BANK OF ENGLAND

Speech

Conditional guidance as a response to supply uncertainty

Appendix to the speech given by

Ben Broadbent, External Member of the Monetary Policy Committee, Bank of England

At the London Business School, Regents Park

Monday 23 September 2013

I would like to thank Alina Barnett, Adrian Chiu and David Barkshire for research assistance, and I am also grateful for helpful comments from other colleagues, especially Spencer Dale, Gabor Pinter, Matt Waldron and Martin Weale. The views expressed are my own and do not necessarily reflect those of the Bank of England or other members of the Monetary Policy Committee.

This technical appendix goes through the details of three points in the main body of the speech: the extent to which flexible pay can account for the productivity shortfall; why (and how) the behaviour of relative prices tells you about losses due to slow reallocation of resources (specifically capital); the sensitivity of optimal policy to unemployment and output growth when both are noisy indicators of the "output gap".

Flexible pay and productivity

The point that productivity is more cyclical when real pay is flexible can be represented with a simple labour demand and supply diagram. For the demand side suppose there's a representative firm that uses CES production with an elasticity of substitution σ. For given output y and a real product wage w its demand for labour is $l = y - (1-\sigma)a - \sigma w$, where a is TFP and all the variables are in logs. Thus, along any given demand curve, (log) labour productivity is $y - l = \sigma w + (1-\sigma)a$.

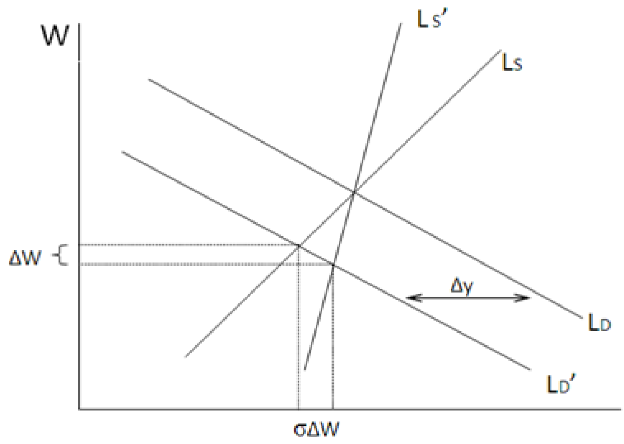

Lower output shifts this demand curve to the left. This might be caused by shift in TFP (the "a" term), which will itself lower output. But it's also possible that y falls independently of a, thanks to a pure demand shock. The Chart represents such a shift and compares two cases: one with a relatively flat labour supply curve (inflexible pay), one where it's steep (flexible pay). In the second case, real pay w will fall further and measured productivity, even for given TFP, will be lower. The size of the difference is $\sigma\Delta w$, where Δw is the extent to which real wages are lower than they otherwise would have been.

While they don't give a precise estimate, Pessoa and Van Reenen (2013) say that real pay has become more flexible since earlier recessions, and is up to 2% weaker (all else equal) as a result. They then assume $\sigma = 1$ (Cobb-Douglas production) and conclude that flexible pay can account for 2% of the measured drop in productivity, relative to past downturns.

My own view is that the effect is smaller. First, a reasonable estimate for the UK suggests σ is around one-half, not one (see, for example, Ellis and Price (2003)). Second, it's not clear that real wages are, in fact, that much lower, given the rate of unemployment, than one might have expected prior to the downturn (see Chart 6 in Broadbent (2012(b)). Third, the maintained hypothesis here is that the entire decline in output is, by assumption, the result of lower aggregate demand, rather than a shift in effective supply (the a term). In other words it's possible that as much as 1% point of the productivity shortfall is the result of more flexible pay and a purely demand-driven downturn (σ times the 2% difference in real wages). But it could also be that the drop in output and wages are themselves the consequences of a drop in effective supply (the a term). Especially if you regard the Δw = 2% figure as a little high, the 1% implied hit to productivity should be as a ceiling, not a central estimate.

