Supply Shocks And the Distribution Of Price Changes

Nathan S. Balke
Associate Professor
Department of Economics
Southern Methodist University

Mark A. Wynne
Senior Economist and Policy Advisor
Federal Reserve Bank of Dallas

This article explores in some detail the dynamics of relative price changes in a simple dynamic general equilibrium model.

In a recent paper, Ball and Mankiw (1995) propose a new measure of supply shocks. Specifically, they advocate a measure of the skewness of price changes across sectors as a superior alternative to existing measures of supply shocks, such as the relative price of oil. Ball and Mankiw begin by showing that various measures of the skewness of the distribution of relative price changes across industries in the producer price index (PPI) are positively correlated with the rate of increase in the overall PPI. They further argue that these measures of skewness are better measures of supply shocks than more traditional measures such as the relative prices of food and energy when used in simple, short-term Phillips curve-type relationships. Their interpretation of the relationship between skewness and aggregate inflation relies heavily on the existence of menu costs associated with changing prices at the firm level. They further argue that since menu cost models were designed to explain monetary nonneutrality, these models "…gain [scientific] credibility from their ability to fit the facts regarding inflation and relative-price changes."

This article builds on the analysis of Ball and Mankiw by exploring in some detail the dynamics of relative price changes in a simple dynamic general equilibrium model. We begin by providing further evidence of a robust statistical relationship between the skewness of the distribution of individual price changes and inflation. We then ask what sort of relationship would we expect to see in a model in which all prices are free to adjust instantaneously. We show that when a simple general equilibrium model with no menu costs is calibrated to match certain features of the real world, it is possible to find a significant relationship between the skewness of individual price changes and aggregate inflation. Thus, our results cast some doubt on Ball and Mankiw's interpretation of the correlation between the skewness of the distribution of price changes and aggregate inflation as supportive of menu cost models.

Relative price changes as aggregate supply shocks

Ball and Mankiw begin their analysis by discussing a simple model in which menu costs associated with changing prices cause firms to adjust nominal prices only in response to large relative price shocks. The existence of menu costs implies the existence of range of inaction over which firms do nothing to change their prices in response to shocks. In Ball and Mankiw's model, the adjustment to large but not to small shocks results in a positive correlation between the skewness and the mean of the (cross section) distribution of price changes. Ball and Mankiw argue that a flexible price model would predict no such relationship. In the flexible price model, a large positive relative price shock is likely to be offset by small declines in the prices of other commodities. Ball and Mankiw examine the cross section distribution of several hundred prices and indeed find a positive correlation between skewness and the mean of the distribution. They interpret this as evidence in favor of the menu cost model.

Yet, as we show below, it is possible for a flexible price model also to generate this positive correlation between skewness and the mean of the distribution of price changes. Ball and Mankiw's assertion that a flexible price model cannot generate this correlation is probably correct for the case in which a large number of sectors are experiencing shocks that are independent of one another and are of the same relative magnitude. This situation would cause relative price changes for the different commodities to be relatively independent of each other also. Yet, in reality, prices across commodities are not independent; shocks in one sector tend to affect prices in other sectors. Furthermore, a few very volatile sectors, such as food and energy, may be responsible for most of the observed volatility in the distribution of price changes. As a result, the correlation between the average inflation rate and the skewness of the distribution of price changes found in the data may arise just because of the importance of a few large price shocks.

Below we consider the implications for the data of a modified version of the general equilibrium model due to Long and Plosser (1983). The model is modified slightly to include a numeraire role for money. The model has complete price flexibility and multiple sectors. Among the key characteristics of this model is that a shock in one sector can spill over to other sectors. We show that as the sectors become more interrelated, it becomes easier for the flexible price model to generate a positive correlation between the skewness of the distribution of price changes and aggregate inflation.

The data

Ball and Mankiw look at the relationship between the distribution of prices in the producer price index (PPI) on an annual basis over the period 1949–89. The advantage of looking at the PPI is that it is available at a high degree of disaggregation. At the four-digit level of disaggregation, the number of component series rises from 213 in 1949 to 343 in 1989. Ball and Mankiw then document the relationship between the distribution of the changes in these several hundred price series and the overall inflation rate (as measured by the PPI).

