

Medium-Term Business Cycles

By DIEGO COMIN AND MARK GERTLER*

Over the postwar period, many industrialized countries have experienced significant medium-frequency oscillations between periods of robust growth versus relative stagnation. Conventional business cycle filters, however, tend to sweep these oscillations into the trend. In this paper we explore whether they may, instead, reflect a persistent response of economic activity to the high-frequency fluctuations normally associated with the cycle. We define as the medium-term cycle the sum of the high- and medium-frequency variation in the data, and then show that these kinds of fluctuations are substantially more volatile and persistent than are the conventional measures. These fluctuations, further, feature significant procyclical movements in both embodied and disembodied technological change, and research and development (R&D), as well as the efficiency and intensity of resource utilization. We then develop a model of medium-term business cycles. A virtue of the framework is that, in addition to offering a unified approach to explaining the high- and medium-frequency variation in the data, it fully endogenizes the movements in productivity that appear central to the persistence of these fluctuations. For comparison, we also explore how well an exogenous productivity model can explain the facts. (JEL E3, O3)

Over the postwar period, many industrialized countries have tended to oscillate between periods of robust growth and relative stagnation. The U.S. economy, for example, experienced sustained high-output growth during the early to late 1960s. From the early 1970s to the early 1980s, however, output growth was low on average. Since the mid-1990s, there has been, for the most part, a return to strong growth. The economies of Europe and Japan have had similar experiences, though the precise details differ.¹ A common feature of these oscillations is that they occur over a longer time frame than is typically considered in conventional business

cycle analysis. Put differently, conventional business cycle detrending methods tend to sweep these kind of oscillations into the trend, thereby removing them from the analysis.

At the same time, these medium-frequency oscillations may be intimately related to the high-frequency output fluctuations normally associated with the business cycle. The long U.S. expansion of the 1960s featured periods of rapid growth and was free of any significant downturn. Similarly, the high growth period over the past ten years has been interrupted by only one recession, considered modest by postwar standards. By contrast, the stagnation of the 1970s and early 1980s featured a number of major recessions, including the two considered the worst of the postwar period. All this raises the possibility that medium-frequency oscillations may, to a significant degree, be the product of business cycle disturbances at the high frequency. Put differently, business cycles may be more persistent phenomena than our conventional measures suggest. It follows that we may need to reexamine how we both identify and model economic fluctuations.

The objective of this paper is twofold: our first goal is to construct an alternative trend/cycle decomposition that includes in the measure of

* Comin: Department of Economics, New York University, 269 Mercer Street, New York, NY 10003 (e-mail: diego.comin@nyu.edu); Gertler: Department of Economics, New York University, 269 Mercer Street, New York, NY 10003 (e-mail: mark.gertler@nyu.edu). We appreciate the helpful comments of two anonymous referees, Olivier Blanchard, Michele Boldrin, Larry Christiano, Jon Faust, Jordi Galí, Pierre-Olivier Gourinchas, Gary Hansen, Louise Keely, Tim Kehoe, Per Krusell, David Laibson, John Leahy, Rodolfo Manuelli, Ed Prescott, Richard Rogerson, and Gianluca Violante. Antonella Trigari and Viktor Tsyrennikov provided outstanding research assistance. Financial assistance from the C.V. Starr Center and the National Science Foundation is greatly appreciated.

¹ See, e.g., Olivier J. Blanchard (1997).

the cycle both the high- and medium-frequency variation in business activity. We refer to this combined high- and medium-frequency variation as the “medium-term business cycle.”² We then present evidence based on a variety of postwar U.S. time-series data that the medium-term cycle differs from the conventionally measured cycle in a number of interesting and important ways. Our second goal is to develop, as a first pass, a quantitative model capable of explaining the facts. Here we attempt to capture both high- and medium-frequency business cycle dynamics within a unified framework.

In Section I, we present evidence on medium-term fluctuations. Roughly speaking, we construct measures of medium-term cycles by running a much smoother trend through the data than is typically used in the conventional high-frequency analysis. We show for a number of key variables that the medium-term cycle is considerably more variable and persistent than the conventionally measured cycle. These findings, further, are statistically significant and not an artifact of a small sample size. In addition, there are a number of interesting patterns. For example, over the medium term there is a strong comovement between output and both disembodied and embodied technological change. Total factor productivity (TFP) moves procyclically over the medium term while the relative price of capital (a measure of embodied technological change) moves countercyclically. R&D moves procyclically. The medium-term cycle, however, is not only a productivity phenomenon. As we show, measures of both the efficiency and intensity of resource utilization also play a role.

In Section II, we develop our model. We consider a simple real business cycle model modified to allow for endogenous productivity, countercyclical markups, and variable factor utilization.³ We incorporate endogenous pro-

ductivity in order to provide a unified explanation for the comovements of TFP, the relative price of capital, and R&D over the medium term. Another consideration leading us in this direction is that many authors have questioned the importance of high-frequency technology shocks based on the view that much of the high-frequency variation in the Solow residual reflects factors such as unmeasured input utilization and imperfect competition as opposed to true technology shifts (e.g., Craig Burnside et al., 1995; Susanto Basu, 1996). Endogenous productivity, however, provides an avenue through which nontechnological shocks that may drive business cycles at the high frequency can generate the strong medium-frequency movements in productivity observed in the data. That is, through this approach, one can capture interesting cyclical variation in both embodied and disembodied technological change without having to rely exclusively on exogenous shifts in technology.

To endogenize productivity dynamics, we use a variation of Paul M. Romer’s (1990) model of R&D expanding the variety of intermediate goods. We extend the Romer framework to allow for an endogenous rate of adoption of new technologies. Julio J. Rotemberg (2003), for example, has argued that in modeling cyclical productivity dynamics, it is important to take into account the significant lag in the diffusion of new technologies.⁴ In this regard, we are able to calibrate our model to allow for realistic steady-state time lags be-

growth model to study business fluctuations that take place over medium-term horizons. They emphasize how their framework may give rise to sunspot fluctuations in the growth rate. In addition to differing considerably in detail, our model has a unique steady-state growth rate: the complementarities in our model then work to magnify the effects of business cycle shocks.

⁴ Rotemberg uses this fact to argue that technology shocks affect only low-frequency movements in output. He then proceeds to develop a framework with exogenous technological change, where nontechnology shocks drive the high-frequency business cycle but have no effect on lower frequencies. Our approach differs by having feedback between the high- and medium-frequency variation in cyclical activity, owing to endogenous productivity. Some support for our approach is that, as we show, measures of the efficiency and intensity of resource utilization show considerable medium-frequency variation and are thus not simply high-frequency phenomena.

² While similar in spirit, our definition differs from Blanchard’s (1997), Ricardo J. Caballero and Mohamad L. Hammour’s (1998) and Robert M. Solow’s (2000), whose notion of the medium term is confined mainly to the medium-frequency variation in the data. We include both the high- and medium-frequency variation because we want to emphasize the interrelation between high- and medium-frequency fluctuations.

³ Our modeling approach is similar to that of George W. Evans et al. (1998), who also develop an endogenous

tween the creation and diffusion of new technologies. At the same time, because adoption as well as R&D intensities vary endogenously over the cycle, the framework can produce the kind of procyclical movement in productivity over the medium term that we document.⁵

As with many conventional quantitative macroeconomic models (e.g., Christiano et al., 2005; Frank Smets and Raf Wouters, 2005), countercyclical movements in both price and wage markups provide the main source of fluctuations at the high frequency. These conventional frameworks typically endogenize markup behavior by introducing nominal rigidities or other types of market frictions. We also endogenize markup behavior but opt for a very simple mechanism that is meant as a stand-in for the kind of richer structure that is typically modeled. Thus, within our framework, an exogenous driving force affects the economy initially by generating countercyclical markup movement. By proceeding with a simple markup structure, however, we are able to avoid additional complexity and keep the focus on the endogenous productivity propagation mechanism, which is the unique feature of our framework.

Section III presents some model simulations and considers how well the framework can capture the broad patterns in the data. Overall, our model does a reasonably good job in characterizing the key features of the medium-term cycle. Even with a nontechnological shock as the main driving force, the model captures most of the cyclical variation in productivity, both at the high and medium frequencies. There are several caveats that we discuss. For comparison, we also explore how well an exogenous technology shock model can explain the data. For this kind of framework to capture movements in both embodied and disembodied technological change, it is necessary to allow for exogenous shocks to the relative price of capital, as well as exogenous shocks to total factor productivity. This approach has some strengths, but also some shortcomings, as we discuss. Concluding remarks are in Section IV.

⁵ Comin and Gertler (2004) provide some evidence that adoption rates are procyclical based on a sample of 22 individual technologies described in Stephen Davies (1979) in the United Kingdom during the postwar period.

I. Measuring the Medium-Term Cycle

In this section we develop some facts about medium-term business fluctuations. To detrend the data, we use a band pass filter, which is basically a two-sided moving average filter, where the moving average depends on the frequencies of the data that one wishes to isolate. A popular alternative is the Hodrick-Prescott (HP) (1997) filter. This filter estimates a trend by maximizing a criterion function that penalizes both deviations of the data from the trend and variation in the trend, depending on the relative weights on each objective in the criterion function. We opt for the band pass filter because it allows us to be precise (in frequency domain terms) about how our measure of the cycle compares with more conventional measures. Our results, however, do not depend on this choice of filters.