Reallocation, lost productivity and the variance of prices

Imagine a firm that uses inputs K and L to produce $y_i = f_i(K_i, L_i)$ and that the price of its output is p_i. Suppose we increase its employment by ΔL_i, holding fixed the other input K. Then the (base-weighted) value of its output will change by

$$p_i^0 \Delta y_i = p_i^0 \int_{L_i^0}^{L_i^0+\Delta L_i} f_{iL}^{'}(K, l)dl \qquad (1)$$

where a zero superscript indicates the starting value and $f_{iL}^{'}$ the marginal product of labour. We want to think about what happens to the change in aggregate output $\Delta Y = \sum_i p_i^0 \Delta y_i$ when shifts in relative demand are met by changes in labour alone. In doing so we assume there is a fixed supply of labour in aggregate (call it L) and that the labour market clears. The first condition means $\sum_i \Delta L_i = 0$. The second means there a common wage w, across all sectors, and that firms are on their labour demand curve

$$w = p_i f_{iL}^{'} \qquad (2)$$

To work out the effect of a reallocation of labour on aggregate output note that the first-order approximation to the integral in (1) is

$$p_i^0 \int_{L_i^0}^{L_i^0+\Delta L_i} f_{iL}^{'}(K, l)dl \approx p_i^0[f_{iL}^{'0} + \tfrac{1}{2}\Delta f_{iL}^{'}]\Delta L_i \approx \tfrac{1}{2} w^0 \left[1 + \frac{\Delta w}{w^0} - \frac{\Delta p_i}{p_i^0}\right] \Delta L_i \qquad (3)$$

where the last term follows from (2). So, aggregating (1) and substituting from (3), the proportionate change in productivity is

$$\frac{\Delta Y_i}{Y} \approx \frac{1}{2}\frac{w^0 L}{Y} \sum_i \left[1 + \frac{\Delta w}{w^0} - \frac{\Delta p_i}{p_i^0}\right] \frac{\Delta L_i}{L} = -\frac{1}{2}\alpha cov_i\left(\frac{\Delta p_i}{p_i^0}, \lambda_i \frac{\Delta L_i}{L_i^0}\right) \qquad (4)$$

where α is the share of wages in national income (wL/Y) and λ_i is employment in sector i relative to the average. Note that, of the three terms in square brackets, the first two aggregate to zero because the wage (and its change) are common to all sectors, so can be taken out of the summation, and we have restricted $\sum_i \Delta L_i = 0$.

The share of wages in GDP is roughly two-thirds. So this relationship says the loss in productivity is (to a first-order approximation) one third the cross-sectoral covariance between inflation and size-weighted employment growth.

Further approximating the relationship between price and employment growth from (2), and using σ as the elasticity of substitution between capital and labour in sector i, and α_i as the share of labour income in that sector, one can re-express this in terms of prices alone:

$$\frac{\Delta Y_i}{Y} \approx -\tfrac{1}{2}\alpha cov_i\left(\frac{\lambda_i\sigma_i}{1-\alpha_i}\frac{\Delta p_i}{p_i^0}, \frac{\Delta p_i}{p_i^0}\right) = -\tfrac{1}{2}\alpha\mu var_i\left(\sqrt{\frac{\mu_i}{\mu}}\frac{\Delta p_i}{p_i^0}\right) \qquad (5)$$

where $\mu_i = \frac{\lambda_i\sigma_i}{1-\alpha_i}$ and $\mu = \frac{\sigma}{1-\alpha}$ is the same quantity for the economy as a whole. Empirical estimates suggest that whole-economy σ is around a half and α two-thirds. On that basis αμ is one and the loss of productivity will be around one half the cross-sectoral variance of (μ-weighted) inflation. This is the relationship we used in the main text.

A few more points:

First, if their derivation is a bit fiddly the intuition behind these relationships is relatively straightforward: if sectors where demand and employment grow (shrink) also see price increases (falls) – i.e. the sectoral supply curves slope upwards – then we know that marginal costs ($w/f_{iL}^{'}$) have become dispersed and aggregate productivity is lower as a result.

Second, other than the assumption that labour demand curves are well defined, there are no restrictions here on production functions: they don't need to be the same in all sectors, nor even do they need to exhibit constant returns to scale (CRS). But: if there are constant returns, and if capital too is mobile, then relative prices don't change in response to demand shocks (relative supply curves are flat) and productivity is invariant to shifts in relative demand.

To see this note that CRS means we can write marginal products in any sector i solely in terms of the capital:labour ratio (call that $k_i \equiv \frac{K_i}{L_i}$): defining the function g as $g_i(x) \equiv f_i(x, 1)$, we have $y_i = L_i g_i(k_i)$, $f_{iK}^{'} = g_i^{'}$ and $f_{iL}^{'} = g_i - k_i g_i^{'}$. So, under full factor mobility, and for given factor prices r and w, the two first-order conditions for optimal employment of labour $= p_i f_{iL}^{'} = p_i(g_i - k_i g_i^{'}$) and capital $r = g_i^{'}$ are enough to solve for the two variables k_i and p_i : the scale of production L_i doesn't matter for prices.

