be restrictive.¹³ Kurmann (2007) proposes a simple method for evaluating the limited-information likelihood under the assumption that the reduced form is a finite-order VAR without imposing determinacy, see section A.2.3 in the Appendix for details. He demonstrates in an empirical application that his method can give very different results from the Fuhrer-Moore approach. This finding does not suffice to infer that the determinacy assumption is incorrect, as the estimates can differ due to sampling uncertainty.¹⁴

The other prominent approach to imposing the VAR assumption on the reduced-form dynamics is VAR-MD, see Sbordone (2005). As explained in section A.2.3 of the Appendix, the difference between VAR-ML and VAR-MD is that the former uses a restricted and the latter an unrestricted estimator for the reduced-form VAR parameters. Therefore, the relationship between VAR-ML and VAR-MD is analogous to the relationship between limited information maximum likelihood (LIML) and two stage least squares (2SLS), respectively, in the linear IV regression model (Fukač and Pagan, 2010).

The assumption that the dynamics of the data can be represented as a finite-order VAR is restrictive. One well-known case when this assumption fails is when there is indeterminacy and sunspots, and the reduced form has moving average (MA) components. When the MA roots are large (i.e., nearly noninvertible), a finite-order VAR may produce an inaccurate representation of the dynamics. Infinite-order VARs may also arise for other reasons, e.g., by omitting relevant variables from a finite-order VAR.¹⁵ Therefore, the VAR approach to identification is less robust than GIV. To gain intuition about the restrictiveness of the VAR assumption, it is useful to think of the analogy to iterated versus direct forecasts, a point made in Magnusson and Mavroeidis (2010). The VAR-GMM moment conditions (15) differ from the GIV moment conditions $E[Y_{t-1}(\pi_t - \beta\pi_{t+1} - \lambda x_t)] = 0$, in that GIV uses direct projections of future inflation on predetermined variables, while VAR-GMM uses iterated multi-step forecasts from a VAR.

Surveys The two alternative survey estimators that we consider here are those that use one-step-ahead forecasts of inflation and those that use lagged two-steps-ahead forecasts. The former

¹³Kurmann (2007) shows that even if the equilibrium of the underlying structural model of the economy is determinate, a limited-information system of equations, where some of the structural equations have been replaced by their reduced form, may have an indeterminate solution. See Kurmann (2007, sec. 3.2) for an example.

¹⁴Kurmann (2007) does not formally test the hypothesis that his approach fits the data significantly better than the Fuhrer-Moore approach. It is not trivial to develop a test for this hypothesis, especially when the structural parameters may be weakly identified. Note that if the model restrictions hold and the true parameters imply determinacy, the Fuhrer-Moore estimator (which imposes a correct restriction) will be more efficient than the Kurmann estimator under conventional asymptotics.

¹⁵See also Fernández-Villaverde et al. (2007). But note that omission of variables from a VAR does not necessarily cause misspecification (Fukač and Pagan, 2007).

substitutes $\pi_{t+1|t}^s$ for $E_t(\pi_{t+1})$ in (8), yielding equations (17)–(18). The second possibility is to substitute $\pi_{t+1|t-1}^s$ for $E_t(\pi_{t+1})$ in (8) to get

$$\pi_t = \beta \pi_{t+1|t-1}^s + \lambda x_t + u_t + \beta \tilde{\varepsilon}_t, \quad (19)$$

$$\tilde{\varepsilon}_t = [E_t(\pi_{t+1}) - E_{t-1}(\pi_{t+1})] + [E_{t-1}(\pi_{t+1}) - \pi_{t+1|t-1}^s]. \quad (20)$$

The first component of $\tilde{\varepsilon}_t$ is orthogonal to $t-1$ information, and the second component has the same interpretation as the one-step survey error in (18). The difference from the previous case is that $\pi_{t+1|t-1}^s$ is certainly predetermined, so it can be more plausibly treated as exogenous.

Some studies treat survey forecasts as exogenous for the estimation of the NKPC (Rudebusch, 2002; Adam and Padula, 2011). This can be justified under very specific assumptions about the timing of expectations and the nature of the disturbance term in the model. For example, if there is no cost-push shock, i.e., $u_t = 0$ in (17), and the disturbance term ε_t defined in (18) is a pure news shock, then $\pi_{t+1|t}^s$ in equation (17) will be exogenous. This will not be true if ε_t is a classical measurement error. When $u_t \neq 0$, exogeneity of $\pi_{t+1|t}^s$ requires that it should be predetermined, i.e., measured *before* π_t , but survey data are actually collected within the quarter. Equation (19) overcomes this problem, because $\pi_{t+1|t-1}^s$ is certainly predetermined, but exogeneity still requires that $\pi_{t+1|t-1}^s$ must be uncorrelated with $\tilde{\varepsilon}_t$ in (20). This will hold if expectations are rational and $[E_{t-1}(\pi_{t+1}) - \pi_{t+1|t-1}^s]$ is a news shock. Of course, even if the survey forecast is exogenous, the forcing variable x_t may still be endogenous.

If survey forecasts or the forcing variable are endogenous, then we need to find instruments. If measurement errors are unsystematic, in the sense that they are unpredictable from information at time $t-1$, and survey forecasts are unbiased, i.e., rational based on their information set, then $E_{t-1}(\varepsilon_t) = 0$ in (18). Therefore, moment conditions for (17) are the same as for the GIV approach (13) where π_{t+1} has been replaced with $\pi_{t+1|t}^s$ and the instruments are predetermined variables, including perhaps lags of $\pi_{t+1|t}^s$. Our view is that it is more robust to treat survey data as endogenous and use instruments for them. Nevertheless, we study the empirical implications of treating survey forecasts as exogenous in section 5.

3.3 Weak identification

Identification of the structural parameter vector ϑ requires that the GMM moment conditions are satisfied only at the true value ϑ_0 . Identification is clearly a necessary condition for obtaining useful estimators of ϑ , but it is not sufficient. *Weak identification* arises when ϑ is almost unidentified, i.e., when the moment conditions are close to being satisfied for all parameters ϑ in a non-vanishing neighborhood of the true value ϑ_0 . In instrumental variable settings, weak identification arises when the instruments are only weakly correlated with the endogenous regressors. When identification is weak, conventional strong instrument asymptotic theory provides a poor approximation to the sampling distribution of GMM estimators and tests, even in large samples. Instead, estimators can be very non-normally distributed and severely biased toward their OLS (or NLS) counterparts, while conventional confidence sets may drastically undershoot their advertised coverage rates. Kleibergen and Mavroeidis (2009) discuss these issues at length in the context of estimation of the NKPC.

As we pointed out in the Introduction, identification of the NKPC is likely to be weak because of the familiar empirical finding that changes in inflation are hard to forecast (Atkeson and Ohanian, 2001; Stock and Watson, 2007), implying that potential instruments which are plausibly exogenous (i.e., lagged) must necessarily be close to irrelevant. Indeed, we demonstrate empirically in section 5 that weak identification is pervasive in U.S. data, confirming a common finding in the literature. Furthermore, there are good theoretical reasons to expect identification to be weak, see Mavroeidis (2005) and Nason and Smith (2008a). For example, it is straightforward to see that when the NKPC is flat, i.e., $\lambda = 0$ in (8), inflation is driven only by cost-push shocks. If these shocks are unpredictable, then so is inflation, and the coefficient on inflation expectations is unidentified because no relevant pre-determined instruments exist. Therefore, the model predicts that identification will become arbitrarily weak as the slope of the NKPC λ gets closer to zero. Another situation in which identification is weak is when monetary policy is very effective in anchoring short-term inflation expectations. If inflation expectations do not vary, their effect on inflation is again unidentified. In other words, effective economic policy is bad for econometric analysis (Mavroeidis, 2010; Cochrane, 2011).

When identification is weak, we show below that GIV and VAR-based estimators of the NKPC can be biased in different directions. This helps explain some of the systematic differences in empirical results that we report in section 5. There we also find that estimates are extremely sensitive to specification choices, which is consistent with the moment conditions being insensitive to the value of the parameter ϑ around the true value ϑ_0 . Recent advances in econometrics have made it possible to do inference that is fully robust to weak identification. Because these weak identification robust methods have been derived using alternative asymptotic approximations that do not assume strong instruments, they are reliable irrespective of the strength of the instruments. Surveys on the consequences of weak identification and methods of inference that are robust to it include Stock et al. (2002), Dufour (2003), Andrews and Stock (2005) and Mikusheva (2013).

Simulation studies Mavroeidis (2005) illustrates the above-mentioned points in the context of GIV estimation of the NKPC, and Kleibergen and Mavroeidis (2009) conduct an extensive set of simulations demonstrating the performance of several alternative GMM methods. The conclusion of these simulation exercises is that the theoretical consequences of weak identification are clearly borne out in empirically realistic tests of the NKPC.

The finite-sample performance of VAR-based estimation methods under weak identification has not received much attention in the literature. Therefore, we provide here simulation results comparing four procedures: GIV estimation (Galí and Gertler, 1999), VAR-MD (Sbordone, 2005), VAR-ML (Kurmann, 2007) and VAR-GMM (introduced above). For simulation purposes, we write the NKPC as

$$\pi_t = c + \gamma_f E_t(\pi_{t+1}) + (1 - \gamma_f)\pi_{t-1} + \lambda x_t + u_t, \quad (21)$$

with accompanying VAR(2) reduced-form dynamics

$$\pi_t = \zeta' Y_{t-1} + v_{\pi t}, \quad x_t = \xi' Y_{t-1} + v_{x t}. \quad (22)$$

Here $Y_t = (\pi_t, x_t, \pi_{t-1}, x_{t-1}, 1)'$, and the reduced-form coefficients are $\zeta = (\zeta_{\pi 1}, \zeta_{x 1}, \zeta_{\pi 2}, \zeta_{x 2}, c_{\pi})'$ and $\xi = (\xi_{\pi 1}, \xi_{x 1}, \xi_{\pi 2}, \xi_{x 2}, c_x)'$. We consider four data generating processes (DGPs) here, as summarized in Table 1. For each DGP, we simulate samples of 200 observations each and calculate

point estimates for each of the four estimators.¹⁶ We execute 10,000 Monte Carlo repetitions per DGP. The GIV estimator that we consider estimates (γ_f, λ, c) by linear GMM, with $\Delta\pi_t$ as the independent variable, $(\pi_{t+1} - \pi_{t-1})$ and x_t (and a constant) as regressors, and Y_{t-1} as instruments.¹⁷ The three VAR-based methods exploit the VAR(2) reduced form for (π_t, x_t) to estimate the same three parameters, as explained in section 3.1.¹⁸ Details about the specifications and estimation procedures, as well as a measure of the strength of identification, are given in section A.3 of the Appendix.

DGPs 1a and *1b* have $\gamma_f = 0.7$, $\lambda = 0.03$ and $c = 0$ as true NKPC parameters. Such parameter values represent typical estimates from the literature, cf. section 4. In *DGP 1a*, our choice of reduced-form parameters ξ for the forcing variable are based on OLS estimates on quarterly U.S. data with the labor share as x_t . The implied reduced-form parameters ζ for inflation feature very limited second-lag dynamics relative to the variances of π_t and x_t , so inflation is hard to predict with lagged variables and identification is weak.¹⁹ *DGP 1b* has ξ set to empirically unrealistic values that yield much better predictability of inflation and thus much stronger identification. The top panels in Figure 2 display the densities of the sampling distributions of the γ_f estimators under *DGPs 1a* and *1b*. Evidently, the four estimators exhibit quite different behaviors in the weakly identified parametrization, *DGP 1a*. The GIV estimates of γ_f are biased downwards toward the probability limit of the OLS estimator, which is close to 0.5 for all DGPs in this paper.²⁰ While the sampling distribution density for GIV has rather fat tails, it is single-peaked and not far from bell-shaped.²¹ In contrast, the three VAR-based estimators all exhibit a distinct bimodal behavior, with a large (or even dominant) share of estimates concentrating around $\gamma_f = 1$. VAR-MD is particularly problematic in this regard. Due to the biased and decidedly non-Gaussian finite-sample distributions of the VAR methods, conventional strong-instrument inference procedures will give spurious results. In the strongly identified parametrization, *DGP 1b*, the situation is entirely different. Here the sampling densities of all four estimators are of the conventional Gaussian shape, and only a slight downward finite-sample bias remains. The VAR estimates do not cluster around $\gamma_f = 1$. As the strong-instrument efficiency comparison in section 3.2 predicts, the sampling densities for the three VAR methods are ever so slightly more narrowly concentrated around the true value $\gamma_f = 0.7$ than the GIV density.

DGPs 2a and *2b* set $\gamma_f = 0.3$, a value at the lower end of estimates reported in the literature, and $\lambda = -0.03$, but are otherwise analogs of *DGPs 1a* and *1b*, respectively. It is striking that the reduced-form parameters for *DGPs 1a* and *2a* are so similar, even though the structural NKPC parameters are completely different, cf. Table 1. The sensitivity of the mapping between reduced-form and structural parameters is a key symptom of weak identification. Because structural estimation works by backing out structural parameters from estimates of reduced-form features of the data, it is clear that weak identification will have serious consequences regardless of the estimation method, which is indeed what we find in our simulations. The bottom panels in Figure 2 display the sam-

¹⁶Because the specifications are overidentified, the VAR estimators do not coincide.

¹⁷Observe that the NKPC (21) may equivalently be written $\Delta\pi_t = c + \gamma_f E_t(\pi_{t+1} - \pi_{t-1}) + \lambda x_t + u_t$.

¹⁸Our VAR-MD implementation differs from that in Sbordone (2005) in the choice of distance function and weight matrix. We focus on Kurmann's (2007) version of VAR-ML for the reasons stated in section 3.2.

¹⁹Mavroeidis (2005) and Nason and Smith (2008a) discuss the role of lag dynamics in identification of the NKPC.

²⁰By "OLS" we mean a simple regression of $\Delta\pi_t$ on $(\pi_{t+1} - \pi_{t-1})$ and x_t .

²¹Simulations with smaller true values of λ yield a stronger bias of γ_f estimates toward 0.5.

pling distributions of the γ_f estimators under *DGPs 2a* and *2b*. The results are similar to those for *DGPs 1a* and *1b* (the slight bimodality of the VAR-MD estimator under *DGP 2b* disappears if the strength of identification is increased further). In particular, the spurious clustering of VAR estimates around $\gamma_f = 1$ under weak identification remains even with the true γ_f value set equal to 0.3. This is interesting, as Sbordone (2005) and Kurmann (2007) both report rather large VAR-based estimates of γ_f , and our empirical VAR-GMM estimates in section 5 similarly concentrate around 1.

Section A.3 in the Appendix refers to an online supplement that provides additional simulation results and Matlab code. Among other things, we find – not surprisingly – that misspecification of the VAR can result in coefficient biases in any direction, irrespective of the strength of identification.

VAR methods To understand the behavior of VAR methods under weak identification, we set $c = \lambda = 0$ for ease of exposition, but our results apply more generally. Equation (21) can then be written as

$$\Delta\pi_t = \delta E_t(\Delta\pi_{t+1}) + \bar{u}_t, \quad (23)$$

where $\delta = \gamma_f/(1 - \gamma_f)$ and $\bar{u}_t = u_t/(1 - \gamma_f)$. Suppose the econometrician bases his analysis on a reduced-form AR(1) in $\Delta\pi_t$.²² Let $\rho_1 = E[\Delta\pi_t \Delta\pi_{t-1}]/E[(\Delta\pi_t)^2]$ denote the first autocorrelation of $\Delta\pi_t$, and let $\hat{\rho}_1$ be its sample analog. Since the specification is just identified (we have one reduced-form parameter to identify γ_f), all VAR estimators coincide, and it is easy to show that the VAR estimator of δ is $1/\hat{\rho}_1$. The implied γ_f estimator is then

$$\hat{\gamma}_f^{\text{VAR}} = \frac{1}{1 + \hat{\rho}_1}. \quad (24)$$

Because the primary cause of weak identification is precisely that $\rho_1 \approx 0$ (inflation changes are nearly unforecastable), we expect to find $\hat{\gamma}_f^{\text{VAR}} \approx 1$. This must be the case for any true value of γ_f that leads to $\rho_1 \approx 0$, i.e., for any empirically realistic DGP. This is different from the weak instrument behavior of the GIV estimator $\hat{\gamma}_f^{\text{GIV}}$. As mentioned earlier, this estimator is biased toward the probability limit of the OLS estimator of γ_f in the regression of $\Delta\pi_t$ on $(\pi_{t+1} - \pi_{t-1})$, which equals $1/2$.²³ We are thus able to explain both biases observed in Figure 2.

Another way to view the weak identification VAR bias is as follows. The VAR estimator of δ is precisely the inverse of the OLS estimator in the regression suggested by equation (23), i.e., the regression of $\Delta\pi_{t+1}$ on $\Delta\pi_t$. There are only two ways in which the AR(1) assumption can hold. The first possibility is that $\bar{u}_t \equiv 0$, i.e., the NKPC is exact. The second possibility is that $\bar{u}_t \neq 0$, but the reduced form for inflation changes is $\Delta\pi_t = \bar{u}_t$.²⁴ This is trivially an AR(1) with coefficient 0. Note, however, that if the DGP for inflation really has this form, γ_f is unidentified. This is the sense in which VAR methods can lead to spurious identification: If the solution is $\Delta\pi_t = \bar{u}_t$, VAR

²²Our weakly identified *DGPs 1a* and *2a* nearly have this reduced form, as their reduced-form coefficient on π_{t-1} in the inflation equation is approximately 1.

²³Under stationarity, the OLS probability limit is $E[(\pi_{t+1} - \pi_{t-1})\Delta\pi_t]/E[(\pi_{t+1} - \pi_{t-1})^2] = \{E[(\Delta\pi_t)^2] + E[\Delta\pi_{t+1}\Delta\pi_t]\}/\{2E[(\Delta\pi_t)^2] + 2E[\Delta\pi_{t+1}\Delta\pi_t]\} = 1/2$.

²⁴Such a reduced form can arise in two ways. If $\gamma_f < 1/2$, the solution is determinate and equation (23) can be iterated forward to yield $\Delta\pi_t = \bar{u}_t$. Alternatively, if $\gamma_f \geq 1/2$, the solution is indeterminate, and the so-called minimum state variable solution $\Delta\pi_t = \bar{u}_t$ satisfies the NKPC.

methods – by being equivalent to OLS estimation of (23) – implicitly select the other possibility $\bar{u}_t \equiv 0$ and obtain a seemingly very precise estimate of γ_f close to 1, even though the parameter is unidentified.

Our discussion has focused on the tractable limits of $\lambda = 0$ and exact identification. While we believe that the intuition translates to empirically realistic settings with $\lambda \approx 0$ and overidentified specifications, there is clearly room for more research on these matters. We stress that it is not obvious from our results, or previous analyses in the literature, that GIV methods perform either better or worse under weak identification than VAR-based methods. Only recently has the econometric literature begun to analyze the consequences of weak identification for estimators other than GMM. Magnusson and Mavroeidis (2010) introduce a robust MD test and apply it to the NKPC. Robust VAR-ML analysis of the NKPC has not been attempted yet in the literature, but weak identification robust procedures for general maximum likelihood analysis are currently under development, cf. Andrews and Mikusheva (2011) and references therein.

3.4 Other issues

A number of other econometric issues have been raised in the literature.

