The Economic Journal, **116** (*April*), 434–456. © Royal Economic Society 2006. Published by Blackwell Publishing, 9600 Garsington Road, Oxford OX4 2DQ, UK and 350 Main Street, Malden, MA 02148, USA.

CLOSED AND OPEN ECONOMY MODELS OF BUSINESS CYCLES WITH MARKED UP AND STICKY PRICES*

Robert J. Barro and Silvana Tenreyro

Shifts in the extent of competition, which affect markups, are possible sources of aggregate fluctuations. Markups are countercyclical; during booms the economy operates more efficiently. In our benchmark model, markups correspond to the prices of differentiated inputs relative to that of undifferentiated final product. If nominal prices of differentiated goods are relatively sticky, unexpected inflation reduces markups, mimicking the effects of increased competition. Similar effects stem from reductions in markups of foreign intermediates and unexpected inflation abroad. The models imply that prices of less competitive goods are more countercyclical. We find support for this hypothesis using data of four-digit manufacturing industries.

An important branch of the macroeconomics literature views variations in markup ratios as major features of business cycles. This literature has recently been surveyed by Rotemberg and Woodford (1999). From the standpoint of generating fluctuations in aggregate economic activity, movements in markups – reflecting shifts in the extent of competition – work similarly to the technological disturbances usually stressed in real business cycle (RBC) models. Hence, shifts in the extent of competition provide another source of real shocks within the RBC framework. In the typical analysis, markups exhibit a countercyclical pattern, and booms are times at which the economy operates more efficiently.

This article begins with a real model in which intermediate inputs are differentiated products that are sold under conditions of imperfect competition. Final product, which can be used for consumption or to create the intermediate goods, is undifferentiated and, hence, competitive. In this model, the markup ratios correspond to the prices of the intermediate goods relative to the price of final product. A reduction in markup ratios spurs the use of intermediates and, thereby, generates an efficient expansion of output and consumption. Labour productivity also rises, and the increase in the marginal product of labour leads to an increase in the real wage rate.

An extended version of the model assumes that the nominal prices of the differentiated intermediate goods are sticky relative to the nominal price of undifferentiated final product. In this environment, unexpected inflation in the price of final product tends to reduce the relative price of intermediates. The expansionary effect on output is the same as that generated from an increase in competition. Hence, some amount of unexpected inflation can look desirable, *ex post*, to the monetary authority.

A further extension allows for trade of the intermediate goods across international borders at non-zero transaction costs. (Final product is assumed to be tradable without transaction costs.) In this model, increases in foreign competition and unexpected

* This research is supported by the National Science Foundation. We are grateful for helpful comments from Daron Acemoglu, Olivier Blanchard, Ricardo Caballero, Pat Kehoe, Nobu Kiyotaki, Greg Mankiw, Ken Rogoff, Julio Rotemberg, Jaume Ventura and two anonymous referees.

inflation in the foreign country tend to be expansionary at home. The model also has implications for the effect of openness on a monetary authority's incentive to inflate.

The models imply that a sector's relative output price will be more countercyclical the less competitive is the sector. In a later Section of the article, we use price data from four-digit manufacturing industries to test this hypothesis. The results support the proposition that less competitive – or, at least, more concentrated – sectors feature more countercyclical movements in their relative output prices.

Elements of our approach are in earlier models of sticky prices and imperfect competition – see, for example, Svensson (1986), Blanchard and Kiyotaki (1987), Romer (1993), Obstfeld and Rogoff (1995), Ireland (1996) and Lane (1997). These models rely on the varieties-type technology or utility function from Spence (1976), Dixit and Stiglitz (1977), and Ethier (1982), and assume monopolistic competition in the provision of final goods or intermediate inputs. This specification of technology or utility, together with the monopolistic competition assumption, implies that the markup is determined by the inverse of the elasticity of substitution between varieties in the production or utility function. Our model differs from these formulations in that the markup is not linked to the elasticity of substitution. This distinction is empirically important because estimated markups tend to be small – see, for example, Basu and Fernald (1997) – implying unrealistically large elasticities of substitution in the existing models. Alternatively, studies that estimate the elasticity of substitution – such as Broda and Weinstein (2004) – tend to find relatively small values, which would require large markups in the existing models. The empirical disconnect between the two parameters calls for a deviation from the standard monopolistic competition assumption present in the literature. We address this issue by modelling imperfect competition as a Bertrand game between the potential providers of each variety of intermediate goods.

1. The Model of a Closed Economy

1.1. *The Real Model*

Competitive firms produce output using a varieties-type production function, which was originated by Spence (1976), Dixit and Stiglitz (1977), and Ethier (1982). The output of firm i is given by

$$Y_i = AL_i^{1-\alpha} \sum_{j=1}^{N} X_{ji}^{\alpha}, \tag{1}$$

where $A > 0$ is a productivity parameter, L_i is firm i's employment of labour, $0 < \alpha < 1$, X_{ji} is the amount of intermediate input of type j used by firm i, and N is the number of types of intermediates available. Everyone has free access to the technology shown in (1). In the basic model, labour is exchanged on a competitive, economy-wide labour market.

We think of the intermediate inputs as specialised goods, such as machine tools and computers. In practice, these goods tend to be durables, so that increases in the X_{ji} require investment outlays. However, to keep things simple, we assume that the intermediate goods are non-durable. This assumption eliminates any dynamic elements but the model can be extended, without changing the basic results, to treat the inputs as capital goods.

Each firm maximises profit, taking as given the economy-wide real wage rate, W, and the price, P_j, of each type of intermediate good. (The prices are all measured in units of final product.) The first-order conditions for the choices of intermediate inputs are

$$A\alpha L_i^{1-\alpha} X_{ji}^{\alpha-1} = P_j, j = 1, \ldots, N. \tag{2}$$

Therefore, every producer of final goods will use all N varieties of the intermediate inputs as long as all of the prices are finite. It can be readily verified that the profit of each firm is zero if the real wage rate equals the marginal product of labour:

$$W = (1-\alpha)Y_i/L_i. \tag{3}$$

Final output is a homogeneous good that can be used for consumption or to produce intermediate goods. All consumer goods are identical. Prices of consumer goods are the same everywhere and are normalised to one.

We use a simple structure to allow for imperfect competition in the exchange of the specialised intermediate inputs. These goods are produced in sectors $j = 1, \ldots, N$. We assume that each sector has a large number of potential firms that have the ability to produce each type of intermediate good, effectively by sticking distinctive labels on the homogeneous final product. However, these firms differ in their costs of production, in the sense of the number of units of final product required to create a unit of intermediate good. We assume that each sector possesses a single leader, who has the lowest costs of production. We normalise so that this lowest-cost provider can produce one unit of intermediate for each unit of final product. If no potential competitors existed, the leader would price at the monopoly level. The constant-elasticity demand function implied by (2) determines the monopoly price of each intermediate good to be $1/\alpha$.[1]

To allow for the potential competition, let $\mu_j > 1$ be the number of units of final good required by the next most efficient producer to create a unit of intermediate good in sector j. The price charged for the good is assumed to be determined from Bertrand competition between the industry leader and the potential entrants. Hence, if $\mu_j \le 1/\alpha$, the leader sets the price just below μ_j and, thereby, obtains the full market. If $\mu_j > 1/\alpha$, the leader prefers to set the lower price $1/\alpha$, which is the monopoly value, and still obtains the full market. Thus, the leader's price is given by

[1] The parameter α also equals the share of payments to intermediate goods in total output. Therefore, the monopoly markup ratio for intermediates is restricted to equal the reciprocal of the factor share of intermediates. This restriction applies because the parameter α in (1) represents two things – factor shares and the degree of substitution across the intermediate inputs. To disentangle these two effects, the production function can be generalised to

$$Y_i = AL_i^{1-\alpha}\left(\sum_{j=1}^{N} X_{ji}^{\sigma}\right)^{\alpha/\sigma},$$

where $0 < \sigma \le 1$. Equation (1) applies when $\alpha = \sigma$. For given α, a higher σ means that the intermediate inputs are closer substitutes, with perfect substitution corresponding to $\sigma = 1$. The monopoly markup ratio can be determined (if $N >> 1$) to equal $1/\sigma$ in the generalised setup, whereas the parameter α still equals the factor share for intermediates. Hence, in this representation, the monopoly markup ratio for intermediates no longer necessarily equals the reciprocal of the income share.

$$P_j = \mu_j \text{ if } 1 < \mu_j < 1/\alpha, \quad P_j = 1/\alpha \text{ if } \mu_j \geq 1/\alpha. \tag{4}$$

To simplify, we assume for now that the structure of competition is the same across sectors, so that $\mu_j = \mu$ for all j, where $1 < \mu < 1/\alpha$. In this case, the parameter μ represents the economy-wide markup ratio.[2]

The quantities of intermediates employed by firm i are given by

$$X_{ji} = (A\alpha/\mu)^{1/(1-\alpha)} L_i, j = 1, \ldots, N. \tag{5}$$

Substitution into (1) and aggregation over the firms determines the aggregate level of output as

$$Y = A^{1/(1-\alpha)} \alpha^{\alpha/(1-\alpha)} L(1/\mu)^{\alpha/(1-\alpha)} N, \tag{6}$$

where L is the economy-wide labour force, which is assumed for now to be constant. A lower μ encourages the use of intermediates (5) and, thereby, raises output, Y, in (6). Since employment, L, is fixed, the increase in output corresponds to a rise in labour's average productivity, Y/L. The rise in productivity occurs for a given form of the production function because the heightened competition corresponding to the lower μ leads to a more efficient – in this case, more intermediated – structure of production.[3] First-best output turns out to correspond to $\mu = 1$ in (6) (see below). The ratio of actual to first-best output is equal to $(1/\mu)^{\alpha/(1-\alpha)} < 1$. Output is also increasing in the productivity parameter, A, and in the number of intermediates, N. In the related growth literature, summarised in Barro and Sala-i-Martin (2004, ch. 6), increases in N are the key to growth. However, N is fixed in the present context.

