(Possibly) Hidden Shocks and the Driving Process for Inflation.

Stephen Wright

December 30, 2013

1 Introduction

A standard New Keynesian Phillips Curve can be written as

$$\pi_t = \beta E_t \pi_{t+1} + \mu_t \tag{1}$$

where μ_t is a composite of all real magnitudes that shift either real marginal costs or the markup.

The data suggest that the following time series representation of inflation captures pretty much all its predictability

$$\pi_t = \left(\frac{1 - \theta L}{1 - \lambda L} \right) \varepsilon_t \tag{2}$$

where Stock and Watson (call them SW1) impose the restriction that $\lambda = 1$, on the assumption that there is a random walk component in trend inflation, with an estimate of θ of at least 0.5, apparently with an upward trend in recent data. In data from 1980 onwards estimates of λ are significantly below unity, but not quantitatively far below, such that Stock and Watson's (fixed coefficient) IMA representation implies an almost identical autocorrelation function for π_t to an unrestricted ARMA(1,1).[1]

[1]SW1 estimate an unobserved components model with a random walk permanent component and a serially uncorrelated transitory component, both with time-varying variances, which reduces to an MA(1) model in differences, with a time-varying (but constrained) MA parameter. The MA parameter appears to rise systematically in recent data.

How can we reconcile the two equations? More generally, what do the time series properties of an observable process, which we believe to be generated by forward-looking agents, tell us about the properties of the driving process (in this case μ_t)? There are two ways of looking at this problem.

The first, and easiest approach, is to assume from the start that price-setting agents' own actions reveal their information set to the econometrician, and thus that shocks to the driving process are fundamental to to the history of inflation itself. So if we start from observing that inflation takes an ARMA(1,1) form as in (2), and that this generates price-setters' expectations, it is straightforward to derive the implied process for μ_t, which must also be an ARMA(1,1), but with a different MA parameter, and which will be likely to be very much closer to (but cannot actually be) white noise. It turns out that this process looks very similar to the "markup shock" in the Smets & Wouters model (call them SW2).

A second approach, allows for the possibility that agents may have information that at least possibly may *not* get revealed to the econometrician. In this approach we start from an assumed ARMA(1,1) process for μ_t, and derive the implied process for π_t. It turns out that this raises the possibility that for some parameter values for μ_t the implied ARMA representation of inflation is nonfundamental, and hence that the shocks to the driving process are *not* revealed to the econometrican (in real time, at least). But I also show that this possibility can be ruled out for a quite wide class of μ_t and π_t processes (including, on a strict interpretation, the Stock-Watson representation).[2]

2 Common Information: SW1 vs SW2

Initially assume that price setters do not have superior information to the econometrician (this turns out to be the implicit assumption that, eg Smets & Wouters (call them SW2) make) and that the only useful information is

[2] One stimulus to this note was the finding, in a parallel piece of research, that it is very difficult to find deep parameter values in the the Smets & Wouters model that violate the "Poor Man's Invertibility Condition" of Fernandez-Villaverde et al, and hence that generate a nonfundamental representation. My suspicion is that this reflects the assumed (multivariate) time series structure of the exogenous driving processes. This is reinforced by the finding tha the only exceptions to this finding are the parameters of the exogenous markup processes, which are both ARMA(1,1)s.

captured by the history of inflation itself, π^t (this is indeed essentially the conclusion of Stock & Watson (SW1)).