Number of instruments Early implementations of the GIV approach to identification of the NKPC used a large number of instruments relative to the sample size: Galí and Gertler (1999) used 4 lags of 6 variables on a sample of 160 observations. This practice is subject to the pitfall of ‘many instrument’ biases, and most subsequent studies have used a significantly smaller number of instruments. There is a large econometric literature on this issue, see e.g., Hansen et al. (2008). It is well known that use of many instruments biases 2SLS estimators towards OLS. Intuitively, in the limit case where the number of instruments is the same as the sample size, the first-stage yields perfect fit, so 2SLS is identical to OLS. The problem becomes more severe if the instruments are many and weak, which is the relevant framework for the NKPC. A recent contribution by Newey and Windmeijer (2009) demonstrates some robustness properties of the weak identification robust methods to many weak instruments. However, Newey and Windmeijer only cover situations in which the instruments are sufficiently informative for the model to be strongly identified. Therefore, their results exclude cases in which the instruments may be *arbitrarily* weak (e.g., completely irrelevant).²⁵ They also ignore the complications arising from uncertainty in the estimation of the long-run variance of the moment conditions, which can be substantial when the number of moment conditions is large. Therefore, we recommend against the use of many instruments in the estimation of the NKPC.

Unit roots Empirically, U.S. inflation appears close to non-stationary in certain subsamples. Fanelli (2008), Mikusheva (2009), Boug et al. (2010) and Nymoen et al. (2010) discuss the implications of inflation having a unit root. This raises issues about the validity of inference when the unit root is left unaccounted for. When the inflation coefficients in the NKPC sum to 1, the model can

---

²⁵ Andrews and Stock (2007) provide conditions under which the weak identification robust tests remain valid under many instruments that may be arbitrarily weak, but these results are for the homoskedastic linear IV regression model, so they do not apply to the NKPC.

be written in terms of changes in inflation. If the instrument set also uses lags of $\Delta\pi_t$, inference will be robust to inflation having a unit root. A number of papers have taken that approach (e.g., Kleibergen and Mavroeidis, 2009) and we study it further in our empirical section.²⁶

Misspecification The various methods described so far may be affected differently by misspecification of the NKPC. Jondeau and Le Bihan (2003, 2008) study some special cases of misspecification in the form of omitted lags of inflation in the NKPC that bias the GIV estimator of the coefficient on future inflation upwards and the VAR-ML estimator downwards. This is consistent with the difference among empirical estimates reported by Fuhrer (1997), Galí and Gertler (1999) and Jondeau and Le Bihan (2005), so Jondeau and Le Bihan argue that misspecification could be the source of disparity in the estimates. However, differences remain when the NKPC is extended to include more lags of inflation.²⁷ Mavroeidis (2005) explores the implications of omitted dynamics for the bias of GIV estimation of the hybrid NKPC.

Cagliarini et al. (2011) and Imbs et al. (2011) discuss another type of misspecification due to aggregation bias. Building on Carvalho (2006), they show that heterogeneity of (Calvo) price rigidity across economic sectors can bias estimates of average price rigidity upwards and bias estimates of the slope of the aggregate NKPC toward zero. Cagliarini et al. (2011) trace this bias to the presence of an additional error term in the aggregate NKPC resulting from sectoral heterogeneity.

Autocorrelated cost-push shocks Autocorrelation of u_t in the NKPC (8) violates the common identifying assumption $E_{t-1}(u_t) = 0$. Because cost-push shocks induce endogenous movements in observables, in general autocorrelated cost-push shocks imply that lagged variables will be correlated with u_t , so that, with the exception of external instruments, all other identifying assumptions listed earlier become invalid. Zhang and Clovis (2010) reiterate this point and perform autocorrelation tests on the residual of the Galí and Gertler (1999) specification. They find evidence of significant residual autocorrelation, which can be removed by including three lags of inflation in the NKPC. Note that the GIV residuals $\tilde{u}_t$ include a future inflation forecast error, so they may exhibit MA(1) autocorrelation even when the structural error u_t is not autocorrelated (Jondeau and Le Bihan, 2005; Mavroeidis, 2005; Eichenbaum and Fisher, 2007). Boug et al. (2010) identify the cost-push shock u_t via VAR-ML and find it to be serially correlated. They also recommend using more lags of inflation in the NKPC. Kuester et al. (2009) show by simulation that GIV estimates of the slope of the NKPC are biased downwards when u_t is autocorrelated, and the Hansen (1982) J test has little power against this misspecification in realistic sample sizes.

Subsample stability Stability tests of the model parameters can be used to test the immunity of the NKPC to the Lucas (1976) critique, as well as to assess the importance of time varying trend inflation and lack of full indexation to it. The standard stability tests of Andrews (1993), Andrews and Ploberger (1994) and Sowell (1996) require strong instruments, but weak identification robust

---

²⁶Note that the restriction that inflation coefficients sum to 1 is not sufficient for inflation to have a unit root. This depends on the dynamics of the forcing variable (Bårdsen et al., 2004).

²⁷As discussed in section 3.2, Kurmann (2007) offers another possible explanation for why VAR-ML estimates that impose determinacy, such as those reported in Fuhrer (1997) and Jondeau and Le Bihan (2005), may differ from GMM estimates.

versions are available (Caner, 2007; Magnusson and Mavroeidis, 2012). Castle et al. (2010) give an extensive discussion of the consequences of structural breaks in the NKPC.

4 Survey of the Empirical Literature

This section surveys the empirical literature on the NKPC. Rather than maintaining a strict chronological order, we have attempted to group the various contributions into the main econometric approaches that were introduced in section 3. Figure 3 and Table 2 present a representative set of results from some of the most frequently cited studies; additional papers are referenced below. The major points of controversy in the literature concern the relative importance of forward- and backward-looking price setting behavior, as well as the degree to which real activity influences inflation dynamics. Although several methodological contributions have been proposed since the beginning of the research program, an empirical consensus is not yet in sight.

4.1 Initial breakthroughs

Limited-information testing of the NKPC was initiated by Fuhrer and Moore (1995) and Roberts (1995). As mentioned by Roberts, previous econometric tests of new Keynesian pricing equations had been based on full-information (system) methods under the assumption of RE. Roberts (1995) shows that three different theoretical frameworks – the staggered contracts model of Taylor (1980), the infinite-horizon staggered pricing model of Calvo (1983) and the quadratic adjustment cost model of Rotemberg (1982) – all lead to (what came to be known as) a difference equation specification of the pure NKPC, with the output gap as forcing variable. He suggests two different limited-information approaches to testing the relationship: first, the use of survey expectations as proxies for the expectation term, and second, McCallum’s (1976) technique of subsuming the RE forecast error into the equation’s error term and instrumenting for next period’s inflation and the output gap with lagged variables (GIV estimation, in the terminology of section 3.1). Using annual U.S. data, Roberts finds a significant role for the output gap.

Seminal contributions by Galí and Gertler (1999) and Sbordone (2002) helped propel the NKPC research agenda into the forefront of empirical macroeconomics. Both papers take the now-standard microfounded RE pricing model to U.S. data and obtain results that are supportive of the model’s fit. Furthermore, both sets of authors exploit the model’s implication that aggregate marginal cost may be proxied by the labor share. Indeed, Galí and Gertler establish that the NKPC only fits U.S. data if the labor share is used as forcing variable instead of the output gap, which may be mismeasured. Galí and Gertler also develop the now-standard *hybrid* NKPC, whose lagged inflation terms introduce intrinsic persistence of the inflation rate on top of the extrinsic persistence imparted by the forcing variable.

4.2 GIV estimation

Using linear and non-linear GIV methods, Galí and Gertler (1999) find that, while the backward-looking inflation term is significant, the forward-looking RE term dominates; they also obtain a significant and correctly signed coefficient on the labor share (unlike the output gap). The

NKPC restrictions are not rejected by overidentification tests or by visual inspection of fitted inflation. Galí et al. (2001) take the model to aggregate Eurozone data, largely confirming the U.S. findings. Benigno and López-Salido (2006) find some heterogeneity in estimated coefficients for major Eurozone countries. Eichenbaum and Fisher (2004, 2007) evaluate a variant of the NKPC with price indexation that was developed by Christiano et al. (2005), and they also introduce variable elasticity of demand, capital adjustment costs and pricing implementation lags. Blanchard and Galí (2007) consider a model with real wage rigidity, which leads to an NKPC featuring the unemployment rate as forcing variable; GIV estimation on U.S. data yields intuitively reasonable coefficients with significant forward-looking behavior. Krause et al. (2008), Ravenna and Walsh (2008) and Blanchard and Galí (2010) explore NKPCs with explicit labor market frictions that lead to alternative expressions for marginal cost. Chowdhury et al. (2006) and Ravenna and Walsh (2006) assume that firms must borrow to pay their wage bill up front each period, which leads to the so-called cost channel of interest rates, i.e., marginal cost is directly influenced by the interest rate. These papers find that, for most countries, the Treasury bill rate enters significantly into an extended NKPC when estimated by GIV, but the coefficients on the forward- and backward-looking inflation terms are not affected much relative to the baseline. Neiss and Nelson (2005) compute the output gap that is implied by a standard new Keynesian model; this theoretically consistent measure turns out to be essentially uncorrelated with quadratically detrended output (which is used by Roberts, 1995, and Galí and Gertler, 1999), and GIV estimates of the slope of the NKPC are even more significant than when using the labor share.²⁸ Gagnon and Khan (2005) extend the NKPC to a more general CES production function and find that structural GIV estimates imply less price stickiness than under the usual Cobb-Douglas specification. In addition to a CES production function, McAdam and Willman (2010) add varying capacity utilization, which decreases the estimated coefficient on the inflation expectation term. Batini et al. (2005) and Rumler (2007) estimate open-economy NKPCs by GIV on data from European countries, finding a significant role for international variables. Gwin and VanHoose (2008) and Shapiro (2008) construct alternative measures of firm marginal costs from micro-level and sectoral data.

While the previously mentioned papers find a significant, and often dominant, role for the forward-looking RE term, a number of papers that use the GIV framework have raised issues with the mainstream analysis (see also the discussion of weak identification below). Bårdsen et al. (2004) point out that the literature has mostly not rejected the homogeneity restriction (i.e., that the coefficients on last and next period's inflation sum to 1), which, under strict exogeneity of the forcing variable, implies that inflation is non-stationary. They show empirically that GIV estimates are quite sensitive to the choice of instrument set and estimator (see also Guay and Pelgrin, 2005), and the Galí et al. (2001) hybrid NKPC is rejected in favor of alternative, encompassing models of inflation. Fuhrer and Olivei (2005) employ a reduced-form VAR to compute expectations of next period's inflation and the output gap; they then use the computed expectations as instruments. Their estimate of the coefficient on the forward-looking expectation term is much smaller than the traditional GIV estimate. A series of papers by Rudd and Whelan (2005, 2006, 2007) contend that the Galí and Gertler (1999) estimation approach yields spurious results. Rudd and Whelan criticize the use of the labor share as a proxy for marginal cost due to its countercyclicality.²⁹ They

²⁸See also Fukač and Pagan (2010) for a critique of “off-model” output detrending.

²⁹This issue is further pursued by Mazumder (2010).

demonstrate that, provided the NKPC leaves out explanatory variables, the use of instruments outside of the model (such as interest rates or wage and commodity price inflation) may bias the estimates in the direction of establishing a high degree of forward-looking behavior. Furthermore, Rudd and Whelan conduct several tests of the incremental explanatory power of the labor share and conclude that it adds essentially no information to inflation forecasting. Estimating the model by GIV in iterated form (cf. section A.2.2 in the Appendix) yields a smaller coefficient on the forward-looking term. Finally, data revisions since 1999 have eroded the significance of the labor share, even in the original Galí and Gertler set-up. Galí et al. (2005) counter that Rudd and Whelan (2005) use a parametrization of the model that does not correspond to the structural parameters in Galí and Gertler (1999) and Galí et al. (2001). Galí et al. (2005) show that if the same parametrization is used, iterated GIV estimation confirms the results in Galí and Gertler (1999) and Galí et al. (2001).

As witnessed by the myriad of parallel sub-models and methods, the literature is still far from producing a consensus set of specifications or empirical conclusions, even within the relatively narrow RE GIV framework.

4.3 VAR estimation

In their seminal paper, Fuhrer and Moore (1995) augment a Taylor (1980) pricing equation with reduced-form VAR equations for the output gap and Treasury bill rate and estimate the resulting system by ML, using the AIM routine from Anderson and Moore (1985) to solve for a RE solution given the parameters. Fuhrer and Moore reject the restrictions implied by the standard pricing model based on a formal likelihood ratio test and inspection of the implied impulse responses, which display too little inflation persistence. The data is more favorable to an alternative real wage contracting model that implies sticky *inflation* instead of just sticky prices. Fuhrer (1997) uses a similar approach to test for the significance of forward-looking rational inflation expectations relative to backward-looking (adaptive) expectations; he finds that the RE component is insignificant. Subsequent papers have applied the AIM-based VAR-ML method to the Galí and Gertler (1999) hybrid NKPC. Fuhrer and Olivei (2005) and Fuhrer (2006) find a small coefficient on the forward-looking term relative to that on lagged inflation. Roberts (2005), who also considers GIV and impulse response matching, estimates a hybrid NKPC with four lags of inflation, obtaining about 50% weight on forward-looking behavior. Jondeau and Le Bihan (2005) estimate the NKPC on data from the U.S. and major European countries. They find that GIV estimates of the coefficient on forward-looking expectations tend to be high, while VAR-ML estimates tend to be lower. Kiley (2007) uses VAR-ML to estimate an NKPC specification with four lags of inflation and expectations of next-period inflation taken with respect to *previous*-period (rather than current-period) information. Here the forward-looking term is dominant, and the Bayesian Information Criterion indicates that the structural model provides as good a fit to U.S. data as a reduced-form VAR.

Kurmann (2007) criticizes the AIM-based approach to ML estimation, as it imposes the extraneous assumption that the RE solution must be unique (determinate), cf. section 3.2. Using an ML method that does not impose uniqueness, he finds evidence of a large share of forward-looking behavior in the Galí and Gertler (1999) U.S. dataset, which contrasts with estimates obtained under the additional uniqueness assumption. It is an open question whether imposing the determinacy assumption matters empirically across other data sets and NKPC specifications. Korenok et al.

(2010) also eschew the AIM algorithm and instead write their model in a form that is amenable to Kalman filtering.³⁰

Sbordone (2002) tests the pure NKPC on U.S. data using a two-step approach akin to that of Campbell and Shiller (1987, 1988). In a first step, she fits a reduced-form VAR to the data. When iterated forward, the pure NKPC implies that inflation is given by an expected present value of future marginal costs, and this quantity may be evaluated using the fitted VAR. The structural parameters of the NKPC can then be estimated by minimizing the squared distance between model-implied and actual inflation. The estimated Calvo (1983) parameter is in line with micro-estimates of price stickiness. Sbordone (2005) refines the estimation approach by interpreting it as minimum distance estimation and accounting for sampling uncertainty of the first-step estimated VAR (see also Kurmann, 2005). She provides estimates of the hybrid NKPC, broadly confirming the conclusions in Galí and Gertler (1999). Tillmann (2008) uses a related MD approach to assess the importance of the cost channel of monetary policy. Sbordone (2006) develops a model of joint price and wage determination, which is estimated by minimum distance. Coenen et al. (2007) construct a model with a general, non-constant hazard rate of price resetting, which they estimate using an indirect inference procedure that matches the model-implied dynamics to the estimated reduced-form VAR. Carriero (2008) rejects the cross-equation restrictions that the one-lag NKPC places on a reduced-form VAR in inflation and the labor share. Guerrieri et al. (2010) develop a microfounded open-economy NKPC in which the relative price of foreign goods enters. To estimate it they use a multi-equation GMM approach that adds reduced-form VAR equations for the labor share and relative foreign goods prices. Their preferred specification yields an insignificant coefficient on lagged inflation. Cornea et al. (2013) use VAR methods to estimate an NKPC with evolutionary switching between forward-looking and backward-looking inflation expectations; they find substantial time-variation and heterogeneity in the type of expectations formation.

Fanelli (2008) and Boug et al. (2010) conduct likelihood-based estimation of the hybrid NKPC, taking into account the possibility that the variables are cointegrated. The likelihood is derived conditional on a reduced-form vector error-correction model for inflation and the output gap, using the Kurmann (2007) approach that does not impose uniqueness of the RE solution. Both papers find that the NKPC restrictions are rejected for the Eurozone. Boug et al. (2010) find some support for the hybrid NKPC in U.S. data, although the residuals are significantly autocorrelated, violating an assumption of the model. The MLE of the coefficient on the inflation expectations term is much larger than that on the lagged inflation term.

While popular in the DSGE literature, Bayesian methods have only been used in a few limited-information analyses of the NKPC. Fuhrer and Olivei (2010) and Fukač and Pagan (2010) compute posteriors for the parameters in versions of the NKPC, where the expectation of next period's inflation is determined by a reduced-form VAR. Cogley and Sbordone (2008) introduce drifting trend inflation into the standard new Keynesian model, which changes the form of the NKPC. They estimate the model using quasi-Bayesian methods, conditional on a reduced-form VAR with drifting parameters and stochastic volatility. Their imputed inflation gap (i.e., the difference between inflation and its trend) is much less persistent than raw inflation, and the quasi-posteriors indicate that once the trend is accounted for, there is no need to allow for backward-looking price

---

³⁰Other papers that use filtering techniques include Nelson and Lee (2007), Kim and Kim (2008) and Kim and Manopimoke (2011).

indexation (see also Sahuc, 2006; Hornstein, 2007). Barnes et al. (2011) and Gumbau-Brisa et al. (2011) argue, however, that this conclusion is sensitive to how the NKPC restrictions are imposed in the estimation. Despite the reservations, the trend inflation research agenda is rapidly becoming one of the most well-cited branches of the NKPC literature.

4.4 Estimation using survey expectations

As mentioned, Roberts (1995) uses survey measures of inflation expectations from the Michigan and Livingston surveys as an alternative to the RE GIV approach.³¹ Roberts (1997) finds that the apparent sluggishness and non-rationality of these survey forecasts generate sufficient inflation persistence in the U.S. NKPC, and the data favors such a specification to the Fuhrer and Moore (1995) sticky inflation model. Rudebusch (2002) estimates the hybrid NKPC on U.S. data by OLS, with Michigan survey data proxying for inflation expectations. He finds a relatively small coefficient on the expectations term but a significantly positive coefficient on the output gap. Adam and Padula (2011) use SPF inflation forecasts and also find the forcing variable to be significant, regardless of whether they use the labor share or output gap, but their OLS estimate of the coefficient on the expectation term is slightly larger than that on lagged inflation. Kozicki and Tinsley (2002) estimate various pricing equations for the U.S. and Canada using survey forecasts and allowing for non-zero trend inflation. Gerberding (2001), Paloviita and Mayes (2005), Paloviita (2006, 2008), Henzel and Wollmershäuser (2008) and Koop and Onorante (2011) estimate NKPCs for European countries using various measures of survey expectations of inflation and various estimation procedures. The estimated extent of forward-looking pricing behavior varies greatly between studies and specifications. Brissimis and Magginas (2008) use SPF forecasts and the Federal Reserve’s Greenbook projections to estimate the U.S. NKPC by GMM. They find a dominant role for forward-looking expectations and a significantly positive coefficient on the labor share. Zhang et al. (2008, 2009) consider both SPF, Greenbook and Michigan survey forecasts. Unlike the RE specification, the survey forecast NKPC gets a positive and significant coefficient on the output gap but its estimates appear more unstable over subsamples (see also Kim and Kim, 2008). Mazumder (2011) uses SPF, Michigan and Greenbook forecasts to test the NKPC with a procyclical measure of marginal costs developed in Mazumder (2010).

Nunes (2010) simultaneously includes rational expectations and SPF forecasts in an NKPC. The GMM estimates point to non-rational expectations only playing a minor role in explaining U.S. inflation dynamics. This conclusion is disputed by Fuhrer and Olivei (2010) and Fuhrer (2012), who proxy for the RE term with expectations from a reduced-form VAR and estimate the NKPC by Bayesian and ML methods. Smith (2009) gives conditions under which it is advantageous to include data on survey forecasts for *statistical* reasons, even if the researcher has a purely rational NKPC in mind.