Households own all of the firms in the economy. The only firms that make profits in equilibrium are the lowest-cost providers of intermediates in each of the N sectors. The ownership rights in these firms are assumed to be distributed evenly across the households. In this case, the model has a representative household, whose net income and consumption correspond to gross output (from (6)) less the total production of intermediates (determined from (5)). The formula for aggregate consumption is

$$C = A^{1/(1-\alpha)} \alpha^{\alpha/(1-\alpha)} L(1/\mu)^{1/(1-\alpha)} (\mu - \alpha) N. \tag{7}$$

It can be verified from (7) that C falls with μ when $\mu \geq 1$ and that C is maximised at $\mu = 1$. That is, consumption of the representative individual is maximised under perfect competition, where the economy-wide markup ratio, μ, equals one.[4] At $\mu = 1$, each type of intermediate good is efficiently utilised up to the point where its marginal product equals unity (the lowest-cost provider's constant cost of production) in (2).

If we had imposed the condition $\mu < 1$, which is inconsistent with the discussion that underlies (4), then (7) implies that C falls if μ decreases. The reason is that interme-

[2] Productivity differences between the leader and the followers might not reflect only differences in firms' technologies. For example, geographical location, taxes and regulations affect the costs of production.

[3] Similar effects occur in the model considered by Basu (1995, section III).

[4] Note that the number N is fixed exogenously in this model. Hence, the economy does not need any monopoly profits to provide incentives for invention, as in Romer (1990) and the rest of the endogenous growth literature summarised in Barro and Sala-i-Martin (2004, chs. 6 and 7).

diates would be utilised too much from an efficiency standpoint if $\mu < 1$. However, the specific results from (7) depend on the unrealistic assumption that the providers of intermediates continue to meet all of the demand even when $\mu < 1$. In fact, if the markup ratio is set below unity, the lowest-cost providers would lose money on each unit produced and sold and would be better off closing down. In this case, output and consumption would collapse to zero. The general lesson – which will be important when we consider unexpected inflation – is that the economy operates inefficiently when $\mu < 1$.

In this real model, business fluctuations could be driven by shocks to the overall productivity parameter, A, and the markup ratio, μ. Movements in A look like the disturbances that are usually stressed in real business cycle models. For given μ, these shocks generate movements in output without any changes in the markup ratio – that is, the markup ratio would be acyclical in this case.

For given parameters of the production function in (1) – or the more general form presented in footnote 1 – shifts in μ in (4) reflect exogenous changes in the economy-wide extent of competition.[5] An increase in μ leads to a decline in output and, if $\mu > 1$, also to a decline in consumption. Hence, these shocks would generate a countercyclical pattern for the markup ratio.

The markup ratios in this model measure the prices of specialised goods that are sold with some degree of monopoly power expressed relative to the prices of competitive goods. In our context, the specialised goods are intermediate inputs, and the undifferentiated goods are final products. However, the results would be similar if, instead, some or all of the specialised goods were final products. In this case, an increase in the extent of competition applicable to the specialised final goods would be expansionary; in particular, households would be better off, and measured real quantities of GDP and consumption would rise.

Our concept of the markup ratio differs from the one stressed in the literature, such as Bils (1987) and Rotemberg and Woodford (1999). That literature focuses on the price of final product expressed relative to the marginal cost of production, which involves variations in inputs such as labour. For example, if labour input is paid the real wage rate W, the Bils markup ratio equals F_L/W. If the production function takes the form of (1), then $F_L = (1-\alpha)(Y/L)$,[6] and, hence, $F_L/W = (1-\alpha)/Sh(L)$, where $Sh(L) = WL/Y$ is labour's share of total gross product. Therefore, to generate a pattern in which this concept of the markup ratio is countercyclical, the labour share has to be procyclical. However, the labour share tends empirically to be countercyclical – see Rotemberg and Woodford (1999, Section 2.1). They argue that the labour share is less countercyclical than it first appears but they are still unable to generate a procyclical pattern.

[5] One possibility that we neglect here is that the state of the business cycle may affect the degree of competition and, hence, the markup ratio. Rotemberg and Saloner (1986) argue that booms tend to intensify competition and, therefore, reduce markup ratios. In the present model, this effect would operate only if the business cycle influenced the production costs of potential competitors relative to the costs of industry leaders.

[6] This result holds more generally if the production function can be written as $Y_i = L_i^{1-\alpha} G(\tilde{X}_{ji}, K_i)$, where $\tilde{X}_{ji}$ represents the vector of intermediate inputs for $j = 1, \ldots, N$, and K_i is capital input.

More generally, it seems that the markup ratio ought to be defined in terms of ratios of prices of goods or inputs that are sold under conditions of monopoly power to prices of goods or inputs that are sold under competitive conditions. The implicit assumption in the Bils (1987) and Rotemberg and Woodford (1999) frameworks[7] seems to be that final products are specialised and sold under conditions of monopoly power, whereas inputs (such as labour and raw materials) are non-specialised and sold under competitive conditions. From an empirical standpoint, it is not obvious that this assumption is reasonable.

Returning to the model, the real wage rate is determined from (3) and (6) to be[8]

$$W = (1-\alpha)A^{1/(1-\alpha)}\alpha^{\alpha/(1-\alpha)}(1/\mu)^{\alpha/(1-\alpha)}N. \tag{8}$$

Therefore, shocks to A and μ cause W to move along with output, which is determined by (6). That is, the real wage rate moves procyclically.

Labour's share of total gross product, WL/Y, is fixed at $1-\alpha$ (see (3)) and is, therefore, acyclical unless α is changing. Correspondingly, the share in payments to intermediate inputs is fixed at α. These payments can be broken down into profits and production costs of intermediate-goods providers. The ratio of profit to output is increasing in μ, whereas the ratio of production costs (for intermediate-goods providers) to output is decreasing in μ. These ratios are invariant with A.

One difficulty with this analysis of shares is that the usual concept of gross product in the national accounts nets out non-durable intermediate inputs, such as the X_{ji} in the model. In the present setting, this concept of product equals consumption, because no durable goods have been introduced.[9] However, as noted before, the treatment of the intermediate inputs as nondurable was a matter of analytical convenience and not an accurate description of the typical differentiated good used as an input to production. We can revise the model to treat the X_{ji}s as durables. If we assume mobility of these durables across types, then the total stock of capital constitutes the single state variable. In this case, the model has dynamics similar to the standard neoclassical growth model. In this setting, the ratio of wL to gross product – which is inclusive of the gross investment outlays on intermediates – is fixed at $1-\alpha$. Therefore, this more realistic version of the model still has the implication that labour's share of gross product would be acyclical.

[7] In this context, some other models that are analogous include Blanchard and Kiyotaki (1987) and Mankiw (1991).

[8] In this model, the Bils markup ratio, F_L/W, always equals one. This result makes sense because, by assumption, the markets for labour and final product are both characterised by perfect competition. We could extend the model to allow some of the labour inputs to be specialised and sold under conditions of monopoly power. In this case, we could consider another markup ratio, involving the wage of specialised labour input expressed relative to that of undifferentiated labour input.

[9] The ratio of payments to labour to consumption is given from (8) and (7) by

$$wL/C = (1-\alpha)\mu/(\mu-\alpha).$$

Therefore, a decrease in μ raises wL/C, that is, the ratio of labour payments to consumption would move procyclically in the model. This pattern does appear in the data for most OECD countries if cyclical patterns are based on Hodrick-Prescott filtering, with wL measured by compensation of employees and C measured by total consumer expenditure or by consumer expenditure on non-durables and services.

A countercyclical pattern for labour's share of gross output tends to emerge if we modify the production function of (1) to have diminishing returns to scale and also introduce a fixed cost of operation for each firm. For example, we could have

$$Y_i = AL_i^{1-\alpha-\beta} \sum_{j=1}^{N} X_{ji}^{\alpha} - \gamma, \tag{9}$$

where $0 < \beta < 1 - \alpha$ and $\gamma > 0$. In this (Marshallian) model, the free-entry condition determines the number of firms so that profit is zero. However, if the number of firms does not respond to temporary variations in A and μ, booms (generated by high A or low μ) have relatively low shares in output of payments to labour and intermediate inputs. The share in output of profits of final-goods producers is procyclical – and is positive in booms and negative in recessions.