Then, using (2) we have

$$E_t \pi_{t+1} | \pi^t = \frac{(\lambda - \theta)\,\varepsilon_t}{1 - \lambda L}$$

and hence, substituting into (1) we must have (letting the superscript F denote "fundamental" - ie, recoverable from the history of inflation)

$$\begin{aligned} \left(\frac{1-\theta L}{1-\lambda L}\right)\varepsilon_t &= \frac{\beta\,(\lambda-\theta)\,\varepsilon_t}{1-\lambda L} + \mu_t^F \\ &\Rightarrow \mu_t^F = \left(\frac{1-\psi L}{1-\lambda L}\right) v_t \\ &\text{where} \\ v_t^F &= (1 - \beta\,(\lambda-\theta))\,\varepsilon_t; \quad \psi_F = \frac{\theta}{(1-\beta\,(\lambda-\theta))} \end{aligned}$$

so the driving process for inflation, thus derived, let's call it μ_t^F, must itself be an ARMA(1,1) but with a different MA parameter. For any representation of inflation that at all resembles that of SW1, $\lambda > \theta$, and hence note that

$$\lambda > \theta \Rightarrow \theta < \psi < \frac{\theta}{1-\beta+\theta} < 1$$

so ψ, the implied MA parameter of μ_t^F, must be distinctly closer to λ than is θ, the MA parameter of inflation, but is strictly less than unity even in the limiting unit root case of SW1.

The figure below plots the MA parameter ψ_F, as a function of λ, assuming $\theta = .5$, roughly consistent with Stock & Watson, and $\beta = 0.99$

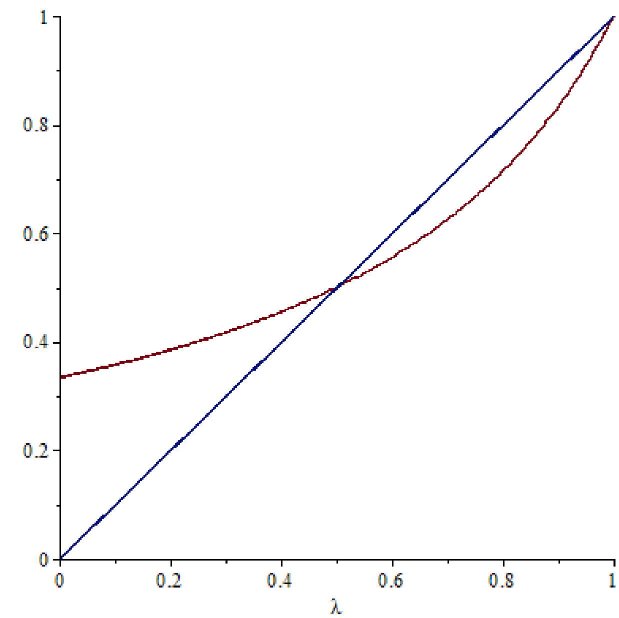

The figure makes it clear that, the larger is λ, the closer μ_t^F must be to a white-noise process, but with MA parameter always strictly less than the AR parameter.

This turns out to be extremely close to the representation of inflation in the Smets & Wouters (SW2) model. If we ignore their backward looking price terms they write

$$\mu_t = mc_t + x_t^p \tag{3}$$

where mc_t is marginal cost, driven by the real economy, and x_t^p is an exogenous ARMA(1,1) "markup" process, with AR parameter =0.9, MA parameter=0.74 (modal estimates). If we simply ignored the real economy and set $\mu_t^F = x_t^p$, then for $\lambda = 0.9$, using the formula above, we would get $\psi_F = 0.83$, which is pretty close to Smets & Wouters' figure. That is, their shock process is very close to being simply what you would need to give you back the time series properties of inflation, under the maintained assumption that (1) was the true model. Thus we can get very close to reconciling SW1 and SW2, just by assuming a markup process that is essentially reverse-engineered from the data for inflation.

3 (Possibly) Hidden Information

Is this the only way to reconcile these two representations of inflation? It turns out that is not, but to show this we have to look at the process the other way around. The shared information assumption I made in the previous section may not be correct. A key question is: if price setters have information which, at least potentially, the econometrican cannot see, will their forward-looking *behaviour* reveal this information to the econometrician?