Survey forecast methods have established a commanding presence in the NKPC literature. So far, the literature has only scratched the surface in terms of providing full-fledged microfoundations, and a detailed understanding of the interplay between non-rational expectation formation and price setting remains elusive.

---

³¹Roberts still instruments for the output gap as it may be correlated with the inflation error term.

4.5 Identification issues and robust inference

The literature’s awareness of the problems associated with weak identification has grown over time. Galí et al. (2001) guide the choice of their instrument set by the first-stage F statistic. Mavroeidis (2004, 2005) provides analytical and simulation evidence that explains why weak identification is likely to be an issue for NKPC estimation. Ma (2002) is the first paper to compute weak identification robust confidence sets (specifically, the Stock and Wright, 2000, S set) for the NKPC, finding the data to be completely uninformative about the structural parameters. Dufour et al. (2006) compute Anderson and Rubin (1949) and Kleibergen (2002) confidence sets for both GIV and survey forecast specifications. The U.S. GIV confidence region is fairly large, while the survey forecast one is empty; no NKPC specification seems to fit Canadian data. Nason and Smith (2008a) reject the hybrid NKPC for both Canada, the U.K. and the U.S. using the Anderson and Rubin (1949) test and the Guggenberger and Smith (2008) GEL test. In contrast, Martins and Gabriel (2009) find very wide robust GEL confidence sets. Using a variety of GMM-based robust tests, Kleibergen and Mavroeidis (2009) conclude that inflation appears to be significantly forward-looking, but the confidence regions are wide. Dufour et al. (2010a,b) carry out robust inference on certain extensions of the NKPC with real wage rigidities and labor market frictions. Magnusson and Mavroeidis (2010) develop a weak identification robust version of Sbordone’s (2005) minimum distance test, finding somewhat smaller confidence regions than when using a robust GIV approach. Kleibergen and Mavroeidis (2013) demonstrate the consequences of ignoring weak identification in Bayesian analyses of the NKPC and propose ways of circumventing the problems.

Some papers have devised methods for improving the strength of identification in GIV estimation of the NKPC. Dees et al. (2009) obtain instruments for individual-country NKPCs by estimating a multi-country cointegrating VAR. Building on Beyer et al. (2008), Kapetanios and Marcellino (2010) and Kapetanios et al. (2011) develop identification robust theory for GMM testing using instruments that have been estimated by principal components from a large set of candidate variables. While this seems to improve identification of the slope of the NKPC, the relative shares of forward- and backward-looking behavior remain very weakly identified. Motivated by the Lucas critique, Magnusson and Mavroeidis (2012) suggest using robust parameter instability tests to improve inference about the NKPC.

The lessons from weak identification analyses have so far only had limited impact on the broader NKPC literature. Papers that do mention the identification issue often either treat it as merely another robustness check or incorrectly dismiss it as a strictly GMM-specific problem. The consequence is that comparison of results across papers is difficult.

5 Empirical Synthesis

In this section we generate estimates of the NKPC corresponding to a wide selection of empirical approaches from the literature.³² Because we use a common data set for all estimates, we are able to highlight the sensitivity of the inference to choices of specification and econometric strategy. While our results largely confirm several isolated results in recent strands of the literature, they also convey the strong message that the *specification* uncertainty surrounding estimation of the NKPC

---

³²Estimation results are obtained using Ox (Doornik, 2007).

is vast. We then show, using a number of benchmark specifications, that even given a model, the *sampling* uncertainty of the estimates tends to be large. Both these conclusions can be explained by the weakness of identification. We also demonstrate that the potential for the data to distinguish between rational and non-rational price setting is limited.

5.1 Data

As in most of the literature, our dataset features U.S. aggregate time series at a quarterly frequency, with the largest possible sample extending from 1947q1 to 2011q4. Most series have been downloaded from the St. Louis Fed’s FRED database. The data consists of alternative series for price and wage inflation, the labor share, output, interest rates and survey measures of inflation expectations. We use the abbreviation “NFB” for the non-farm business sector. See section A.4 in the Appendix for a detailed description of the data and transformations.

A few of our data series deserve mention here. Survey forecasts of inflation are taken from the Survey of Professional Forecasters (SPF) and the Federal Reserve’s Greenbooks (GB). We consider both one-quarter-ahead inflation forecasts made at time t , $\pi_{t+1|t}^s$, and two-quarters-ahead inflation forecasts made at time $t-1$, $\pi_{t+1|t-1}^s$.

Inflation gaps are calculated as the raw inflation rate minus a measure of trend inflation. Our two model-based measures of trend inflation are the smoothed (two-sided) and filtered (one-sided) permanent components of inflation from the UC-SV model of Stock and Watson (2007, 2010). For CPI inflation, 10-year CPI inflation forecasts serve as an additional measure of trend inflation (this series starts in 1991).

Real-time data on inflation and output is obtained from the Philadelphia Fed’s website. We have compiled a unique dataset on real-time changes in the labor share (real unit labor cost), for use as instruments, by combining internal records from the Bureau of Labor Statistics with figures from the bureau’s historical news releases.³³

Our output gaps include the official estimate from the Congressional Budget Office (CBO) as well as various detrended output series. We also compute labor share gaps. This is done to remove trends such as the recent dramatic decline in the labor share, which may arguably be attributed to secular changes outside of the new Keynesian model.³⁴ In addition to full-sample gaps, we use real-time output data or current-vintage labor share data to compute *one-sided* gaps, for which the trend is determined using only data points up to time t . Because such series do not estimate the trend from future data, they (or their lags) can more plausibly be treated as exogenous for estimation purposes.³⁵ Another stationary analog of the labor share is the cointegrating relationship between real wages and labor productivity found by Sbordone (2005, fn. 19). Like most of the literature, we consider non-detrended labor share series as well.

In our empirical analysis we ignore measurement error in the estimates of the trends of inflation and forcing variables.

---

³³We are grateful to Shawn Sprague for assisting us in obtaining the real-time labor share data.

³⁴See Gwin and VanHoose (2008) for a discussion of the need to detrend marginal cost measures.

³⁵Unfortunately, we cannot use our real-time labor share data to construct actual real-time labor share gaps. Because the BLS base year changes over time, we can only compute real-time *changes* in the (log) labor share, not levels. Our one-sided labor share gaps therefore rely on current-vintage data.

5.2 Specification sensitivity

We take the specification of Galí and Gertler (1999) and Galí et al. (2001) to be our benchmark: a hybrid NKPC (9) with one lag of inflation and the labor share as forcing variable, estimated by GIV under the RE assumption. As discussed in section 4, GIV analyses typically find point estimates of the coefficient on expectations γ_f in the 0.5–0.7 range, and the coefficient on lagged inflation γ_b , the measure of intrinsic persistence, is often significantly positive and not significantly different from $1 - \gamma_f$. The coefficient on the labor share λ is generally estimated to be positive but borderline significant (using the usual strong-instrument inference). In Table 3 we replicate these findings using data of the same vintage as Galí and Gertler (1999) but with the Galí et al. (2001) instrument set.³⁶ Later papers have mostly obtained insignificant λ estimates, and like Rudd and Whelan (2007) we find that this is even true on the Galí and Gertler (1999) sample if revised data (as of 2012) is used. Using the output gap as forcing variable also typically yields an insignificant estimate of λ , and early papers in the literature tended to find negative point estimates.

The estimation results reported in the literature differ in terms of the choice of data series, estimation sample and various other aspects of the specification, such as the number of inflation lags, any additional regressors, the measurement of inflation expectations, and the identification assumptions, including the set of instruments and other identifying restrictions. As we showed in Figure 3, estimates of λ and γ_f reported in various papers differ markedly, but the key message is that all highly cited papers obtain a positive slope coefficient ($\lambda > 0$), and, with the exception of Fuhrer (2006), generally find forward-looking behavior to be dominant ($\gamma_f > 0.5$). The results presented in Figure 3 are a tiny subset of possible specifications. Table 4 presents various dimensions of the specification choice that have been considered in the literature.³⁷ These combinations of choices produce a very large number specifications that are not objectionable on a priori grounds.

To gauge the sensitivity of the results about the importance of forward-looking behavior to variations in data, sample and identification assumptions, we obtain estimates of the coefficients (λ, γ_f) in the baseline NKPC (9) for various combinations of the specification choices listed in Table 4. We then plot the point estimates in (γ_f, λ) -space. These plots do not convey any information about sampling uncertainty, i.e., they are not confidence sets. Confidence sets for a subset of those specifications are analyzed in section 5.3 below. However, these plots, which we refer to as “clouds”, do give a useful visual impression of the specification uncertainty. We study the specifications with the labor share and output gap as forcing variable separately, because the coefficient λ on the forcing variable is not comparable across these cases. As we are only able to report a limited number of results here, we invite interested readers to explore the myriad of possible clouds using our interactive Matlab plotting tool, available in the online supplement.³⁸

We first look at the specification settings that have been used in the literature (i.e., not using real-time data or survey expectations as instruments). Figures 4 and 5 report the results for the

³⁶We obtained the 1998 vintage data from Adrian Pagan. We use CUE rather than 2-step GMM (cf. section A.2.1 in the Appendix) because the former is invariant to reparametrization of the moment conditions. The results are comparable to the bottom two rows of Table 2 in Galí et al. (2001).

³⁷The only components of the table that have not been explored extensively in the literature are some of the real-time data series (but see Paloviita and Mayes, 2005, Dufour et al., 2006, and Wright, 2009) and the use of survey expectations as instruments (but see Wright, 2009, and Nunes, 2010). The latter is motivated by evidence that surveys typically forecast inflation better than most alternatives, see Ang et al. (2007).

³⁸<https://sites.google.com/site/sophoclesmavroeidis/research/working-papers/online-supplement-for-nkpc-review>

labor share and output gap as forcing variable, respectively. Figure 4 also contains the Galí and Gertler (1999) vintage point estimate and associated Wald confidence ellipse from Table 3 for comparison. These plots contain more than 600,000 estimates combined. Observe that the plotted parameter space $(\gamma_f, \lambda) \in [-1, 2] \times [-0.3, 0.3]$ is much larger than that of Figure 3. Table 5 reports summary statistics for the point estimates in Figures 4 and 5.

The main messages from the figures are that (i) estimates of the coefficient on the forcing variable are symmetrically dispersed around zero, and (ii) estimates of the coefficient on expectations are on average around 3/4 and very dispersed, though the vast majority (around 90%) of those are positive. Importantly, only about half of the estimates lie in the positive orthant $\lambda > 0, \gamma_f > 0$. Moreover, the fraction of cases in which λ and γ_f both appear statistically significantly positive using (one-sided) 5%-level individual t-tests is quite small, while most of the reported estimates in the literature appear to fall in that category. It is interesting that the frequency of significantly positive coefficients for the output gap specifications is almost double the frequency for the labor share ones. This is not in line with the view that NKPC specifications with the output gap as forcing variable more frequently have estimates with the ‘wrong sign’ than do specifications using the labor share as forcing variable (Galí and Gertler, 1999). It is important to stress that the results based on t-tests are reported for the comparison with the literature, and they do not yield reliable evidence on the significance of the coefficients. In the next subsection we report results that are robust to weak identification.

To shed some light on the issue of weak identification, the penultimate row of Table 5 reports the median value of the heteroskedasticity and autocorrelation robust first-stage F statistic of Montiel Olea and Pflueger (2013), denoted F_{HAR} . A low value of this statistic can be thought of as a warning sign for weak instruments.³⁹ We see that instruments are quite strong for forecasting the forcing variable (the median F statistic is 63.7 for labor share specifications and 166.5 for output gap specifications) but rather weak for forecasting the inflation expectation proxy (median F is 3.1 and 4.2, respectively).⁴⁰ Even though this is not a formal test of weak instruments, and we do not recommend the use of pre-tests in place of weak identification robust inference, these results reinforce the intuition that changes in inflation are hard to forecast and we should therefore worry about weak identification. Figure 6 displays smoothed density estimates of our first-stage F statistics for forecasting the expectations proxy, treating RE GIV separately from time- t dated SPF/GB forecasts, and using all instrument sets in Table 4. As one might expect, time- t dated survey forecasts are much better predicted by the various instrument sets than is next period’s realized inflation. The median F for forecasting next period’s inflation is 2.7 across all labor share specifications (3.6 across output gap specifications). In comparison, the median is 12.8 (12.5) for SPF/GB forecasts, and if we restrict attention to the instrument set that includes lagged survey forecasts, the median F is even higher at 42.1 (43.9). This suggests that survey forecast specifications of the NKPC may be more strongly identified than their GIV counterparts, and the

³⁹The well-known rule of thumb of $F > 10$ is a commonly used benchmark (Stock and Yogo, 2002), although Montiel Olea and Pflueger (2013) show that this condition is neither necessary nor sufficient for instruments to be strong in the presence of heteroskedasticity and autocorrelation.

⁴⁰Low values of the F statistic for forecasting the forcing variable do arise in specifications that use the real-time (RT) instrument set. For labor share specifications, the median value of the F statistic is 9.5 for the RT regressions, whereas the median is 63.1 for all other instrument sets in Table 4. For output gap specifications, the corresponding medians are 69.3 and 170.6, respectively.

evidence reported in section 5.3 corroborates this conjecture.

The final row of Table 5 reports the rejection frequencies of a weak identification robust version of Hansen’s (1982) J test of overidentifying restrictions, see section A.2.6 in the Appendix for details. The rejection frequencies are just over 3% at the 5% level, so there is no systematic evidence against the validity of the overidentifying restrictions. Notice, however, that this test is less powerful than the standard J test because it uses larger critical values.

We now take a closer look at the different dimensions of the specification choice. In the following we do not exclude estimates that use the real-time or survey instrument sets. In the remainder of this subsection, the discussion is organized in self-contained paragraphs that can be skipped without affecting the readability of the rest of the article. Additional details are provided in section A.5 of the Appendix.

CUE versus 2-step GMM We generate GMM estimates using both the efficient 2-step estimator (2S) and the continuous updating estimator (CUE) of Hansen et al. (1996). These are described in section A.2.1 in the Appendix. Table 6 compares summary statistics of point estimates based on 2S and CUE GMM for the various specifications listed in Table 4.⁴¹ The 2S and CUE are very similar for λ , and CUE is typically larger than 2S for γ_f . Moreover, the 2S estimates are closer to the corresponding OLS estimates than CUE. This finding is consistent with the well-known bias of GMM estimators towards the OLS probability limit, which is stronger for 2S than for CUE (Stock et al., 2002). The relatively better bias properties of the CUE come at the cost of greater dispersion, which is confirmed by the 90% interquantile ranges: the ones for the CUE are more than double the corresponding ones for 2S. Bårdsen et al. (2004) and Guay and Pelgrin (2005) also report large sensitivity of NKPC estimates to the choice of GMM estimator, as well as to the set of instruments.

VAR assumption Our VAR-GMM estimates are based on the moment condition (15). The reduced form evolution of inflation is thus restricted to be a linear function of the variables in the instrument set. Table 7 reports summary statistics comparing GIV and VAR-GMM estimates, while Figure 7 plots clouds for estimates that impose the VAR assumption and those that do not. There is no noticeable difference in the estimates of λ between the VAR and GIV methods, but there is a substantial difference in γ_f : in the vast majority of cases (about 80%), imposing the VAR assumption increases the estimate of γ_f and the median estimate is actually 1. This is consistent with the results reported in Sbordon (2005) that use VAR-MD and find no role for intrinsic persistence, as well as with the VAR-ML results in Kurmann (2007). It is inconsistent with Fuhrer (2006, 2012), who additionally imposes determinacy, cf. section 3.2. As we pointed out in section 3.3, weak identification can cause VAR estimates of γ_f to be biased toward 1. Imposing the additional restrictions that coefficients on inflation in the NKPC sum to one (i.e., $\gamma(1) = \gamma_f$ in equation (9)) and that inflation enters the VAR in first differences (thus using lags of $\Delta\pi$ as instruments) causes γ_f estimates to concentrate even more tightly around 1, but also increases the dispersion of the λ estimates.

---

⁴¹In a small number of cases the CUE either failed to converge or produced estimates that were very large in absolute value. This is consistent with the well-known property that the finite sample distribution of the CUE has fat tails (Hansen et al., 1996).

Survey forecasts There are large and systematic differences in the effect of using survey inflation forecasts relative to RE GIV across labor share and output gap specifications, sample period, inflation series (GDP deflator versus CPI) and forecast source (SPF versus GB). Survey forecasts typically increase the estimate of λ across most specifications and sample periods, especially when the output gap is used as forcing variable. The estimate of γ_f moves in different directions across specifications: it is typically much lower than GIV in labor share specifications and either the same or higher in output gap specifications. This is illustrated in Figure 8, which plots the post-1984 cloud for RE GIV estimates against that for time- $(t - 1)$ dated exogenous SPF forecasts (results are similar for other choices of survey forecasts), treating GDP deflator and CPI specifications separately. Further details are given in Table 9 in section A.5 of the Appendix.

Subsample variation is also quite striking. The reduction in γ_f relative to GIV is much more evident in the post-1981 sample (SPF CPI forecasts are only available from 1981q3). Survey specifications with CPI inflation typically yield much larger estimates of γ_f than those with GDP deflator inflation. SPF and GB forecasts do not yield systematically different full-sample estimates, though there are some systematic difference to the estimates of γ_f before and after 1984. Treating surveys as endogenous or exogenous does not seem to make much difference to the central tendency of the estimates, though it does make a difference to dispersion (the latter estimates are a lot less dispersed, as expected).

Instruments The last two rows of Table 6 give median differences for specifications using the Galí and Gertler (1999) instrument set (GG), which is considerably larger than the rest. This instrument set produces estimates for γ_f that are typically lower than average. The GG estimates are also less dispersed and more concentrated around the OLS estimates. Other than GG, the choice of instrument set does not substantially change the central tendency of the estimates.

Number of inflation lags in the NKPC Estimates of γ_f are very sensitive to the number of inflation lags included in the model, while estimates of λ seem to be unaffected, on average. Specifically, adding lags of inflation to the NKPC tends to reduce the estimate of γ_f by about 0.25 when we add 1 lag to the pure NKPC, and by a similar amount when we add three more lags. This corroborates results reported by Rudd and Whelan (2005), but it need not be due to misspecification of the more restrictive NKPC specifications, as was suggested by Rudd and Whelan (2005) and Mavroeidis (2005). The direction of the movement in the point estimates is entirely consistent with the possibility that specifications with more inflation lags are more weakly identified, in which case estimates of γ_f would exhibit a larger bias towards $\gamma_f = 1/2$.⁴² Indeed, the median first-stage F_{HAR} statistics for inflation expectations are 24.5, 3.2, and 2.3, for the 0, 1 and 4-lag NKPC models, respectively, across all specifications that use the labor share as the forcing variable.⁴³ This is further corroborated by the size of the robust confidence regions reported in the next subsection: they get progressively larger as we move from 0 lags to 4 lags.

⁴²As explained in section 3.3, the probability limit of the OLS γ_f estimator is often close to 0.5 for empirically realistic NKPC parametrizations.

⁴³The corresponding numbers for output gap specifications are 25.2, 4.1 and 3.1.

Inflation series Figure 8 indicates that survey forecast specifications are more sensitive to the choice of inflation series than GIV estimates are. Figure 9 compares GDP deflator and CPI estimates, pooling across all GIV and survey forecast specifications and all subsamples in Table 4. Estimates of both parameters are considerably more dispersed in GDP deflator specifications, but the median difference across these inflation series is very small. This is partly a result of the general decrease in the dispersion of estimates from the pre-1984 to the post-1984 sample, since CPI specifications are under-represented in samples that contain data before 1981 due to the lack of CPI survey forecasts. The bottom row of plots in Figure 8 compares CPI versus GDP deflator estimates for a common post-1984 sample, and it is apparent that the dispersion of the estimates is generally smaller and not substantially different across inflation series.