Their data analysis reveals a number of interesting facts. First, there is considerable variation in the distribution of price changes over time. For example, in 1987 the distribution is fairly symmetric, while in 1973 it is skewed sharply to the right and in 1986 it is skewed sharply to the left.[1] Not surprisingly, both 1973 and 1986 were also years in which there were significant oil price shocks, with oil prices rising dramatically in 1973 and falling dramatically in 1986.

Ball and Mankiw document a statistically significant relationship between various measures of skewness and the overall inflation rate.[2] They also show that the skewness of the distribution of price changes tends to dominate the standard deviation of the distribution as an explanatory variable for inflation. This result is robust to their use of any of three measures of skewness.

In the analysis presented below we will examine the relationship between the distribution of price changes and aggregate inflation in the context of a multisectoral model that is calibrated to match certain characteristics of the U.S. economy. Considerations of tractability prevent us from considering a model with more than a small number of sectors. In fact, we work with a version of the real business cycle model proposed by Long and Plosser (1983) that has only six sectors. Before proceeding, then, it is important to verify that the empirical regularities observed by Ball and Mankiw in the prices that make up the PPI are also present when we consider more aggregated measures of prices.

The six sectors Long and Plosser use to calibrate their model are agriculture, mining, construction, manufacturing, transportation and trade, and services and miscellaneous. Table 1 presents some summary statistics for inflation rates (as measured by the implicit gross domestic product (GDP) deflators for these sectors) over the period 1949–93. The table reveals a number of interesting facts about the time series behavior of sectoral inflation rates.[3] First, there are notable differences in the average rates of inflation across the six sectors over the sample period, ranging from a low of just under 2 percent per year in agriculture to a

Table 1
Statistics on Inflation Rates by Sector, 1948–93
(Annual data)

	Mean	Variance
Agriculture	1.99	114.34
Mining	3.28	8.39
Construction	4.01	210.74
Manufacturing	5.23	21.53
Transportation and trade	3.41	7.02
Services and miscellaneous	4.94	4.45

high of nearly 5.25 percent per year in manufacturing. Second, there are dramatic differences in the variances of the individual inflation rates across sectors, from a low of 4.45 in services and miscellaneous to a high of 210.74 in construction.

Table 2 presents some simple regression results for the relationship between the rate of inflation as measured by the fixed-weight GDP deflator and measures of the distribution of prices across six sectors of the U.S. economy. The first column shows the results of regressing the rate of inflation on its own lagged value, while the second column shows the results of adding an unweighted measure of skewness to this basic regression. Comparing columns 1 and 2, we see that skewness has significant explanatory power for the inflation rate: the $\bar{R}^2$ increases from 0.52 to 0.66, and all of the coefficients in the second regression are significant at the 1-percent level. Column 3 shows what happens if we use a weighted measure of skewness instead. We obtain an even higher $\bar{R}^2$ and again all coefficient estimates are significant at the 1-percent level. The results in this table compare favorably with the results reported in Tables IIIA and IIIB of Ball and Mankiw. In fact, working with our more aggregated price data, we find an even stronger statistical relationship (in the sense of higher $\bar{R}^2$) between the skewness of the distribution of price changes and the aggregate inflation rate.

To summarize, it is clear that we can obtain the same strong statistical relationship between the skewness of the distribution of price changes and aggregate inflation looking at only six prices as Ball and Mankiw do looking at several hundred prices. The relationship uncovered by Ball and Mankiw seems robust, and our decision to focus on just six sectors does not seem to be a gross violation of the spirit of their analysis. Our objective in what follows is to see to what extent we can replicate the facts about the relationship between skewness and inflation as documented here in the context of a simple dynamic general equilibrium model.

A simple dynamic general equilibrium model with multiple sectors

A logical starting point for an investigation of the relationship between the distribution of price changes across sectors and aggregate inflation is the equilibrium business cycle model of Long and Plosser (1983). A great virtue of this model is that it has multiple sectors, but more importantly, the calculation of decision rules is simplified because of restrictions on preferences and the rate of depreciation of capital. The original version of this model was a "real"model in every sense of the word, in that there was no role for money.

For our purposes we would like to extend the model to include money as a numeraire. Benassy (1995) has recently proposed a version of the Long and Plosser model that incorporates money by including real balances in the utility function. While introducing money into the model in this way is not entirely satisfactory, it is well-known that this specification is functionally equivalent in certain circumstances to the more popular cash-in-advance and shopping-time formulations of the demand for real balances. However, Benassy works with a single-sector variant of the Long and Plosser model, and it is far from straightforward to extend his analysis to a multiple-sector setting (the essence of the problem that arises in this regard is the absence of a single correct measure of the price level in a multisector environment). An alternative is to introduce money via some sort of cash-in-advance constraint on either purchases of consumption goods (or some subset thereof) or purchases of capital goods, or both. However, it rapidly becomes apparent that it is no longer possible to calculate simple closed form decision rules in either of these cases.