This question is not, of course, new. The "news" literature has already established that the answer to this question is not always yes. However, as far as I am aware it has not focussed on the necessary links between the time series properties of the driving process and those of the observables. I show below that some restrictions on the time series properties of the driving process imply that price-setters *must* reveal their information to the econometrician.

To answer the question,we need to solve the problem from the other end: start with a process for μ_t and derive the implied process for π_t. In any rational expectations model, time series properties of any observable that involves forward-looking terms in some driving process (here, the composite variable, μ_t) must be driven by the time series properties of the driving process.

A standard critique of many macro models is that these driving processes are usually exogenous, and that the time series properties of the observables involve few if any endogenous sources of dynamics (indeed the analysis of the previous section suggests that this is close to holding in the Smets & Wouters model). However while it must be true that the times series properties of the driving process *determine* those of the observables, this does not imply that they must be the *same.*

Here simple models can be misleading. A standard textbook implementation of (1), for example, assumes that μ_t is an AR(1). In this special case it is straightforward to show (see below) that π_t will also be an AR(1), with the same AR(1) parameter, with an innovation that is simply a scaling of the shock to μ_t, and is therefore fundamental for the history of π_t. However, this result does not fully generalise. In particular, once MA terms enter the picture, nonfundamental representations may arise.

4 Time Series Properties of Observables and Driving Processes: the ARMA(1,1) case

By generalising this simple example by the smallest possible amount, to the case where the driving process is an ARMA(1,1), it is possible to see what does, and does not, generalise.[3]

Proposition 1 *Assume an observable process y_t is forward-looking in some driving process, x_t (as for example in the New Keynesian Phillips Curve (1)) and thus has the forward solution*

$$y_t = E\sum_{i=0}^{\infty}\beta^i E\left(x_{t+i}|I_t\right) \tag{4}$$

for some $t-$dated information set, I. Assume further that x_t is an ARMA(1,1) process

$$x_t = c\left(L\right)v_t = \left(\frac{1-\psi L}{1-\lambda L}\right)v_t \tag{5}$$

where $|\psi| < 1$, $|\lambda| < 1$ so the process is stationary and v_t is fundamental for x^t, and $I_t = x^t$, the history of x_t (either an infinite history, or the finite joint history (x^t, v^t)). Then:

1. *The observable y_t process will also have a fundamental ARMA(1,1) representation*

$$y_t = \left(\frac{1-\theta L}{1-\lambda L}\right)\varepsilon_t \tag{6}$$

with associated nonfundamental representation

$$y_t = \left(\frac{1-\theta^{-1} L}{1-\lambda L}\right)\eta_t \tag{7}$$

2. *Letting*

$$\psi^* = \frac{1}{1+\beta\left(1-\lambda\right)} \in \left(\lambda, 1\right)$$

[3] In Appendix C I attempt an extension to a more general analysis. While it is possible to draw some conclusions on the links between the processes for x_t and y_t, it is not obvious how to derive general implications on fundamentalness.

then

$$\begin{aligned}\psi &\leq \psi^* \Rightarrow \theta = \frac{\psi(1-\lambda\beta)}{1-\psi\beta}, v_t = \psi\theta^{-1}\varepsilon_t \\ \psi &> \psi^* \Rightarrow \theta = \frac{1-\psi\beta}{\psi(1-\lambda\beta)}, v_t = (\psi\theta)^{-1}\eta_t\end{aligned}$$

where ε_t *and* η_t *are the shocks to (6) and (7) respectively.*

3. *Using the Beveridge Nelson decomposition we can write (5) as*

$$x_t = c(1)v_t + c^*(L)\Delta v_t$$

where $c^*(L) = c(L) - c(1)$. *A sufficient condition for* v_t *to be fundamental for* y^t *is* $c(1) \geq 1$, *or equivalently,* $\lambda \geq \psi$. *In particular, any driving process that is an AR(1) with* $\lambda \geq 0$ *will have innovations that are fundamental for* y^t.