Conventional decompositions of the data with the band pass filter associate the business cycle with the frequencies between 2 and 32 quarters (e.g., Marianne Baxter and Robert G. King, 1999, and Christiano and Terry J. Fitzgerald, 2003). This decomposition produces a representation of the cycle that is very similar to the one that the standard HP filter generates. As we show, however, in each of these cases, the trend exhibits considerable variation about a simple linear trend, reflecting the presence of significant cyclical activity at the medium frequencies. Our measure of fluctuations, accordingly, adds back in this medium-frequency variation. We do this, simply put, by applying a filter that identifies a much smoother trend in the data than does the conventional filter.

In particular, we define the medium-term cycle as including frequencies between 2 and 200 quarters. The trend accordingly consists of variation in the data at frequencies of 200 quarters and below. We allow for a very smooth nonlinear trend, as opposed to simply employing a linear trend for several reasons. First, over the sample period there have been a number of factors, such as demographics, that are likely to have introduced variation in the data, mainly at the very low frequencies. Rather than try to incorporate all these factors in our model, we concentrate mainly on explaining the high- and medium-frequency variation in the data. Second, with a linear trend, estimates of some of

the key moments in the data can become very imprecise. We chose a cutoff for the trend in the frequency domain of 200 quarters (roughly the sample size) because we found that with this decomposition we obtain, for the most part, both reasonably smooth trends and estimates of the variability of the filtered data that have a reasonable degree of precision.

We decompose the medium-term cycle as follows: we refer to the frequencies between 2 and 32 quarters (the standard representation of cycles) as the high-frequency component of the medium-term cycle, and frequencies between 32 and 200 quarters as the medium-frequency component. Two points are worth keeping in mind: first, in general it is not correct to think of the medium-frequency variation in the data as orthogonal to the high-frequency variation. For this reason, we focus on understanding the overall fluctuations in the data between frequencies of 2 and 200 quarters and not simply the medium-frequency component. Second, it is important to be careful about the mapping between the frequency domain and the time domain: even though our measure of the cycle includes frequencies up to 50 years, as we show, its representation in the time domain leads to cycles on the order of a decade, reflecting the distribution of the mass of the filtered data over the frequency domain. For the conventional business cycle decomposition, as we show, the cycles are much shorter.

We proceed as follows. Because most of our series are nonstationary, we first convert all the data into growth rates by taking log differences. We then apply the band pass filter to the growth rate data, obtaining a measure of trend growth rates that corresponds to frequencies 200 quarters and below. We then cumulate the trend growth rates to obtain measures of the trends in log levels.⁶ We also perform several robustness checks. Because the band pass filter provides an approximation only for finite series, we “pad” the series by forecasting and backcasting, to minimize biases that are likely to arise, especially at sample endpoints. In addition, as we discuss, we also perform a simple Monte Carlo

⁶ For series where the growth rate or the log level are stationary, such as per capita output, we obtain virtually identical results by filtering the data in log levels (after first removing a linear trend).

exercise to assess the statistical significance of the trends we measure.

The data are quarterly from 1948:1 to 2001:2, except noted otherwise. We consider two sets of variables. The first includes “standard” business cycle series, including output, consumption, investment, hours, and labor productivity.⁷ The second includes “other” variables useful for characterizing the movements in both productivity and resource utilization over the medium term. These include TFP, the quality adjusted relative price of capital, private R&D, the markup of price over social marginal cost, and capacity utilization. The series on TFP, the relative price of capital, and private R&D are available only at an annual frequency.⁸

We measure the markup as the ratio of the marginal product of labor to the household’s marginal rate of substitution between consumption and leisure (where the latter reflects the social marginal cost of labor). To do so, we assume that production is Cobb-Douglas, and that both the coefficient of relative risk aversion and the Frisch labor supply are unity. As Jordi Galí et al. (2002) note, this total markup can be decomposed into the product of a price markup (the ratio of the marginal product of labor to the real wage) and a wage markup (the ratio of the real wage to the household’s marginal rate of substitution).⁹ For capacity utilization we use

⁷ Output, hours, and labor productivity are for nonfarm business. Consumption includes nondurables and services. Investment is nonresidential. We normalize all variables by the working-age population (ages 16–65).

⁸ TFP is from the Bureau of Economic Analysis (BEA). R&D is nonfederally funded R&D expenditures as reported by the National Science Foundation (NSF). The relative price of capital is the quality adjusted price of equipment and structures from Jason G. Cummins and Giovanni L. Violante (2002) (which is, in turn, based on Robert J. Gordon, 1990) divided by the BEA price of nondurable consumption and services. While this series conventionally applied in the literature, it is the subject of some controversy (e.g., Bart Hobijn, 2001). We thus also considered the BEA measure and found that it exhibits nearly identical cyclical properties to the Gordon series, though it is only about 60 percent as volatile overall. We return to this latter point in Section III. Finally, another issue is that the Cummins-Violante series is based on extrapolating the Gordon series after 1983. We find, however, that the pre-1983 cyclical properties are very similar to those of the whole series.

⁹ See also Robert E. Hall (1997) for an analysis of the cyclical properties of the labor market distortion over the postwar period.




FIGURE 1. NONFARM BUSINESS OUTPUT PER PERSON 16-65

the Federal Reserve Board's measure. Given some limitations of this series, we also considered an alternative measure of utilization based on the ratio of electricity usage to the capital stock, and obtained very similar results.¹⁰ Finally, we convert all variables to logs.

Figure 1 plots the medium-term cycle for per capita output. The line with circles gives the percent deviation of per capita output from trend for the medium-term cycle. The solid line gives the medium-term component (i.e., the variation in the data at frequencies between 32 and 200 quarters). The difference between the two lines is the high-frequency component (the variation at frequencies between 2 and 32 quarters). Put differently, this difference is the measure of the cycle in conventional analysis.

Overall, the medium-term cycle appears to capture the sustained swings in economic activity that we described in the introduction. There is sustained upward movement in output relative to trend over the 1960s. Output reverses course in the 1970s through the early 1980s. There is again a sustained upward movement beginning in the mid-1990s. As the figure

makes clear, these persistent movements reflect medium-frequency variation in the data that are large relative to the high-frequency variation. In this regard, the medium-term cycle appears much larger in overall magnitude than the conventionally measured cycle. The percentage rise in output relative to trend over the 1960s, for example, is roughly 14 percent and is mirrored by a roughly similar percentage drop during the period of the productivity slowdown. Another good example involves the period around the 1980-1982 recession: the medium-term cycle suggests a much larger and more protracted drop relative to trend, than does the conventional high-frequency measure.

Figure 2 confirms that our measure of trend per capita output grows at a relatively smooth pace. There is a steady modest decline in the trend growth rate over the postwar period until the mid-1990s, where it gently reverses course. From a more formal perspective, the ratio of the standard deviation of output growth at the low frequencies (below 200 quarters) to the standard deviation of output growth at frequencies between 2 and 200 quarters is only 0.08. Thus, our measure of trend growth is very smooth relative to the overall variation in the growth rate.

Table 1 presents a set of formal statistics to confirm the visual impression left by Figure 1 that the medium-term cycle in output is consid-

¹⁰ In principle, the Board series includes variable labor utilization as well as capital utilization. We expect that the former, however, may be less important over the medium frequencies.




FIGURE 2. LONG-RUN TREND ANNUAL GROWTH RATE IN NONFARM BUSINESS OUTPUT PER PERSON 16–65

TABLE 1—STANDARD DEVIATIONS: QUARTERLY FREQUENCIES

Frequency	Medium-term cycle 2–200	High-frequency component 2–32	Medium-frequency component 32–200
Output per capita	4.12 (2.71–5.52)	2.39 (1.93–2.85)	3.25 (2.41–4.09)
Hours	2.8 (2.38–3.62)	1.73 (1.49–2.01)	2.2 (1.98–2.8)
Labor productivity	2.4 (1.93–2.87)	1.13 (1.03–1.23)	2.07 (1.7–2.44)
Consumption	2.28 (1.72–2.84)	0.93 (0.87–0.99)	2.2 (1.7–2.7)
Investment	10.08 (2.17–17.99)	4.88 (3.04–6.72)	8.64 (2.19–15.09)

erably more volatile than the conventionally measured high-frequency cycle.¹¹ It also presents a similar set of statistics for each of the traditional quarterly business cycle variables. For each variable, the table reports the percent standard deviation over the medium-term cycle, and also the standard deviations of both the

high- and medium-frequency components. Below each number are 95-percent confidence intervals.¹² As the first row indicates, the standard deviation of per capita output over the medium term is roughly twice that of the high-frequency component (4.12 versus 2.39). Moreover, the standard deviation of the medium-frequency component (3.25) is large in comparison to the

¹¹ Interestingly, by examining data only from the first half of the sample, it would not be possible to detect a medium-term cycle that is distinct from the high-frequency cycle. For the period 1948:1–1973:1, for example, the two different filters yield very similar cycles. Only by examining data from the full sample does a distinct medium-term cycle emerge.