Using inflation gaps to account for trend inflation tends to produce somewhat lower estimates of γ_f irrespective of whether the labor share or output gap is used as the forcing variable. For λ , there is a small positive difference only in output gap specifications. The reason why these differences are not large may be that the sum of the coefficients on inflation are close to one, thus mitigating the impact of any trend inflation, as discussed in section 2.2. The inflation gap point estimates do, however, cluster much tighter around $\lambda = 0$.

The remaining inflation series yield results that are similar to either GDP deflator or CPI estimates. Using the chain-type GDP price index gives very similar results to those for GDP deflator inflation. For GIV, PCE estimates are similar to CPI results, although in output gap specifications λ tends to be estimated higher with CPI compared to PCE. There is little difference between using CPI/PCE inflation and their core inflation equivalents, except that the core estimates are less dispersed.

Output gap and labor share series There is generally very little systematic difference in the results based on alternative labor share and output gap series, except that use of detrended labor share series (labor share “gaps”), using either pseudo-real-time or full-sample trends, increases the dispersion of the estimates of λ , without much change in central tendency. This could be due to the fact that the detrended series are harder to forecast, thus making identification somewhat weaker.⁴⁴ A striking conclusion is that the addition of a good decade’s worth of data (and data revisions) since Galí and Gertler (1999) completely overturns their conclusion that labor share specifications yield markedly different results from output gap specifications.

Sample There is little systematic difference in the central tendency of estimates before and after 1984, cf. Table 10 in the Appendix. As Figure 8 suggests, survey estimates are, however, sensitive to sample choice. This is consistent with Zhang et al. (2008). Figure 10 reports pre- and post-1984 estimates in the survey specifications with GDP deflator inflation. For RE GIV specifications, the central tendency of estimates does not depend much on the choice of sample, but post-1984 estimates are more tightly concentrated around $\lambda = 0$.

Other specification choices The restriction that coefficients sum to 1 does not matter much except for VAR specifications, as discussed above. Use of oil prices or interest rates in the NKPC

---

⁴⁴The median of the first-stage F_{HAR} statistic for the labor share is 71 for the levels and 50 for the gap specifications, respectively.

does not affect the central tendency of the point estimates. This is consistent with Chowdhury et al. (2006) and Ravenna and Walsh (2006).

5.3 Sampling uncertainty

The previous subsection focused on specification sensitivity, characterized by the variation in point estimates across specifications. We now turn to sampling uncertainty, which we measure conventionally using confidence sets for selected specifications based on methods that are robust to weak identification. Our robust confidence sets, called S sets, are based on the S test of Stock and Wright (2000), described in section A.2.5 of the Appendix. This is a test of the validity of the model's identifying restrictions at a hypothesized value of the structural parameters. Other weak identification robust methods, such as conditional likelihood ratio or score tests (Moreira, 2003; Kleibergen, 2005), are more powerful than the S test under strong identification, but they are technically more involved and computationally more demanding. We do not report results based on those tests because in all the cases that we considered they gave similar results to the S test.

S sets are obtained by inverting the S test, i.e., by performing an S test for each candidate value of the parameters in the parameter region and collecting all the points that are not rejected at the given significance level. Unlike Wald sets, which are elliptical and can be computed analytically, S sets need to be computed by grid search over the parameter space and they can be disjoint. In this exercise, we use the same parameter region as the one that was used for the cloud plots (which includes over 90% of all point estimates), namely, $\lambda \in [-0.3, 0.3]$ and $\gamma_f \in [-1, 2]$. For each specification, we evaluate the test at over 1000 grid points. Because this procedure is computationally intensive, we consider only a subset of all the specifications listed in Table 4, consisting of about 1400 specifications, see Table 8.⁴⁵ The cloud of point estimates for the specifications in Table 8 is qualitatively similar to that for the full set of specifications in Table 4. Perhaps not surprisingly, the union of the joint 90% S sets for all specifications in Table 8 covers the entire parameter region in our plots.⁴⁶ These findings are detailed in section A.5 of the Appendix.

To get a sense of the impact of different specification choices on sampling uncertainty, we compare the average size of 90% and 95% S sets across different specification choices (see section A.5 in the Appendix for details). The S sets are generally quite large, covering on average between 1/3 (90% level) and 1/2 (95% level) of the parameter space for both labor share and output gap specifications.⁴⁷ However, there is systematic variation in size across specification choices. With regards to the impact of adding lags of inflation to the NKPC, the size of the S sets becomes progressively larger as we move from 0 to 4 lags. The difference between the pure and 1-lag hybrid NKPC is small, but adding three more lags of inflation to the hybrid NKPC roughly doubles the S sets, on average. The size of the S sets is smaller over the full sample than over pre- and post-1984 subsamples, as expected, but pre-1984 S sets are smaller than post-1984 S sets. More striking differences arise when we compare RE GIV to survey inflation expectations and when we compare different instrument sets. S sets for RE GIV are on average much larger than for

⁴⁵Computation of the S sets for VAR-GMM takes about 100 times longer than using the other single-equation methods. Therefore, we only consider 16 specifications that impose the VAR assumption.

⁴⁶The union of the S sets may be formally interpreted as a projection-based grand S set that projects over a latent hyperparameter which indexes the different specifications.

⁴⁷Additionally, the S sets are, on average, between 3 to 7 times larger than the corresponding Wald ellipses.

surveys, as anticipated in the discussion of first-stage F statistics above, and it looks like most of the difference arises from using GB forecasts. With regards to the different instruments, RT (external) instruments yield the largest S sets covering 50–80% of the parameter space. These are almost double the size of the S sets for exactly identified models, which are the smallest. Use of lagged survey forecasts as instruments produces on average smaller S sets than using lags of realized inflation, as conjectured by Wright (2009).

It is interesting to compute how often the S sets for (λ, γ_f) lie entirely in the positive orthant, as would be required to find significant evidence of forward-looking behavior. First, recall that S sets can be empty, which would indicate violation of the model’s overidentifying restrictions, but the frequency of empty S sets (for the overidentified specifications) is considerably below the nominal significance level, so there is no systematic evidence against the validity of the identifying restrictions. Jointly significantly positive coefficients λ and γ_f occur in a very small fraction of the specifications considered (less than 5% at the 10% level).⁴⁸ This happens more frequently when the output gap rather than the labor share is used as the forcing variable. Interestingly, when the forcing variable is the output gap, we obtain significantly positive coefficients only when survey forecasts are used to proxy for inflation expectations, whereas when it is the labor share, the occurrence of positive S sets is equally (un)likely for survey and RE GIV specifications. Significantly positive coefficients almost never arise when 4 lags of inflation are included in the NKPC, or when real-time instruments are used. Detailed results are provided in section A.5 in the Appendix.

Next, we draw 90% S sets and Wald confidence ellipses (based on the CUE) for (λ, γ_f) in the NKPC for a number of different specifications. The complete collection of robust confidence sets can be accessed using our interactive Matlab plotting tool in the online supplement (cf. footnote 38). Figure 11 reports the results for specifications based on GDP deflator inflation using either the labor share (NFB) or output gap (CBO) as forcing variables, imposing the restriction that inflation coefficients sum to 1, and using the “small” instrument set (three lags of $\Delta\pi_t$ and x_t). Three samples are considered: the full available sample, and the pre-1984 and post-1984 subsamples. The confidence sets are not completely uninformative, and they are particularly tight along the λ axis over the full sample, but rather wide across the γ_f axis. All S sets (and most Wald ellipses) contain $\lambda = 0$. For most specifications, identification is sufficiently weak for the results to be consistent both with the view that there is no forward-looking behavior, i.e., no role for expectations in price setting, as well as with the view that expectations matter a lot. Martins and Gabriel (2009) and Kleibergen and Mavroeidis (2009) reach similar conclusions. Regarding subsample variation, even though the point estimates differ considerably across the pre- and post-1984 samples, the sampling uncertainty is so large that we cannot infer that the coefficients have changed over time.⁴⁹

Figure 12 reports confidence sets for the full sample based on the assumption that the reduced form is a VAR(3) in the change in inflation and the forcing variable. The point estimate of γ_f is larger than 1 when either the labor share or the output gap are used as forcing variables, and the confidence sets are considerably tighter than for the corresponding GIV specification – compare with the top row of Figure 11. Hence, the VAR assumption appears to be informative in these specifications, consistent with the results in Magnusson and Mavroeidis (2010). However, we should

⁴⁸Specifically, this is the fraction of the S sets that are non-empty and lie entirely in the area $(\lambda, \gamma_f) \in (0, .3] \times (0, 2]$.

⁴⁹Using weak identification robust stability tests, Kleibergen and Mavroeidis (2009) find some evidence of instability before 1984 when they use a shorter pre-1984 sample, but no evidence of instability using their full 1960–2008 sample, which is consistent with Magnusson and Mavroeidis (2012).

stress that, due to computational limitations, we have only looked at very few VAR specifications, so this result should be viewed as tentative. A more thorough investigation is needed in order to assess the validity of the VAR assumption.

Figure 13 reports confidence sets based on survey specifications. Results are reported for SPF and GB GDP deflator inflation forecasts, as well as SPF CPI inflation forecasts. For SPF GDP deflator forecasts, the results are quite similar to the corresponding RE GIV specifications, given in the bottom row of Figure 11 (post-1984 sample). However, when we use the GB forecasts, confidence sets become considerably smaller. S sets based on SPF CPI forecasts are comparable in size and have considerable overlap with those based on SPF GDP deflator forecasts. Results for SPF GDP deflator inflation specifications over the pre- and post-1984 samples look very different. In particular, the 90% Wald ellipses do not overlap, which, if identification were strong, would suggest time-variation in the coefficients of the NKPC, as suggested by Zhang et al. (2008). However, the S sets do overlap considerably over the two subsamples, so it is not clear whether the survey-based NKPC is unstable.

To assess the empirical success of the external instruments approach, we plot robust confidence sets for GDP deflator inflation using the real-time (RT) instrument set in Figure 14. These figures are comparable to the top row in Figure 11, although the sample starts in 1971 due to data availability. Figure 14 demonstrates the unfortunate fact that the most plausibly exogenous instrument set also results in very weak identification, as the 90% robust confidence sets contain all reasonable γ_f values. Post-1984 confidence sets (not reported) are even larger.

5.4 Nesting RE and survey expectations

Finally, we assess the relative importance of rational and survey expectations in the NKPC, as studied by Nunes (2010), Fuhrer and Olivei (2010) and Fuhrer (2012), cf. section 3.1. Figure 15 reports CUE estimates and 90% S sets for the coefficients of future inflation (RE) and time- t dated one-quarter-ahead forecasts of inflation in the model

$$\pi_t = \lambda x_t + \gamma_{RE} \pi_{t+1} + \gamma_s \pi_{t+1|t}^s + \gamma_b \pi_{t-1} + \tilde{u}_t. \quad (25)$$

The coefficient λ here is treated as well-identified, and it is concentrated out.⁵⁰ We consider both SPF and GB GDP deflator inflation forecasts over the full available samples as well as a sample that starts in 1984q1. The instrument set is the same as in Nunes (2010), i.e., GGLS (see Table 4) plus two lags of survey inflation forecasts. The point estimates generally indicate a dominant role for RE, consistent with the evidence in Nunes (2010), and different from the preferred estimates in Fuhrer and Olivei (2010) and Fuhrer (2012). However, as acknowledged by Nunes (2010), sampling uncertainty is very large, and there is considerable sensitivity to data and estimation sample. Only when we use the labor share as forcing variable and the full available sample can we conclude that the RE term is dominant. Interestingly, all the confidence sets exclude $\gamma_{RE} = \gamma_s = 0$.⁵¹

⁵⁰This is a reasonable assumption since in all cases the first-stage F_{HAR} statistic for x_t is over 100.

⁵¹A number of authors have argued that non-rational expectations can account for the intrinsic persistence in the NKPC (Roberts, 1997; Brissimis and Magginas, 2008; Nunes, 2010). Estimation results that impose $\gamma_b = 0$ (not reported) are quite similar to the ones reported.

6 Conclusion

Based on the foregoing comparison of more than 100 papers from the literature with our analysis of thousands of a priori reasonable new Keynesian Phillips curve specifications estimated on U.S. data, we reach six main conclusions.

First, estimation of the NKPC using macro data is subject to a severe weak instruments problem. Consequently, seemingly innocuous specification changes lead to big differences in point estimates. The specification sensitivity is even larger than what has been reported in the literature. Moreover, given a choice of specification, sampling uncertainty is typically large, as weak identification robust confidence sets often cover a substantial part of the parameter space. While these findings are purely empirical, there are good theoretical explanations for why identification of the NKPC is weak.

Second, we do not reject the NKPC – far from it. However, we are unable to pin down the role of expectations in the inflation process sufficiently accurately for the results to be useful for policy analysis. The evidence is consistent both with the view that expectations matter a lot, as well as with the opposite view that they matter very little.

Third, because standard inference methods and efficiency comparisons are unreliable, weak identification robust methods should be used when possible. Weak identification is not a GMM-specific problem.

Fourth, estimation methods that rely on the assumption that inflation expectations can be proxied by a reduced-form vector autoregression (VAR) typically point toward a much greater role for forward-looking expectations in price determination than do less restrictive estimators. We demonstrate that VAR-based inference can be spurious when identification is weak. Because the VAR assumption is not innocuous, we recommend that VAR estimates be compared to non-VAR estimates when possible.

Fifth, it is hard to interpret the empirical results from specifications that use survey forecasts to proxy for inflation expectations. They often appear to be more strongly identified than other types of specifications, but they are particularly sensitive to the choice of forecast source, sample and inflation series. Moreover, the survey forecast specification of the NKPC is not microfounded unless the forecasts are rational, which does not seem to hold empirically. It is an interesting topic for future research to develop an internally consistent framework for analyzing inflation dynamics under non-rational expectation formation.

Sixth, researchers should be aware of the large and frequent revisions to NKPC data. We have proposed an estimation method that uses revisions as external instruments. While its assumptions are appealingly unrestrictive, it does not yield informative empirical results.

The evidence we present in this paper leads us to conclude that identification of the NKPC is too weak to warrant research on conceptually minor extensions. Issues related to the choice of explanatory variables, instruments, alternate data constructions and small modifications of the model are likely to be dwarfed by identification problems. Instead, we think it will be more fruitful to explore fundamentally new sources of identification, such as micro/sectoral data, cross-country models, information from large data sets and stability restrictions. Some recent papers have taken up this challenge, and we hope more will follow. The onus is not purely on applied researchers; theoretical macroeconomists can help by developing models that can be taken to the data in ways that directly address the identification issue.

A Appendix

A.1 Calibration of impulse responses

The model used to generate the impulse responses in Figure 1 is based on the canonical three-equation new Keynesian framework as described by Galí (2008):

$$\begin{aligned}\pi_t &= \gamma_f E_t(\pi_{t+1}) + (1 - \gamma_f)\pi_{t-1} + \lambda mc_t, \\ x_t &= -\frac{1}{\sigma}(i_t - E_t(\pi_{t+1}) - \rho) + E_t(x_{t+1}), \quad mc_t = \left(\sigma + \frac{\varphi + \alpha}{1 - \alpha}\right)x_t, \\ i_t &= \rho + \phi_\pi \pi_t + \phi_x x_t + v_t^i, \quad v_t^i = \rho_v v_{t-1}^i + \varepsilon_t^v.\end{aligned}\tag{26}$$

The first equation is a hybrid NKPC, the second is the dynamic IS curve, the third relates log real marginal cost (in deviation from the zero inflation steady state) mc_t to the log output gap x_t , the fourth is a Taylor rule for the nominal interest rate i_t , and the fifth equation specifies that the Taylor rule disturbance v_t^i follows an AR(1) process. We call ε_t^v the monetary policy shock.

When calibrating the structural parameters of the model, we use the benchmark values in Galí (2008, p. 52). The rate of time preference is $\rho = -\log(0.99)$, the elasticity of intertemporal substitution is $1/\sigma = 1$, the Frisch elasticity of labor supply is $\varphi = 1$, the labor exponent in the production function is $1 - \alpha = 2/3$, the Taylor rule coefficients are $(\phi_\pi, \phi_x) = (1.5, 0.5/4)$, and the AR(1) coefficient for the Taylor rule disturbance is $\rho_v = 0.5$.

A.2 Econometric methods

A.2.1 GMM estimation and optimal instruments

GMM estimation can be briefly described as follows. Let $f_T(\vartheta)$ denote sample moments, whose expectation vanishes at the true value of the parameters. For example, for the moment conditions (13) we set $f_T(\vartheta) = T^{-1} \sum_{t=1}^T Z_t h_t(\vartheta)$. Define the GMM objective function

$$S_T(\vartheta, \bar{\vartheta}) = f_T(\vartheta)' W_T(\bar{\vartheta}) f_T(\vartheta),\tag{27}$$

where $\bar{\vartheta}$ is some preliminary estimator of ϑ , and W_T is a weighting matrix that may depend on the data and on ϑ . A GMM estimator is the minimizer of $S_T(\vartheta, \bar{\vartheta})$ with respect to ϑ , if it exists. Given the particular choice of moments $f_T(\vartheta)$, efficient GMM estimation requires $W_T(\bar{\vartheta})$ to be a consistent estimator of the inverse of the variance of $\sqrt{T}f_T(\vartheta)$ – the long-run variance of the moment conditions. The most commonly used GMM estimator is a 2-step estimator, where the preliminary estimator $\bar{\vartheta}$ is obtained using some weight matrix that does not depend on ϑ . When the moment conditions are linear, $\bar{\vartheta}$ may be obtained using two-stage least squares.

Setting $\bar{\vartheta} = \vartheta$, so the efficient weight matrix estimator $W_T(\vartheta)$ is evaluated at the same parameters as the sample moments $f_T(\vartheta)$, yields the continuously updated estimator (CUE), which was proposed by Hansen et al. (1996). 2-step GMM and CUE are asymptotically equivalent under strong identification, but the latter has certain advantages under weak identification (see, e.g., Stock et al., 2002).

Optimal instruments When identification is given by conditional moment restrictions of the form $E_{t-1}[h_t(\vartheta_0)] = 0$, where $h_t(\cdot)$ is an $s \times 1$ vector-valued function, there is an infinite number of predetermined variables Z_t that can be used as instruments to form unconditional moment restrictions $E[Z_t h_t(\vartheta_0)] = 0$. Efficiency (under strong identification) in the class of all GMM estimators amounts to choosing the instruments in a way that minimizes the asymptotic variance of the GMM estimator among all possible instruments $Z_t \in \mathcal{I}_{t-1}$, where $\mathcal{I}_{t-1}$ denotes the information set at time $t - 1$. If the residual function $h_t(\vartheta_0)$ is a martingale difference sequence (MDS), the optimal instruments are given by

$$Z_t^o = \left[E_{t-1} \left(\frac{\partial h_t(\vartheta_0)}{\partial \vartheta'} \right) \right]' \{ E_{t-1} [h_t(\vartheta_0) h_t(\vartheta_0)'] \}^{-1}, \quad (28)$$

see Chamberlain (1987).

When estimating the NKPC by VAR-GMM, the two-dimensional residual vector (16) satisfies the conditional moment restriction $E_{t-1}[\tilde{h}_t(\vartheta_0)] = 0$. The residual vector is a MDS because it is adapted to the information set at time t . Moreover, the VAR assumption implies that $E_{t-1}[\partial \tilde{h}_t(\vartheta_0) / \partial \vartheta']$ is spanned by Y_{t-1} . So, applying the formula for the optimal instruments (28), we see that under conditional homoskedasticity, the optimal instruments are spanned by Y_{t-1} .