Proof. See Appendix A ■

Part 1 of the proposition shows that one feature of the AR(1) case does generalise to the ARMA(1,1) case: the *order* of the ARMA process of the observable process is the same as that of the driving process. This result appears to generalise for any ARMA process (see Appendix C).

Part 2 shows that, while the AR parameter λ of x_t is equal to that of y_t the MA parameter θ is *not* equal to ψ except in the special case that $\psi = 0$, and hence x_t is a pure AR(1). Furthermore, for sufficiently large values of ψ, the structural shock v_t is nonfundamental for y^t - i.e., it cannot be recovered from the history of the observable, but is instead a scaling of the shock to its nonfundamental representation (7).

Part 3 exploits the feature that the critical value, ψ^*, that determines whether v_t is fundamental for y^t is strictly greater than λ, implying that at this value, the Beveridge-Nelson permanent component for x_t, $c(1)v_t = (1-\psi^*)/(1-\lambda)v_t$ must have lower variance than v_t; equivalently, Cochrane's (1988) variance ratio for x_t slopes downwards. Thus, only if the variance ratio for x_t is sufficiently downward-sloping will the problem of nonfundamentalness arise. This immediately rules out any nonfundamentalness problem for any x_t process that is an AR(1) with $\lambda \geq 0$ (thus subsuming the case where x_t is white noise). But it also rules out nonfundamentalness for a considerably wider class of ARMA(1,1) driving processes for which $c(1) > 1$.

5 SW1 vs SW2 again

Not all, however. In the inflation example, in Section 2 above, it should now be evident that I used the ARMA properties of inflation to reveal the required properties of *one* possible driving process for inflation, μ_t^F, with shock v_t^F, which arises if the New Keynesian Phillips Curve (1) has generated the data for inflation, *and* if econometricians have the same information as price setters - thus for the case where the driving shock is fundamental for inflation. But the Proposition reveals that, there is another possible driving process, call it μ_t^N,that in principle could match the data for inflation and the NKPC equally well, but for which the driving shock v_t^N is not fundamental for π^t

To examine this case we need to exploit the second case of Part 2 of the Proposition, and find the second value of ψ, above the critical value ψ^*, that is consistent with an observed MA parameter for inflation of around $\theta = 0.5$. To simplify, again assume the same value for the AR parameter $\lambda = 0.9$, as in Smets & Wouters, and $\beta = 0.99$. This would imply that μ_t^N the (hidden) ARMA(1,1,) driving process for inflation, would have an MA parameter $\psi_N = 0.96$, as compared to the value $\psi_F = 0.83$ derived above for the fundamental driving process μ_t^F.

Nothing in the data for inflation allows us to distinguish between these two driving processes: for a given value of λ, they both match the data equally well, under the maintained assumption that the New Keynesian Phillips Curve has generated the data for inflation.But the properties of inflation do however at least allow us to rule out certain attributes of the driving process.

One striking feature of the two candidate driving processes is that they are both very close to white noise: both have univariate R^2s very close to zero.[4] However, they do differ nontrivially in one crucial respect: they have very different values of $c(1)$, and hence they differ significantly in their long-run predictability. It turns out that this allows both to match the ARMA properties of inflation equally well. To see this, note that, to a good approx-

[4] μ_t^F has a univariate R^2 of 0.025, μ_t^N has a univariate R^2 of 0.019.

imation, with $\beta \approx 1$,

$$\begin{aligned}\psi &\leq \psi^* \Rightarrow \theta \approx \psi\left(\frac{1-\lambda}{1-\psi}\right) = \frac{\psi}{c(1)} \\ \psi &> \psi^* \Rightarrow \theta \approx \frac{1-\psi}{\psi(1-\lambda)} = \frac{c(1)}{\psi}\end{aligned}$$

This brings out two special cases, both of which we can rule out on the basis of the time series properties of inflation.