We opt instead to introduce money in a somewhat novel manner. Specifically, we assume that consumers are obliged to hold some fraction υ of their consumption purchases during each period in the form of cash at the end of the period. Thus, we posit the following constraint on household choices:

$$M_t \geq \upsilon \sum_{i=1}^{N} P_{i,t} C_{i,t}, \tag{1}$$

where M_t denotes the stock of nominal money balances held at the *end* of period t and $\sum_{i=1}^{N} P_{i,t} C_{i,t}$ denotes nominal consumption expenditures *during* period t, with $P_{i,t}$ denoting the price of good i at date t, and $C_{i,t}$ denoting the quantity of good i purchased for consumption purposes at date t. The existence of this constraint can be thought of as arising due to the need to, say, maintain some minimum level of cash balances in a bank account to facilitate consumption purchases made with inside money. It will turn out that money does not play a very important role in our economy.

The rest of the model is relatively standard.

Households. The economy is populated by a large number of identical consumers, each of whom has preferences summarized by the following utility function:

$$U = \sum_{i=0}^{\infty} \beta^t u(C_t, L_t), \tag{2}$$

where $1 > \beta > 0$ is the discount factor, $C_t = (C_{1,t}, C_{2,t}, \ldots, C_{N,t})'$ is an $N \times 1$ vector of commodities consumed at date t, and L_t denotes leisure at date t. The point-in-time utility function is furthermore assumed to have the following specific functional form:

$$u(C_t, L_t) = \theta_0 \log(L_t) + \sum_{i=1}^{N} \theta_i \log(C_{i,t}), \tag{3}$$

where $\theta_i \geq 0$, $\forall$ i. If $\theta_i = 0$ for some $i \geq 1$ then that commodity has no utility value to the consumer.

The budget constraint of the representative consumer is given by

$$\begin{aligned} &\sum_{i=1}^{N} W_{i,t} H_{i,t} + \sum_{i=1}^{N} \sum_{j=1}^{N} R_{i,j,t} K_{i,j,t-1} + \mu_t M_{t-1} \\ &\geq \sum_{i=1}^{N} P_{i,t} C_{i,t} + \sum_{j=1}^{N} P_{j,t} \sum_{i=1}^{N} K_{i,j,t} + M_t, \end{aligned} \tag{4}$$

where $W_{i,t}$ denotes the (nominal) wage in sector i at date t (which in equilibrium will be the same in all sectors), $H_{i,t}$ denotes hours worked in sector i at date t, $R_{i,j,t}$ denotes the nominal rental rate on type j capital employed in sector i in period t, $K_{i,j,t-1}$ represents the quantity of type j capital employed in sector i during period t, (that is, capital in place at the end of period $t-1$) and μ_t represents the gross rate of increase in the money stock at date t. The sources of funds each period are wage income $\sum_{i=1}^{N} W_{i,t} H_{i,t}$, income from capital $\sum_{i=1}^{N} \sum_{j=1}^{N} R_{i,j,t} K_{i,j,t-1}$ and a transfer from the government that is directly proportional to nominal money holdings held at the end of the previous period, $\mu_t M_{t-1}$. The uses of funds each period are consumption expenditures, $\sum_{i=1}^{N} P_{i,t} C_{i,t}$, purchases of new capital equipment, $\sum_{j=1}^{N} P_{j,t} \sum_{i=1}^{N} K_{i,j,t}$, and funds

Table 2
Inflation and the distribution of price changes, 1948–93

	(1)	(2)	(3)
Constant	.010*	.013**	.014**
	(.005)	(.004)	(.004)
Lagged inflation	.715**	.642**	.609**
	(.102)	(.088)	(.080)
Unweighted skewness		.011**	
		(.003)	
Weighted skewness			.004**
			(.001)
$\bar{R}^2$	.52	.66	.72
Durbin–Watson statistic	1.54	1.78	2.02

*,** denotes significance at the 5-percent and 1-percent levels, respectively.

NOTE: Standard errors are in parentheses.

held over to the next period, M_t.