In the case of GIV estimation of the NKPC, the residuals are not adapted to $\mathcal{I}_t$ since $h_t(\vartheta_0) = \tilde{u}_t \in \mathcal{I}_{t+1}$, see equation (12). Under the assumption $E_{t-1}(u_t) = 0$, $h_t(\vartheta_0)$ can be represented as a moving average of order 1, e.g., $h_t(\vartheta_0) = v_t - \varphi v_{t+1} = \varphi(L^{-1})v_t$, say, where v_t is an MDS with $E_{t-1}(v_t) = 0$. Following Hayashi and Sims (1983), the optimal instruments can be obtained as follows. First, forward-filter $h_t(\vartheta_0)$ to get $v_t = \varphi(L^{-1})^{-1}h_t(\vartheta_0)$. Then compute the optimal instruments Z_t^o for $E_{t-1}(v_t) = 0$ using (28). Finally, transform these instruments to the optimal instruments $\tilde{Z}_t^o$ for $E_{t-1}[h_t(\vartheta_0)] = 0$, which are given by $\tilde{Z}_t^o = \varphi(L)^{-1}Z_t^o = \sum_{j=0}^{\infty} \varphi^j Z_{t-j}^o$.

A.2.2 GIV estimation with iterated instruments

Rudd and Whelan (2005) suggested the following alternative to the Galí and Gertler (1999) approach. Iterating equation (8) q periods forward using $E_t(u_{t+j}) = 0$, $j > 0$, and the law of iterated expectations, we get

$$\pi_t = \beta^{q+1} E_t(\pi_{t+q+1}) + \lambda \sum_{j=0}^q \beta^j E_t(x_{t+j}) + u_t. \quad (29)$$

Rudd and Whelan (2005) use the GIV approach to estimate the above relation. We now point out how the iterated method relates to the previously described Galí and Gertler (1999) procedure. Using the definition of the residual $h_t(\vartheta)$ in (12), equation (29) can be equivalently written as

$$E_t \left[\sum_{j=0}^q \beta^j h_{t+j}(\vartheta) \right] = u_t.$$

The identifying restriction $E_{t-1}(u_t) = 0$ then implies the unconditional moment restrictions

$$E \left[\sum_{j=0}^q \beta^j Y_{t-1} h_{t+j}(\vartheta) \right] = 0, \quad (30)$$

where Y_{t-1} is a vector of lags of π_t , x_t and any other variables used in the analysis. If we further assume that the distribution of the data is stationary, so that $E[Y_{t-1} h_{t+j}(\vartheta)] = E[Y_{t-j-1} h_t(\vartheta)]$, then equation (30) is equivalent to

$$E \left[\left(\sum_{j=0}^q \beta^j Y_{t-j-1} \right) h_t(\vartheta) \right] = 0. \quad (31)$$

This makes it clear that the only difference between the iterated moment conditions (30) and the difference equation moment conditions (13) is in the choice of instruments. That is, the underlying identifying assumption $E_{t-1}(u_t) = 0$ is the same, but each method uses a different subset of all admissible instruments.

A.2.3 Alternative VAR estimators

VAR-MD This approach was introduced by Campbell and Shiller (1987) for the estimation of asset pricing models and was popularized in the NKPC literature by the work of Sbordone (2002, 2005, 2006). It can be described briefly as follows. The structural model (8) implies restrictions on the reduced-form VAR coefficients A in (14). These restrictions can be written as $g(A, \vartheta) = 0$, where g is a vector-valued “distance” function. Typically, the number of restrictions exceeds the number of structural parameters, so the minimum distance estimator is defined as the minimizer of the objective function

$$g(\hat{A}, \vartheta)' W g(\hat{A}, \vartheta),$$

where $\hat{A}$ is a consistent first-step estimator of the reduced-form parameters (such as the OLS estimator), and W is a possibly random weight matrix. The optimal choice of W is a consistent estimator of the inverse of the asymptotic variance of $g(\hat{A}, \vartheta)$.

Define e_π and e_x to be the unit vectors with 1 in the position of π_t and x_t in Y_t , respectively. If we take time- $(t-1)$ expectations on both sides of the difference equation specification (8) and use the VAR implication $E_t(Y_{t+1}) = AY_t$, we obtain the parameter restrictions

$$g_1(A, \vartheta)' \equiv e'_\pi A - \beta e'_\pi A^2 - \lambda e'_x A = 0. \quad (32)$$

Note that $A = E(Y_t Y'_{t-1}) E(Y_{t-1} Y'_{t-1})^{-1} = E(Y_{t+1} Y'_t) E(Y_t Y'_t)^{-1}$, so that (32) can be equivalently written as

$$E[Y_{t-1}(\pi_t - \beta Y'_t \zeta - \lambda x_t)] = 0, \quad (33)$$

where $\zeta = A'e_\pi$ are the coefficients of the projection of π_{t+1} on Y_t . These are exactly the moment conditions for VAR-GMM given in (15).⁵²

The VAR-MD distance function is not unique. If we iterate the pure NKPC (8) forward an infinite number of times, we obtain the so-called “closed-form” solution

$$\pi_t = \lambda \sum_{j=0}^{\infty} \beta^j E_t(x_{t+j}) + u_t, \quad (34)$$

provided the series converges and the terminal condition $\lim_{\tau \rightarrow \infty} E_t(\beta^\tau \pi_{t+\tau}) = 0$ holds. Using $E_t(Y_{t+j}) = A^j Y_t$, we can write (34) as

$$\pi_t = \lambda (I - \beta A)^{-1} x_t + u_t.$$

The assumption $E_{t-1}(u_t) = 0$ implies the restrictions

$$g_2(A, \vartheta)' = e'_\pi A - \lambda e'_x A (I - \beta A)^{-1} = 0. \quad (35)$$

The distance function $g_1(A, \vartheta)$ defined in (32) satisfies $g_1(A, \vartheta) = (I - \beta A)' g_2(A, \vartheta)$. Because $(I - \beta A)$ is nonsingular, the two sets of restrictions are equivalent. However, the MD estimator is not invariant to nonlinear transformations of the distance function, so in finite samples the choice of distance function matters.⁵³

VAR-ML Suppose

$$z_t = \sum_{j=1}^l A_j z_{t-j} + v_t \quad (36)$$

denotes the l -th order VAR of z_t , whose companion representation was given in equation (14) above, where A_j are $n \times n$ coefficient matrices, and v_t is a $n \times 1$ vector of reduced-form errors. We have omitted deterministic terms for simplicity. An alternative to MD estimation is to maximize the likelihood function of the finite-order VAR (36) subject to the cross-equation restrictions (32) implied by the structural NKPC (8).

ML estimation of the constrained VAR is typically implemented by solving out the equality constraints (32) to express some of the reduced-form parameters in the likelihood in terms of the structural parameters, ϑ , and the remaining reduced-form parameters. Denoting the latter as ψ , the restricted reduced-form coefficients can be expressed as $A_j(\vartheta, \psi)$, $j = 1, \dots, l$. Assume, as in most of the literature, that the VAR errors are i.i.d. Gaussian and homoskedastic, i.e., $v_t \sim N(0, \Omega)$, where Ω is a $l \times l$ positive definite variance matrix. After concentrating with respect to Ω , the log-likelihood function can be written as

$$L(\vartheta, \psi) = \text{constant} - \frac{T}{2} \log \det [\hat{\Omega}(\vartheta, \psi)], \quad (37)$$

---

⁵²The VAR-GMM estimator reduces to the VAR-MD estimator for a particular (inefficient) choice of block diagonal GMM weight matrix. Thus, VAR-MD can be viewed as a variant of VAR-GMM, which imposes that the OLS moment conditions $E[(Y_t - AY_{t-1})Y'_{t-1}] = 0$ for the VAR companion matrix A hold exactly.

⁵³This was briefly discussed in Sbordone (2005), and in more detail in Barnes et al. (2011). An exception occurs when the model is just-identified and the equations $g_1(\hat{A}, \vartheta) = 0$ and $g_2(\hat{A}, \vartheta) = 0$ can be solved for ϑ as a function of $\hat{A}$, in which case the VAR-MD estimator does not depend on the choice of distance function.

where $\hat{\Omega}(\vartheta, \psi) = T^{-1} \sum_{t=l+1}^T \hat{v}_t(\vartheta, \psi) \hat{v}_t(\vartheta, \psi)'$ and $\hat{v}_t(\vartheta, \psi) = z_t - \sum_{j=1}^l A_j(\vartheta, \psi) z_{t-j}$. The choice of ψ is not unique, i.e., there are several ways of imposing the restrictions (32) on the likelihood. The pioneering approach by Fuhrer and Moore (1995) chooses ψ to be all the coefficients in the VAR except those corresponding to the equation for inflation. Computation of $A_j(\vartheta, \psi)$ then requires solving for the reduced-form coefficients in the inflation equation as functions of all other structural and reduced-form parameters. There are generically multiple solutions to this problem, so this mapping is not unique, and evaluating the likelihood (37) at all of the possible VAR solutions can be impractical, see Kurmann (2007). Fuhrer and Moore (1995) circumvent this issue by restricting the parameter space to the *determinacy region*, which by definition contains the parameter combinations for which there is a unique stable VAR solution.

An alternative approach, proposed by Kurmann (2007), is to set ψ equal to all the reduced-form VAR coefficients except those corresponding to the equation for the forcing variable x_t . He shows that the mapping $A_j(\vartheta, \psi)$ is then unique, except on a set of measure zero, and so evaluation of the likelihood (37) is straightforward on the entire parameter space, also outside the determinacy region. Inside the determinacy region the method gives the same likelihood as the Fuhrer-Moore approach. The following example from Kurmann (2007) illustrates. Suppose the reduced form is a VAR(1) in $(\pi_t, x_t)'$:

$$\begin{aligned}\pi_t &= a_{\pi\pi}\pi_{t-1} + a_{\pi x}x_{t-1} + v_{\pi t}, \\ x_t &= a_{x\pi}\pi_{t-1} + a_{xx}x_{t-1} + v_{xt},\end{aligned}\tag{38}$$

where $v_{\pi t}$ and v_{xt} are i.i.d. reduced-form shocks. The restrictions (32) can be expressed as

$$a_{x\pi}(\lambda + \beta a_{\pi x}) = (1 - \beta a_{\pi\pi})a_{\pi\pi}, \quad a_{xx}(\lambda + \beta a_{\pi x}) = (1 - \beta a_{\pi\pi})a_{\pi x}.$$

If we solve these equations for $a_{\pi\pi}, a_{\pi x}$ as functions of $\vartheta = (\beta, \lambda)'$ and $\psi = (a_{x\pi}, a_{xx})'$, then it can be shown that there are generically three solutions, see Kurmann (2007, sec. 2).⁵⁴ If we instead set $\psi = (a_{\pi\pi}, a_{\pi x})'$ and solve for the reduced-form parameters $a_{x\pi}, a_{xx}$, then there is a unique solution unless $\lambda + \beta a_{\pi x} = 0$.

Relationship between VAR methods Let $\hat{A}$ be the OLS estimator of the VAR coefficients. VAR-ML can be thought of as minimizing the distance between $\hat{A}$ and $A(\vartheta, \psi)$ with respect to ϑ and ψ . VAR-MD instead sets ψ equal to its OLS estimator and only minimizes this distance with respect to ϑ . Thus, the relationship between VAR-MD and VAR-ML is analogous to the relationship between 2SLS and LIML, respectively, in the textbook linear IV model (Fukač and Pagan, 2010). The analogy suggests that VAR-ML and VAR-MD should be asymptotically equivalent under strong identification, but not so under weak identification. Moreover, computation of VAR-MD (like 2SLS) is easier than VAR-ML (like LIML). Another difference is that VAR-ML is invariant to nonlinear transformations, i.e., it gives the same results in finite samples whether we specify the model as a difference equation (8) or in closed form (34).

An advantage of the VAR-GMM estimator relative to VAR-MD and VAR-ML is that it is easy to add zero restrictions to the coefficients of the reduced-form VAR so as to avoid many instrument issues. For instance, if you want to use four lags of inflation but only two lags of x_t and other

⁵⁴The multiplicity of solutions increases with the dimension of the VAR.

variables in the VAR, as in Galí et al. (2001), you just need to include only those variables in Y_{t-1} in the moment conditions (15)–(16). Hence, it is straightforward to check the implications of imposing the VAR assumption given any choice of instruments, as we do in our empirical section.

A.2.4 External instruments

Consider a generalized version of the model (9) that does not place any exclusion restrictions on the lags (we assume w_t in eq. (9) is part of Y_{t-1}):

$$\pi_t = \lambda x_t + \gamma_f \pi_{t+1}^e + \delta' Y_{t-1} + u_t. \quad (39)$$

We use π_{t+1}^e to denote inflation expectations so as to allow for the possibility that these may not be rational. Define Y_t^s to be the vintage- s observation, i.e., the statistical agency's estimate of Y_t published at time s . Variables without superscripts denote the latent true values of the series. Re-arrange (39):

$$\pi_t = \lambda x_t + \gamma_f \pi_{t+1} + \delta' Y_{t-1} + \tilde{u}_t, \quad \tilde{u}_t = u_t + \gamma_f (\pi_{t+1}^e - \pi_{t+1}).$$

Suppose that

$$E(u_t | Y_{t-1}, \{Y_{t-1}^r, \text{all } r\}) = 0. \quad (40)$$

This assumption can be interpreted as saying that the only way that data revisions enter the model is through their use in forming expectations. For Y_{t-1}^r to be valid instruments it must be the case that

$$E(\tilde{u}_t | Y_{t-1}, Y_{t-1}^r) = 0. \quad (41)$$

If $\gamma_f = 0$, then $\tilde{u}_t = u_t$ and data revisions are exogenous by (40). Whether they are relevant is an empirical issue, and depends on the extent to which expectations are formed using published data.

If instead $\gamma_f \neq 0$, things are more complicated. Let π_{t+1}^* denote the rational expectation of π_{t+1} , and suppose that

$$\pi_{t+1}^e = \pi_{t+1}^* + \zeta_t, \quad \text{where } E(\zeta_t | Y_{t-1}^{t-1}) = 0, \quad \text{and} \quad (42)$$

$$E(\pi_{t+1} - \pi_{t+1}^* | Y_{t-1}^{t-1}) = 0. \quad (43)$$

Condition (42) holds if agents have rational expectations, in which case $\zeta_t = 0$, but it also holds under departures from RE, in which case ζ_t is some “opinion” that is orthogonal to observable vintage- $(t-1)$ data. Condition (43) says that the information set used to compute π_{t+1}^* contains Y_{t-1}^{t-1} . Under these conditions, it can be shown that $E(\tilde{u}_t | Y_{t-1}^{t-1}) = 0$. In other words, Y_{t-1}^{t-1} is an exogenous instrument in the model (39). But this treats Y_{t-1} as endogenous, which leaves the model underidentified (there are two more endogenous variables, π_{t+1} and x_t , than instruments). This identification problem can be “solved” by imposing some exclusion restrictions on elements of Y_{t-1} , though this goes against the idea of external instruments. Alternatively, if we replace Y_{t-1}^{t-1} in conditions (42) and (43) with $\{Y_{t-1}^r\}_{r \leq t-1}$, then we could use several vintages of Y_{t-1} as instruments. This would satisfy the order condition for identification, but those instruments are likely to be weak.

A.2.5 S test

The S statistic for testing the null hypothesis $H_0 : \vartheta = \vartheta_0$ is given by T times the value of the continuous updated GMM objective function (27) at ϑ_0 , i.e., $TS_T(\vartheta_0, \vartheta_0)$. Under H_0 and some regularity conditions, this statistic is asymptotically $\chi^2(k)$ with degrees of freedom equal to the number of moment restrictions (or instruments), irrespective of whether the model is identified or not. A $(1 - a)\%$ level S set is obtained by collecting all points ϑ_0 for which $S_T(\vartheta_0, \vartheta_0)$ does not exceed the $(1 - a)$ percentile of $\chi^2(k)$. When the model also contains exogenous and predetermined variables, e.g., w_t and π_{t-j} in (9), their coefficients are concentrated out in order to improve the power of the test, see Stock and Wright (2000, Theorem 3).

A.2.6 Weak identification robust Hansen test

The minimum value of the S statistic, $\min_{\vartheta} TS_T(\vartheta, \vartheta)$, coincides with Hansen's J test of overidentifying restrictions that is based on the continuous updated GMM objective function, see Hansen et al. (1996). Its strong-instruments asymptotic distribution under the null of correct specification is the usual $\chi^2(k - p)$, where k is the number of identifying restrictions and p is the total number of estimated parameters. Under weak instruments, the asymptotic distribution of this statistic is bounded by $\chi^2(k - q)$, where q is the number of coefficients on exogenous regressors (cf. Stock and Wright, 2000, Theorem 3). Hence, since $q < p$, a robust version of the test can be obtained using the larger critical values associated with quantiles of $\chi^2(k - q)$. The robust test is less powerful than the standard one, because it uses the same test statistic but larger critical values.

A.3 Simulation study

Here we give details on the simulations presented in section 3.3. The main parameter choices for our four DGPs are listed in Table 1. All DGPs set the intercepts in the reduced-form VAR equal to 0. The innovations $v_t = (v_{\pi t}, v_{xt})'$ are distributed i.i.d. Gaussian with mean zero and covariance matrix $\Omega = ((0.07, 0.03)', (0.03, 0.70)')$, a typical reduced-form estimate on quarterly U.S. data from 1960–2011. The last column in Table 1 lists the smallest eigenvalue of the population concentration matrix for the GIV specification. This is a measure of the strength of identification, for which higher values mean stronger identification. It can loosely be thought of as an analog of 2 times the smallest first-stage F statistic in homoskedastic linear IV (there are two endogenous regressors).

For DGPs 1a and 2a, the reduced-form coefficients ξ are set to values that are close to the OLS estimates on the above-mentioned sample, with x_t equal to the labor share. DGPs 1a–b are indeterminate, since none of the Blanchard and Kahn (1980) generalized eigenvalues are outside the unit circle. DGPs 2a–b are determinate.

The four estimators we consider are implemented as follows. GIV estimation uses efficient two-step linear GMM with instruments Y_{t-1} and the Newey and West (1987) HAC long-run variance estimator. VAR-GMM is based on efficient two-step GMM with heteroskedasticity robust weight matrix. Because the moment conditions (15) are not linear, we resort to numerical optimization, although we only have to optimize over the scalar parameter γ_f . VAR-MD uses a distance function of the difference equation type (32) and a two-step efficient procedure. The estimator is available in closed form. VAR-ML uses the Kurmann (2007) approach, described in subsection A.2.3, to optimize over the parameters $(\gamma_f, \lambda, c, \zeta)$. This requires numerical optimization, which we carry out

using Matlab’s `fmincon` routine. The optimizer is provided with analytical first derivatives of the log likelihood, and we consider eight different initial values per estimation.

Matlab code and a full documentation of our approach are available in the online supplement (cf. footnote 38). The documentation also provides a more comprehensive set of results, including additional DGPs (some with VARMA reduced form such that the VAR assumption does not hold), the behavior of λ estimators, and rejection frequencies for t-tests, overidentification tests and the S test.

A.4 Data description

Most series mentioned in Table 4 are either self-explanatory or described in section 5.1. Unless otherwise noted, the series are from the St. Louis Fed FRED database. All growth rates are logarithmic and quarterly. Here we give details on some of the more involved constructions. Complete data and transformation files are available in the online supplement (cf. footnote 38).

Wage and commodity price inflation, which we use as instruments, refer to the growth in business sector compensation per hour and commodity PPI inflation from the Bureau of Labor Statistics (BLS), respectively (the latter series is not seasonally adjusted). Interest rates are U.S. Treasury rates.