Third, note that I've assumed throughout that the production functions are constant: all we're doing is simulating pure (cross-sectoral) demand shocks, and the formula for lost productivity (roughly one half the cross-sectoral variance of prices) depends on this assumption. But if there were such things as TFP shocks, at a sectoral level, these would obviously have effects on relative prices too. By the same token they would also have a (more direct) impact on aggregate TFP.

Extracting a signal about the output gap from noisy growth, unemployment

We begin with a general result due to Svensson and Woodford (2002). They analyse how optimal policy, in a general linear model of the economy, with quadratic preferences, is affected by incomplete information about economically relevant variables, in particular the "output gap". They demonstrate that the sensitivity of policy to the estimate of the output gap is invariant to its degree of accuracy. Specifically, suppose that, when you know the output gap x with complete certainty, optimal policy involves some reaction function akin to a Taylor rule:

$$R = \beta x + inflation\ terms$$

If you then introduce some noise into the observation of x, the optimal policy is just

$$R = \beta \hat{x} + inflation\ terms$$

where $\hat{x}$ is the best possible estimate of the output gap. Everything else, including the sensitivity parameter β, is the same as in the complete-information case. This means the relative sensitivities of optimal policy to employment and output growth are the same as in the formula for $\hat{x}$.

We now turn, in a particular (and very simple) setting, to what that "signal extraction" formula looks like. The economy has an output gap x subject to persistent demand shocks d, unemployment u that depends on a distributed lag of x and a white-noise disturbance u*, and output growth that's (by definition) a combination of changes in the output gap and an underlying supply shock, also assumed to be white noise, Δs (Δ denotes the change in a variable so s itself is a random walk). Thus:

$$x_t = \alpha x_{t-1} + d_t \tag{1}$$
$$U_t = -\gamma_1 x_t - \gamma_2 x_{t-1} + u_t^* \tag{2}$$
$$\Delta y_t = \Delta x_t + \Delta s_t \tag{3}$$

The standard deviations of the variances of the three shocks are σ_d, σ_u and σ_s respectively. We observe Δy and u. Our task is to find the best possible estimate of x.

The easiest way to do this is to re-express (1)-(3) in state space form and use a Kalman filter. So the state variables are x and its lagged value and evolve according to the state equation

$$\begin{bmatrix} x_t \\ x_{t-1} \end{bmatrix} = \begin{bmatrix} \alpha & 0 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} x_{t-1} \\ x_{t-2} \end{bmatrix} + \begin{bmatrix} d_t \\ 0 \end{bmatrix} \tag{1)'$$

The observed variables growth and unemployment are given by the measurement equation.

$$\begin{bmatrix} u_t \\ \Delta y_t \end{bmatrix} = \begin{bmatrix} -\gamma_1 & -\gamma_2 \\ 1 & -1 \end{bmatrix} \begin{bmatrix} x_t \\ x_{t-1} \end{bmatrix} + \begin{bmatrix} u_t^* \\ \Delta s_t \end{bmatrix} \qquad (2)'$$

The assumed parameter values are derived from simple regression estimates on quarterly UK data: $\alpha = 0.8$, $\gamma_1 = 0.1$, $\gamma_2 = 0.4$, $\sigma_d^2 = 0.2$, $\sigma_u^2 = 0.1$, $\sigma_s^2 = 0.3$. We then solve the problem numerically. What drops out is an updating rule for the estimated output gap:

$$\hat{x}_t = \hat{x}_{t-1} + \kappa_y(\Delta y_t - \Delta\hat{y}_t) + \kappa_u(u_t - \hat{u}_t),$$

where $\Delta\hat{y}_t$ and $\hat{u}_t$ are prior expectations of growth and unemployment respectively. Chart 9 in the main text (reproduced below) plots the values of the κ coefficients as you vary σ_s.