A labour–leisure choice could be introduced, so that L would be variable. We can think of each individual's work effort as depending on W with the usual types of substitution and income effects. For example, an economy-wide, temporary decline in μ would raise output in each sector, (6), and also raise the economy-wide real wage rate, W, in (8). Since the disturbance is temporary, the income effect from the higher real wage rate would be weak; therefore, the dominant impact on current labour supply would derive from the substitution effect that favours work over leisure. Hence, L would tend to rise, implying that employment would be procyclical.

Most of the results are similar if we modify the basic model to allow each sector to have a different degree of competition and, hence, a different markup ratio, μ_j. (We assume that $\mu_j \leq 1/\alpha$ applies for all j.) The solution for aggregate output is then a generalisation of (6):

$$Y = A^{1/(1-\alpha)}\alpha^{\alpha/(1-\alpha)} L \left[\sum_{j=1}^{N} (1/\mu_j)^{\alpha/(1-\alpha)}\right]. \tag{10}$$

Therefore, each sector is weighted inversely to its markup ratio, μ_j. The formula for aggregate consumption is now modified from (7) to

$$C = A^{1/(1-\alpha)}\alpha^{\alpha/(1-\alpha)} L \left[\sum_{j=1}^{N} (1/\mu_j)^{\alpha/(1-\alpha)} - \alpha \sum_{j=1}^{N} (1/\mu_j)^{1/(1-\alpha)}\right]. \tag{11}$$

The consumption maximising value for each of the markup ratios is again $\mu_j = 1$.

A common generalisation of the production function, mentioned in footnote 1, is

$$Y = AL_i^{1-\alpha} X_i^{\alpha}, \text{ with } X_i = \left(\sum_{j=1}^{N} X_i^{\sigma}\right)^{1/\sigma}. \tag{12}$$

Our specification in (1) corresponds to $\alpha = \sigma$. In a setup with constant N, it is easy to show that the symmetric equilibrium in the more general specification will be identical to ours, up to a constant.

The main results do not change if some intermediate goods are provided under perfect competition ($\mu_j = 1$) or if some final goods are provided under conditions of imperfect competition. The main point is that the relevant markup is the price of goods sold under conditions of imperfect competition compared to goods sold under competitive conditions.

1.2. *Nominal Price Stickiness*

To introduce nominal elements and a possible role for monetary policy, we use a simple setting in which the nominal prices of the intermediate goods involve some stickiness, whereas the prices of the final goods are flexible.[10] More generally, the assumption is that the more specialised and, hence, less competitive products – which, in our model, are the intermediate inputs – tend to feature less flexibility in their nominal prices. This specification accords with the theoretical model of Rotemberg and Saloner (1987), who find that the cost of having the wrong price tends to be greater for duopolists than for monopolists. Hence, with some fixed costs for changing prices, the prices of less competitive goods would tend to adjust less often. Empirical support for this specification is provided by Carlton (1986), who finds in the data of Stigler and Kindahl (1970) that less competitive industries (as gauged by concentration ratios) have more rigid prices (as measured by the frequency of zero month-to-month changes). Basu (1995) refers to this evidence and uses it to motivate an assumption of relative rigidity in the prices of intermediate goods. More recently, Bils and Klenow (2004, Table 3) find, for goods contained in the US consumer price index that the frequency of price adjustment is negatively and significantly related to the concentration ratio.

Let p_j be the nominal price of the *jth* intermediate good and p the nominal price of final goods (and, hence, consumer goods), all of which sell at one price. If all nominal prices were flexible, the preceding analysis would go through, with the relative price of each intermediate good, p_j/p, set to equal the markup ratio, μ_j, in accordance with (4). If the degree of competition were the same in each sector, $\mu_j = \mu$ for all j would again apply. The nominal wage rate would then equal pW, where W is given in (8).

Suppose that the nominal price of final product, p, is determined through some stochastic process by the country's monetary authority. That is, nominal monetary aggregates or nominal interest rates – which we do not model explicitly – are assumed to adjust to achieve a target nominal price of final goods.[11]

We assume now that the lowest-cost provider of intermediates in sector j sets the nominal price p_j one period in advance. That is, the industry leader effectively offers a contract to its buyers (who are producers of final product) in which the nominal price of the intermediate good is guaranteed for the next period. We assume that the other part of the contract is that the leader commits to meet the demand for the intermediate

[10] Our setting has synchronised price setting for the various goods. That is, the prices of intermediate goods all adjust together with a one-period lag (and the price of the single type of final product and the nominal wage rate adjust with no lag). Some alternative models assume staggered price adjustment – see, for example, Calvo (1983) and Chari *et al.* (2002). Staggered price adjustment may be empirically realistic and important for the model's detailed dynamics. However, this specification adds complexity without affecting the main results of our analysis.

[11] In much of the related business-cycle literature, such as Blanchard and Kiyotaki (1987) and Mankiw (1991), real effects from monetary stimuli depend on movements in real money balances. In the present model, the real effects from nominal shocks derive, instead, from changes in the prices of intermediates goods relative to the price of final product.

good that the customers turn out to express next period. (We consider later that the leader might, instead, opt not to meet the demand when the price fails to cover the cost of production.)

Other potential providers of intermediates in sector j can be thought of as offering similar fixed-nominal-price contracts. However, in the equilibrium, the lowest-cost provider will again price so that the next most efficient firm (and, moreover, all of the other firms) will be motivated not to participate.

To find the leader's nominal price, the only new element that we need is the probability distribution of p. As a first approximation, the industry leader in sector j will set the price as

$$p_j \approx \mu_j \mathrm{E}p, \tag{13}$$

where μ_j is the markup ratio given in (4), and Ep is each producer's one-period-ahead expectation of p. (All agents are assumed to have the same information and, therefore, the same value for Ep.) In the present case, μ_j will represent the target markup ratio, which will not be realised exactly if p departs from Ep. Equation (13) implies $p_j/p = \mu_j$, as before, if p is known with certainty one period in advance. When p is uncertain, the entire probability distribution of p would generally matter for the leader's optimal choice of p_j.[12] However, for present purposes, we assume that (13) is a satisfactory approximation.[13]

If p exceeds Ep, $p_j/p = \mu_j(\mathrm{E}p/p)$ falls correspondingly below the intended markup level, μ_j, in all sectors, and the demand for intermediates rises. We continue to assume, for now, that the lowest-cost provider of intermediate goods in each sector always meets the demand even when the real price is lower than intended.[14] If we also assume that all sectors have the same markup ratio – $\mu_j = \mu$ for all j – the expression for aggregate output from (6) is modified by replacing the parameter μ by $\mu(\mathrm{E}p/p)$ to get

$$Y = A^{1/(1-\alpha)}\alpha^{\alpha/(1-\alpha)}L[(1/\mu)(p/\ \mathrm{E}p)]^{\alpha/(1-\alpha)}N. \tag{14}$$

[12] For a monopolist, the value of p_j that maximises expected profit is given by

$$p_j = (1/\alpha)\frac{\int_0^\infty p^{1/(1-\alpha)}f(p)\ \mathrm{d}p}{\int_0^\infty p^{\alpha/(1-\alpha)}f(p)\mathrm{d}p},$$

where $f(\cdot)$ is the one-period-ahead probability density function for p. If $\alpha = 1/2$, this expression simplifies to

$$p_j = \frac{\mathrm{E}p}{\alpha}(1+s^2),$$

where s is the coefficient of variation of p. Hence, in this case, (13) holds (with $\mu_j = 1/\alpha$) if $s << 1$. If $\log(p)$ is normally distributed with variance σ^2, then

$$p_j = \frac{\mathrm{E}p}{\alpha}\exp[(\frac{\alpha}{1-\alpha})\sigma^2].$$

Hence, (13) holds here (with $\mu_j = 1/\alpha$) if $\sigma^2 << 1$.

[13] The subsequent analysis would not change materially if we modified the right-hand side of (13) to include higher moments of the distribution of p.

[14] If p falls short of Ep, the markup ratio rises above the target level, μ_j. In this case, higher cost firms in each sector might find it profitable to produce and sell intermediate goods. However, these firms would not have been willing, *ex ante*, to offer a fixed-nominal-price contract in which they were willing to meet whatever demands were realised. Since we are considering only these types of contracts, we assume that these competitor firms do not enter the market, *ex post*.

Hence, unexpected inflation raises output. Moreover, because of the distortion from markup pricing of the intermediate goods, this expansion of output is efficient over a range of unexpected inflation. If $\mu > 1$, the outcome $p/\mathrm{E}p = \mu > 1$ would generate the efficient level of production. That is, this amount of unexpected inflation would exactly offset the distortion from markup pricing.