All forcing variables are in logs. “Output” refers to real GDP per capita. We estimate trends by fitting linear or quadratic polynomials in time, or by the HP or Baxter-King filters. The Baxter-King filtered gaps retain cycles of duration between 6 and 32 quarters. For the HP filter, we use a smoothing parameter (commonly referred to as λ) of 1,600 for output and 10,000 for the labor share. Our computation of *real-time* (one-sided) output gaps proceeds as follows. For every quarter, the series of real-time (i.e., then-current estimate of) output per capita is loaded. An AR(6) in changes is fitted to this series and used to generate forecasts several quarters ahead. The detrending routines are then applied to the concatenation of the real-time series and the generated forecasts. Pseudo-real-time labor share gaps are calculated somewhat differently from the output gaps. First, the data used is not actually real-time but is instead based on the latest vintage, as explained in the main text. Second, to better capture the marked decrease in the labor share in the latter half of the sample, the forecasting regression is an AR(15) in second differences.⁵⁵

The real-time labor share data set is gleaned from the BLS’s Productivity and Cost “Preliminary” news releases on nominal unit labor costs and the implicit price deflator. The data corresponds approximately to what was known around the middle of each quarter, like in the Philadelphia Fed’s real-time dataset. Data vintages from 1971q2 to 1993q4 have been manually typed in from scanned PDFs of the BLS news releases, available in the St. Louis Fed’s FRASER document database. Vintages from 1994q1 to 2001q2 are parsed from electronic news releases available on the BLS website. Finally, vintages from 2001q3 and onward are parsed from vintages of the BLS’ internal “edit 60” flat text file.

A.5 Additional empirical results

Table 9 lists median differences in estimates of λ and γ_f across different survey specifications and over different sub-samples. Table 10 reports a number of similar pairwise comparisons across other

---

⁵⁵For the business sector log labor share, the AIC selects 15 lags on the full sample.

specification choices. The results are discussed in section 5.2.

Figure 16 displays point estimates for all specifications listed in Table 8. For the collection of labor share specifications, the union of the joint 90% S sets (not shown) covers the entire plotted parameter space; the same is true for the collection of output gap specifications.

Table 11 reports the average size of 90% and 95% S sets as a fraction of the plotted parameter space $(\gamma_f, \lambda) \in [-2, 1] \times [-0.3, 0.3]$ for various specification choices and samples. Here ‘all’ refers to the different options listed in Table 8. The sample end date varies by series: for GB data the sample ends in 2005q4, while for all other series it stretches to 2011q4.

Table 12 reports some additional statistics associated with the S sets corresponding to the specifications of Table 8. Row “% empty S set” gives the fraction of the overidentified specifications for which the S sets are empty at the specified significance level. The rest of the rows give the frequency of non-empty and positive S sets for all specifications and for various subcategories.

References

- Adam, K. and Padula, M. (2011), ‘Inflation dynamics and subjective expectations in the United States’, *Economic Inquiry* **49**(1), 13–25.
- Adjemian, S., Bastani, H., Juillard, M., Mihoubi, F., Perendia, G., Ratto, M. and Villemot, S. (2011), ‘Dynare: Reference Manual, Version 4’, *Dynare Working Papers* (1).
- An, S. and Schorfheide, F. (2007), ‘Bayesian Analysis of DSGE Models’, *Econometric Reviews* **26**(2–4), 113–172.
- Anderson, G. and Moore, G. (1985), ‘A linear algebraic procedure for solving linear perfect foresight models’, *Economics Letters* **17**(3), 247–252.
- Anderson, T. W. and Rubin, H. (1949), ‘Estimation of the Parameters of a Single Equation in a Complete System of Stochastic Equations’, *The Annals of Mathematical Statistics* **20**(1), 46–63.
- Andrews, D. W. K. (1993), ‘Tests for Parameter Instability and Structural Change With Unknown Change Point’, *Econometrica* **61**(4), 821–856.
- Andrews, D. W. K. and Ploberger, W. (1994), ‘Optimal Tests when a Nuisance Parameter is Present Only Under the Alternative’, *Econometrica* **62**(6), 1383–1414.
- Andrews, D. W. and Stock, J. H. (2005), ‘Inference with Weak Instruments’, *National Bureau of Economic Research Technical Working Paper Series* (313).
- Andrews, D. W. and Stock, J. H. (2007), ‘Testing with many weak instruments’, *Journal of Econometrics* **138**(1), 24–46.
- Andrews, I. and Mikusheva, A. (2011), ‘Maximum Likelihood Inference in Weakly Identified Models’. Working paper, December version.
- Ang, A., Bekaert, G. and Wei, M. (2007), ‘Do macro variables, asset markets, or surveys forecast inflation better?’, *Journal of Monetary Economics* **54**(4), 1163–1212.
- Angeletos, G.-M. and La’O, J. (2009), ‘Incomplete information, higher-order beliefs and price inertia’, *Journal of Monetary Economics* **56**(Supplement), S19–S37.
- Ascari, G. (2004), ‘Staggered Prices and Trend Inflation: Some Nuisances’, *Review of Economic Dynamics* **7**(3), 642–667.
- Ascari, G. and Sbordone, A. (2013), ‘The macroeconomics of trend inflation’. Mimeo.
- Atkeson, A. and Ohanian, L. E. (2001), ‘Are Phillips Curves Useful for Forecasting Inflation?’, *Federal Reserve Bank of Minneapolis Quarterly Review* **25**(1), 2–11.
- Barnes, M. L., Gumbau-Brisa, F., Lie, D. and Olivei, G. P. (2011), ‘Estimation of Forward-Looking Relationships in Closed Form: An Application to the New Keynesian Phillips Curve’, *Federal Reserve Bank of Boston Working Papers* (11-3). June version.

- Batini, N., Jackson, B. and Nickell, S. (2005), ‘An open-economy new Keynesian Phillips curve for the U.K.’, *Journal of Monetary Economics* **52**(6), 1061–1071.
- Benigno, P. and López-Salido, J. D. (2006), ‘Inflation Persistence and Optimal Monetary Policy in the Euro Area’, *Journal of Money, Credit and Banking* **38**(3), 587–614.
- Bernanke, B. S. (2007), ‘Inflation Expectations and Inflation Forecasting’. Speech at the Monetary Economics Workshop of the National Bureau of Economic Research Summer Institute.
- Bernanke, B. S. (2008), ‘Outstanding Issues in the Analysis of Inflation’. Speech at the Federal Reserve Bank of Boston’s 53rd Annual Economic Conference.
- Beyer, A., Farmer, R. E. A., Henry, J. and Marcellino, M. (2008), ‘Factor analysis in a model with rational expectations’, *Econometrics Journal* **11**(2), 271–286.
- Blanchard, O. and Galí, J. (2007), ‘Real Wage Rigidities and the New Keynesian Model’, *Journal of Money, Credit and Banking* **39**, 35–65.
- Blanchard, O. and Galí, J. (2010), ‘Labor Markets and Monetary Policy: A New Keynesian Model with Unemployment’, *American Economic Journal: Macroeconomics* **2**(2), 1–30.
- Blanchard, O. J. and Kahn, C. M. (1980), ‘The Solution of Linear Difference Models under Rational Expectations’, *Econometrica* **48**(5), 1305–1311.
- Boug, P., Cappelen, r. and Swensen, A. R. (2010), ‘The new Keynesian Phillips curve revisited’, *Journal of Economic Dynamics and Control* **34**(5), 858–874.
- Bårdsen, G., Jansen, E. S. and Nymoen, R. (2004), ‘Econometric Evaluation of the New Keynesian Phillips Curve’, *Oxford Bulletin of Economics and Statistics* **66**, 671–686.
- Brissimis, S. N. and Magginas, N. S. (2008), ‘Inflation Forecasts and the New Keynesian Phillips Curve’, *International Journal of Central Banking* **4**(2), 1–22.
- Cagliarini, A., Robinson, T. and Tran, A. (2011), ‘Reconciling microeconomic and macroeconomic estimates of price stickiness’, *Journal of Macroeconomics* **33**(1), 102–120.
- Calvo, G. A. (1983), ‘Staggered prices in a utility-maximizing framework’, *Journal of Monetary Economics* **12**(3), 383–398.
- Campbell, J. and Shiller, R. (1988), ‘The dividend-price ratio and expectations of future dividends and discount factors’, *Review of Financial Studies* **1**(3), 195–228.
- Campbell, J. Y. and Shiller, R. J. (1987), ‘Cointegration and Tests of Present Value Models’, *Journal of Political Economy* **95**(5), 1062–1088.
- Caner, M. (2007), ‘Boundedly pivotal structural change tests in continuous updating GMM with strong, weak identification and completely unidentified cases’, *Journal of Econometrics* **137**(1), 28–67.

- Carriero, A. (2008), ‘A simple test of the New Keynesian Phillips Curve’, *Economics Letters* **100**(2), 241–244.
- Carvalho, C. (2006), ‘Heterogeneity in price stickiness and the real effects of monetary shocks’, *Frontiers in Macroeconomics* **6**(1).
- Castle, J. L., Doornik, J. A., Hendry, D. F. and Nymoen, R. (2010), ‘Testing the Invariance of Expectations Models of Inflation’. Working paper, November 9 version.
- Chamberlain, G. (1987), ‘Asymptotic efficiency in estimation with conditional moment restrictions’, *Journal of Econometrics* **34**(3), 305–334.
- Chowdhury, I., Hoffmann, M. and Schabert, A. (2006), ‘Inflation dynamics and the cost channel of monetary transmission’, *European Economic Review* **50**(4), 995–1016.
- Christiano, L. J., Eichenbaum, M. and Evans, C. L. (2005), ‘Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy’, *Journal of Political Economy* **113**(1), 1–45.
- Cochrane, J. (2011), ‘Determinacy and Identification with Taylor Rules’, *Journal of Political Economy* **119**(3), 565–615.
- Coenen, G., Levin, A. T. and Christoffel, K. (2007), ‘Identifying the influences of nominal and real rigidities in aggregate price-setting behavior’, *Journal of Monetary Economics* **54**(8), 2439–2466.
- Cogley, T. and Sbordone, A. M. (2008), ‘Trend Inflation, Indexation, and Inflation Persistence in the New Keynesian Phillips Curve’, *American Economic Review* **98**(5), 2101–26.
- Cornea, A., Hommes, C. H. and Massaro, D. (2013), ‘Behavioral Heterogeneity in U.S. Inflation Dynamics’. Working paper, January 4 version.
- Dees, S., Pesaran, M. H., Smith, L. V. and Smith, R. P. (2009), ‘Identification of New Keynesian Phillips Curves from a Global Perspective’, *Journal of Money, Credit and Banking* **41**(7), 1481–1502.
- Doornik, J. A. (2007), *Object-Oriented Matrix Programming Using Ox*, third edn, Timberlake Consultants Press, London.
- Dotsey, M., King, R. G. and Wolman, A. L. (1999), ‘State-Dependent Pricing And The General Equilibrium Dynamics Of Money And Output’, *The Quarterly Journal of Economics* **114**(2), 655–690.
- Dufour, J.-M. (2003), ‘Identification, weak instruments, and statistical inference in econometrics’, *Canadian Journal of Economics* **36**(4), 767–808.
- Dufour, J.-M., Khalaf, L. and Kichian, M. (2006), ‘Inflation dynamics and the New Keynesian Phillips Curve: An identification robust econometric analysis’, *Journal of Economic Dynamics and Control* **30**, 1707–1727.

- Dufour, J.-M., Khalaf, L. and Kichian, M. (2010a), ‘Estimation uncertainty in structural inflation models with real wage rigidities’, *Computational Statistics and Data Analysis* **54**(11), 2554–2561.
- Dufour, J.-M., Khalaf, L. and Kichian, M. (2010b), ‘On the precision of Calvo parameter estimates in structural NKPC models’, *Journal of Economic Dynamics and Control* **34**(9), 1582–1595.
- Eichenbaum, M. and Fisher, J. D. (2004), ‘Evaluating the Calvo Model of Sticky Prices’, *National Bureau of Economic Research Working Paper Series* (10617).
- Eichenbaum, M. and Fisher, J. D. (2007), ‘Estimating the frequency of price re-optimization in Calvo-style models’, *Journal of Monetary Economics* **54**(7), 2032–2047.
- Estrella, A. and Fuhrer, J. C. (2002), ‘Dynamic Inconsistencies: Counterfactual Implications of a Class of Rational-Expectations Models’, *The American Economic Review* **92**(4), 1013–1028.
- Fanelli, L. (2008), ‘Testing the New Keynesian Phillips Curve Through Vector Autoregressive Models: Results from the Euro Area’, *Oxford Bulletin of Economics and Statistics* **70**(1), 53–66.
- Fernández-Villaverde, J., Rubio-Ramírez, J. F., Sargent, T. J. and Watson, M. W. (2007), ‘ABCs (and Ds) of Understanding VARs’, *The American Economic Review* **97**(3), 1021–1026.
- Fischer, S. (1977), ‘Wage indexation and macroeconomic stability’, *Carnegie-Rochester Conference Series on Public Policy* **5**, 107–147.
- Fisher, I. (1973), ‘I Discovered the Phillips Curve: “A Statistical Relation between Unemployment and Price Changes”’, *The Journal of Political Economy* **81**(2), 496–502.
- Friedman, M. (1968), ‘The Role of Monetary Policy’, *The American Economic Review* **58**(1), 1–17.
- Fuhrer, J. C. (1997), ‘The (Un)Importance of Forward-Looking Behavior in Price Specifications’, *Journal of Money, Credit, and Banking* **29**(3), 338–350.
- Fuhrer, J. C. (2006), ‘Intrinsic and Inherited Inflation Persistence’, *International Journal of Central Banking* **2**(3), 49–86.
- Fuhrer, J. C. (2012), ‘The Role of Expectations in Inflation Dynamics’, *International Journal of Central Banking* **8**, 137–165.
- Fuhrer, J. C. and Moore, G. (1995), ‘Inflation Persistence’, *The Quarterly Journal of Economics* **110**(1), 127–159.
- Fuhrer, J. C. and Olivei, G. (2005), Estimating Forward-Looking Euler Equations with GMM and Maximum Likelihood Estimators: An Optimal Instruments Approach, in J. Faust, A. Orphanides and D. Reifschneider, eds, ‘Models and Monetary Policy: Research in the Tradition of Dale Henderson, Richard Porter, and Peter Tinsley’, Board of Governors of the Federal Reserve System.
- Fuhrer, J. C. and Olivei, G. P. (2010), ‘The Role of Expectations and Output in the Inflation Process: An Empirical Assessment’, *Federal Reserve Bank of Boston Public Policy Briefs* (10-2).

- Fukač, M. and Pagan, A. R. (2007), ‘Commentary on “An estimated DSGE model for the United Kingdom”’, *Federal Reserve Bank of St. Louis Review* **89**(4), 233–240.
- Fukač, M. and Pagan, A. R. (2010), ‘Limited information estimation and evaluation of DSGE models’, *Journal of Applied Econometrics* **25**(1), 55–70.
- Gagnon, E. and Khan, H. (2005), ‘New Phillips curve under alternative production technologies for Canada, the United States, and the Euro area’, *European Economic Review* **49**(6), 1571–1602.
- Galí, J. (2008), *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework*, Princeton University Press.
- Galí, J. and Gertler, M. (1999), ‘Inflation dynamics: A structural econometric analysis’, *Journal of Monetary Economics* **44**(2), 195–222.
- Galí, J., Gertler, M. and López-Salido, J. (2001), ‘European inflation dynamics’, *European Economic Review* **45**(7), 1237–1270.
- Galí, J., Gertler, M. and López-Salido, J. D. (2005), ‘Robustness of the estimates of the hybrid New Keynesian Phillips curve’, *Journal of Monetary Economics* **52**(6), 1107–1118.
- Gerberding, C. (2001), ‘The information content of survey data on expected price developments for monetary policy’, *Deutsche Bundesbank discussion paper* (9/01).
- Gordon, R. J. (1990), U.S. Inflation, Labor’s Share, and the Natural Rate of Unemployment, in H. König, ed., ‘Economics of Wage Determination’, New York: Springer Verlag, pp. 1–34.
- Gordon, R. J. (2011), ‘The History of the Phillips Curve: Consensus and Bifurcation’, *Economica* **78**(309), 10–50.
- Guay, A. and Pelgrin, F. (2005), ‘The U.S. New Keynesian Phillips Curve: An Empirical Assessment’. Working paper, September 7 version.
- Guerrieri, L., Gust, C. and López-Salido, J. D. (2010), ‘International Competition and Inflation: A New Keynesian Perspective’, *American Economic Journal: Macroeconomics* **2**(4), 247–280.
- Guggenberger, P. and Smith, R. J. (2008), ‘Generalized empirical likelihood tests in time series models with potential identification failure’, *Journal of Econometrics* **142**(1), 134–161.
- Gumbau-Brisa, F., Lie, D. and Olivei, G. P. (2011), ‘A Response to Cogley and Sbordone’s Comment on “Closed-Form Estimates of the New Keynesian Phillips Curve with Time-Varying Trend Inflation”’, *Federal Reserve Bank of Boston Working Papers* (11-4). June version.
- Gwin, C. R. and VanHoose, D. D. (2008), ‘Alternative measures of marginal cost and inflation in estimations of new Keynesian inflation dynamics’, *Journal of Macroeconomics* **30**(3), 928–940.
- Hansen, C., Hausman, J. and Newey, W. (2008), ‘Estimation With Many Instrumental Variables’, *Journal of Business & Economic Statistics* **26**, 398–422.

- Hansen, L. P. (1982), ‘Large Sample Properties of Generalized Method of Moments Estimators’, *Econometrica* **50**(4), 1029–1054.
- Hansen, L. P., Heaton, J. and Yaron, A. (1996), ‘Finite-Sample Properties of Some Alternative GMM Estimators’, *Journal of Business & Economic Statistics* **14**(3), 262–80.
- Hansen, L. P. and Singleton, K. J. (1982), ‘Generalized Instrumental Variables Estimation of Non-linear Rational Expectations Models’, *Econometrica* **50**(5), 1269–1286.
- Hayashi, F. and Sims, C. A. (1983), ‘Nearly Efficient Estimation of Time Series Models with Predetermined, but not Exogenous, Instruments’, *Econometrica* **51**(3), 783–798.
- Henry, S. G. B. and Pagan, A. R. (2004), ‘The Econometrics of the New Keynesian Policy Model: Introduction’, *Oxford Bulletin of Economics and Statistics* **66**, 581–607.
- Henzel, S. and Wollmershäuser, T. (2008), ‘The New Keynesian Phillips curve and the role of expectations: Evidence from the CESifo World Economic Survey’, *Economic Modelling* **25**(5), 811–832.
- Hornstein, A. (2007), ‘Evolving inflation dynamics and the New Keynesian Phillips curve’, *Federal Reserve Bank of Richmond Economic Quarterly* **93**(4), 317–339.
- Hume, D. (1752), Of money, in ‘Essays’, George Routledge and Sons.
- Imbs, J., Jondeau, E. and Pelgrin, F. (2011), ‘Sectoral Phillips curves and the aggregate Phillips curve’, *Journal of Monetary Economics* **58**(4), 328–344.
- Jondeau, E. and Le Bihan, H. (2003), ‘ML vs GMM Estimates of Hybrid Macroeconomic Models (With an Application to the New Phillips Curve)’, *Banque de France notes d’études et de recherche* (103).
- Jondeau, E. and Le Bihan, H. (2005), ‘Testing for the New Keynesian Phillips Curve. Additional international evidence’, *Economic Modelling* **22**(3), 521–550.
- Jondeau, E. and Le Bihan, H. (2008), ‘Examining bias in estimators of linear rational expectations models under misspecification’, *Journal of Econometrics* **143**(2), 375–395.
- Kapetanios, G., Khalaf, L. and Marcellino, M. (2011), ‘Factor based identification-robust inference in IV regressions’. Working paper, March 22 version.
- Kapetanios, G. and Marcellino, M. (2010), ‘Factor-GMM estimation with large sets of possibly weak instruments’, *Computational Statistics and Data Analysis* **54**(11), 2655–2675.
- Kiley, M. T. (2007), ‘A Quantitative Comparison of Sticky-Price and Sticky-Information Models of Price Setting’, *Journal of Money, Credit and Banking* **39**(1), 101–125.
- Kim, C.-J. and Kim, Y. (2008), ‘Is the Backward-Looking Component Important in a New Keynesian Phillips Curve?’, *Studies in Nonlinear Dynamics & Econometrics* **12**(3). Article 5.
- Kim, C.-J. and Manopimoke, P. (2011), ‘Trend Inflation and the New Keynesian Phillips Curve’. Working paper, July 11 version.