The remaining constraint that the consumer faces is on the allocation of available time,

$$L_t + \sum_{i=1}^{N} H_{i,t} = 1, \tag{5}$$

which states that the sum of leisure and time worked in each sector cannot exceed the total amount of time available, which we normalize to 1.

The consumer's problem is to maximize the objective function given in equation 2 above subject to the budget constraint (equation 4), the cash constraint (equation 1), and the constraint on the allocation of time (equation 5), taking as given the prices at which he or she can purchase consumption and capital goods and the wage and rental rates at which labor and capital services are sold to the business sector.

Firms. Production possibilities in the i'th sector are given by the following production function:

$$Y_{i,t} = Z_{i,t} H_{i,t}^{b_i} \prod_{j=1}^{N} K_{i,j,t-1}^{a_{i,j}}, \tag{6}$$

where $Y_{i,t}$ denotes output of the i'th good at date t, $Z_{i,t}$ is a random variable or productivity shock that denotes the state of technology in the i'th sector at date t, $H_{i,t}$ denotes hours worked in the i'th sector at date t, and $K_{i,j,t}$ denotes the quantity of output of the j'th industry employed in the i'th industry at date t. The parameters of the production function, b_i and $a_{i,j}$ are assumed to satisfy $b_i > 0$, $a_{i,j} > 0$ and $b_i + \sum_{j=1}^{N} a_{i,j} = 1$ for i = 1, 2, …, N. The production side of this model can be viewed in two ways. We can think of each sector as producing both consumption and capital goods that are used in every sector, with the capital depreciating at a 100-percent rate. Or, we can think of each sector as producing consumption goods and goods that are used as intermediate inputs in the production of other goods. The two interpretations are equivalent.

The firm's optimization problem is to maximize profits, taking as given the available technology, the price at which output can be sold and the prices or rental rates of the labor and capital inputs.

Equilibrium. Straightforward manipulation of the first order conditions for the household and firm maximization problems allows us to obtain the following decision rules:

$$(7) \qquad C_{i,t} = \left(\frac{\theta_i}{\gamma_i}\right) Y_{i,t},$$

$$(8) \qquad L_t = \frac{\theta_0(1+\nu(1-\beta))}{\theta_0(1+\nu(1-\beta)) + \sum_{i=1}^{N} \gamma_i b_i},$$

$$(9) \qquad H_{i,t} = \frac{\gamma_i b_i}{\theta_0(1+\nu(1-\beta)) + \sum_{j=1}^{N} \gamma_j b_j}, \text{ and}$$

$$(10) \qquad K_{i,j,t} = \left(\frac{\beta \gamma_i a_{i,j}}{\gamma_j}\right) Y_{j,t},$$

where $\gamma_j = \theta_j + \beta \sum_{i=1}^{N} \gamma_i a_{i,j}$. The derivation of these rules is explained in more detail in the Appendix.

Some comments are in order. As noted, the simple form of these decision rules stems from the particular assumptions we have made about preferences, production possibilities and the durability of capital. Equation 7 shows that consumption of each type of good is simply a constant fraction (θ_i/γ_i) of the available output of that type of good, with the fraction being a complicated function of the parameters of the underlying preferences and technology. Likewise, equation 10 shows that the amount of each sector's output that is allocated for use in production in other sectors is a constant fraction ($\beta\gamma_i a_{i,j}/\gamma_j$) of the available output. Perhaps more surprisingly, equations 8 and 9 show that total leisure and the allocation of effort to each sector are independent of realizations of the productivity shocks and are also independent of the endogenously chosen "state" of each sector, as summarized by the level of output produced in each sector.

To better understand why the allocation of labor across sectors (and total labor or leisure) is independent of the current state of the economy, consider the condition determining the equilibrium allocation of labor.[4] This condition states that the value of the marginal product of labor in each of its alternative employments should equal the wage rate, where all prices and wages are denominated in utility units. The wage rate is simply the marginal utility of leisure. The question then is, Given an initially optimal allocation of labor across sectors, would a change in either the available capital stock or the state of technology change either side of this equation? Given the specification of preferences we are working with, the marginal utility of leisure at any point in time depends only on the labor–leisure allocation at that time, so any effect of the capital stock or technology on the optimal allocation must come about through changes in the value of the marginal product. Consider the effect of a higher than expected realization of the technology. For a given allocation of capital to a particular sector, one effect of the technology shock would be to raise the marginal physical product of labor. However, working against this, the technology shock will put downward pressure on the price of the sector's output, lowering the value of the marginal product of labor. It just so happens in this case that these two effects offset each other, leaving the value of the marginal product unchanged. In other words, the optimal allocation of labor to the sector is unaffected by realizations of the technology shock. Similar reasoning applies to determining the effects of greater availability of capital in a sector.