1. If $\psi = 0$ and hence μ_t was a pure AR(1), then, as noted above, inflation would also be an AR(1). The data quite strongly reject this.

2. If $\psi = \lambda \iff c(1) = 1$ and hence μ_t was white noise, inflation would also be white noise. The data also (very) strongly reject this: inflation *is* predictable in terms of its own past.

It is on the face of it therefore somewhat surprising that both driving processes consistent with the inflation data should be so close to white noise, if white noise is so clearly rejected in the data. But the explanation is in the quite marked degree of nonlinearity of $c(1)$ when both λ and ψ are large: for μ_t^F, $c(1) \approx 1.7$, whereas for μ_t^N, $c(1) \approx 0.4$.

What is the intuition behind the result that a sufficient degree of "variance compression", or long-run predictability of the driving process for inflation, can result in a hidden shock problem? The explanation here is in the way optimal forecasts by price-setters should respond to shocks. In both cases, the first term in the forward expansion in (4) setting $x_t = \mu_t$, is affected symmetrically by a shock to μ_t. But the impact on forecasts is distinctly different. A current shock to μ_t^N will be expected to partially unwind, dampening the impact of the shock, whereas a shock to μ_t^F will be reinforced, accentuating it. But the lagged impact on μ_t of the shock in the previous period is very similar (since ψ^N is close to ψ^F), so that the *relative* impact of the lagged shock is much larger for μ_t^N, thus leading to an MA coefficient greater than unity, hence a nonfundamental representation.

Note that if we interpret SW1's results as implying that λ is *precisely* equal to unity, then we can still derive the implied properties of μ_t, which itself becomes an MA(1) in differences, but Proposition 1 implies that on this literal interpretation of their results there can be no hidden shock, since, by inspection of part 2 of the proposition, $\lambda = 1 \Rightarrow \psi^* = 1$, so as long as we assume that price-setters use a representation of μ_t that is fundamental to μ^t, there can be no μ_t^N. However, it is not clear whether we should necessarily take such a literal interpretation.

Is there any way to discriminate between these two explanations of inflation on grounds, say, of plausibility? The only other possible way (at least that I can think of) is by asking how much better price-setters must be at forecasting inflation than the econometrician in the nonfundamental case. If the model required them to forecast significantly better, this might be questioned on grounds of plausibility. But the answer comes straight out of Robertson & Wright: namely, they would forecast better, but not very much better. By assumption, price-setters get to use the history of the shocks to the nonfundamental representation to predict inflation, but Robertson & Wright show that the implied improvement in predictive power is quite limited - and especially so in recent years.

Furthermore, we do not actually have to assume that price-setters can forecast this well. Robertson & Wright also show that the nonfundamental representation of inflation is the best *possible* predictive model consistent with inflation data. Price setters might, arguably, not be able to forecast this well, but might have *some* information that the econometrician does not have. In Appendix B I consider a particularly simple case, which assumes that price-setters have enough information to derive the persistent-transitory decomposition of inflation, a special case of which is the representation used by Stock & Watson. It is easy to show that this in turn implies an equivalent decomposition of the driving process μ_t, but which shares the feature, noted above, that μ_t is (much) closer to being, but is not quite, white noise.

Finally, does all of this matter? I would argue that, yes it does. While both the assumed driving processes derived above would appear on the face of it to be pretty similar in one characteristic - that they are both very close to white noise - this does *not* mean that they would be similar in the data. Their shocks would not be particularly strongly correlated. For an ARMA(1,1) process the correlation between the fundamental and nonfundamental shocks is simply equal to $|\theta|$, hence around 0.5, so looking at inflation shocks will tell us something, but not very much about what the shocks to the nonfundamental driving process, μ_t^N look like.

More precisely, it will not tell us very much in real time. With the benefit of hindsight, however, it will tell us a lot more, since we *can* recover these shocks from the two-sided history of inflation. This is worth investigating.