¹² To compute confidence intervals, we make use of the following. Let σ_x^2 be the variance of a stochastic process $\{X_t\}_{t=0}^T$. If X_t is stationary, the sample variance $\hat{\sigma}_x^2$ is asymptotically distributed $N(\sigma_x^2, \Sigma)$. We compute Σ by truncating the lags in the theoretical expression using the Newey-West truncation parameter.

TABLE 2—STANDARD DEVIATIONS: ANNUAL FREQUENCIES

Standard variables	Medium-term cycle 0–50	High-frequency component 0–8	Medium-frequency component 8–50
Output per capita	3.87 (3.06–4.67)	2.2 (1.91–2.48)	3.15 (2.55–3.75)
Hours	2.84 (2.42–3.26)	1.59 (1.44–1.74)	2.33 (2.03–2.63)
Labor productivity	2.37 (2.04–2.7)	1.01 (0.96–1.07)	2.11 (1.83–2.4)
Consumption	2.2 (1.83–2.57)	0.8 (0.73–0.87)	2.01 (1.72–2.3)
Investment	9.59 (4.5–14.68)	4.5 (3.34–5.66)	8.29 (3.64–12.94)
Other variables			
TFP	2.39 (2.07–2.7)	1.25 (1.15–1.36)	2.01 (1.76–2.25)
Relative price of capital (Gordon)	4.34 (3.12–5.39)	1.54 (1.27–1.49)	4.12 (3.05–5.14)
R&D	7.65 (4–8.99)	3.4 (2.65–4.13)	5.95 (3.64–8.26)
Total markup	4.35 (3.27–5.43)	2.38 (2.01–2.74)	3.65 (2.74–3.56)
Capacity utilization	4.26 (3.28–5.24)	3.11 (2.38–3.84)	2.81 (2.3–3.32)

latter. A similar pattern exists for hours, labor productivity, consumption, and investment. There are a few natural qualifications, however: as one might expect, a priori, the ratio of variation over the medium-term cycle to variation over the high frequency is somewhat lower for hours (2.8 to 1.73) and somewhat higher for labor productivity (2.4 to 1.13).

Importantly, as the last column of the table shows, the medium-frequency variation in each variable is statistically significant. This can be seen by noting that the confidence interval for the standard deviation of each variable lies safely above zero. In this regard, our measure of the medium-term cycle differs in a statistically significant way from the conventional high-frequency measure.

While Table 1 presents standard deviations of quarterly data, Table 2 presents a similar set of statistics for annual data. It does so for both the standard business cycle series listed in Table 1 and for the other variables we consider (TFP, the relative price of capital, R&D, and the markup). We consider annual data for two reasons. First, because we are attempting to model fluctuations that are considerably more persistent than are typically considered, the quantitative framework we develop in the next section

aims to explain variation in the annual as opposed to the quarterly data. Second, some of the nonstandard series we consider are available only at the annual frequency.

Two central conclusions emerge from Table 2. First, for the standard business cycle series, the standard deviations of the annual data are quite similar to those of the quarterly data, i.e., the message from annual data about the relative importance of the medium-term cycle is virtually the same as from quarterly data. Second, the “other” variables listed in Table 2 similarly exhibit significantly greater volatility over the medium term than over the high-frequency cycle: in each case the standard deviation is more than double. For the relative price of capital, the standard deviation is nearly three times larger than in the high frequency.¹³

Next, Table 3 reports (again, for the annual data) the first-order autocorrelation of each variable over the medium term and over the high frequency. Interestingly, for each series, the

¹³ For the BEA measure of the relative price of capital, the standard deviation is 2.59 for the medium term, 0.82 for the high frequency, and 2.44 for the medium frequency. Thus, while the overall volatility is lower, the breakdown across frequencies is similar to the Gordon series.

TABLE 3—FIRST-ORDER AUTOCORRELATIONS: ANNUAL FREQUENCIES

Standard variables	Medium-term cycle 0–50	High-frequency component 0–8
Output per capita	0.65	0.16
Hours	0.65	0.11
Labor productivity	0.79	0.19
Consumption	0.84	0.4
Investment	0.76	0.22
Other variables		
TFP	0.7	0.16
Relative price of capital (Gordon)	0.86	0.33
R&D	0.83	0.07
Total markup	0.66	0.09
Capacity utilization	0.51	0.19

high-frequency component shows relatively little persistence. In this case, for example, annual output being one percentage point above trend in year t implies that it is likely to be only 0.16 percent above trend in year $t + 1$.¹⁴ Put differently, conventional detrending methods yield output fluctuations at the annual frequency that exhibit low persistence. The same is true for the other series. In sharp contrast, the medium-term cycle exhibits significantly greater persistence for each series. The first-order autocorrelation for output is 0.65, for example, and is similarly large for the other series.

Table 4 presents statistics on comovements among the variables, again for both the standard business cycle series and the nonstandard ones. In the first column is the cross correlation of each annual variable with annual output over the medium-term cycle, and in the second column is the same statistic for the conventionally measured cycle. By and large, the positive comovements among the standard variables that are a hallmark of the conventional business cycle are also a feature of the medium-term cycle. Similarly, the nonstandard variables for the most part display a similar comovement with output across the different frequencies. One notable exception is the relative price of capital. At the high frequency there is only a

TABLE 4—CONTEMPORANEOUS CORRELATION WITH OUTPUT: ANNUAL FREQUENCIES

Standard variables	Medium-term cycle 0–50	High-frequency component 0–8
Hours	0.92	0.85
Labor productivity	0.67	0.77
Consumption	0.75	0.87
Investment	0.67	0.84
Other variables		
TFP	0.8	0.83
Relative price of capital (Gordon)	−0.56	−0.24
R&D	0.31	0.3
Total markup	−0.82	−0.65
Capacity utilization	0.67	0.93

relatively weak negative correlation between annual movements in the relative price of capital (−0.24) and output. This negative relation is clearly stronger over the medium-term cycle (−0.56).¹⁵ In addition, while for R&D the contemporaneous relationship with output is similarly positive across the different measures of the cycle, the lead/lag pattern is distinct. As we show later (Figure 5) there appears a lead of R&D over output over the medium-term cycle. This pattern is absent at the high frequency.

To fully appreciate the behavior of the productivity-related variables in the medium-term cycle, it is useful to examine the respective detrended time series. The three panels in Figure 3 plot the medium-term cycle along with the associated medium-term component for TFP, the relative price of capital, and R&D, respectively. The TFP series is clearly procyclical over the medium term, exhibiting about two-thirds of the volatility of output. Interestingly, the medium-term cycle in the relative price of capital is almost the mirror image of that of TFP. To the extent that declines in the relative price of capital reflect embodied technological change, the figure suggests that there is a strong positive comovement of both disembodied and embodied technological change. The productivity slowdown of the 1970s and early 1980s, for example, featured a slowdown of both types of

¹⁴ This point estimate is slightly smaller than the point estimate for HP-filtered output (around 0.4), but this difference is not statistically significant.

¹⁵ For the BEA measure, the cross correlation over the medium-term cycle is −0.46 and over the high frequency is −0.17, suggesting a very similar cyclical pattern to the Gordon series.