The result for aggregate consumption is now a modification of (7):

$$C = A^{1/(1-\alpha)}\alpha^{\alpha/(1-\alpha)}L\left\{[(1/\mu)(p/\ \mathrm{E}p)]^{1/(1-\alpha)}(\frac{\mu\ \mathrm{E}p}{p} - \alpha)N\right\}. \tag{15}$$

We can show that C rises with $p/\mathrm{E}p$ if $p/\mathrm{E}p < \mu$ and is maximised when $p/\mathrm{E}p = \mu$. This result corresponds to the efficient use of intermediates for production when $p/\mathrm{E}p = \mu$.

We can also show that C falls with $p/\mathrm{E}p$ when $p/\mathrm{E}p \geq \mu$. In this range, intermediates would be overutilised from an efficiency standpoint (if the lowest-cost providers still meet the demand). Hence, while gross output continues to rise with unexpected inflation, net output and consumption decline. Thus, although some amount of unexpected inflation would be attractive – because it offsets existing distortions – too much unexpected inflation would be undesirable because it creates net new distortions.

If $p/\mathrm{E}p > \mu$, the real price of intermediates falls short of the lowest-cost provider's cost of production. As in the case discussed before where $\mu < 1$, the lowest-cost producers of intermediates would then do better, *ex post*, by shutting down. Of course, this failure to meet demand violates the form of the fixed-nominal-price contract that we had assumed, that is, the willingness to meet whatever demand materialised, *ex post*, at the set nominal price. In any case, if the leader in each sector were to shut down, too much unexpected inflation would result in a drastic decline of output and consumption. This result reinforces the conclusion that too much unexpected inflation would have adverse consequences.

From the standpoint of a policy maker, the model rationalises a loss function in which some amount of unexpected inflation (for prices of final product), $\pi - \pi^e$, reduces the loss. This kind of effect is often assumed in monetary models of rules versus discretion, such as Kydland and Prescott (1977) and Barro and Gordon (1983). In the present model and in some of the rules-versus-discretion literature, the negative effect of $\pi - \pi^e$ on the loss diminishes with the size of $\pi - \pi^e$, eventually becomes nil, and subsequently changes sign. The amount of the initial loss reduction and the size of the interval over which unexpected inflation is beneficial depends on the extent of the existing distortion. In the present model, the distortion increases with the markup ratio, μ. Thus, the policy maker would value unexpected inflation more when μ was higher, that is, when the extent of competition was smaller. Therefore, in the rules-versus-discretion setting, a higher μ would result in a higher equilibrium rate of inflation.

The positive effect of unexpected inflation on output in (14) reflects an increase in the use of intermediates and, thereby, a rise in the marginal product of labour. Therefore, unexpected inflation increases the real wage rate, which is now given as a modification of (8) by

$$W = (1-\alpha)A^{1/(1-\alpha)}\alpha^{\alpha/(1-\alpha)}[(1/\mu)(p/\mathrm{E}p)]^{\alpha/(1-\alpha)}N. \tag{16}$$

This result means that unexpected inflation causes the real wage rate to move procyclically. Labour's share of the total gross product, WL/Y, is, however, still constant at $1 - \alpha$.

If the total labour supply, L, is fixed, employment is constant and, hence, acyclical. However, as before, a positive response of labour supply to the increased real wage rate would generate a procyclical pattern for employment, as observed in the data. This positive response is particularly likely for unexpected price-level changes, which have to be temporary. In this case, the income effect from a higher real wage rate would be minor, and the main influence would be the substitution effect that favoured work over leisure.

If the markup ratios, μ_j, are heterogeneous across the sectors, (10) applies to aggregate output if $1/\mu_j$ is replaced by $(p/\mathrm{E}p)(1/\mu_j)$. The same change applies to the expression for aggregate consumption in (11). However, these formulas are again valid only if the lowest-cost provider in each sector always meets the demand for intermediates. As $p/\mathrm{E}p$ rises to reach the various μ_j, the corresponding sectors become unprofitable, and the providers would have the incentive, *ex post*, to shut down.

A new element with heterogeneity in the markup ratios is that the sectors do not all become unprofitable at the same time – in the earlier context, when $p/\mathrm{E}p$ reached the common markup ratio, μ. Suppose, without loss of generality, that the sectors are ordered so that $\mu_1 < \mu_2 < \mu_3$ etc. Then sector 1 would be motivated to close down when $p/\mathrm{E}p$ reached μ_1, sector 2 would also be motivated to close when $p/\mathrm{E}p$ reached μ_2 and so on. Thus, this adverse effect of excessive inflation now sets in only gradually.

2. The Model with Two Open Economies

2.1. *The Real Model*

To consider two or more open economies, we use a variant of the framework developed in Krugman (1980), Helpman and Krugman (1985, chs. 10–11), Alesina *et al.* (2000), and Chari *et al.* (2002). Suppose that there are two countries, where country I produces the intermediates $j = 1, \ldots, N^I$ and country II the intermediates $j = N^I + 1, \ldots, N^I + N^{II}$. (Generalisation to more than two countries is straightforward.) We assume that the countries do not overlap in the types of intermediate goods that they produce, so that there still exists a single lowest-cost provider for each variety of intermediate. Hence, domestic and foreign producers do not compete directly in the provision of a particular type of intermediate input.

Within each country, there is assumed to be free trade and no transaction costs for shipping goods. The shipping of an intermediate good across country borders entails transaction costs, which can reflect transport expenses and trade barriers. Specifically, we assume an iceberg technology, whereby, for each unit of intermediate good shipped from country I to country II or the reverse, $1 - b$ units arrive, with $0 < b < 1$. Note that the trading cost, b, reflects the using up of real resources, not a transfer from one party to another (as would be true for a tariff). We assume that transaction costs for shipping final product (and, therefore, consumer goods) abroad are nil.

The production function for a producer of final goods is now modified from (1) to

$$Y_i = AL_i^{1-\alpha}\left(\sum_{j=1}^{N^I+N^{II}} X_{ji}^{\alpha}\right). \tag{17}$$

Hence, for a producer in country I, there are now N^I domestic and N^{II} foreign types of intermediate goods available. Each type of intermediate in country I (and, analogously, in country II) is assumed to feature a single real price, P_j, which applies at the point of origin for domestic and foreign purchasers. Since foreigners receive only $1 - b$ units for each unit bought, their effective price per unit of j-type intermediate good employed in production is $P_j/(1 - b)$. Thus, domestic purchasers of intermediates face markup pricing, whereas foreign purchasers face markup pricing and shipping costs. This price differential will impart a home bias in the demand for intermediate inputs.

The first-order conditions for the choices of intermediate inputs by the producers of final product in country I are now

$$\begin{aligned} A\alpha L_i^{1-\alpha}X_{ji}^{\alpha-1} &= P_j, j = 1, \ldots, N^I, \\ A\alpha L_i^{1-\alpha}X_{ji}^{\alpha-1} &= (\frac{P_j}{1-b}), j = N^I + 1, \ldots, N^I + N^{II}. \end{aligned} \tag{18}$$

The new element is that the price relevant for foreign goods is $P_j/(1 - b)$.

The determination of the markup ratios for each of the intermediate goods produced in country I again follows from Bertrand competition as

$$\begin{aligned} &P_j = \mu_j \text{ if } 1 < \mu_j < 1/\alpha, \\ &P_j = 1/\alpha \text{ if } \mu_j \geq 1/\alpha, \text{ for } j = 1, \ldots, N^I. \end{aligned} \tag{19}$$

As before, we interpret μ_j as reflecting the degree of competition, in the sense of the gap between the production costs of the most efficient and next most efficient providers of intermediates in each sector. If the structure of competition is the same across sectors of country I, the same markup ratio, $\mu_j = \mu^I$, applies in each sector of country I. Extension to the case of heterogeneity across the sectors is straightforward and follows the analysis for the one-country case.

Pricing solutions of the form of (19) also apply to the intermediate goods produced in country II. If the structure of competition is the same across these sectors, the single markup ratio, μ^{II}, applies to all of the sectors in country II.

Substitution of $P_j = \mu^I$ for $j = 1, \ldots, N^I$ and $P_j = \mu^{II}$ for $j = N^I + 1, \ldots, N^I + N^{II}$ into (18) determines the quantities of intermediates employed by firm i in country I:

$$\begin{aligned} X_{ji} &= (A\alpha/\mu^I)^{1/(1-\alpha)}L_i, j = 1, \ldots, N^I, \\ X_{ji} &= [(A\alpha/\mu^{II})(1-b)]^{1/(1-\alpha)}L_i, j = N^I + 1, \ldots, N^I + N^{II}. \end{aligned} \tag{20}$$

Substitution of the results from (20) into (17) and aggregation over the firms determines the level of aggregate output in country I:

$$Y^I = A^{1/(1-\alpha)}\alpha^{\alpha/(1-\alpha)}L^I\left[\left(\frac{1}{\mu^I}\right)^{\alpha/(1-\alpha)}N^I + \left(\frac{1-b}{\mu^{II}}\right)^{\alpha/(1-\alpha)}N^{II}\right], \tag{21}$$

where L^I is the aggregate labour in country *I*. The new element, relative to the closed-economy result in (6), is the term involving the N^{II} foreign types of intermediate goods. These inputs count with the weight $[(1-b)/\mu^{II}]^{\alpha/(1-\alpha)} < 1$, which tends to be less than that for the N^I domestic types because of the shipping cost term, $1-b$. From the perspective of incentives to use the intermediate inputs, markup pricing ($\mu^{II} > 1$) and trading costs ($b > 0$) have similar and reinforcing effects.