We can use the equations above to write dollar-denominated prices in our extended model as

$$(11) \qquad P_{i,t} = \frac{1}{\nu} \frac{\gamma_i}{\sum_{j=1}^{N} \theta_j} \frac{M_t}{Y_{i,t}}.$$

That is, nominal prices are directly proportional to the nominal money stock.[5] It is straightforward to show that these prices are also directly proportional to the utility-denominated prices calculated by Long and Plosser.

Dynamics. The dynamic behavior of this economy is implied by the technology as summarized by the production functions, along with the decision rules for the inputs to the production processes. It is convenient to write the system in logarithmic form as follows,

$$(12) \qquad y_t = k_y + Ay_{t-1} + z_t,$$

where we adopt the convention that lower case letters denote the logarithms of the corresponding upper case variable. Thus, y_t is the $N \times 1$ vector $(\log(Y_{1,t}), \log(Y_{2,t}), \ldots, \log(Y_{N,t}))'$, k_y is an $N \times 1$ vector of constants, and z_t is the $N \times 1$ stochastic vector $(\log(Z_{1,t}), \log(Z_{2,t}), \ldots, \log(Z_{N,t}))'$. Since our primary focus in this article is on the evolution of the distribution of prices, we will also need to specify a stochastic process for the log of the nominal money stock, m_t.

The evolution of prices is given by

$$(13) \qquad p_t = k_p + \iota_N m_t + y_t,$$

where $p_t = (\log(P_{1,t}), \log(P_{2,t}), \ldots, \log(P_{N,t}))'$, k_p is a vector of constants and ι_N is an $N \times 1$ vector of ones. An important point to note from this expression is that the money stock only affects the mean of the distribution of prices across sectors and not any of the higher moments (such as the standard deviation or skewness).

Measures of the aggregate price level. We can easily calculate a variety of measures of the aggregate price level that correspond to the measures commonly used to gauge inflation in the real world. Three such measures are defined in the Appendix. We will concentrate on just one of them, a fixed-weight measure of the aggregate price level that corresponds to the fixed-weight GDP deflator. We construct a fixed-weight GDP deflator starting from the definition

$$(14) \qquad PGDPF_t = \frac{\sum_{i=1}^{N} P_{i,t} Y_{i,b}}{\sum_{i=1}^{N} P_{i,b} Y_{i,b}}.$$

That is, the value of the fixed-weight GDP deflator at date t, $PGDPF_t$, equals the ratio of the cost of purchasing the base-year market basket of final output at current-year prices, $\sum_{i=1}^{N} P_{i,t} Y_{i,b}$, to the cost of purchasing the base-year market basket of final output at base-year prices, $\sum_{i=1}^{N} P_{i,b} Y_{i,b}$. Making substitutions using equation 11 above, we obtain

$$(15) \qquad PGDPF_t = \frac{M_t}{M_b} \frac{1}{\sum_{i=1}^{N} \gamma_i} \sum_{i=1}^{N} \frac{\gamma_i Y_{i,b}}{Y_{i,t}},$$

which can also be rewritten in logs as

$$(16) \qquad pgdpf_t \equiv \log(PGDPF_t) = m_t - \overline{m} - \log\left(\sum_{i=1}^{N} \gamma_i\right) + \log\left(\sum_{i=1}^{N} \gamma_i \exp(\overline{y}_i - y_{i,t})\right),$$

where bars over variables denote steady-state values, which we pick as the base year.

Calibration. Long and Plosser calibrate their model to six sectors of the U.S. economy using a consolidated version of the twenty-three sector input–output table for 1967 (U.S. Department of Commerce, 1975). This yields an estimate of the A matrix for the model above. Given an estimate of the A matrix, the vector of coefficients b is recovered from the assumption of constant returns to scale, *i.e.,* $b_i = 1 - \sum_{j=1}^{N} a_{i,j}$. To calibrate the vector θ, we note that the decision rules for consumption of each type of good imply that $\theta_i = \gamma_i \frac{C_{i,t}}{Y_{i,t}}$. We can obtain estimates of the share (γ_i) of each sector's output in aggregate output from the 1967 input–output table. The same table also allows us to estimate the fraction of each sector's output that was allocated to consumption that year, which together with the estimate of γ_i allows us to obtain an estimate of θ_i. Finally, we set the discount factor β equal to 0.95 and the parameter ν equal to 1.