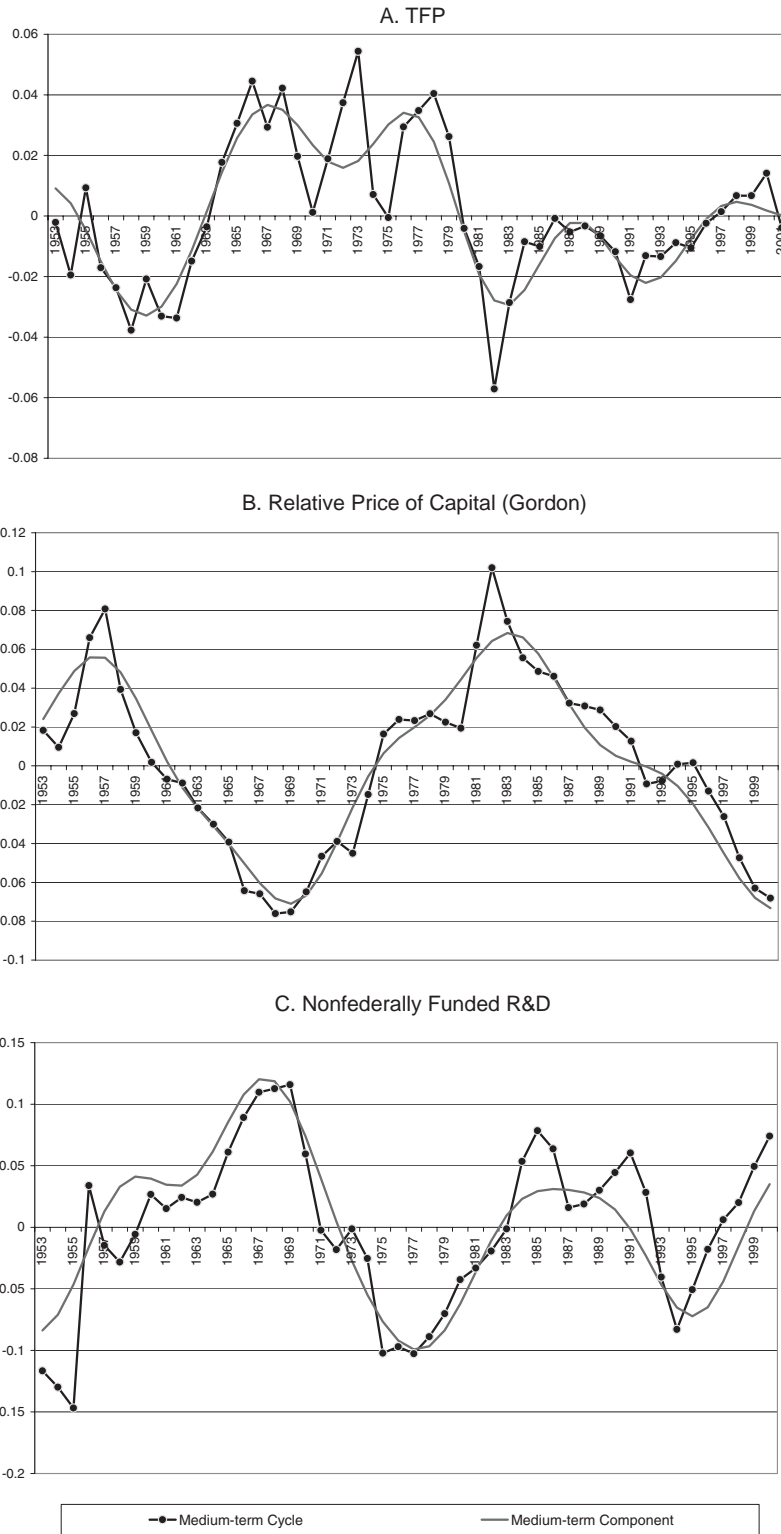


FIGURE 3. A. TFP B. RELATIVE PRICE OF CAPITAL (GORDON) C. NONFEDERALLY FUNDED R&D

technological change. Conversely, both types of technological change were features of the productivity boom in the latter part of the 1990s. Finally, as the bottom panel shows, the movement in R&D over the medium term is highly procyclical. Note the strong inverse relation with the relative price of capital. This is a significant finding, since the R&D data are measures of the resources spent to develop new manufacturing goods, which should enhance embodied technological change.

Thus, overall, procyclical movement of disembodied and embodied technological change and of R&D, particularly at the medium frequencies, appears to be a salient feature of the medium-term cycle. As we noted, however, the medium-term cycle is not only a productivity phenomenon, since medium-frequency movements in the markup and capacity utilization appear important as well, as Tables 2, 3, and 4 suggest.

II. Model

We now develop a model of medium-term business fluctuations. The model is annual as opposed to quarterly, as we noted earlier, because we are interested in capturing fluctuations over a longer horizon than is typically studied. To this end, we abstract from a number of complicating factors that otherwise might be useful for understanding quarterly dynamics, such as money and nominal rigidities.

We consider a two-sector version of a reasonably conventional real business cycle model, modified to allow for, among other things, endogenous productivity and endogenous countercyclical markups.¹⁶ Procyclical entry and exit by final goods firms induces procyclical variation in the degree of competition and, hence, countercyclical variation in (final goods) price markups, in the spirit of Rotemberg and Michael Woodford (1995). The precise formulation we use, though, is based on Galí and Fabrizio Zilibotti (1995).

Final goods producers use intermediate goods as inputs, along with capital and labor. As in Romer (1990), creation of new specialized intermediate goods stemming from R&D is the

source of technological change. As we noted in the introduction, we modify the Romer framework to allow for an endogenous rate of adoption of new technologies, along with endogenous R&D. By doing so, we are able to allow for empirically reasonable diffusion lags but still generate endogenous medium-term swings in productivity. Roughly speaking, endogenous R&D permits disturbances to have permanent effects on productivity movements, while endogenous adoption rates accelerate the transition of productivity to the new steady state. Because the model has two sectors, consumption goods versus capital goods, we are able to distinguish between embodied and disembodied technological change.

Households are conventional, except we allow them a bit of market power in labor supply in order to motivate a wage markup. We then allow for exogenous variation in this market power. This mechanism combined with the entry/exit mechanism in the goods market permits the model to generate countercyclical movements in both price and wage markups that are consistent with the data. As we alluded to earlier, this simple structure is meant as a stand-in for a richer model of markup variation.

We first describe final goods firms, and also characterize the entry and exit/markup mechanism. We next characterize the creation and adoption of new intermediate goods. Then we turn to households, and finally characterize the complete equilibrium.

A. Final Goods Output

Final Output Composite.—There are two sectors: a capital goods sector (k) and a consumption goods sector (c). Within each sector there is a final output composite, $Y_{x,t}$; $x = k, c$. We take the price of the consumption good composite as the numeraire.

The composite for each sector is a CES aggregate of the output of $N_{x,t}$ final goods firms, each producing a differentiated product. Let $Y_{x,t}^j$ be the output of final goods firm j in sector x . Then:

$$(1) \quad Y_{x,t} = \left(\int_0^{N_{x,t}} (Y_{x,t}^j)^{1/\mu_{x,t}} dj \right)^{\mu_{x,t}},$$

¹⁶ For a recent characterization of the real business cycle framework and its variants, see Sergio Rebelo (2005).

with $\mu_{x,t} > 1$. The time-varying parameter $\mu_{x,t}$ is inversely related to the price elasticity of substitution between the differentiated goods. In the symmetric equilibrium, it is the gross markup of price over marginal cost that final goods producers charge. Individual firms take as given the time paths of $N_{x,t}$ and $\mu_{x,t}$. As we discuss shortly, the entry/exit process jointly determines these two variables in equilibrium.

Each final goods firm j produces a specialized good using the following inputs: capital services, $U_{x,t}^j K_{x,t}^j$, where $K_{x,t}^j$ is capital (of firm j in sector x) and $U_{x,t}^j$ is the capital utilization rate; labor, $L_{x,t}^j$; and an intermediate good composite, $M_{x,t}^j$. The production function is Cobb-Douglas, as follows:

$$(2) \quad Y_{x,t}^j = [(U_{x,t}^j K_{x,t}^j)^\alpha (L_{x,t}^j)^{1-\alpha}]^{1-\gamma} [M_{x,t}^j]^\gamma,$$

where γ is the intermediate goods share and α is the capital goods share of value added. Following Jeremy Greenwood et al. (1988), we assume that the depreciation rate of capital is increasing in the utilization rate, i.e., for each firm the depreciation rate is given by $\delta(U_{x,t}^j)$, with $\delta'(\cdot) > 0$.

Intermediate Goods Composite.—The intermediate good composite used by firm j , in turn, is the following CES aggregate of differentiated intermediate goods, where $A_{x,t}$ is the total number of specialized goods in use within each sector:¹⁷

$$(3) \quad M_{x,t}^j = \left(\int_0^{A_{x,t}} (M_{x,t}^{j,k})^{1/\vartheta} dk \right)^\vartheta,$$

with $\vartheta > 1$. Within each sector, each producer k is a monopolistic competitor who has acquired the right to market the respective good via the product development and adoption process that we describe below. In the symmetric equilibrium, ϑ is the producer's gross markup. The CES formulation implies gains from expanding variety. Thus, as in Romer (1990) and others, creation and adoption of new intermediate prod-

ucts (i.e., increasing $A_{x,t}$) provides the ultimate source of productivity growth.

Observe that with fixed values of $N_{x,t}$ and $A_{x,t}$, the model dynamics would mimic those of a (two-sector) real business cycle model. In our framework, however, entry/exit behavior will induce stationary procyclical fluctuations in $N_{x,t}$ that will induce countercyclical fluctuations in $\mu_{x,t}$. In turn, endogenous technology development and adoption will induce permanent adjustments of $A_{x,t}$ inducing long procyclical swings in both embodied and disembodied productivity.

We now turn to the entry/exit and endogenous productivity mechanisms.

B. Entry/Exit and Markups

To characterize the joint determination $N_{x,t}$ and $\mu_{x,t}$, for each sector, we use a simple approach based on Galí and Zilibotti (1995) (GZ) where procyclical competitive pressures associated with endogenous procyclical net entry induce countercyclical movements in the markup.¹⁸ Within the GZ framework, competitive pressures are increasing in the number of firms in the market. In particular, the elasticity of substitution among final goods is increasing in the number of active firms. Because the markup varies inversely with the elasticity of substitution, the model implies the following reduced-form relation between the markup and the total number of firms:

$$(4) \quad \mu_{x,t} = \mu(N_{x,t}); \quad \mu'(\cdot) < 0.$$

As in GZ, we assume that final output firms in sector x must pay a per-period operating cost $b_x \Psi_t$, which the firm takes as given. Let $\Pi(\mu_{x,t}, P_{x,t} Y_{x,t}^j)$ be firm profits (which depend on the markup and output value) where $P_{x,t}$ is the price of the output composite in sector x (with $P_{c,t} = 1$). Then, for any firm j to be in the market, it must be the case that profits are not less than operating costs:

$$(5) \quad \Pi(\mu_{x,t}, P_{x,t} Y_{x,t}^j) \geq b_x \Psi_t.$$

¹⁸ It is also possible to generate countercyclical markups by allowing for money and nominal price rigidities (see, e.g., the discussion in Galí et al., 2002). To keep the model tractable, however, we abstract from money and nominal rigidities.