The real wage rate in country *I* again equals the marginal product of labour and can be calculated from (17) and (21) as[15]

$$W^I = (1-\alpha)A^{1/(1-\alpha)}\alpha^{\alpha/(1-\alpha)}\left[\left(\frac{1}{\mu^I}\right)^{\alpha/(1-\alpha)} N^I + \left(\frac{1-b}{\mu^{II}}\right)^{\alpha/(1-\alpha)} N^{II}\right]. \tag{22}$$

Labour's share of the country's total gross product, $W^I L^I/Y^I$, is still the constant $1-\alpha$.

The second part of (20) determines the quantity of intermediate goods produced in country *II* and used by final-goods producers in country *I*. The value of these imported goods, gross of shipping costs, is determined by multiplying the quantity of intermediates by $\mu^{II}/(1-b)$. The resulting expression for imports, which is gross of the iceberg losses on the intermediate goods shipped from country *II* to country *I*, is

$$\begin{aligned}&\text{Value of imports of intermediates to country } I\\ &= (A\alpha)^{1/(1-\alpha)}\left(\frac{1-b}{\mu^{II}}\right)^{\alpha/(1-\alpha)} N^{II}L^I.\end{aligned} \tag{23}$$

An expression analogous to the second part of (20) determines the quantity of country *I*'s intermediates used by final-goods producers in country *II*. The corresponding value of the exports of intermediate goods from country *I* to country *II* can be calculated, after multiplication by $\mu^I/(1-b)$, as

$$\begin{aligned}&\text{Value of exports of intermediates from country } I\\ &= (A\alpha)^{1/(1-\alpha)}\left(\frac{1-b}{\mu^{I}}\right)^{\alpha/(1-\alpha)} N^{I}L^{II}.\end{aligned} \tag{24}$$

This expression is gross of the iceberg losses on the intermediate goods shipped from country *I* to country *II*.

Balanced trade in intermediate goods results if

$$\frac{N^I}{\mu^I L^I} = \frac{N^{II}}{\mu^{II}L^{II}}. \tag{25}$$

Alternatively, if the left-hand side of (25) exceeds (or falls short of) the right-hand side, country *I* has a net surplus (or deficit) in the trade in intermediate goods with country *II*. If there are no internationally traded financial assets, as we assume, the net trades of goods and services across country borders must be balanced. Therefore, any net surplus or deficit in the trade of intermediate goods is balanced by an equal-size net deficit or surplus in the trade of final products.

[15] If $b > 0$, the real wage rate in country *II*, W^{II}, generally differs from W^I if $N^I \neq N^{II}$ and $\mu^I \neq \mu^{II}$. We assume here that labour can move freely within a country but cannot move from one country to another. Therefore, a single real wage rate applies within a country, but different rates can apply across countries.

The formulas for imports and exports of intermediate goods resemble gravity-type equations, in the sense of depending on country sizes. However, imports to country I depend on the product $N^{II}L^{I}$, whereas exports from country I depend on the product $N^{I}L^{II}$. Thus, from the standpoint of exports of intermediates, the relevant size variable is the number of varieties, N^{I} or N^{II}, that a country knows how to produce. In a more general context, this concept of size would depend on a country's level of technological advancement and might be proxied by per capita GDP. From the standpoint of imports of intermediates, the relevant size variable is the quantity of labour, L^{I} or L^{II}. This concept relates to a country's level of GDP.

The results in (23) and (24) also resemble familiar gravity models in the sense of predicting that higher trading costs, b, reduce the overall volume of trade. Empirically, the parameter b might relate to distance, other measures of transport costs, the nature of monetary systems, and the extent of similarities in language, legal systems, culture, colonial heritage, and other variables.

In the present model, the degree of monopoly power reduces the volume of trade in a manner similar to that for trading costs. Specifically, a higher markup ratio at home, μ^{I}, reduces exports of intermediate goods, and a higher ratio abroad, μ^{II}, reduces imports of intermediate goods. An increase in markup ratios in both countries lowers the overall volume of trade.

Aggregate consumption in country I now equals the country's output of final goods less its production of intermediates plus the country's net surplus in intermediate trade with country II.[16] The resulting formula is

$$C^{I} = A^{1/(1-\alpha)}\alpha^{\alpha/(1-\alpha)}\times \left\{ \begin{array}{c} (1/\mu^{I})^{1/(1-\alpha)}N^{I}\left[(\mu^{I}-\alpha)L^{I}+\alpha(1-b)^{\alpha/(1-\alpha)}(\mu^{I}-1)L^{II}\right] \\ +(1-\alpha)\left(\frac{1-b}{\mu^{II}}\right)^{\alpha/(1-\alpha)}N^{II}L^{I} \end{array} \right\}. \tag{26}$$

It follows immediately that C^{I} is diminishing in country II's markup ratio, μ^{II}. However, this result assumes that producers of intermediate goods in country II always meet the demand. If, instead, these producers would shut down in the face of losses, $\mu^{II}=1$ would be the preferred markup ratio from the standpoint of country I.

For given μ^{II}, the effect of μ^{I} on C^{I} involves two considerations. First, with respect to home purchases of domestically produced intermediates, we again find that the consumption maximising markup ratio is $\mu^{I}=1$, so that the use of these intermediates would not be distorted. However, for foreign buyers, the maximisation of C^{I} dictates a markup ratio of $\mu^{I}=1/\alpha$, the monopoly value. This result corresponds to the usual monopoly tariff, which applies because the residents of country I do not internalise the benefits of competition for residents of country II. If it were possible for domestic

[16] This equality holds because there is, by assumption, no net borrowing or lending between the two countries. Otherwise, some disturbances – such as a temporary shock to the markup ratio in one country – might motivate net borrowing or lending between the countries. The introduction of these international capital flows would not change any of the main results.

sellers of intermediates to discriminate between domestic and foreign buyers,[17] the consumption maximising choices would be markup ratios of 1 for domestic buyers and $1/\alpha$ for foreign buyers. If a single markup ratio applies to all buyers – as we have been assuming – the consumption maximising choice of $1/\mu^I$ turns out to be a weighted average of 1 and α. The proportionate weights on these values are given, respectively, by L^I and $(1-b)^{\alpha/(1-\alpha)}L^{II}$. Hence, a larger trading cost, b, and a lower ratio of L^{II} to L^I make the monopoly tariff effect less important. The actual markup ratio that the producers of intermediates determine – from (19) – may be higher or lower than the consumption maximising value, depending on the extent of competition that prevails in country *I*.

In a recent paper, Kose and Yi (2001) investigate the role of transport costs in the transmission of business cycles. In their model, lower transport costs lead to higher trade, which tends to increase business-cycle co-movement. However, another force is that lower trading costs intensify a 'resource-shifting channel', whereby capital and other resources move to a country that receives a favourable productivity shock. Kose and Yi show that, without vertical specialisation, their model cannot replicate the empirical finding that higher trade induces higher business-cycle co-movement. In particular, the resource-shifting channel dominates, leading to a negative correlation between trade and business-cycle comovement. Our model allows for the 'back-and-forth' trade called for by these findings.

To our knowledge, there is no evidence that trade in final goods is less costly than trade in intermediates. However, our main conclusions do not depend on this assumption. The model can be easily extended to allow for trading costs in the final goods.

2.2. *Nominal Price Stickiness*

We now introduce nominal price stickiness and, hence, possible roles for each country's monetary policy into the open-economy setting. For country *I*, let p_j again be the nominal price of the *jth* intermediate good and p the nominal price of final goods (and, hence, consumer goods), all of which sell at one price. Country *II* uses a different currency and denominates its prices, p_j^* and p^*, in units of that currency. If all nominal prices were flexible, the preceding analysis would go through, with the relative prices of each intermediate good, p_j/p and p_j^*/p^*, set at the markup level, μ_j. This markup depends on the sectoral competition in country *I* for $j = 1, \ldots, N_I$, and in country *II* for $j = N_I + 1, \ldots, N_I + N_{II}$. We again assume that the degree of competition is the same across sectors within each country.

Suppose that p and p^* are determined through some independent stochastic processes by each country's monetary authority. We assume that the nominal exchange rate, ϵ, is flexible and adjusts so that the standard *PPP* condition holds:

$$\epsilon = p/p^*. \tag{27}$$

[17] An individual producer has an incentive to discriminate only if the elasticities of demand differ across the groups. However, a policy maker who cares about the welfare of the representative domestic individual has the incentive to discriminate even if the elasticities of demand are all the same. A tariff or other levy on international trade could generate this price discrimination.