Of course, all the above analysis has been predicated on the assumption that only data about inflation can tell us anything about the process for inflation. But this is more or less what Stock and Watson tell us. Their conclusion is that it is very difficult to find a predictive model for inflation that does better than their simple IMA model. It is puzzling, and not a little worrying for macroeconomists, that such a crucial macroeconomic variable appears to have almost no multivariate properties. But analysing the necessary characteristics of any driving process for inflation consistent with forward-looking behaviour at least should give us some idea of what we should be looking for.[5]

[5] Of course it is possible that price setters do have some information of a multivariate nature that assists them in predicting the driving process, and hence inflation, ie, that $I_t \supset x^t$. But it would be surprising, were this the case, if this did not result in *some* degree of multivariate predictability for the econometrician.

Appendix

A Proof of Proposition

Combining (4) and (5) we have

$$
\begin{aligned}
y_t &= x_t + \sum_{i=1}^{\infty} \beta^i E\left(x_{t+i}|x^t\right) \\
&= \left[\left(\frac{1-\psi L}{1-\lambda L}\right) + \sum_{i=1}^{\infty} \beta^i \lambda^{i-1} \frac{(\lambda-\psi)}{1-\lambda L}\right] \ v_t \\
(1-\lambda L)\, y_t &= \left[1 - \psi L + \frac{(\lambda-\psi)}{\lambda} \sum_{i=1}^{\infty} \beta^i \lambda^i\right] \ v_t \\
&= \left[1 - \psi L + \frac{(\lambda-\psi)}{\lambda} \frac{\lambda\beta}{1-\lambda\beta}\right] \ v_t \\
&= \left[1 + \frac{(\lambda-\psi)\,\beta}{1-\lambda\beta} - \psi L\right] \ v_t \\
&= \left[\frac{1-\psi\beta}{1-\lambda\beta} - \psi L\right] \ v_t = \left[c\,(\beta) - \psi L\right] = \left[1 - \gamma L\right] \frac{\psi}{\gamma} \ v_t
\end{aligned}
$$

hence y_t is a "structural" ARMA(1,1) (ie, defined in terms of the structural shock v_t) with $\gamma\,(\psi) = \psi\left(\frac{1-\lambda\beta}{1-\psi\beta}\right) = \psi/c\,(\beta)\,,$ with $\gamma'_\psi > 0$, proving part 1 of the proposition.

We then have, letting

$$
\psi^* = \frac{1}{1+\beta\,(1-\lambda)}
$$

$$
\psi \gtreqless \psi^* \Longleftrightarrow \gamma\,(\psi^*) \gtreqless 1
$$

hence we have two cases:

a) for $\psi < \psi^*$, $\theta = \gamma \Rightarrow v_t = \frac{\theta}{\psi}\varepsilon_t$

b)for $\psi > \psi^*$, $\theta = \gamma^{-1} \Rightarrow \eta_t = \frac{\gamma}{\psi} v_t = \frac{1}{\theta\psi} v_t \Rightarrow \eta_t = \theta\psi v_t$ hence, equating (6) to (7) we have $\varepsilon_t = \left(1-\theta^{-1}L\right)/\left(1-\lambda L\right)\theta\psi v_t$, proving part 2 of the proposition.

Finally note that we have

$$c(1) > 1 \Longleftrightarrow \lambda > \psi$$

since

$$c(1) = \frac{1-\psi}{1-\lambda} = 1 + \frac{\lambda - \psi}{1-\lambda}$$

and, given $|\lambda| < 1$, as assumed in the proposition,

$$\psi^* > \lambda$$

since

$$\psi^* - \lambda = \frac{1 - \lambda(1 + \beta(1-\lambda))}{1 + \beta(1-\lambda)} = \frac{(1-\lambda)(1-\beta\lambda)}{1+\beta(1-\lambda)} > 0$$

hence

$$c(1) > 1 \Rightarrow \psi < \lambda < \psi^*$$

and hence, case b) above is ruled out, so v_t is fundamental for y^t, proving part 3 of the proposition. The reverse does not apply since for $\psi \in (\lambda, \psi^*)$, $c(1) < 1$ but v_t is still fundamental for y^t, thus $c(1) > 1$ is sufficient but not necessary.■

B Deriving μ_t directly from the Stock-Watson Decomposition

The SW1 decomposition allows a particularly straightforward derivation of the time series properties of the implied driving variable, call it μ_t^{SW}.