¹⁷ Production is “roundabout,” in the respect that specialized intermediate goods are made by costlessly transforming final consumption goods. See, e.g., Basu (1995).

In a symmetric competitive equilibrium, equation (5) will hold with equality. This zero-profit condition along with equation (4) will then jointly determine the markup and the quantity of firms. As we show, in the general equilibrium, entry will vary procyclically within each sector and the markup, countercyclically.

In order to ensure a balanced growth equilibrium with stationary markups in each sector, we assume that the time-varying component of operating costs, Ψ_t , drifts up over time proportionately with the social value of the aggregate capital stock, as follows:

$$(6) \quad \Psi_t = P_t^I K_t,$$

where P_t^I is the social price of the capital (i.e., the competitive price) and K_t is the aggregate capital stock. Note that an individual firm takes Ψ_t as given, since it depends only on aggregate factors. One possible interpretation of this formulation is that operating costs are proportionate to the sophistication of the economy, as measured by the social value of the capital stock. It is important to note that, because Ψ_t drifts quite slowly in relative terms, its variation does not significantly affect medium-term dynamics.

C. R&D and Adoption

As we discussed earlier, we modify Romer (1990) to allow for two stages: R&D and adoption. In addition, we distinguish between innovation in the capital goods sector versus the consumption goods sector. Doing so allows us to capture the medium-term variation in the relative price of capital. We first characterize R&D and then adoption.

R&D (Creation of New Intermediate Goods).—Innovators within each sector develop new intermediate goods for the production of final output. They then sell the rights to this good to an adopter who converts the idea for the new product into an employable input, as we describe in the next section.

We assume a technology for creating new specialized goods that permits a simple decentralization of the innovation process, as in Romer (1990). We introduce some modifications, however, that permit us to calibrate the model in a manner that is reasonably consistent with the

data and also that enhance the ability of the model to generate procyclical movements in R&D activity.

Specifically, each innovator p in sector x conducts R&D by using the final consumption good composite as input into developing new products. He finances this activity by obtaining loans from households. Let $S_{x,t}^p$ be the total amount of R&D by innovator p in sector x and let $Z_{x,t}^p$ be his total stock of innovations. Let $\varphi_{x,t}$ be a productivity parameter that the innovator takes as given, and let $1 - \phi$ be the probability that any existing specialized good becomes obsolete in the subsequent period (whether or not it has been adopted). Then the innovation technology is given as follows:

$$(7) \quad Z_{x,t+1}^p = \varphi_{x,t} S_{x,t}^p + \phi Z_{x,t}^p,$$

where ϕ is the implied product survival rate. Allowing for a (small) rate of obsolescence of technologies is not only realistic, it also gives us flexibility in calibrating the model to match the balanced growth evidence.

We assume that the technology coefficient $\varphi_{x,t}$ depends on aggregate conditions which are taken as given by individual innovators. In particular, we assume

$$(8) \quad \varphi_{x,t} = \chi_x Z_{x,t} [\Psi_t^p \cdot S_{x,t}^{1-\rho}]^{-1},$$

with $0 < \rho \leq 1$ and where χ_x is a scale parameter. As in the Romer model, there is a positive spillover of the aggregate stock of innovations, $Z_{x,t}$. In addition, we relax the restriction of a unit elastic response of new technology development to R&D by introducing a congestion externality via the factor $[\Psi_t^p \cdot S_{x,t}^{1-\rho}]^{-1}$. This congestion effect raises the cost of developing new products as the aggregate level of R&D intensity increases (since the factor is decreasing in $S_{x,t}$). It is straightforward to show that in equilibrium the R&D elasticity of new technology creation becomes ρ under this formulation. We are thus free to pick this elasticity to match the evidence. As a way to ensure that the growth rate of new intermediate product is stationary, we also assume that the congestion effect depends positively on the scaling factor Ψ_t . This is tantamount to assuming that, everything else equal, the marginal gain from R&D declines as the economy becomes more

sophisticated, as measured by the aggregate capital stock.

Let $J_{x,t}$ be the value of an “unadopted” specialized intermediate good. This is the price at which an innovator can sell a new product to an adopter. Given the linearity of the R&D technology (as perceived by the individual product developer) and given free entry, each new product developer must exactly break even. As a consequence, the following arbitrage conditions must be satisfied in any equilibrium with an interior solution for innovation:

$$(9) \quad 1/\varphi_{x,t} = \phi E_t \{ \Lambda_{t+1} J_{x,t+1} \},$$

where Λ_{t+1} is the stochastic discount factor.

The right side of equation (9) is the discounted marginal benefit from developing a new product, while the left side is the marginal cost.

Equation (9) strongly hints at how the framework generates procyclical R&D. During a boom, for example, the expected value of a new “unadopted” product, $E_t J_{x,t+1}$, increases. That is, since the profit flow from specialized intermediate goods rises, the benefit to creating new types of these goods goes up. R&D spending will increase in response.

Adoption (Conversion of Z to A).—We next characterize how newly developed intermediate goods are adopted over time (i.e., the process of converting $Z_{x,t}$ to $A_{x,t}$). Our goal is to capture the notion that adoption takes time on average, but allow for adoption intensities to vary procyclically, consistent with the evidence described in Comin and Gertler (2004). In addition, we would like to characterize the diffusion process in a way that minimizes complications from aggregation. In particular, we would like to avoid having to keep track, for every available technology, of the fraction of firms that have or have not adopted it.

These considerations lead us to the following formulation: a competitive set of “adopters” converts available technologies into use. As we discussed earlier, they buy the rights to the technology from the original innovator. They then convert the technology into usable form by employing a costly and potentially time-consuming process. To fund adoption expenses, they obtain loans from households. Once the

new product is in usable form, the adopter is able to manufacture it and sell it to final goods producers who use it as part of the intermediate goods composite in the production function (see equations (2) and (3)).

The pace of adoption depends positively on the level of adoption expenditures in the following simple way: whether or not an adopter succeeds in making a product usable in any given period is a random draw with (success) probability $\lambda_{x,t}$ that is increasing in the amount the adopter spends, $H_{x,t}$ (in units of the final consumption goods composite). If the adopter is not successful, he may try again in the subsequent period. Thus, under our formulation there is slow diffusion of technologies on average (as they are slow on average to become usable). The diffusion rate, however, varies positively with the intensity of adoption expenditures. At the same time, aggregation is simple since once a technology is in usable form, all firms may employ it.

In particular, let $V_{x,t}$ be the value to an adopter of successfully bringing a new product into use, i.e., the present value of profits the adopter would receive from marketing the specialized intermediate good, given by:

$$(10) \quad V_{x,t} = \Pi_{m,x,t} + \phi E_t \{ \Lambda_{t+1} V_{x,t+1} \}$$

where $\Pi_{m,x,t}$ is profits at t in sector x earned on a specialized intermediate input. Then we may express the value of an “unadopted” product to the adopter as

$$(11) \quad J_{x,t} = \max_{H_{x,t}} \{ -H_{x,t} + \phi E_t (\Lambda_{t+1} [\lambda_{x,t} V_{x,t+1} + (1 - \lambda_{x,t}) J_{x,t+1}] \}$$

where $\lambda_{x,t}$ is the following increasing function of adoption expenditures, $H_{x,t}$, and a scaling parameter parameter $\Gamma_{x,t}$ that is exogenous to the adopter:

$$(12) \quad \lambda_{x,t} = \lambda \left(\frac{A_{x,t}}{\Psi_t} \cdot H_{x,t} \right),$$

with $\lambda' > 0$, $\lambda'' < 0$, and where the adopter takes the scaling factor $A_{x,t}/\Psi_t$ as given. We

introduce the (slow-moving) scaling factor to keep the adoption rate stable along a balanced-growth path. Otherwise, the pace of adoption could increase without bound as the economy grows. Our formulation assumes a spillover effect from aggregate adoption, $A_{x,t}$, to individual adoption (i.e., one can learn from the adoption activities of others) that diminishes as the economy becomes more complex, as measured by the value of aggregate capital stock Ψ_t .