This condition is consistent with the assumption that final product is homogeneous and internationally tradeable with zero transaction costs.

Assume again that, in country *I*, the nominal price p_j for $j = 1, \ldots, N^I$ is set one period in advance by the lowest-cost producer of each type of intermediate good. We assume that all buyers – domestic and foreign – pay the same price at the point of origin for each intermediate good. That is, we assume that nominal prices are sticky in the units of the producer's currency. (We shall make a parallel assumption about price setting in country *II*.) The nominal price of each of country *I*'s intermediate goods in the nominal currency unit of country *II* is given from (27) by $p_j/\epsilon = p_j(p^*/p)$. Hence, the relative price faced by a buyer in country *II* is given, after division by p^*, as p_j/p, just as in country 1. The quantity demanded of this intermediate good by producers of final product in both countries will again be a constant-elasticity function of this common relative price.

The nominal prices of intermediate goods will now be given in country *I* by

$$p_j \approx \mu^I \mathrm{E}p, \tag{28}$$

for $j = 1, \ldots, N^I$, and in country *II* by

$$p_j^* \approx \mu^{II}\ \mathrm{E}\ p^*, \tag{29}$$

for $j = N^I + 1, \ldots, N^I + N^{II}$. If p exceeds Ep, the relative price p_j/p falls correspondingly below the intended markup level, μ^I. Therefore, the demand for country *I* intermediates by final goods producers rises in both countries. Analogously, an excess of p^* above Ep^* raises the demand for country *II* intermediates in both countries. We assume, for now, that the producers of intermediate goods in each country meet the demands that are forthcoming at these reduced real prices.

From the standpoint of output in country *I*, the parameter μ^I in (21) is replaced by $\mu^I \mathrm{E}p/p$ for the N^I sectors of country *I*. Analogously, the parameter μ^{II} is replaced by $\mu^{II}\mathrm{E}p^*/p^*$ for the N^{II} sectors of country *II*. Therefore, country *I*'s output is now given by

$$Y^I = A^{1/(1-\alpha)}\alpha^{\alpha/(1-\alpha)}L^I \times \left\{[(1/\mu^I)(p/\mathrm{E}p)]^{\alpha/(1-\alpha)}N^I + [(\frac{1-b}{\mu^{II}})(p^*/\mathrm{E}p^*)]^{\alpha/(1-\alpha)}N^{II}\right\}. \tag{30}$$

Hence, unexpected inflation in either country raises output in country *I*. (The results are analogous for country *II*.) The effect from foreign inflation is attenuated by the trading cost term, $1 - b$. The relative strengths of domestic and foreign unexpected inflation on domestic output depend also on the size of the home country, measured by N^I, relative to the size of the foreign country, N^{II}. Because of the distortion from the markup pricing of intermediate goods, unexpected inflation tends to offset the distortion and leads, thereby, to an efficient expansion of output. The outcomes $p/\mathrm{E}p = \mu^I > 1$ and $p^*/\mathrm{E}p^* = \mu^{II} > 1$ would generate the efficient levels of production in both countries.

The real wage rate in country *I* is now

$$W^I = (1-\alpha)A^{1/(1-\alpha)}\alpha^{\alpha/(1-\alpha)} \times \left\{[(1/\mu^I)(p/\mathrm{E}p)]^{\alpha/(1-\alpha)}N^I + \left[\left(\frac{1-b}{\mu^{II}}\right)(p^*/\ \mathrm{E}\ p^*)\right]^{\alpha/(1-\alpha)}N^{II}\right\}. \tag{31}$$

Therefore, unexpected inflation in either country raises the real wage rate in country *I*. The mechanism is that unexpected inflation spurs increased use of intermediates in country *I* – domestic in the case of domestic inflation and foreign in the case of foreign inflation – and, thereby, raises the marginal product of labour. Thus, the real wage rate moves procyclically in response to nominal stimuli. Labour's share of the total gross product, $W^I L^I / Y^I$, is, however, still fixed at $1 - \alpha$.

We can evaluate how different outcomes for unexpected inflation, say $p/\mathrm{E}p$ for country *I*, affect consumption, C^I.[18] The optimal choice of $p/\mathrm{E}p$ (*ex post*) is analogous to the consumption maximising choice of country *I*'s markup ratio, μ^I, as considered in the previous Section. If $\mu^I = 1/\alpha$ (the monopoly value), C^I is increasing in $p/\mathrm{E}p$ when $p/\mathrm{E}p = 1$, but this effect diminishes toward zero as $p/\mathrm{E}p$ approaches a value that lies between 1 and $1/\alpha$.[19] Consumption, C^I, decreases with $p/\mathrm{E}p$ at still higher values of $p/\mathrm{E}p$. If $1 < \mu^I < 1/\alpha$, the range in which unexpected inflation is valued is narrower, and if $\mu^I < 1/\alpha$, the optimal choice of $p/\mathrm{E}p$ may be less than one. Therefore, as in the closed-economy model, the policy maker would value unexpected inflation only over some range. Moreover, because some of the benefit from unexpected inflation now goes to foreigners, the range for which the marginal valuation is positive is narrower in the two-country model than in the one-country setting. In particular, if $\mu^I = 1/\alpha$, it is no longer optimal for $p/\mathrm{E}p$ to be as high as $1/\alpha$.

This model has some surprising implications about how trade and monetary union affect a policy maker's incentive to inflate, *ex post*.[20] If country *I* is a closed economy, the monetary authority would value surprise inflation as long as $p/\mathrm{E}p < \mu^I$. Since $\mu^I > 1$, the preferred inflation surprise is always positive. Moreover, the higher the distortion, μ^I, the greater is the incentive to inflate. Therefore, in a discretionary equilibrium of the type considered by Barro and Gordon (1983), the inflation rate would tend to be higher the higher is μ^I.

If country *I* opens up to trade, say by entering into a trading union with country *II*, the incentive to inflate diminishes because part of the benefit from inflation surprises goes to residents of country *II*. (The assumption is that these foreign benefits are not internalised in some way by the policy maker of country *I*.) This effect is more important the lower is the trading cost, *b*, and the higher is the size of country *II*, as measured by L^{II}. Therefore, in a discretionary equilibrium, the inflation rate would tend to be lower than under autonomy. Moreover, the more open the economy the lower the equilibrium inflation rate. Empirical support for these predictions is provided by Romer (1993), who finds a negative effect of openness on inflation, and Neiss (2001), who finds that countries with higher average markups have higher inflation rates.[21]

Now suppose that the two countries go further by adopting a monetary union. The key assumption here is that the common monetary authority takes into account the

[18] We again omit any international capital markets, so that each country's consumption equals its net income.

[19] By analogy to the results in the previous Section, the consumption maximising value of $p/\mu^I \mathrm{E}p$ is a weighted average of 1 and α, where the proportionate weights are given by L^I and $(1 - b)^{\alpha/(1-\alpha)} L^{II}$, respectively.

[20] We are grateful to Jaume Ventura for these ideas (which we have hopefully not misinterpreted).

[21] Lane (1997) constructs an alternative model in which greater international openness leads to a lower inflation rate. He also provides additional empirical support for this negative relationship.

beneficial effects of surprise inflation in both countries. In this case, the benefits from inflation surprises are similar to those that arose under autonomy. Hence, the equilibrium inflation rate tends to be higher than that without the monetary union.[22] Of course, monetary union can have other effects that favour low inflation, for example, if a client country effectively obtains the policy commitment possessed by the anchor country (Alesina and Barro, 2002). The point here is that the present model identifies one reason why monetary union would be inflationary.

3. A Little Empirical Evidence

In the various versions of the model, the critical variables are markup ratios, measured as prices of specialised, imperfectly competitive products expressed relative to undifferentiated, competitive products. These relative output prices move countercyclically, either because of shifts in the extent of competition or because of nominal disturbances that mimic the effects of changes in competition. In the model, the specialised products were intermediate goods, and the undifferentiated ones were final goods and labour inputs. However, these identifications are not crucial for the general approach. The basic hypothesis is that the relative prices of less competitive goods move countercyclically. Therefore, an important test of the theory is that appropriately measured relative prices move in the hypothesised manner during business cycles.

We begin with some existing empirical evidence that bears on the model's predictions about relative prices. Rotemberg and Woodford (1991, Table 8) find that markups are more negatively correlated with real GNP in sectors with higher four-firm concentration ratios. Thus, if the concentration ratio is a satisfactory measure of imperfect competition, the conclusion is that markups are more countercyclical in less competitive industries. However, their analysis depends on inferring markup ratios from an estimated model of the production function and we are unsure about the proper interpretation of these constructed measures. We have similar misgivings about constructed estimates of markup ratios in the analyses of Bils (1987) and Hall (1988). Therefore, we find it preferable to rely on empirical evidence that uses movements in observed relative prices.