Assume that price-setters have an information set that allows them to derive a permanent-transitory decomposition of inflation of the form

$$\begin{aligned}
\pi_t &= \pi_t^P + \pi_t^T = \frac{v_{1t}}{1-\lambda L} + v_{2t}, \\
E(v_{1t}v_{2t}) &= 0;\ E(v_{it}v_{jt-k}) = 0, i = 1,2;\ j = 1,2, \forall k > 0
\end{aligned}$$

where SW1 impose $\lambda = 1$ but we can consider this as limiting case (we can also in principle follow SW1 and allow the two shocks to be drawn from time-varying distributions, but it will not matter for the results here). Then $E_t\pi_{t+1} = \lambda\pi_t^P$, hence, substituting into the NKPC we have

$$\begin{aligned}
\pi_t &= \beta E_t \pi_{t+1} + \mu_t^{SW} \\
\pi_t^P + \pi_t^T &= \beta\lambda\pi_t^P + \mu_t^{SW} \\
&\Rightarrow \mu_t^{SW} = (1-\beta\lambda)\,\pi_t^P + \pi_t^T = \mu_t^P + \mu_t^T
\end{aligned}$$

so μ_t must have an identical transitory component to inflation, but the weight on the permanent component will be (much) lower, hence μ_t^{SW} will be very much closer to white noise. But, as in the main case of the paper, μ_t^{SW} can never be pure white noise, since it is this component that is the source of any predictability of inflation. (Equivalently, any shock to μ_t^P is multiplied up by the factor $(1-\beta\lambda)^{-1}$ because of its predictive power for future marginal costs, and thus matters much more for the time series properties of inflation than it does for marginal costs).

The assumption that the SW1 decomposition is known is of course just an assumption. This representation of inflation is *not* fundamental for π^t (indeed *no* representation with two shocks can be fundamental to a single time series). But Robertson & Wright show that it is a special case of a predictive model for inflation, with an implied R^2 that lies roughly halfway between the lower bound given by the fundamental ARMA representation, and the upper bound given by the nonfundamental representation. As such it arguably represents a plausible representation of price-setting behaviour using an information set that is superior to the econometrician's but is still some way from providing the best possible predictor.

C A general derivation of the link between time series properties of driving and observables processes

Consider a general process for the driving variable in Proposition 1,

$$x_t = c(L)\,v_t = \sum_{i=0}^{\infty} c_i L^i v_t$$

with $c_0 = 1$. Then, substituting into the forward sum, we have

$$
\begin{aligned}
y_t &= \sum_{i=0}^{\infty} \beta^i E_t c(L)\, v_{t+i} \\
&= c_0 v_t + c_1 v_{t-1} + c_2 v_{t-2} + ... \\
&\quad + \beta \left(c_1 v_t + c_2 v_{t-1} + c_3 v_{t-2} + ...\right) \\
&\quad + \beta^2 \left(c_2 v_t + c_3 v_{t-1} + c_4 v_{t-2} + ...\right) \\
&\quad + \beta^3 \left(c_3 v_t + c_4 v_{t-1} + c_5 v_{t-2} + ...\right) \\
&\quad + ... \\
&= \sum_{i=0}^{\infty} \gamma_i L^i v_t = \gamma(L)\, v_t
\end{aligned}
$$

where, $\gamma(L)$ is a structural representation, i.e., defined in terms of the structural shock v_t, and thus may in principle be nonfundamental.