In period t the adopter spends $H_{x,t}$. With probability $\lambda(\Gamma_{x,t}H_{x,t})$ he succeeds in making the product usable in the subsequent period, enabling him to generate a profit stream worth $R_{t+1}^{-1}\phi V_{x,t+1}$ in current value. With probability $1 - \lambda(\Gamma_{x,t}H_{x,t})$, he fails and must start over. At the margin, each adopter adjusts conversion expenditures until the marginal expenditure (one unit of final goods consumption output) equals the discounted marginal benefit:

$$(13) \quad 1 = \frac{A_{x,t}}{\Psi_t} \lambda' \cdot \phi E_t \{ \Lambda_{t+1} [V_{x,t+1} - J_{x,t+1}] \}.$$

It is now easy to see why adoption expenditures will move procyclically. During booms, the value of having a marketable product rises relative to one still in the development phase, i.e., $V_{x,t+1}$ rises relative to $J_{x,t+1}$. In this case, $H_{x,t}$ will increase since it becomes worth more at the margin to invest in adoption. The reverse, of course, will happen in downturns.

Overall, the following partial adjustment equation characterizes the adoption rate by any adopter q :

$$(14) \quad A_{x,t+1}^q = \lambda_{x,t} \phi [Z_{x,t}^q - A_{x,t}^q] + \phi A_{x,t}^q.$$

The term in brackets is the stock of products at t that the adopter owns but has not yet converted. He is able to successfully adopt the fraction $\lambda_{x,t}$ (which depends on his expenditure intensity $H_{x,t}$).¹⁹ The business cycle affects this evolution in two ways. First, the stock of new products created moves procyclically due to procyclical movement in R&D

intensity, raising $Z_{x,t}^q$ on average across adopters. Second, adoption expenditures vary procyclically, causing the procyclical variation in adoption rates.

D. Households and Government

Households.—Our formulation of the household sector is reasonably standard. One distinctive feature, however, is that we allow each household to have a bit of market power in labor supply as a way to motivate a wage markup. As we noted earlier, time variation in the wage markup will provide the source of exogenous disturbances in our baseline model.

There is a continuum of households of measure unity. Each household supplies a differentiated type of labor L_t^h . Labor input L_t is the following CES aggregate of all the different types of household labor:

$$(15) \quad L_t = \left(\int_0^1 (L_t^h)^{1/\mu_{w,t}} dh \right)^{\mu_{w,t}},$$

where $\mu_{w,t}$ governs the elasticity of substitution across different types of labor and, in equilibrium, will correspond to a markup of wages over each household's cost of supply labor (in perfect analogy to the price markup for each sector $\mu_{w,t}$). Let W_t^h be the wage of household h and W_t the wage index. Then cost minimization by firms implies that each household faces the following demand curve for its labor:

$$(16) \quad L_t^h = \left(\frac{W_t^h}{W_t} \right)^{-\mu_{w,t}/(1-\mu_{w,t})} L_t$$

with

$$(17) \quad W_t = \left[\int_0^1 (W_t^h)^{1/(1-\mu_{w,t})} dh \right]^{1-\mu_{w,t}}.$$

In addition to supplying labor, each household consumes and saves. We allow for perfect consumption insurance (implying in our case that all households will consume the same amount). A household may save by either accumulating capital or lending to innovators and adopters. It also has equity claims on all monopolistically

¹⁹ Note that the total (per period) investment in adoption undertaken by the q^{th} adopter is $H_{x,t} \cdot [Z_{x,t}^q - A_{x,t}^q]$.

competitive firms. It makes one-period loans to innovators and adopters and also rents capital that it has accumulated directly to firms.

Let C_t be consumption. Then the household maximizes the present discounted utility as given by the following expression:

$$(18) \quad E_t \sum_{i=0}^{\infty} \beta^i \left[\ln C_{t+i} - \frac{(L_{t+i}^h)^{1+\zeta}}{1+\zeta} \right],$$

with $0 < \beta < 1$ and $\zeta > 0$. The budget constraint is as follows:

$$(19) \quad C_t = W_t^h L_t^h + \Pi_t + [D_t + P_t^k] K_t \\ - P_t^k K_{t+1} + R_t B_t - B_{t+1} - T_t,$$

where Π_t reflects the profits of monopolistic competitors paid out fully as dividends to households, B_t is total loans the household makes at $t - 1$ that are payable at t , and T_t reflects lump-sum taxes. R_t is the (possibly state-contingent) payoff on the loans. The household's decision problem is simply to choose the wage rate, consumption, labor supply, capital, and bonds to maximize equation (18) subject to (19), as well as the labor demand curve (16).

Observe that because a household has some market power, the wage it chooses is a markup $\mu_{w,t}$ over its marginal rate of substitution between labor and consumption, $(L_t^h)^\zeta C_t$:

$$(20) \quad W_t^h = \mu_{w,t} \cdot (L_t^h)^\zeta C_t.$$

Further, in the symmetric equilibrium, all households will charge the same wage and supply the same amount of labor. In equilibrium, thus, $\mu_{w,t}$ is the markup of the wage index over the household marginal rate of substitution.

As we noted in the introduction, a number of authors have argued that countercyclical movements in (something like) $\mu_{w,t}$ provides the main source of high-frequency business fluctuations, although there is considerable disagreement as to what truly underlies the fluctuations in this variable. As we noted earlier, various labor market frictions that give rise to wage rigidity can generate countercyclical movements in $\mu_{w,t}$. Rather than adding further to the complexity, however, we simply explore the implications of exogenous shifts in $\mu_{w,t}$, keep-

ing in mind that these shifts may reflect labor market frictions that we have not explicitly modeled.

Government.—Finally, government spending is financed with lump-sum taxes:

$$(21) \quad G_t = T_t.$$

E. Symmetric Equilibrium Conditions

The economy has a symmetric sequence of markets equilibrium. The endogenous state variables are the aggregate capital stock, K_t , the total stocks of intermediate goods invented in each sector, $Z_{c,t}$ and $Z_{k,t}$, and the total stock of intermediate goods adopted in each sector, $A_{c,t}$ and $A_{k,t}$. The following system of equations characterizes the equilibrium.

Resource Constraints and Technology.—Let Y_t denote aggregate net value added output. This quantity equals the sum over sectors of the value of gross output, $P_{x,t} Y_{x,t}$, net expenditures on specialized intermediate goods, $(A_{x,t})^{1-\vartheta} M_{x,t}$, and operating costs, $\psi_{x,t}$:

$$(22) \quad Y_t = \sum_{x=c,k} [P_{x,t} Y_{x,t} - (A_{x,t})^{1-\vartheta} M_{x,t} - \psi_{x,t}].$$

(Recall, $P_{c,t} = 1$.) Note that net expenditures on intermediate goods in a sector is the total produced $M_{x,t}$ normalized by the per-unit cost of producing $(A_{x,t})^{\vartheta-1}$. The latter reflects the productivity gains from diversity implied by the CES aggregator (equation (3)).

The uses of output, in turn, are divided between consumption, C_t , investment, $P_{k,t} Y_{k,t}$, government spending, G_t , and the total costs of R&D and adoption, $\sum_{x=c,k} [S_{x,t} + (Z_{x,t-j} - A_{x,t}) H_{x,t}]$:

$$(23) \quad Y_t = C_t + P_{k,t} Y_{k,t} + G_t \\ + \sum_{x=c,k} [S_{x,t} + (Z_{x,t-j} - A_{x,t}) H_{x,t}].$$

In addition, capital evolves over time according to:

$$(24) \quad K_{t+1} = (1 - \delta(U_t)) K_t + Y_{k,t}.$$

The gross final output composite in each sector, $Y_{x,t}$, is given by

$$(25) \quad Y_{x,t} = (N_{x,t})^{\mu_{x,t}-1} \left[\left(\frac{U_t K_t}{L_t} \right)^\alpha (L_{x,t}) \right]^{1-\gamma} \cdot (M_{x,t})^\gamma,$$

where the term $(N_{x,t})^{\mu_{x,t}-1}$ reflects the efficiency gains to diversity implied by the CES aggregator for final goods in each sector.

Factor Markets.—The labor market in each sector satisfies the requirement that the marginal product of labor equals the product of the (sector-specific) markup and the household's marginal rate of substitution between leisure and consumption:

$$(26) \quad (1 - \alpha)(1 - \gamma) \frac{P_{x,t} Y_{x,t}}{L_{x,t}} = \mu_{x,t} \mu_{w,t} L_t^\zeta C_t,$$

where aggregate labor input L_t equals the sum of labor input across sectors:

$$(27) \quad L_t = L_{c,t} + L_{k,t}.$$

Observe that the total markup of the marginal product of labor over the household's marginal rate of substitution in each sector corresponds to the product of the price and wage markups: $\mu_{x,t} \mu_{w,t}$. This latter expression, in turn, corresponds to the labor market wedge that we used earlier to compute the time series for the markup.