Basu (1995, Table 5) found from US sectoral data in manufacturing that the ratios of prices of materials inputs – defined to include all intermediate goods and services – to wages tended to move countercyclically.[23] This evidence supports our basic model, in which the specialised, imperfectly competitive products were identified with the intermediate inputs. However, Basu's evidence can be viewed as a restatement of the familiar observation that real wages are procyclical. From the perspective of our theory, it would be more interesting to examine the behaviour of prices of intermediate goods expressed relative to prices of final product. Moreover, the Basu analysis relies on the

[22] Similarly, in Rogoff's (1985) paper, cooperation among monetary authorities may raise the benefit attached to inflation surprises and, thereby, increase the equilibrium inflation rate.

[23] From the perspective of the present theory, materials input is a heterogeneous category that includes raw materials, which are likely to be highly competitive, and manufactured goods, which may resemble the specialised intermediate inputs that enter into the model. However, the empirical results for the aggregate of materials input may nevertheless be relevant for the model because, as Basu (1995) observes, 'raw materials and energy are actually only a small fraction of total intermediate inputs. In a modern economy, by far the largest share of these inputs is devoted to purchases of goods manufactured by other firms.'

identification of the less competitive goods with the intermediate inputs, and this constraint is unnecessarily restrictive.

Kraay and Ventura (2002, Table 3) used data for a sample of OECD countries to examine the cyclical behaviour of prices of goods of varying capital intensity. They found that the relative prices of capital-intensive products were countercyclical. This evidence supports the present model if, as seems plausible, more capital-intensive products tend to be more specialised and, hence, less competitive. However, the extent of the association between capital intensity and the degree of competition is unclear.

In the model with nominal rigidities, a key assumption is that the prices of differentiated goods, which are sold under conditions of imperfect competition, are sticky, whereas prices of undifferentiated goods, sold under perfect competition, are flexible. Direct evidence on this assumption is provided by Bils and Klenow (2004), who use the micro data underlying the US consumer price index to examine the frequency of price change for 350 categories of goods and services covering 70% of consumer spending. They find a strong negative correlation between the four-firm concentration ratio (a proxy for imperfect competition) and the frequency of price change. They also find that raw goods, which tend to be more homogeneous and, presumably, more competitive, exhibit a much higher frequency of price adjustment.

To find evidence that bears more directly on our theory, we examined the behaviour of price deflators for industry shipments in US manufacturing at the four-digit level. The data, assembled by the National Bureau of Economic Research, are annual from 1958 to 1997 and cover over 400 industries.[24] We constructed the price ratio p_{jt}/p_t for each sector j and year t, where p_{jt} is the price deflator for shipments from industry j and p_t is the overall GDP deflator.

To see the implications of our theory for this relative output price, consider the expression for output in (10), which applies to a closed economy but allows for a different markup ratio, μ_j, in each sector. This formula works in the context of sticky output prices if we replace μ_j by a modified term $\tilde{\mu}_j$, which is the product of μ_j and the price-surprise term, $(\mathrm{E}p/p)_j$, which we now allow to vary across sectors:[25]

$$\tilde{\mu}_j = \mu_j(\mathrm{E}\ p/p)_j. \tag{32}$$

Suppose, first, that the price-surprise term, $(\mathrm{E}p/p)_j$, is the same for all sectors and that some exogenous change in the extent of competition moves the target markup ratios, μ_j, in the same proportion, η, in all sectors. In this case, the growth rate of aggregate output would be a constant times η. Or, to put things differently, the cyclical pattern in $\log(\tilde{\mu}_j)$ would be the same in all sectors.

Suppose now that the target markup ratios, μ_j, are constant but that less competitive sectors – those with higher μ_j – exhibit more price stickiness, in the sense that the delay in adjustment of p_j to changes in p is greater. In this case, if p rises unexpectedly over

[24] The data originate from the Annual Surveys of Manufactures and the Censuses of Manufactures of the US Bureau of the Census. See Bartelsman and Gray (1996) for a discussion. Updates of the data are compiled by Eric Bartlesman, Randy Becker and Wayne Gray and are available from the website of the National Bureau of Economic Research (NBER.org).

[25] The formula also applies to an open economy, where the terms for the imported intermediate goods involve the product of the foreign target markup ratio, μ_j^*, the trading cost term, $1/(1 - b)$, and the foreign price surprise, $(\mathrm{E}p/p)_j^*$.

some period – which will cause aggregate output to expand temporarily in the model – the less competitive sectors will tend to have lower values of the price-surprise term, $(\mathrm{E}p/p)_j$. Hence, $\tilde{\mu}_j/\mu_j$ will typically be lower the higher is μ_j. Therefore, unexpected inflation will cause $\log(\tilde{\mu}_j)$ to be more counter-cyclical the less competitive the sector (that is, the higher is μ_j).

Now assume that the price surprise is nil, so that $(\mathrm{E}p/p)_j = 1$ in all sectors. Suppose that some real disturbance reduces the extent of competition economy-wide, so that the target markup ratios, μ_j, tend to rise in each sector. The key question for our analysis is how the proportionate changes in μ_j relate to μ_j. In general, we have no presumption about this relationship. However, if some sectors are competitive, so that $\mu_j = 1$ always applies, the presumption is that a general reduction in the extent of competition would have a greater proportionate effect on the less competitive sectors (because the effect on the most competitive sectors is zero). In this case, $\log(\tilde{\mu}_j)$ tends to be more counter-cyclical the less competitive the sector (that is, the higher is μ_j). Therefore, under these conditions, we would expect this cyclical pattern to apply whether the underlying shock was monetary or to the extent of competition.

Returning to the data, we used as an indicator of the business cycle the first difference in annual data of the logarithm of real per capita GDP. The first finding is that relative prices of manufacturing products are countercyclical overall. An OLS regression with 13,845 observations of the first difference of $\log(p_{jt}/p_t)$ on the first difference of the log of real per capita GDP yields an estimated coefficient of -0.108 (s.e. = 0.021).[26] The result that manufacturing relative prices are countercyclical would support the underlying model if manufacturing were generally less competitive than the rest of the economy. However, for our purposes, the more interesting issue is how the cyclical pattern within the manufacturing sector relates to an industry's degree of competition.

We measured an industry's extent of competition by using the Herfindahl-Hirschman index of firm concentration for 1982.[27] The assumption, as in Rotemberg and Woodford (1991), is that greater industrial concentration signals less competition. We then ran OLS regressions in which the data were stratified either into two halves or into deciles in accordance with the Herfindahl indexes. The results are in Table 1.

For the two-way division by the Herfindahl indexes, the estimated coefficient in the lower half is -0.026 (0.029), whereas that in the upper half is -0.189 (0.029). The t-statistic for the difference in these two coefficients is 4.0 and is significant at the 0.00 level. Therefore, we conclude that relative output prices were more countercyclical in the more concentrated manufacturing sectors.[28]

For the decile breakdown, Table 1 shows that the estimated coefficients tend to be more negative as the Herfindahl index rises. However, the pattern is not uniform. A test

[26] This conclusion is similar if we instead relate Hodrick-Prescott filtered values of $\log(p_{jt}/p_t)$ to the H-P filtered value of the log of real per capita GDP. We used a standard smoothing parameter for annual data of 100 to compute all of the filtered values. The OLS estimate of the slope coefficient was then -0.22 (s.e. = 0.05).

[27] These data are available for 1982, 1987, 1992, and 1997. The 1982 figures are in US Bureau of the Census, *US Economic Census*, 1982. Our results do not change greatly if we base the groupings for the Herfindahl indexes on data for the other available years.

[28] Using H-P filtered values, we get coefficients of -0.107 (0.058) and -0.330 (0.075), respectively. The t-statistic here for the difference in coefficients is 2.4.

Table 1

Cyclical Coefficients as a Function of Industrial Concentration

Half of Herfindahl index	Range of log(Herfindahl)	Estimated coefficient (s.e.)	Number of observations
1	1.6–6.1	−0.026 (0.029)	6,903
2	6.1–8.0	−0.189 (0.029)	6,904
Decile for Herfindahl index	Range of log(Herfindahl)	Estimated coefficient (s.e.)	Number of observations
1	1.6–4.6	−0.129 (0.065)	1,365
2	4.6–5.1	−0.029 (0.065)	1,402
3	5.1–5.4	−0.056 (0.066)	1,365
4	5.4–5.8	0.047 (0.065)	1,404
5	5.8–6.1	0.049 (0.067)	1,365
6	6.1–6.4	−0.250 (0.065)	1,404
7	6.4–6.7	−0.208 (0.067)	1,365
8	6.7–7.0	−0.022 (0.065)	1,404
9	7.0–7.3	−0.417 (0.066)	1,365
10	7.3–8.0	−0.072 (0.065)	1,404

Note. The estimated coefficients come from OLS regressions of the first difference of log (p_{jt}/p_t) on the first difference of the log of real per capita GDP, where p_{jt} is the deflator for shipments by industry j and p_t is the GDP deflator. Each regression is run for observations corresponding to the indicated range for the Herfindahl indexes. Separate constant terms, not shown, are included for. each of the ranges. See the text for sources of data.

of the hypothesis that all of these coefficients are equal yields an F-statistic (with 9 and 13,825 degrees of freedom) of 4.8. This result is again significant at the 0.00 level.[29] One possible problem with our procedure is that a measure of industrial concentration, such as the Herfindahl-Hirschman index, need not be an accurate gauge of the extent of competition. Another potential problem is that sectors that exhibit greater concentration may have characteristics aside from less competition that cause their relative prices to be more countercyclical. It may be possible to hold constant some of these other characteristics in an extended analysis.