Experiment 1. Our first experiment examines the behavior of inflation and the distribution of prices in an economy with six sectors but with no input–output relations between the sectors and with each sector subject to *i.i.d.* shocks of equal variance. Thus, we set

$$A = \begin{pmatrix} 0.33 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0.33 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0.33 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0.33 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0.33 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0.33 \end{pmatrix}$$

and $\theta = (0.167, 0.167, 0.167, 0.167, 0.167, 0.167)'$. We assume that the technology innovations hitting each sector are *i.i.d.* with zero mean and unit variance. A priori, we expect that there will be no relationship between measures of the cross-section distribution of prices and aggregate inflation in this economy. We think that Ball and Mankiw have an economy such as this in mind when they question the ability of a flexible price model to generate the correlations between skewness and inflation that are found in the data.

Experiment 2. For our second experiment, we calibrate the A matrix and the θ vector along the same lines as Long and Plosser, but retain the assumption of *i.i.d.* shocks. Thus, we set

$$A = \begin{pmatrix} 0.4471 & 0.0033 & 0.0146 & 0.2093 & 0.0999 & 0.1591 \\ 0.0000 & 0.0935 & 0.0427 & 0.1744 & 0.0549 & 0.4854 \\ 0.0029 & 0.0104 & 0.0003 & 0.4189 & 0.1209 & 0.0893 \\ 0.0618 & 0.0340 & 0.0050 & 0.4576 & 0.0611 & 0.1267 \\ 0.0017 & 0.0004 & 0.0166 & 0.1246 & 0.1040 & 0.3249 \\ 0.0174 & 0.0212 & 0.0595 & 0.1998 & 0.0871 & 0.3805 \end{pmatrix}$$

and $\theta = (0.003465, 0.000162, 0.022046, 0.139968, 0.089811, 0.15012)'$. The economy of this experiment differs from that of the first experiment mainly in allowing for complicated input–output type relations between the various sectors.

Experiment 3. For our third experiment, we retain the specifications of the A matrix and θ vector used in the second experiment but calibrate the technology shocks to the actual postwar data. Thus, we estimate Solow residuals for each sector as

$$\begin{aligned} z_{i,t} \equiv \log(Z_{i,t}) &= \log(y_{i,t}) \\ &\quad - (1-\alpha_i)\log(k_{i,t-1}) - \alpha_i \log(n_{i,t}) \end{aligned}$$

where $\log(y_{i,t})$ is the BP-filtered log of output in sector i produced during period t, $\log(n_{i,t})$ is the BP-filtered log of full time equivalent employees in the i'th sector during period t, $\log(k_{i,t-1})$ is the BP-filtered log of the net (real) capital stock in the i'th sector as of the *end* of period $t - 1$ (*i.e.,* capital available in the i'th sector at the *beginning* of period t), and α_i is the average value over the sample period (1947–94) of labor's share in the i'th sector (defined as the ratio of nominal compensation of employees in the i'th sector to nominal GDP in that sector). For the BP filter, we used the parameter values $up = 2$, $dn = 8$ and $K = 3$.[6]

We assume that total factor productivity in the model evolves according to

$$z_t = Pz_{t-1} + \epsilon_t$$

where we use OLS to estimate the matrix P as

$$P = \begin{pmatrix} 0.109 & 0.0 & 0.0 & 0.0 & 0.0 & 0.0 \\ 0.0 & 0.024 & 0.0 & 0.0 & 0.0 & 0.0 \\ 0.0 & 0.0 & 0.458 & 0.0 & 0.0 & 0.0 \\ 0.0 & 0.0 & 0.0 & 0.127 & 0.0 & 0.0 \\ 0.0 & 0.0 & 0.0 & 0.0 & 0.106 & 0.0 \\ 0.0 & 0.0 & 0.0 & 0.0 & 0.0 & 0.329 \end{pmatrix}$$

The innovations $\epsilon_t = (\epsilon_{1,t}, \epsilon_{2,t}, \epsilon_{3,t}, \epsilon_{4,t}, \epsilon_{5,t}, \epsilon_{6,t})'$ are assumed to be *i.i.d.* with variances $\sigma^2 = (1.01, 1.26, 0.25, 0.54, 0.19, 0.03) \times 10^{-3}$.