It is straightforward to show that the capital/labor ratios and the capital utilization rates are uniform across sectors. The equilibrium conditions for the sectoral allocations of capital, the sectoral utilization rates, and the sectoral allocations of the intermediate goods are, then, respectively:

$$(28) \quad \alpha(1 - \gamma) \frac{P_{x,t} Y_{x,t}}{K_{x,t}} = \mu_{x,t} [D_t + \delta(U_t) P_t^K],$$

$$(29) \quad \alpha(1 - \gamma) \frac{P_{x,t} Y_{x,t}}{U_t} = \mu_{x,t} \delta'(U_t) P_t^K K_{x,t},$$

$$(30) \quad \gamma \frac{P_{x,t} Y_{x,t}}{M_{x,t}} = \mu_{x,t} P_{x,t}^M,$$

where $P_{x,t}^M$ is the relative price of the intermediate goods composite in sector x (see below), and with

$$(31) \quad K_{x,t} = (K_t/L_t) L_{x,t}.$$

Consumption/Saving.—We can express the intertemporal Euler equation as

$$(32) \quad E_t \left\{ \Lambda_{t+1} \left[\frac{D_t + P_{t+1}^K}{P_t^K} \right] \right\} = 1,$$

with $\Lambda_{t+1} = \beta C_t / C_{t+1}$.

Arbitrage between acquisition of capital and loans to adopters and innovators implies

$$(33) \quad \{E_t \Lambda_{t+1} R_{t+1}\} \\ = E_t \left\{ \Lambda_{t+1} \left[\frac{D_t + P_{t+1}^K}{P_t^K} \right] \right\}.$$

We now characterize the variation in final goods firms, $N_{x,t}$, and in specialized intermediate goods, $A_{x,t}$, beginning with the former.

Entry and Exit.—Free entry by final goods producers in each sector yields the following inverse link between the sector's markup, $\mu_{x,t}(N_{x,t})$, and per-firm output, $(N_{x,t})^{1-\mu_{x,t}} Y_{x,t} / N_{x,t} = (N_{x,t})^{-\mu_{x,t}} Y_{x,t}$.

$$(34) \quad \frac{\mu_{x,t}(N_{x,t}) - 1}{\mu_{x,t}(N_{x,t})} (N_{x,t})^{-\mu_{x,t}} Y_{x,t} = b_x \Psi_t$$

where the scaling factor, Ψ_t , drifts slowly over time according to equation (6).

Evolution of Productivity.—The R&D technology (equations (7) and (8)) implies that in the symmetric equilibrium the gross growth rate of new intermediate products in each sector depends positively on R&D (scaled relative to the social value of the capital stock), as follows:

$$(35) \quad \frac{Z_{x,t+1}}{Z_{x,t}} = \chi_x \left(\frac{S_{x,t}}{\Psi_t} \right)^\rho + \phi.$$

Similarly, the optimality condition for adoption, along with the adoption technology (equations (12) and (14)), imply that the gross rate of adoption of new intermediate goods in equilibrium obeys the following partial adjustment mechanism, where the rate of adjustment depends positively on adoption expenditures, $H_{x,t}$:

$$(36) \quad \frac{A_{x,t+1}}{A_{x,t}} = \lambda \left(\frac{A_{x,t}}{\Psi_t} H_{x,t} \right) \phi \left[\frac{Z_{x,t}}{A_{x,t}} - 1 \right] + \phi.$$

Free entry into research and development (see equations (7), (8), and (9)) implies that the discounted value of the gain from R&D must equal the total cost

$$(37) \quad \phi E_t \{ \Lambda_{t+1} \cdot J_{x,t+1} \} \cdot (Z_{x,t+1} - \phi Z_{x,t}) = S_{x,t}.$$

Similarly, free entry into adoption (see equations (12) and (13)) implies that adoption expenditures will be increasing in the difference between the value of an adopted project $V_{x,t+1}$, versus the value of an unadopted one $J_{x,t+1}$:

$$(38) \quad \frac{A_{x,t}}{\Psi_t} \lambda \cdot \phi E_t \{ \Lambda_{t+1} [V_{x,t+1} - J_{x,t+1}] \} = 1,$$

with $V_{x,t}$ given by the discounted stream of earnings from a specialized intermediate good:

$$(39) \quad V_{x,t} = \left[\mu_{x,t}^{-1} (1 - 1/\vartheta) \gamma \frac{P_{x,t} Y_{x,t}}{A_{x,t}} \right] + \phi E_t \{ \Lambda_{t+1} V_{x,t+1} \},$$

where the first term in brackets is profits at time t and where the value of an unadopted technology, $J_{x,t}$, obeys

$$(40) \quad J_{x,t} = -H_{x,t} + \phi E_t \{ \Lambda_{t+1} [J_{x,t+1} + \lambda (\Gamma_{x,t} H_{x,t}) (V_{x,t+1} - J_{x,t+1})] \}.$$

Relative Prices and Productivity Dynamics.—Finally, we present expressions for the equilibrium-relative prices of intermediate goods and of capital. We also draw the connection of these prices to movements in both disembodied and embodied technological change.

From the CES aggregator for intermediate goods (equation (3)), the relative price of the intermediate good composite in the symmetric equilibrium is given by

$$(41) \quad P_{x,t}^M = \vartheta (A_{x,t})^{1-\vartheta}.$$

Observe that $P_{x,t}^M$ declines as the number of specialized intermediate products, $A_{x,t}$, increases, since ϑ exceeds unity. Technological progress is thus manifested in falling prices of intermediate goods.

To see the link between intermediate goods prices and disembodied productivity, consider the expression for net value added output within each sector, $Y_{x,t}^v$ (obtained by netting out materials and overhead costs):

$$(42) \quad Y_{x,t}^v = \theta_{x,t} K_{x,t}^\alpha L_{x,t}^{1-\alpha}$$

where $\theta_{x,t}$ is the Solow residual in sector x and is given by

$$(43) \quad \theta_{x,t} = \left(1 - \frac{\gamma}{\mu_{x,t} \vartheta} \right) N_{x,t}^{\mu_{x,t}-1} \times \left(\frac{\gamma}{\mu_{x,t} P_{x,t}^M} \right)^{\gamma(1-\gamma)} U_t^\alpha - \frac{\psi_{x,t}}{K_{x,t}^\alpha L_{x,t}^{1-\alpha}}.$$

Observe that $\theta_{x,t}$ increases as $P_{x,t}^M$ declines. The strength of the effect, further, depends positively on the intermediate goods share parameter γ .

Movements in $P_{x,t}^M$ generate variation in $\theta_{x,t}$ that reflects true technological change. We should expect the influence of $P_{x,t}^M$ on $\theta_{x,t}$ to play out mainly over the medium and low frequencies, since the force underlying this relationship is the variation of $A_{x,t}$, which, in turn, is a product of the endogenous research and technology adoption processes. Observe also that underlying the Solow residual are several others factors that do not reflect true technological change and that are likely to be important at the high frequency. The capital utilization rate, U_t , enters $\theta_{x,t}$ because it reflects unmeasured intensity of input usage. The term $N_{x,t}^{\mu_{x,t}-1}$ reflects the agglomeration effect present in the CES composite for final output. The markup $\mu_{x,t}$ also enters $\theta_{x,t}$ due to the roundabout input-output nature of production. On balance, U_t , $N_{x,t}^{\mu_{x,t}-1}$, and $\mu_{x,t}$ contribute to procyclical variation in $\theta_{x,t}$, particularly at the high frequencies, while $P_{x,t}^M$ contributes to procyclical variation mainly at the medium frequencies and below.

Lastly, we consider the relative price of capital and embodied technological change. In equilibrium, P_t^K obeys

$$(44) \quad P_t^K = \left(\frac{\mu_{k,t}}{\mu_{c,t}} \right) \cdot \left(\frac{N_{k,t}^{1-\mu_{k,t}}}{N_{c,t}^{1-\mu_{c,t}}} \right) \cdot P_t^I,$$

where P_t^I is the social price of capital (i.e., the price that would arise in competitive markets), given by

$$(45) \quad P_t^I = \left(\frac{A_{k,t}}{A_{c,t}} \right)^{(1-\theta)\gamma} = \left(\frac{P_{k,t}^M}{P_{c,t}^M} \right)^\gamma.$$

Movements in the social price of capital depend on the state of technology in the capital goods sector relative to the consumption sector, as measured by the ratio of intermediate goods sector prices. Note that P_t^I is likely to vary countercyclically over the medium term because the pace of endogenous technological change in the capital goods sector relative to the consumption goods sector is likely to vary procyclically: in particular, R&D and adoption exhibit stronger procyclical variation in the investment sector as a consequence of the stronger procyclical variation of profitability in this sector.

As with the effect of $P_{x,t}^M$ on $\theta_{x,t}$, the impact of $P_{k,t}^M/P_{c,t}^M$ on P_t^K reflects true technological progress and is also likely to play out over the medium and low frequencies. There are similarly nontechnological influences on P_t^K that are likely to be particularly relevant at the high frequency. Specifically, P_t^K depends on the relative markups in the two sectors and also on the relative number of firms operating in each market (the latter reflects the relative strength of the agglomeration effect from the CES aggregators). Note that the high-frequency variation in P_t^K is also likely to be countercyclical because entry exhibits greater procyclical variation in the capital goods sector and, conversely, the markup exhibits greater countercyclical variation. Thus, forces at both the high and medium frequency contribute to countercyclical variation in P_t^K over the medium term.

III. Model Simulations

In this section we explore the ability of the model to generate medium-term cycles. We solve the model by loglinearizing around the deterministic balanced growth path and then

employing the Anderson-Moore code, which provides numerical solutions for general first-order systems of difference equations. We first describe the calibration and then turn to some numerical exercises.