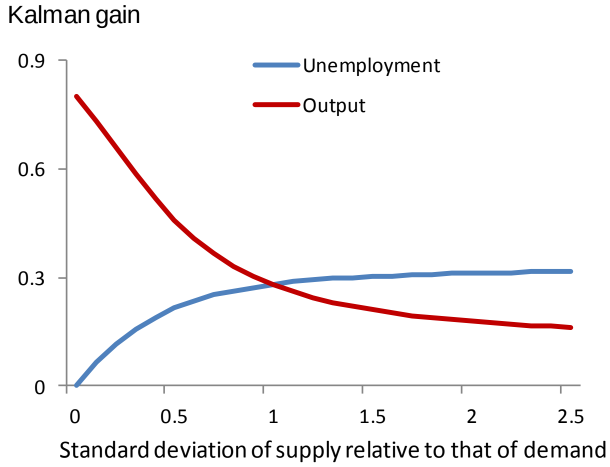

Note that these coefficients depend on other parameters in the model. In particular, the weight on unemployment rises, relative to that on output if either (i) demand shocks are more persistent (higher α) and/or (ii) the sooner unemployment reacts to demand (γ1 rises relative to γ2) – see the graphs below. The intuition for (i) is that, if demand shocks are more persistent, so are any policy errors one makes by (wrongly) reacting to a supply shock Δs. So the weight on output falls. And the sooner the one reacts to the other, the better unemployment serves as an indicator of the output gap.

Sensitivity analysis

Sensitivity of the relative gain to changes in α

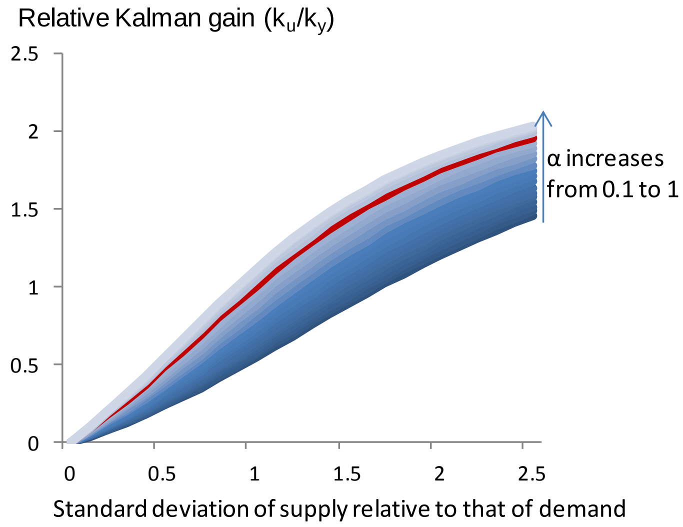

Note: Red line indicates the base-line simulation gain when $\alpha = 0.8$

Sensitivity of the relative gain to changes in γ_2/γ_1

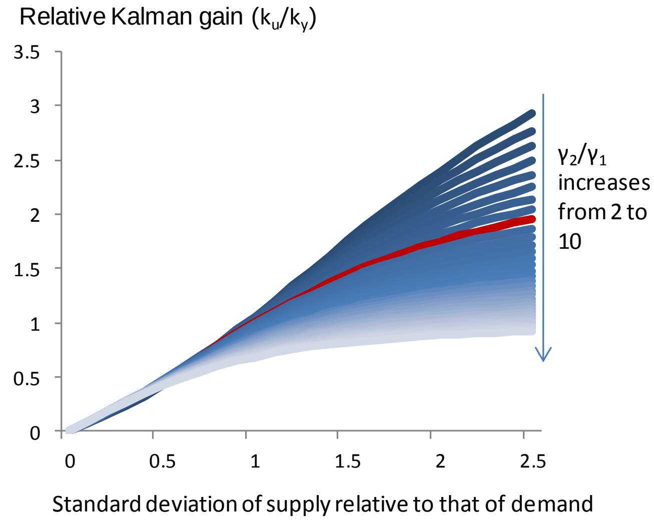

Note: Red line indicates the base-line simulation gain when $\gamma_1/\gamma_1 = 4$

References

Broadbent, B., 2012(b), "Productivity and the allocation of resources", Speech given at Durham Business School available at http://www.bankofengland.co.uk/publications/Documents/speeches/2012/speech599.pdf

Ellis, C. and S. Price, 2003, "UK business investment: long-run elasticities and short-run dynamics", Bank of England Working Paper Series no. 196

Pessoa, J.P and J Van Reenen, 2013, "The UK productivity and jobs puzzle: does the answer lie in the labour market flexibility?", CEP working paper.

Lippi, F, 2003, "Monetary policy with unobserved potential output", BIS working paper no. 19.

Svensson, L. and M. Woodford, 2000, " Indicator variables for optimal policy", European Central Bank, Working Paper Series no 0012.