Experiment 4. For our fourth experiment, we estimate a simple VAR for the technology innovations in each sector. Again, we assume total factor productivity follows the process

$$z_t = Pz_{t-1} + \epsilon_t$$

where now the matrix P is given by

$$P = \begin{pmatrix} 0.231 & 0.285 & 0.205 & 1.031 & 0.872 & 1.831 \\ 0.006 & 0.116 & 0.948 & 0.468 & 0.344 & 2.334 \\ 0.044 & 0.031 & 0.515 & 0.458 & 0.438 & 0.022 \\ 0.188 & 0.127 & 0.438 & 0.074 & 0.188 & 0.230 \\ 0.083 & 0.068 & 0.348 & 0.075 & 0.020 & 0.092 \\ 0.065 & 0.004 & 0.075 & 0.066 & 0.166 & 0.079 \end{pmatrix}$$

and again the innovations ϵ_t are assumed to be *i.i.d.* with variances $\sigma^2 = (0.75, 0.87, 0.20, 0.41, 0.13, 0.02) \times 10^{-3}$.

Each of these experiments introduces progressively more interaction between the sectors and allows for greater diversity in the shocks hitting the sectors. In the first experiment, there is no interaction and the shocks hitting each sector are completely independent of one another. The second experiment allows for interaction through input–output relationships but retains the assumption of independent shocks. The third experiment allows for interaction through input–output relationships and allows for serially correlated shocks in each sector. The fourth experiment allows for input–output type interaction between sectors and for serial correlation in the state of technology across sectors.

One final comment on the experiments. In each of the four experiments reported here, we hold the money stock constant, so that the only source of fluctuations in the model economies are technological innovations. This technique allows us to completely isolate the effects of what Ball and Mankiw call "supply shocks" on the relationship between the distribution of price changes and aggregate inflation.

Results

For each experiment, we calibrate the model as described above and simulate it for fifty periods one hundred times. We use the time series on prices generated in each of the one hundred simulations to run the regressions described in Table 2. Table 3 reports the average value over all one hundred simulations of the regression coefficients on the weighted and unweighted measures of the skewness of the distribution of prices at each date for an inflation regression in which we use the rate of change of a fixed-weight measure of the GDP deflator as the measure of inflation. For comparison, we also report the coefficients estimated using actual data.

Moving down the rows of Table 3, we see that in the basic economy with no interaction between sectors and *i.i.d.* shocks of equal variance hitting each sector, we are unable to generate a significant role for skewness in explaining the inflation rate. Note that the coefficient estimates are the same for the weighted and unweighted measures of skewness as all sectors are identical by construction. Moving to the economy of experiment 2, there is still no role for unweighted skewness in explaining inflation, but the weighted measure is now significant. Recall that this economy differs from that in experiment 1 only in that it allows for input–output type relationships between all of the sectors. For the economies of experiment 3 and experiment 4, the weighted measure of skewness helps explain inflation. However, in none of our experiments is the unweighted measure correlated with inflation in a statistically significant sense.

Summarizing our results, it is clear that we are able to replicate to a surprising degree the key correlation between skewness and inflation Ball and Mankiw find in the data. Most importantly, we are able to do so in the context of a simple equilibrium model with fully flexible prices, thus raising questions about Ball and Mankiw's interpretation that this correlation results from the existence of menu costs. In Balke and Wynne (1995), we also document other aspects of the relationship between the distribution of price changes and the overall inflation rate and show that skewness seems to be a leading indicator of aggregate inflation. Our model has less success in replicating this feature of the data.

Table 3
Estimated coefficients on skewness in inflation regression

	Unweighted skewness	Weighted skewness
Data	.011**	.004**
	(.003)	(.001)
Experiment 1	.038	.038
	(.182)	(.182)
Experiment 2	–.037	.297**
	(.222)	(.102)
Experiment 3	.000	.005**
	(.003)	(.002)
Experiment 4	.001	.007**
	(.004)	(.003)

*,** denotes significance at the 5-percent and 1-percent levels, respectively.

NOTE: Standard errors are in parentheses.

SOURCE: Authors' calculations.