By summing all terms in v_t we have

$$
\gamma_0 = \sum_{i=0}^{\infty} c_i \beta^i = c(\beta)
$$

and we also have

$$
\begin{aligned}
\gamma(1) &= c_0 + c_1 + c_2 + ... \\
&\quad + \beta \left(c_1 + c_2 + c_3 + ...\right) \\
&\quad + \beta^2 \left(c_2 + c_3 + c_4 + ...\right) \\
&\quad + \beta^3 \left(c_3 + c_4 + c_5 + ...\right) \\
&= \sum_{i=0}^{\infty} c_i \sum_{j=0}^{i} \beta^j = \sum_{i=0}^{\infty} c_i \frac{1-\beta^{i+1}}{1-\beta} \\
&= \frac{\sum_{i=0}^{\infty} c_i}{1-\beta} - \frac{\beta \sum_{i=0}^{\infty} c_i \beta^i}{1-\beta} \\
&= \frac{c(1) - \beta c(\beta)}{1-\beta}
\end{aligned}
$$

I can only find a recursive definition for the individual γ_i. We have

$$
\begin{aligned}
\gamma_1 &= c_1 + \beta c_2 + \beta^2 c_3 + ... \\
&= \beta^{-1} \left(c_1 \beta + c_2 \beta^2 + c_3 \beta^3 + ...\right) \\
&= \beta^{-1} \left(\gamma_0 - c_0\right)
\end{aligned}
$$

$$\begin{aligned}\gamma_2 &= c_2 + \beta c_3 + \beta^2 c_4 + \ldots. \\ &= \beta^{-1}\left(c_2\beta + c_3\beta^2 + c_4\beta^3 + \ldots.\right) \\ &= \beta^{-1}\left(\gamma_1 - c_1\right)\end{aligned}$$

hence by by iteration

$$\gamma_{k+1} = \beta^{-1}\left(\gamma_k - c_k\right)$$

which is not very intuitive but can at least be programmed.

Finally, note that $\gamma(L)$ needs to be re-scaled such that the impact effect is unity, define a new polynomial and shock process

$$y_t = d(L)\,\omega_t$$

where

$$\begin{aligned}\omega_t &= \gamma_0 v_t = c(\beta)\,v_t \\ d(L) &= \frac{\gamma(L)}{c(\beta)}\end{aligned}$$

and hence

$$d_{k+1} = \beta^{-1}\left(d_k - \frac{c_k}{c(\beta)}\right)$$

which then gives

$$d(1) = \frac{\gamma(1)}{\gamma(0)} = \frac{\frac{c(1)}{c(\beta)} - \beta}{1-\beta}$$

Hence

$$c(1) > c(\beta) \Rightarrow d(1) > 1$$

eg, for an ar(1) this gives

$$d(1) = \frac{\frac{1-\lambda\beta}{1-\lambda} - \beta}{1-\beta} = \frac{1}{1-\lambda}$$

For the arma(1,1) with ψ the MA parameter of the driving process, and γ the (structural) MA parameter of the y_t process, it gives

$$d(1) = \frac{\frac{(1-\psi)(1-\lambda\beta)}{(1-\lambda)(1-\psi\beta)} - \beta}{1-\beta} = \frac{1-\psi\left(\frac{1-\lambda\beta}{1-\psi\beta}\right)}{1-\lambda} = \frac{1-\gamma}{1-\lambda}$$

which also checks out (somewhat miraculously). We also know already for this case that the white noise case, $\psi = \lambda$ for the driving process maps to white noise, and hence $d(1) = 1$ for the observable. There may be a similarly straightforward dividing line for more general arma models.

There is however no obvious way to see for the general case under what conditions the $d(L)$ process is fundamental, hence whether, or under what conditions, $\omega_t = \varepsilon_t$ or η_t.