- Kleibergen, F. (2002), ‘Pivotal Statistics for Testing Structural Parameters in Instrumental Variables Regression’, *Econometrica* **70**(5), 1781–1803.
- Kleibergen, F. (2005), ‘Testing Parameters in GMM without Assuming That They Are Identified’, *Econometrica* **73**(4), 1103–1123.
- Kleibergen, F. and Mavroeidis, S. (2009), ‘Weak Instrument Robust Tests in GMM and the New Keynesian Phillips Curve’, *Journal of Business & Economic Statistics* **27**(3), 293–311.
- Kleibergen, F. and Mavroeidis, S. (2013), ‘Identification issues in limited-information Bayesian analysis of structural macroeconomic models’. Working paper, May 21 version.
- Koop, G. and Onorante, L. (2011), ‘Estimating Phillips Curves in Turbulent Times using the ECBs Survey of Professional Forecasters’. Working paper, February version.
- Korenok, O., Radchenko, S. and Swanson, N. R. (2010), ‘International evidence on the efficacy of new-Keynesian models of inflation persistence’, *Journal of Applied Econometrics* **25**(1), 31–54.
- Kozicki, S. and Tinsley, P. (2002), ‘Alternative Sources of the Lag Dynamics of Inflation’, *Federal Reserve Bank of Kansas City Research Working Paper* (02-12).
- Krause, M. U., Lopez-Salido, D. and Lubik, T. A. (2008), ‘Inflation dynamics with search frictions: A structural econometric analysis’, *Journal of Monetary Economics* **55**(5), 892–916.
- Kuester, K., Müller, G. J. and Stölting, S. (2009), ‘Is the New Keynesian Phillips curve flat?’, *Economics Letters* **103**(1), 39–41.
- Kurmann, A. (2005), ‘Quantifying the uncertainty about the fit of a new Keynesian pricing model’, *Journal of Monetary Economics* **52**(6), 1119–1134.
- Kurmann, A. (2007), ‘VAR-based estimation of Euler equations with an application to New Keynesian pricing’, *Journal of Economic Dynamics and Control* **31**(3), 767–796.
- Kurz, M. (2011), ‘A New Keynesian Model with Diverse Beliefs’. Working paper, September 7 version.
- Lindé, J. (2005), ‘Estimating New-Keynesian Phillips curves: A full information maximum likelihood approach’, *Journal of Monetary Economics* **52**(6), 1135–1149.
- Lubik, T. A. and Schorfheide, F. (2004), ‘Testing for Indeterminacy: An Application to U.S. Monetary Policy’, *American Economic Review* **94**(1), 190–217.
- Lucas, R. E. (1976), ‘Econometric policy evaluation: A critique’, *Carnegie-Rochester Conference Series on Public Policy* **1**, 19–46.
- Ma, A. (2002), ‘GMM estimation of the new Phillips curve’, *Economics Letters* **76**(3), 411–417.
- Magnusson, L. M. and Mavroeidis, S. (2010), ‘Identification-Robust Minimum Distance Estimation of the New Keynesian Phillips Curve’, *Journal of Money, Credit and Banking* **42**(2–3), 465–481.

- Magnusson, L. M. and Mavroeidis, S. (2012), ‘Identification using stability restrictions’. Working paper, December 4 version.
- Mankiw, N. G. (2001), ‘The Inexorable and Mysterious Tradeoff between Inflation and Unemployment’, *The Economic Journal* **111**(471), C45–C61.
- Mankiw, N. G. and Reis, R. (2011), Imperfect Information and Aggregate Supply, *in* B. M. Friedman and M. Woodford, eds, ‘Handbook of Monetary Economics’, Vol. 3, Elsevier, chapter 5, pp. 183–229.
- Mankiw, N. G., Reis, R. and Wolfers, J. (2004), Disagreement about Inflation Expectations, *in* M. Gertler and K. Rogoff, eds, ‘NBER Macroeconomics Annual 2003’, Vol. 18, The MIT Press, pp. 209–270.
- Martins, L. F. and Gabriel, V. J. (2009), ‘New Keynesian Phillips Curves and potential identification failures: A Generalized Empirical Likelihood analysis’, *Journal of Macroeconomics* **31**(4), 561–571.
- Mavroeidis, S. (2004), ‘Weak Identification of Forward-looking Models in Monetary Economics’, *Oxford Bulletin of Economics and Statistics* **66**(supplement), 609–635.
- Mavroeidis, S. (2005), ‘Identification Issues in Forward-Looking Models Estimated by GMM, with an Application to the Phillips Curve’, *Journal of Money, Credit and Banking* **37**(3), 421–48.
- Mavroeidis, S. (2010), ‘Monetary Policy Rules and Macroeconomic Stability: Some New Evidence’, *American Economic Review* **100**(1), 491–503.
- Mazumder, S. (2010), ‘The new Keynesian Phillips curve and the cyclicalities of marginal cost’, *Journal of Macroeconomics* **32**(3), 747–765.
- Mazumder, S. (2011), ‘The empirical validity of the New Keynesian Phillips curve using survey forecasts of inflation’, *Economic Modelling* **28**(6), 2439–2450.
- McAdam, P. and Willman, A. (2010), ‘Technology, Utilization and Inflation: Re-assessing the New Keynesian Fundamental’. Working paper, October 15 version.
- McCallum, B. T. (1976), ‘Rational Expectations and the Natural Rate Hypothesis: Some Consistent Estimates’, *Econometrica* **44**(1), 43–52.
- Mikusheva, A. (2009), ‘Comment’, *Journal of Business & Economic Statistics* **27**(3), 322–323.
- Mikusheva, A. (2013), ‘Survey on statistical inferences in weakly-identified instrumental variable models’, *Applied Econometrics* **29**(1), 117–131.
- Montiel Olea, J. L. and Pflueger, C. (2013), ‘A Robust Test for Weak Instruments’, *Journal of Business & Economic Statistics*. Forthcoming.
- Moreira, M. J. (2003), ‘A Conditional Likelihood Ratio Test for Structural Models’, *Econometrica* **71**(4), 1027–1048.

- Nakamura, E. and Steinsson, J. (2013), ‘Price Rigidity: Microeconomic Evidence and Macroeconomic Implications’, *National Bureau of Economic Research Working Paper Series* (18705).
- Nason, J. M. and Smith, G. W. (2008a), ‘Identifying the new Keynesian Phillips curve’, *Journal of Applied Econometrics* **23**(5), 525–551.
- Nason, J. M. and Smith, G. W. (2008b), ‘The New Keynesian Phillips Curve: Lessons From Single-Equation Econometric Estimation’, *Federal Reserve Bank of Richmond Economic Quarterly* **94**(4), 361–395.
- Neiss, K. S. and Nelson, E. (2005), ‘Inflation Dynamics, Marginal Cost, and the Output Gap: Evidence from Three Countries’, *Journal of Money, Credit and Banking* **37**(6), 1019–1045.
- Nelson, C. R. and Lee, J. (2007), ‘Expectation horizon and the Phillips Curve: the solution to an empirical puzzle’, *Journal of Applied Econometrics* **22**(1), 161–178.
- Newey, W. K. and McFadden, D. (1994), Large sample estimation and hypothesis testing, in R. F. Engle and D. L. McFadden, eds, ‘Handbook of Econometrics’, Vol. 4, Elsevier, pp. 2111–2245.
- Newey, W. K. and West, K. D. (1987), ‘A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix’, *Econometrica* **55**(3), 703–708.
- Newey, W. K. and Windmeijer, F. (2009), ‘Generalized Method of Moments With Many Weak Moment Conditions’, *Econometrica* **77**(3), 687–719.
- Nunes, R. (2010), ‘Inflation Dynamics: The Role of Expectations’, *Journal of Money, Credit and Banking* **42**(6), 1161–1172.
- Nymoen, R., Swensen, A. R. and Tveter, E. (2010), ‘The New Keynesian Phillips Curve: A meta-analysis’. Working paper, October 7 version.
- Ólafsson, T. T. (2006), ‘The New Keynesian Phillips Curve: In Search of Improvements and Adaptation to the Open Economy’, *Central Bank of Iceland Working Papers* (31).
- Paloviita, M. (2006), ‘Inflation Dynamics in the Euro Area and the Role of Expectations’, *Empirical Economics* **31**, 847–860.
- Paloviita, M. (2008), ‘Comparing alternative Phillips curve specifications: European results with survey-based expectations’, *Applied Economics* **40**(17), 2259–2270.
- Paloviita, M. and Mayes, D. (2005), ‘The use of real-time information in Phillips-curve relationships for the euro area’, *The North American Journal of Economics and Finance* **16**(3), 415–434.
- Phelps, E. S. (1967), ‘Phillips Curves, Expectations of Inflation and Optimal Unemployment over Time’, *Economica* **34**(135), 254–281.
- Phillips, A. W. (1958), ‘The relation between unemployment and the rate of change of money wage rates in the United Kingdom, 1861–1957’, *Economica* **25**, 283–299.

- Preston, B. (2005), ‘Learning about Monetary Policy Rules when Long-Horizon Expectations Matter’, *International Journal of Central Banking* **1**(2), 81–126.
- Ravenna, F. and Walsh, C. E. (2006), ‘Optimal monetary policy with the cost channel’, *Journal of Monetary Economics* **53**(2), 199–216.
- Ravenna, F. and Walsh, C. E. (2008), ‘Vacancies, unemployment, and the Phillips curve’, *European Economic Review* **52**(8), 1494–1521.
- Roberts, J. M. (1995), ‘New Keynesian Economics and the Phillips Curve’, *Journal of Money, Credit and Banking* **27**(4), 975–984.
- Roberts, J. M. (1997), ‘Is inflation sticky?’, *Journal of Monetary Economics* **39**(2), 173–196.
- Roberts, J. M. (2005), ‘How Well Does the New Keynesian Sticky-Price Model Fit the Data?’, *Contributions to Macroeconomics* **5**. Issue 1, Article 10.
- Rotemberg, J. J. (1982), ‘Sticky Prices in the United States’, *Journal of Political Economy* **90**(6), 1187–1211.
- Rudd, J. and Whelan, K. (2005), ‘New tests of the new-Keynesian Phillips curve’, *Journal of Monetary Economics* **52**(6), 1167–1181.
- Rudd, J. and Whelan, K. (2006), ‘Can Rational Expectations Sticky-Price Models Explain Inflation Dynamics?’, *The American Economic Review* **96**(1), 303–320.
- Rudd, J. and Whelan, K. (2007), ‘Modeling Inflation Dynamics: A Critical Review of Recent Research’, *Journal of Money, Credit and Banking* **39**, 155–170.
- Rudebusch, G. D. (2002), ‘Assessing Nominal Income Rules for Monetary Policy with Model and Data Uncertainty’, *The Economic Journal* **112**(479), 402–432.
- Rumler, F. (2007), ‘Estimates of the Open Economy New Keynesian Phillips Curve for Euro Area Countries’, *Open Economies Review* **18**(4), 427–451.
- Sahuc, J.-G. (2006), ‘Partial indexation, trend inflation, and the hybrid Phillips curve’, *Economics Letters* **90**(1), 42–50.
- Samuelson, P. A. and Solow, R. M. (1960), ‘Analytical Aspects of Anti-Inflation Policy’, *The American Economic Review* **50**(2), 177–194.
- Sargent, T. and Wallace, N. (1975), ‘“Rational” Expectations, the Optimal Monetary Instrument, and the Optimal Money Supply Rule’, *Journal of Political Economy* **83**(2), 241–254.
- Sbordone, A. M. (2002), ‘Prices and unit labor costs: a new test of price stickiness’, *Journal of Monetary Economics* **49**(2), 265–292.
- Sbordone, A. M. (2005), ‘Do expected future marginal costs drive inflation dynamics?’, *Journal of Monetary Economics* **52**(6), 1183–1197.

- Sbordone, A. M. (2006), ‘U.S. Wage and Price Dynamics: A Limited-Information Approach’, *International Journal of Central Banking* **2**(3), 155–191.
- Schorfheide, F. (2011), ‘Estimation and Evaluation of DSGE Models: Progress and Challenges’, *National Bureau of Economic Research Working Paper Series* (16781).
- Shapiro, A. H. (2008), ‘Estimating the New Keynesian Phillips Curve: A Vertical Production Chain Approach’, *Journal of Money, Credit and Banking* **40**(4), 627–666.
- Sheedy, K. D. (2010), ‘Intrinsic inflation persistence’, *Journal of Monetary Economics* **57**(8), 1049–1061.
- Smith, G. W. (2009), ‘Pooling forecasts in linear rational expectations models’, *Journal of Economic Dynamics and Control* **33**(11), 1858–1866.
- Sowell, F. (1996), ‘Optimal Tests for Parameter Instability in the Generalized Method of Moments Framework’, *Econometrica* **64**(5), 1085–1107.
- Stock, J. H. and Watson, M. W. (2007), ‘Why Has U.S. Inflation Become Harder to Forecast?’, *Journal of Money, Credit and Banking* **39**(1), 3–33.
- Stock, J. H. and Watson, M. W. (2010), ‘Modeling Inflation After the Crisis’, *National Bureau of Economic Research Working Papers* (16488). Prepared for the Federal Reserve Bank of Kansas City Symposium at Jackson Hole.
- Stock, J. H. and Wright, J. (2000), ‘GMM with Weak Identification’, *Econometrica* **68**(5), 1055–1096.
- Stock, J. H., Wright, J. H. and Yogo, M. (2002), ‘A Survey of Weak Instruments and Weak Identification in Generalized Method of Moments’, *Journal of Business & Economic Statistics* **20**(4), 518–529.
- Stock, J. H. and Yogo, M. (2002), ‘Testing for Weak Instruments in Linear IV Regression’, *National Bureau of Economic Research Working Papers* (0284).
- Taylor, J. B. (1979), ‘Staggered Wage Setting in a Macro Model’, *American Economic Review* **69**(2), 108–113.
- Taylor, J. B. (1980), ‘Aggregate Dynamics and Staggered Contracts’, *Journal of Political Economy* **88**(1), 1–23.
- Thomas, L. B. (1999), ‘Survey Measures of Expected U.S. Inflation’, *Journal of Economic Perspectives* **13**(4), 125–144.
- Tillmann, P. (2008), ‘Do interest rates drive inflation dynamics? An analysis of the cost channel of monetary transmission’, *Journal of Economic Dynamics and Control* **32**(9), 2723–2744.
- Tsoukis, C., Kapetanios, G. and Pearlman, J. (2011), ‘Elusive Persistence: Wage and Price Rigidities, the New Keynesian Phillips Curve and Inflation Dynamics’, *Journal of Economic Surveys* **25**(4), 737–768.

- White, H. (1980), ‘A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity’, *Econometrica* **48**(4), 817–838.
- Woodford, M. (2003), *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton University Press.
- Wright, J. H. (2009), ‘Comment’, *Journal of Business & Economic Statistics* **27**(3), 323–326.
- Yao, F. (2011), ‘Monetary Policy, Trend Inflation and Inflation Persistence’. SFB 649 Discussion Paper 2011-008.
- Yun, T. (1996), ‘Nominal price rigidity, money supply endogeneity, and business cycles’, *Journal of Monetary Economics* **37**(2), 345–370.
- Zhang, C. and Clovis, J. (2010), ‘The New Keynesian Phillips Curve of rational expectations: A serial correlation extension’, *Journal of Applied Economics* **13**(1), 159–179.
- Zhang, C., Osborn, D. R. and Kim, D. H. (2008), ‘The New Keynesian Phillips Curve: From Sticky Inflation to Sticky Prices’, *Journal of Money, Credit and Banking* **40**(4), 667–699.
- Zhang, C., Osborn, D. R. and Kim, D. H. (2009), ‘Observed Inflation Forecasts and the New Keynesian Phillips Curve’, *Oxford Bulletin of Economics and Statistics* **71**(3), 375–398.

DGP s FOR MONTE CARLO SIMULATIONS

DGP	NKPC		Red. form: forcing var.				Red. form: inflation				Conc.
	γ_f	λ	$\xi_{\pi 1}$	$\xi_{x 1}$	$\xi_{\pi 2}$	$\xi_{x 2}$	$\zeta_{\pi 1}$	$\zeta_{x 1}$	$\zeta_{\pi 2}$	$\zeta_{x 2}$	
<i>1a</i>	0.7	0.03	0.20	0.80	−0.10	0.10	1.00	−0.07	0.01	−0.01	1.8
<i>1b</i>	0.7	0.03	0.10	0.70	0.10	−0.10	0.79	0.25	0.05	−0.05	108.4
<i>2a</i>	0.3	−0.03	0.20	0.80	−0.10	0.10	0.98	−0.06	0.01	−0.01	0.7
<i>2b</i>	0.3	−0.03	0.70	1.05	0.08	−0.08	0.91	−0.07	−0.01	0.01	30.1

Table 1: List of our DGPs. Columns 2 and 3 list the true values of (γ_f, λ) . The following 8 columns list the VAR(2) reduced form coefficients in the equations for the forcing variable (ξ) and inflation (ζ). The last column lists the minimum eigenvalue of the population concentration matrix (“Conc.”), cf. section A.3 in the Appendix. Numbers in columns 8–11 have been rounded off to 2 decimal points, while numbers in the last column have been rounded off to 1 decimal point.

OVERVIEW OF ESTIMATION APPROACHES IN THE LITERATURE

Papers	Estimation approach	Expectation vs. lags	Slope	Rejection of model?
Galí and Gertler (1999), Galí et al. (2001, 2005)	RE GIV.	Forward-looking behavior dominant, but backward-looking term significant.	Significantly positive for labor share.	No, based on over-ID test and visual fit.
Fuhrer and Moore (1995), Fuhrer (1997, 2006)	RE VAR-ML (AIM).	Price setting not very forward-looking; need large intrinsic persistence.	Positive for both labor share and output gap, but significance varies.	Pure NKPC rejected based on LR test and IRFs.
Roberts (1995, 1997, 2005)	GIV, VAR-ML, IRF matching; RE and survey forecasts.	Sluggish survey forecasts impart necessary persistence. For RE, need more than 1 lag of inflation.	Positive for both labor share and output gap, but significance varies.	No.
Sbordone (2002, 2005)	RE VAR-MD.	Forward-looking behavior clearly dominant, but lag is significant.	Positive but marginally insignificant in hybrid model.	No, based on over-ID test and visual fit.
Rudd and Whelan (2005, 2006, 2007)	RE GIV (iterated).	Lagged inflation very significant.	Neither labor share nor output gap adds explanatory power.	Yes, forcing variable doesn't help explain inflation.
Rudebusch (2002)	OLS; survey forecasts.	Four-quarter MA of lagged inflation receives larger weight than forecast.	Output gap coefficient positive and significant.	No.
Ravenna and Walsh (2006)	RE GIV, interest rate added to NKPC.	(Pure NKPC.)	(Not directly estimated.)	No, based on over-ID test.
Cogley and Sbordone (2008)	Bayesian estimation using VAR with drifting parameters and stochastic volatility.	Backward-looking term insignificant once trend inflation is accounted for.	(Not directly estimated.)	No, based on visual fit and magnitude of forecast errors.

Table 2: Overview of well-cited papers in the limited-information empirical NKPC literature. Papers have been grouped according to authorship. Each series of papers has been ranked by the number of Google Scholar citations (as of mid-September 2012) for the single most cited paper within the series.