Conclusions

In this article, we explore the relationship between shifts in the distribution of prices and the aggregate inflation rate in the context of a simple dynamic general equilibrium model with multiple sectors. The idea that changes in the distribution of relative price changes might have implications for the overall inflation rate was first proposed by Ball and Mankiw (1995). A crucial part of the story that they tell is that firms face significant adjustment costs associated with changing nominal prices. The existence of these so-called menu costs means that firms respond (in the sense of adjusting their prices) to large shocks but not to small shocks. We show that in the context of a simple dynamic general equilibrium model with no costs of adjusting prices it is possible to observe the same correlation between the skewness of the distribution of price change and the overall inflation rate. Our model is driven solely by supply shocks in the form of technological innovations.

The analysis in this article leaves a number of issues open for further research. First, it would be interesting to document in a more thorough fashion the behavior of the distribution of price changes over the business cycle and its relationship to aggregate activity and aggregate inflation. It would also be interesting to extend the analysis above to a model with more sophisticated dynamics and a more important role for money. Finally, it would be interesting to extend the model outlined above to allow for a limited degree of price stickiness (say along the lines of Ohanian, Stockman, and Kilian 1994) and see how much nominal rigidities can contribute to explaining the relationship between the distribution of price changes and inflation in an equilibrium model. Some of these issues are addressed in Balke and Wynne (1995).

Notes

We thank Evan Koenig and Finn Kydland for comments on an earlier draft. Whitney Andrew provided excellent research assistance.

1 The skewness of a distribution is defined as $E[(x-\mu)^3]/\sigma^3$, where μ is the mean of the distribution of X and σ is the standard deviation.

2 Ball and Mankiw look at two measure of skewness in addition to the conventional measure defined in note 1 above. The first, intended to measure the mass in the

upper tail of the distribution relative to the mass in the lower tail, is defined as

$$AsymX = \int_{-\infty}^{-x} rh(r)dr + \int_{x}^{\infty} rh(r)dr,$$

where r is defined as the relative price *change* defined as the four-digit industry inflation rate minus the average of the four-digit industry inflation rates (*i.e.*, $r_{it} = \Delta\log(p_{i,t}) - \frac{1}{N}\sum_{i=1}^{N}\Delta\log(p_{i,t})$ and $h(r)$ is the density of r. The tails are defined as relative price changes greater than X percent or less than $-X$ percent. Their second alternative measure of skewness is

$$Q = \int_{-\infty}^{\infty} \left| r \right| \bullet rh(r)dr.$$

This variable measures the average of the product of each relative price change and its own absolute value.

[3] Formal tests for nonstationarity indicate that all of the sectoral inflation rates with the exception of agriculture are nonstationary, meaning that in samples of infinite size the variances of these series will be undefined. However, this is not necessarily a problem from our perspective as the data in Table 1 are simply presented to illustrate differences in the rates of change of prices in different sectors.

[4] See also Long and Plosser (1983, 49–50).

[5] Note that in the basic Long and Plosser model, nominal GDP (denominated in utility terms) is a constant. In our extended model, nominal GDP (denominated now in terms of dollars) is directly proportional to the money stock.

[6] The BP filter is explained in Baxter and King (1995).

References

Balke, Nathan S., and Mark A. Wynne, 1995, "A general equilibrium analysis of the relationship between relative price changes and aggregate inflation," manuscript in progress, Federal Reserve Bank of Dallas.

Ball, Laurence and N. Gregory Mankiw, 1995, "Relative-Price Changes as Aggregate Supply Shocks," *Quarterly Journal of Economics*, February, 161–193.

Baxter, Marianne, and Robert G. King, 1995, "Measuring Business Cycles: Approximate Band-Pass Filters for Economic Time Series, NBER Working Paper No. 5022, February.

Benassy, Jean-Pascal, 1995, "Money and wage contracts in an optimizing model of the business cycle," *Journal of Monetary Economics*, 35, 2, April, 303–315.

Long, John B., and Charles I. Plosser, 1983, "Real Business Cycles," *Journal of Political Economy*, vol. 91, no. 1, 39–69.

Ohanian, Lee E., Alan C. Stockman, and Lutz Kilian, 1994, "The effects of real and monetary shocks in a business cycle model with some sticky prices," mimeo University of Pennsylvania.

U.S. Department of Commerce: Bureau of the Census, 1975, Historical Statistics of the United States: Colonial Times to 1970, Washington, D.C.: U.S. Government Printing Office.