4. Summary of Major Results

In the basic model, intermediate inputs are specialised and, hence, imperfectly competitive, whereas final product is undifferentiated and, therefore, competitive. An increase in the extent of competition encourages use of the intermediate goods and leads, thereby, to an expansion of output, labour productivity, and consumption. The increase in the marginal product of labour implies a rise in the real wage rate. The likely positive effect of the real wage rate on labour supply generates an expansion of employment. Similar effects from increases in competition would result if some or all of the specialised goods were final products.

[29] For the H-P filtered values, the ten estimated coefficients are −0.14 (0.10), −0.04 (0.12), −0.09 (0.13), −0.14 (0.13), −0.13 (0.15), −0.47 (0.19), −0.40 (0.16), −0.16 (0.17), −0.42 (0.13) and −0.23 (0.18). Hence, the magnitude of the coefficient tends again to rise with the Herfindahl index. However, in this case, the F-statistic for the hypothesis of equal coefficients is only 1.0.

The basic model treats the intermediate goods as non-durables. However, a more realistic identification of the intermediate inputs is with investment goods, such as machine tools and computers. Hence, in a more general setting, the increased investment in these goods would be the key channel that connects the underlying disturbances to the responses of output.

The link with nominal variables and monetary policy arises in the model because the specialised intermediate inputs are assumed to feature relatively sticky nominal prices. Under these conditions, nominal expansion – in the form of an unexpected increase in the price of final product – tends to reduce the relative price of the intermediate goods. Hence, nominal expansion tends to mimic the real effects of an increase in the extent of competition. Specifically, the model predicts increases in output, consumption, labour productivity, the real wage rate, and employment. These effects would arise in a more general model as long as the less competitive goods tended to have more rigid nominal prices.

The extension to an open economy is straightforward if trade in the specialised intermediate inputs involves transaction costs, whereas trade in undifferentiated final product does not entail these costs. The latter assumption implies a standard *PPP* condition for final goods. In this environment, increases in the extent of foreign competition reduce the real cost of foreign-produced intermediate inputs and are, therefore, expansionary at home. Similarly, if the foreign nominal prices of intermediate goods are sticky relative to the price of final product, unexpected inflation abroad tends to lower the real cost of the foreign intermediate goods. Unexpected foreign inflation is, therefore, expansionary for the home country.

The models predict that a sector's relative output price will be more countercyclical the less competitive the sector. This hypothesis was supported by empirical evidence on the cyclical behaviour of prices from four-digit manufacturing industries.

Harvard University
London School of Economics

Submitted: 22 May 2003
Accepted: 21 February 2005

References

Alesina, A. and Barro, R. J. (2002). 'Currency unions', *Quarterly Journal of Economics*, vol. 117 (May), pp. 409–36.

Alesina, A., Spolaore, E. and Wacziarg, R. (2000). 'Economic integration and political disintegration', *American Economic Review*, vol. 90 (September), pp. 1276–96.

Barro, R. J. and Gordon, D. B. (1983). 'Rules, discretion, and reputation in a model of monetary policy', *Journal of Monetary Economics*, vol. 12 (July), pp. 101–21.

Barro, R. J. and Sala-i-Martin, X. (2004). *Economic Growth*, 2nd edn., Cambridge MA: MIT Press.

Bartelsman, E. J. and Gray, W. (1996). 'The NBER manufacturing productivity database', NBER technical working paper no. 205, October.

Basu, S. (1995). 'Intermediate goods and business cycles: implications for productivity and welfare', *American Economic Review*, vol. 85 (June), pp. 512–31.

Basu, S. and Fernald, J. G. (1997). 'Returns to scale in US production: estimates and implications', *Journal of Political Economy*, vol. 105(2), pp. 249–83.

Bils, M. (1987). 'The cyclical behavior of marginal cost and price', *American Economic Review*, vol. 77 (December), pp. 838–57.

Bils, M. and Klenow, P. (2004). 'Some evidence on the importance of sticky prices', *Journal of Political Economy*, vol. 112 (October), pp. 947–85.

Blanchard, O. J. and Kiyotaki, N. (1987). 'Monopolistic competition and the effects of aggregate demand', *American Economic Review*, vol. 77 (September), pp. 647–66.

Broda, C. and Weinstein, D. (2004). 'Are we underestimating the gains from globalization for the United States', unpublished, Federal Reserve Bank of New York.

Calvo, G. (1983). 'Staggered prices in a utility-maximizing framework', *Journal of Monetary Economics*, vol. 12 (September), pp. 383–98.

Carlton, D.W. (1986). 'The rigidity of prices', *American Economic Review*, vol. 76 (September), pp. 637–58.

Chari, V. V., Kehoe, P. J. and McGrattan, E. R. (2002). 'Can sticky price models generate volatile and persistent exchange rates?', *Review of Economic Studies*, vol. 69 (August), pp. 533–63.

Dixit, A. and Stiglitz, J. (1977). 'Monopolistic competition and optimum product diversity', *American Economic Review*, vol. 67 (June), pp. 297–308.

Ethier, W. J. (1982). 'National and international returns to scale in the modern theory of international trade', *American Economic Review*, vol. 72 (June), pp. 389–405.

Hall, R. E. (1988). 'The relation between price and marginal cost in US industry', *Journal of Political Economy*, vol. 96 (October), pp. 921–47.

Helpman, E. and Krugman, P. R. (1985). *Market Structure and Foreign Trade*, Cambridge MA: MIT Press.

Ireland, P. (1996). 'The role of countercyclical monetary policy', *Journal of Political Economy*, vol. 104 (August), pp. 704–23.

Kose, A. and Yi, K. (2001). 'International trade and business cycles: is vertical specialization the missing link?', *American Economic Review, Papers and Proceedings*, vol. 91 (May), pp. 371–5.

Kraay, A. and Ventura, J. (2002). 'Product prices and the OECD cycle', *Advances in Macroeconomics, The B.E Journals in Macroeconomics*, vol. 2(1), pp. 1–16.

Krugman, P. (1980). 'Scale economies, product differentiation, and the pattern of trade', *American Economic Review*, vol. 70 (December), pp. 950–9.

Kydland, F. E. and Prescott, E. C. (1977). 'Rules rather than discretion: the inconsistency of optimal plans', *Journal of Political Economy*, vol. 85 (June), pp. 473–92.

Lane, P. (1997). 'Inflation in open economies', *Journal of International Economics*, vol. 42 (May), pp. 327–47.

Mankiw, N. G. (1991). 'A general equilibrium example', in (N. G. Mankiw and D. Romer, eds.), *New Keynesian Economics*, pp. 36–40, Cambridge MA: MIT Press.

Neiss, K. (2001). 'The markup and inflation: evidence in OECD countries', *Canadian Journal of Economics*, vol. 34 (May), pp. 570–87.

Obstfeld, M. and Rogoff, K. (1995). 'Exchange rate dynamics redux', *Journal of Political Economy*, vol. 103 (June), pp. 624–60.

Rogoff, K. (1985). 'Can international monetary policy cooperation be counterproductive?' *Journal of International Economics*, vol. 18 (May), pp. 199–217.

Romer, D. (1993). 'Openness and inflation: theory and evidence', *Quarterly Journal of Economics*, vol. 108 (November), pp. 870–903.

Romer, P. M. (1990). 'Endogenous technological change', *Journal of Political Economy*, vol. 98 (October), part II, pp. S71–102.

Rotemberg, J. and Saloner, G. (1986). 'A supergame-theoretic model of price wars during booms', *American Economic Review*, vol. 76 (June), pp. 390–407.

Rotemberg, J. and Saloner, G. (1987). 'The relative rigidity of monopoly pricing', *American Economic Review*, vol. 77 (December), pp. 917–26.

Rotemberg, J. and Woodford, M. (1991). 'Markups and the business cycle', *NBER Macroeconomics Annual* vol. 1991, pp. 63–129.

Rotemberg, J. and Woodford, M. (1999). 'The cyclical behavior of prices and costs', in (J. Taylor and M. Woodford, eds.), *Handbook of Macroeconomics*, pp. 297–346, Amsterdam: North-Holland Elsevier.

Spence, A. M. (1976). 'Product selection, fixed costs, and monopolistic competition', *Review of Economic Studies*, vol. 43 (June), pp. 217–35.

Stigler, G. J. and Kindahl, J. K. (1970). *The Behavior of Industrial Prices*, New York: National Bureau of Economic Research.

Svensson, L. (1986). 'Sticky good prices, flexible asset prices, monopolistic competition, and monetary policy', *Review of Economic Studies*, vol. 53 (July), pp. 385–405.