BASELINE GIV ESTIMATES USING DIFFERENT DATA VINTAGES

Data vintage	const	λ	γ_f	γ_b	Hansen test
1998	0.041 (0.030)	0.026 (0.013)	0.615 (0.057)	0.340 (0.058)	5.263 [0.628]
2012	-0.049 (0.040)	0.018 (0.012)	0.719 (0.099)	0.240 (0.095)	9.816 [0.199]

Table 3: Comparison of GIV estimates of the hybrid NKPC based on 1998 and 2012 vintages of data. The estimation sample is 1970q1 to 1998q1. Inflation: GDP deflator. Labor share: NFB. Instruments: four lags of inflation and two lags of the labor share, wage inflation, and quadratically-detrended output. Estimation method: CUE GMM. Weight matrix: Newey and West (1987) with automatic lag truncation (4 lags). Standard errors in parentheses and p-values in square brackets.

NKPC SPECIFICATION COMBINATIONS

Specification settings	Options
Inflation (π_t)	GDP deflator, CPI, chained GDP def., GNP def., chained GNP def., NFB GDP def., PCE, core PCE, core CPI, filtered GDP def. gap, smoothed GDP def. gap, filt. CPI gap, sm. CPI gap, SPF-based CPI gap, filt. core CPI gap, sm. core CPI gap, filt. PCE gap, sm. PCE gap, filt. core PCE gap, sm. core PCE gap
Labor share (ls)	NFB, NFB coint. relation, HP filtered NFB gap, Baxter-King filt. NFB gap, linearly detrended NFB gap, quadratically detrended NFB gap, real-time NFB HP gap, real-time NFB BK gap, real-time NFB lin. detr. gap, real-time NFB quadr. detr. gap
Output gap (ygap)	CBO, HP filt., BK filt., lin. detr., quadr. detr., real-time HP filt., real-time BK filt., real-time lin. detr., real-time quadr. detr.
Reduced form	Unrestricted, VAR
Survey forecasts ($\pi_{t \tau}^s$)	SPF CPI, SPF GDP def., GB GDP def.
Expectations	π_{t+1} (endogenous), $\pi_{t+1 t}^s$ (endog.), $\pi_{t+1 t}^s$ (exogenous) $\pi_{t+1 t-1}^s$ (endog.), $\pi_{t+1 t-1}^s$ (exog.)
Instruments	GG: 4 lags of π_t , ls, ygap, 10y–90d yield spread, wage infl., commodity price infl. GGLS: 4 lags of π_t and 2 lags of ls, ygap, wage infl. small: 4 lags of π_t and 3 lags of forcing variable exact: 1 extra lag of each endog. regr. (just-identified) RT: 2 real-time lags of GDP def. inflation, Δ ls, ygap survey: 2 lags of 1-quarter SPF/GB forecasts, forcing variable Extra regressors (e.g., oil) added to instruments (if endog., use 2 lags)
Inflation lags	0 lags (pure NKPC), 1 lag, 4 lags
Parameter restrictions	No restrictions, $\gamma(1) = \gamma_f$ (inflation coef's sum to 1) With $\gamma(1) = \gamma_f$, use lags of $\Delta\pi_t$ instead of π_t as instr's
Oil shocks	None, log change of WTI spot price divided by GDP def.
Interest rate	None, 90-day Treasury rate
Sample	Full available, 1960–1997, 1968–2005, 1968–2008, 1971–2008, 1981–2008, 1984–end of sample
GMM estimator	2-step, CUE

Table 4: List of the specification options that we consider when estimating the NKPC (9). The efficient GMM weight matrix is computed using the Newey and West (1987) heteroskedasticity and autocorrelation consistent estimator with automatic lag truncation, except for VAR specifications, which use the White (1980) heteroskedasticity consistent estimator.

SUMMARY STATISTICS FOR ESTIMATION RESULTS

Parameter	labor share		output gap	
	λ	γ_f	λ	γ_f
median	0.004	0.753	0.004	0.760
5th percentile	-0.068	-0.648	-0.070	-0.771
95th percentile	0.135	1.814	0.133	1.831
fraction both positive	0.525		0.505	
... and signif. (5% one-sided t-test)	0.102		0.179	
... and $\gamma_f > 0.5$	0.087		0.155	
median F_{HAR}	63.73	3.079	166.46	4.154
fraction rejections by 5% Hansen test	0.033		0.032	

Table 5: Summary statistics of estimation results across specifications listed in Table 4, excluding real-time and survey instrument sets. Hansen test is evaluated at the CUE using larger critical values that are robust to weak identification, cf. section A.2.6 of the Appendix; results for this statistic do not include VAR-GMM specifications.

PAIRWISE COMPARISON OF POINT ESTIMATES: 2-STEP VS. CUE GMM

Parameter	labor share				output gap			
	λ		γ_f		λ		γ_f	
Method	2S	CUE	2S	CUE	2S	CUE	2S	CUE
median	0.003	0.004	0.676	0.798	0.004	0.002	0.670	0.784
90% IQR	0.141	0.273	1.213	2.855	0.123	0.264	1.250	3.272
med. diff 2S-CUE	-0.000		-0.030		0.000		-0.024	
med. diff from OLS	0.000	0.002	0.168	0.239	-0.002	-0.003	0.156	0.213
med. diff 2S-CUE (GG)	-0.001		-0.057		0.002		-0.041	
med. diff from OLS (GG)	0.001	0.002	0.124	0.208	-0.002	-0.005	0.129	0.195

Table 6: Comparison of 2-step and CUE GMM estimates for the specifications in Table 4 (excluding VAR specifications, for which we only computed 2-step GMM). “90% IQR” is the difference between the 95th and 5th percentiles. Rows labeled “GG” focus on results for the GG instrument set.

PAIRWISE COMPARISON OF POINT ESTIMATES: IMPACT OF VAR ASSUMPTION

Parameter	labor share				output gap			
	λ		γ_f		λ		γ_f	
Method	VAR	GIV	VAR	GIV	VAR	GIV	VAR	GIV
median	0.002	0.002	0.997	0.747	0.002	0.001	0.997	0.750
90% IQR	0.137	0.087	1.912	1.124	0.198	0.106	3.207	1.116
median diff	-0.001		0.178		0.000		0.176	
fraction positive diff	0.481		0.844		0.513		0.797	

Table 7: Effect of imposing the VAR identification assumption.

NKPC SPECIFICATION COMBINATIONS USED FOR ROBUST CONFIDENCE REGIONS

Specification settings	Options
Inflation (π_t)	GDP deflator, CPI
Labor share (ls)	NFB, real-time NFB BK gap
Output gap (ygap)	CBO, real-time BK filter
Reduced form	Unrestricted, VAR
Survey forecasts ($\pi_{t \tau}^s$)	SPF CPI, SPF GDP def., GB GDP def.
Expectations	π_{t+1} (endogenous), $\pi_{t+1 t}^s$ (endogenous)
Instruments	GGLS, small, exact, RT, survey
Inflation lags	0, 1, 4
Parameter restrictions	No restrictions, $\gamma(1) = \gamma_f$ (inflation coef's sum to 1)
Sample	Full available, pre-1984, 1984-end of sample

Table 8: Different specifications of the NKPC for which we compute robust confidence sets.

PAIRWISE COMPARISON OF POINT ESTIMATES: RE GIV VS. SURVEY FORECASTS

Parameter	labor share		output gap	
	λ	γ_f	λ	γ_f
SPF vs GIV: post-81	0.002	-0.200	0.042	0.092
SPF vs GIV: CPI post-81	-0.002	0.030	0.049	0.286
SPF vs GIV: GDP def. post-81	0.004	-0.326	0.038	-0.043
SPF vs GIV: GDP def. post-68	0.002	-0.200	0.042	0.092
GB vs GIV: post-68	0.039	-0.093	0.047	-0.000
endogenous vs exogenous	-0.000	-0.000	-0.001	-0.045
GB vs SPF	0.003	-0.016	-0.007	-0.063
GB vs SPF: pre-1984	-0.010	0.069	-0.017	0.007
GB vs SPF: post-1984	0.009	-0.076	0.003	-0.103

Table 9: Effect of using observed inflation forecasts (SPF or GB) to proxy for inflation expectations in the NKPC. Numbers are median pairwise differences in estimates across specifications that differ by one characteristic, keeping all other specification aspects constant. For example, “SPF vs GIV” is the median difference of coefficient estimates in SPF specifications from the corresponding RE GIV specifications.

PAIRWISE COMPARISONS OF POINT ESTIMATES: OTHER SPECIFICATION CHOICES

Parameter	labor share		output gap	
	λ	γ_f	λ	γ_f
RT vs small instr.	-0.003	0.008	-0.003	-0.012
survey vs small instr.	0.001	0.001	-0.001	-0.018
1 vs 0 lags	-0.001	-0.256	-0.000	-0.256
4 vs 1 lags	-0.004	-0.223	0.009	-0.277
post- vs pre-84 sample	-0.008	-0.024	-0.002	0.000
GDP def. vs CPI inflation	0.001	0.011	0.004	0.018
gap vs level of inflation	-0.001	-0.054	0.013	-0.072
oil vs no controls	0.000	0.001	-0.001	-0.001
Treas. rate vs no controls	-0.000	0.005	0.001	0.000

Table 10: Median pairwise differences in estimates across specifications that differ by one characteristic, keeping all other specification aspects constant. For example, row “RT vs small instr.” compares the specifications that use “small” and “RT” instrument sets; row “1 vs 0 lags” gives the median difference of estimates in the hybrid NKPC with 1 lag from the corresponding estimates in the pure NKPC.

SIZE OF ROBUST CONFIDENCE REGIONS

		labor share		output gap	
		90%	95%	90%	95%
spec.	Sample				
all	all	0.33	0.47	0.30	0.47
0 lags	all	0.19	0.34	0.18	0.34
1 lag	all	0.24	0.38	0.19	0.34
4 lags	all	0.49	0.63	0.48	0.67
all	full	0.25	0.36	0.24	0.39
all	pre-1984	0.35	0.50	0.25	0.41
all	1984q1–end of sample	0.41	0.56	0.41	0.59
all RE	1984q1–2011q3	0.49	0.65	0.46	0.65
all survey	1984q1–end of sample	0.35	0.50	0.38	0.55
SPF only	1984q1–2011q4	0.37	0.53	0.44	0.62
GB only	1984q1–2005q4	0.31	0.45	0.26	0.42
GGLS instr.	1984q1–end of sample	0.33	0.48	0.43	0.60
small instr.	1984q1–end of sample	0.46	0.66	0.43	0.59
exact instr.	1984q1–end of sample	0.29	0.42	0.32	0.52
RT instr.	1984q1–end of sample	0.64	0.78	0.54	0.73
survey instr.	1984q1–end of sample	0.32	0.47	0.35	0.52

Table 11: Size of 90% and 95% S sets corresponding to the different specifications of the NKPC in Table 8, as a fraction of the parameter space shown in Figure 16 (using 1025 grid points). “RE” refers to all RE GIV and VAR-GMM specifications.

ANATOMY OF ROBUST CONFIDENCE REGIONS

	labor share		output gap	
Level	90%	95%	90%	95%
No. of specs	698	698	698	698
fraction empty S set	0.038	0.013	0.016	0.004
fraction non-empty and positive	0.034	0.010	0.049	0.020
all RE	0.034	0.010	0.000	0.000
all survey	0.035	0.010	0.085	0.035
0 lags	0.072	0.014	0.094	0.036
1 lag	0.046	0.018	0.074	0.032
4 lags	0.004	0.000	0.000	0.000
GGLS instr.	0.062	0.021	0.042	0.000
small instr.	0.035	0.014	0.035	0.021
exact instr.	0.000	0.000	0.079	0.043
RT instr.	0.007	0.007	0.000	0.000
survey instr.	0.069	0.008	0.092	0.038

Table 12: Anatomy of 90% and 95% S sets for specifications listed in Table 8. S sets are computed on the parameter space shown in Figure 16 using 1025 grid points. “Fraction empty” lists fraction of cases with empty S set. “Fraction non-empty and positive” lists fraction of cases for which S set is non-empty and lies entirely in positive orthant; rows “all RE” to “survey instr.” give the latter fraction for different subcases. “RE” refers to all RE GIV and VAR-GMM specifications.

IMPULSE RESPONSES TO MONETARY POLICY SHOCK

Figure 1: Impulse responses of inflation and the output gap to a 25 basis point monetary policy shock in a standard three-equation new Keynesian model with $\gamma_f = 0.3, 0.4, \dots, 0.8$ and $\lambda = 0.01, 0.03, 0.05$. The other parameters are calibrated to the benchmark values listed in Appendix A.1. More sluggish responses correspond to lower values of γ_f . The figure was generated using Dynare (Adjemian et al., 2011).

SAMPLING DISTRIBUTION OF γ_f ESTIMATORS

Figure 2: Kernel-smoothed density estimates of the sampling distribution of γ_f estimators in the hybrid NKPC model (21) for the DGPs listed in Table 1. The dotted vertical line marks the true parameter value.

POINT ESTIMATES REPORTED IN THE LITERATURE

Figure 3: Point estimates of λ (vertical axis) and γ_f (horizontal axis) reported in the literature. Only estimates that use U.S. data and the labor share as forcing variable are plotted. For some papers the semi-structural point estimates have been imputed from point estimates of deeper parameters. The dotted blue lines indicate 95% confidence intervals for λ where available. We include papers with readily available estimates and more than 25 Google Scholar citations as of mid-September 2012: Galí and Gertler (1999), Galí, Gertler and López-Salido (2001), Fuhrer and Olivei (2005), Gagnon and Khan (2005), Guay and Pelgrin (2005), Jondeau and Le Bihan (2005), Roberts (2005), Sbordone (2005), Dufour, Khalaf and Kichian (2006), Fuhrer (2006), Kiley (2007), Kurmann (2007), Rudd and Whelan (2007), Brissimis and Magginas (2008), Adam and Padula (2011) and Henzel and Wollmershäuser (2008).

POINT ESTIMATES: LABOR SHARE SPECIFICATIONS

Figure 4: Point estimates of λ, γ_f from the various specifications listed in Table 4 that use the labor share as forcing variable, excluding real-time and survey instrument sets. The black dot and ellipse represent the point estimate and 90% joint Wald confidence set from the 1998 vintage results in Table 3.

POINT ESTIMATES: OUTPUT GAP SPECIFICATIONS

Figure 5: Point estimates of λ, γ_f from the various specifications listed in Table 4 that use the output gap as forcing variable, excluding real-time and survey instrument sets.

FIRST-STAGE F STATISTICS

Figure 6: Smoothed density estimates of robust first-stage F statistics for forecasting the expectation proxy, using realized next-period inflation (solid line, left axis) or time- t dated endogenous SPF/GB survey forecasts of inflation (dotted line, right axis). The left and right panels show results for all labor share and output gap specifications, respectively, that are listed in Table 4.

POINT ESTIMATES: VAR (RED) VS. UNRESTRICTED (BLUE)

Figure 7: Point estimates of the coefficient on the forcing variable and inflation expectations for the specification combinations listed in Table 4. The red points correspond to estimates that impose the VAR assumption, while the blue points do not impose the assumption. The left and right panels plot specifications with the labor share and output gap as forcing variable, resp.

POINT ESTIMATES: SURVEY FORECASTS (RED) VS. RE GIV (BLUE)

Figure 8: Point estimates of the coefficient on the forcing variable and inflation expectations using time- $(t - 1)$ dated exogenous SPF forecasts (red) versus RE GIV (blue), for the various specification settings listed in Table 4. The top row of figures use GDP deflator inflation, while the bottom row uses CPI inflation. In each row, the left and right panels correspond to labor share and output gap specifications, resp. Sample: 1984q1–2011q4.

POINT ESTIMATES: CPI (RED) VS. GDP DEFLATOR (BLUE) INFLATION

Figure 9: Point estimates of the coefficient on the forcing variable and inflation expectations using CPI (red) versus GDP deflator (blue) inflation, for the various specification settings listed in Table 4. The top row gives results for all available samples, while the bottom row selects only the post-1984q1 sample. The left and right panels plot specifications with the labor share and output gap as forcing variable, resp.

POINT ESTIMATES: PRE-1984 (RED) AND POST-1984 (BLUE) SAMPLES

Figure 10: Point estimates of the coefficient on the forcing variable and inflation expectations in survey forecast specifications (SPF and GB) on pre-1983q4 (red) and post-1984q1 (blue) samples. The inflation series is GDP deflator inflation. The left and right panels plot specifications with the labor share and output gap as forcing variable, resp.

Figure 11: 90% S set (grey), 90% Wald ellipse and CUE GMM point estimate (bullet) of the coefficients of the labor share and future inflation in the hybrid NKPC specification with one lag of inflation, where inflation coefficients sum to 1. Inflation: GDP deflator. Forcing variable: NFB labor share (left panels), CBO output gap (right panels). Instruments: three lags of $\Delta\pi_t$ and the forcing variable. Sample: starts 1948q2 (labor share), 1949q4 (output gap), ends 2011q3; full sample (top row), pre-1983q4 (middle row), post-1984q1 (bottom row). Weight matrix: Newey-West with automatic lag truncation.

ROBUST CONFIDENCE REGIONS: VAR SPECIFICATIONS

Figure 12: 90% S set (grey), 90% Wald ellipse and CUE VAR-GMM point estimate (bullet) of the coefficients of the labor share and future inflation in the hybrid NKPC specification with one lag of inflation, where inflation coefficients sum to 1. Inflation: GDP deflator. Forcing variable: NFB labor share (left panel), CBO output gap (right panel). Instruments: three lags of $\Delta\pi_t$ and the forcing variable, implying VAR(3) reduced form. Sample: starts 1948q2 (labor share), 1949q4 (output gap), ends 2011q3. Weight matrix: White (1980) heteroskedasticity-consistent.

ROBUST CONFIDENCE REGIONS: SURVEY FORECAST SPECIFICATIONS

Figure 13: 90% S set (grey), 90% Wald ellipse and CUE GMM point estimate (bullet) of the coefficients of the labor share and future inflation in the hybrid NKPC specification with one lag of inflation, where inflation coefficients sum to 1. Forcing variable: NFB labor share (left panels), CBO output gap (right panels). Instruments: three lags of $\Delta\pi_t$ and the forcing variable. Specifications: SPF GDP deflator 1984q1–2011q4 (top row), GB GDP deflator 1984q1–2005q4 (second row), SPF CPI 1984q1–2011q4 (third row), SPF GDP deflator 1968q3–1983q4 (bottom row), all time- t dated and endogenous. Weight matrix: Newey-West with automatic lag truncation.

ROBUST CONFIDENCE REGIONS: REAL-TIME INSTRUMENTS

Figure 14: 90% S set (grey), 90% Wald ellipse and CUE GMM point estimate (bullet) of the coefficients of the labor share and future inflation in the hybrid NKPC specification with one lag of inflation, where inflation coefficients sum to 1. Inflation: GDP deflator. Forcing variable: NFB labor share (left panels), CBO output gap (right panels). Instruments: 2 lags of GDP deflator inflation, the output gap and the change in the labor share, all measured in real time. Sample: 1971q1-2011q2. Weight matrix: Newey-West with automatic lag truncation.

ROBUST CONFIDENCE REGIONS: SPECIFICATIONS WITH BOTH RATIONAL AND SURVEY EXPECTATIONS

Figure 15: 90% S set (grey), 90% Wald ellipse and CUE GMM point estimate (bullet) of the coefficients of future inflation (γ_{RE} , vertical axis) and 1-quarter forecasts (γ_s , horizontal axis) in the nesting NKPC specification (25). Inflation: GDP deflator. Forcing variable: NFB labor share (left two columns), CBO output gap (right two columns). Instruments: GGLS plus two lags of survey forecasts. Sample: full available (top row), 1984q1–end of sample (bottom row). Survey forecasts: SPF (1st and 3rd columns), Greenbook (2nd and 4th columns), all time- t dated and endogenous. Weight matrix: Newey-West with automatic lag truncation.

POINT ESTIMATES FOR SPECIFICATIONS IN TABLE 8

Figure 16: CUE point estimates of the coefficient on the forcing variable and inflation expectations for the various specification settings listed in Table 8. The left and right panels plot specifications with the labor share and output gap as forcing variable, resp.