A. Model Calibration

The calibration we present here is meant as a benchmark. We have found that our results are robust to reasonable variations around this benchmark. To the extent possible, we use the restrictions of balanced growth to pin down parameter values. Otherwise, we look for evidence elsewhere in the literature. There is a total of 18 parameters. Ten appear routinely in other studies; the eight others relate to the R&D and adoption processes and also to the entry/exit mechanism.

We begin with the standard parameters. We set the discount factor β equal to 0.95 to match the steady state share of nonresidential investment to output. Based on steady-state evidence, we also choose the following number: (the capital share) $\alpha = 0.33$; (the materials share) $\gamma = 0.50$; (government consumption to output) $G/Y = 0.2$; (the depreciation rate) $\delta = 0.08$; and (the steady-state utilization rate) $U = 0.8$. We set the inverse of the Frisch elasticity of labor supply ζ at unity, which represents an intermediate value for the range of estimates across the micro and macro literature. Similarly, we set the elasticity of the change in the depreciation rate with respect to the utilization rate, $(\delta'/\delta')U$ at 0.33, which lies within the range that other studies suggest.²⁰ Finally, based on evidence in Basu and John G. Fernald (1997), we fix the steady-state gross valued added markup in the consumption goods sector, μ^c , equal to 1.1 and the corresponding markup for the capital goods sector, μ^k , at 1.2.

We next turn to the “nonstandard” parameters. The steady-state growth rates of GDP and the relative price of capital in the model are functions of the growth rate of new technologies developed in the consumption and capital goods sectors. By using the balanced growth restric-

²⁰ We set U equal to 0.8 based on the average capacity utilization level in the postwar period as measured by the Board of Governors. We set $(\delta'/\delta')U$ equal to $1/3$ based on the range of values used in the literature that vary from 0.1 used by Robert G. King and Rebelo (1999) to unity used by Baxter and Dorsey Farr (2001).

tions and matching the average growth rate of nonfarm business output per working-age person (0.024), the average growth rate of the Gordon quality adjusted price of capital relative to the BEA price of consumption goods and services (-0.026), and the average share of private research expenditures in GDP during the postwar period (0.009),²¹ we can identify the annual survival probability of new technologies, ϕ , and the productivity parameters in the technologies for creating new intermediate goods in each sector, χ_c and χ_k . Accordingly, we set $\phi = 0.97$ (implying an annual obsolescence rate of three percent), $\chi_c = 2.27$, and $\chi_k = 32.36$.

There is no direct evidence on the gross markup ϑ for specialized intermediate goods. Given the specialized nature of these products, it seems that an appropriate number would be at the high range of the estimates of markups in the literature for other types of goods. Accordingly, we choose a value of 1.6, but emphasize that our results are robust to reasonable variations around this number.

There is also no simple way to identify the parameter ρ , the elasticity of new intermediate goods with respect to R&D. Zvi Griliches (1990) presents some estimates using the number of new patents as a proxy for technological change. The estimates are noisy and range from about 0.6 to 1.0, depending on the use of panel versus cross-sectional data. We accordingly set ρ equal to 0.8, the midpoint of these estimates.

We now consider the parameters that govern the adoption process. To calibrate the steady-state adoption frequency λ , we use information on diffusion lags. Given that λ is the probability a technology is adopted in any given period, the average time to adoption for any intermediate good is $1/\lambda$. Mansfield (1989) examines a sample of embodied technologies and finds a median time to adoption of 8.2 years. Comin and Gertler (2004) examine a sample of British data that includes both disembodied and embodied

innovations. They find median diffusion lags of 9.8 and 12.5 years, respectively. Assuming that the median is not too different from the mean, a reasonable value for λ is 0.1, implying an average time to adoption of ten years.

Another parameter we need to calibrate is the elasticity of the frequency of adoption with respect to adoption expenditures. Doing so is difficult because we do not have good measures of adoption expenditures, let alone adoption rates. One partial measure of adoption expenditures we do have is development costs incurred by manufacturing firms trying to make new capital goods usable (which is a subset of the overall measure of R&D that we used earlier). A simple regression of the rate of decline in the relative price of capital (the relevant measure of the adoption rate of new embodied technologies in the context of our model) on this measure of adoption costs and a constant yields an elasticity of 0.95. Admittedly, this estimate is crude, given that we do not control for other determinants of the changes in the relative price of capital. On the other hand, we think it provides a plausible benchmark value.

We next turn to the entry/exit mechanism. We set the elasticity of the price markup with respect to the number of firms in each sector (i.e., $\partial \log \mu_{x,t} / \partial \log N_{x,t}$). By doing so, the overall medium-term variation in the number of firms is roughly consistent with the data. In the data, the percent standard deviation of the number of firms relative to the standard deviation of the total markup over the medium-term cycle is 0.65. With our parameterization, the model produces a ratio of 0.15. Given, however, that the data do not weight firms by size, and given that most of the cyclical variation is due to smaller firms, the true size-weighted variation in the number of firms (relative to the markup) is probably closer to 0.15 than 0.65. As a check, we show that our results are largely unaffected by varying the price markup elasticity over the range from 0.5 to 1.5.

Finally, we fix the autocorrelation of the preference/wage markup shock so that the model generates an autocorrelation that approximately matches that of the overall gross markup, $\mu_t \cdot \mu_t^w$, as measured by Galí et al. (2002). This results in a value of 0.6. We then choose a standard deviation of the shock, so that the unconditional variance of the overall gross markup matches the data (for the medium-term cycle).

²¹ The NSF divides R&D into three components: basic, applied, and development. The first two are clearly included in our notion of R&D. Development is defined by the NSF as expenses devoted to improving existing products or processes. Therefore, some of what the NSF classifies as development expenses may belong to our investments in technology adoption. We take this view in the calibration and assume that only half of the private expenses in development as reported by the NSF are R&D. That brings the average private R&D expenses in the postwar period to $\%_{10}$ of 1 percentage point.

B. Some Numerical Experiments

To be clear, the exercises that follow are meant simply as a first pass at exploring whether the mechanisms we emphasize have potential for explaining the data: they are not formal statistical tests. As we discussed earlier, we treat innovations to the wage markup, μ_t^w , stemming from exogenous fluctuations in household labor market power as the main source of disturbances to the economy. We keep in mind, though, that this simple mechanism is meant as a shortcut for a richer description of countercyclical wage markup behavior.²²

Ideally, we would like to evaluate the model against moments of the data that are conditional on a shock to μ_t^w . There are several difficulties with this approach, however. First, while identified VAR methods may be useful for measuring the near-term effects of shocks to the economy, they usually are very imprecise at measuring the kind of medium-term effects that are the focus of our analysis.²³ Second, identification of orthogonalized movements in μ_t^w is highly problematic, particularly given that in practice there is likely a strong endogenous component to this variable.

Rather than attempt to match imprecisely estimated conditional moments, we follow the lead of the traditional real business cycle literature by exploring how well a single shock, presumed to be the principle driving force, can explain the unconditional patterns in the data. We view this kind of exercise as a check on the plausibility of the scenario we offer, as opposed to a clear test against alternatives. While the RBC literature focused on the ability of technology shocks to account for the high-frequency variation in the data, we instead consider the ability of our model driven by shocks to μ_t^w to explain the combined high- and medium-frequency variation, as well as each of the components. We also explore how well an exogenous technology shock model can explain medium-term fluctuations, as compared to our benchmark model.

²² As discussed in Galí et al. (2002), a model with either nominal or real wage rigidities can generate a countercyclical wage markup.

²³ In particular, long-run effects of impulse response functions from identified VARs are measured much less precisely than are short-run effects.

The Baseline Model with Endogenous Productivity.—Before confronting the data, we first gain some intuition for the workings of the model by examining the model impulse responses to a shock to μ_t^w . In Figure 4, the solid line in each panel is the response of the respective variable within our full-blown model. The dotted line corresponds to the simple benchmark with neither endogenous productivity nor endogenous markup variation (though allowing for endogenous factor utilization). The period length is a year.

The rise in μ_t^w effectively raises the price of labor, reducing labor demand and output. Both the initial impact and the impact over time on output is larger in our full-blown model than in the benchmark model. Within the full-blown model, the decline in output leads to a rise in both the consumption and capital goods price markups (due to a decline in entry within each sector), enhancing the initial drop in output. Over time, output climbs back to trend but does not make it back all the way due to the endogenous permanent decline in productivity, which we describe shortly. After a decade or so, it levels off at a new steady state. The permanent decline in output is about one-half of the initial decline. In the benchmark model without endogenous productivity, output simply reverts to its initial steady state.

Even though the shock is not technological, the rise in μ_t^w generates a persistent decline in TFP and labor productivity and a persistent increase in the relative price of capital. The initial decline in measured TFP and labor productivity results mainly from variable factor utilization (in conjunction with overhead costs) and from the rise in the price markups.²⁴ Over time, there is a decline in true productivity relative to trend. In particular, the initial contraction in economic activity induces a drop in both R&D and the rate of technology adop-

²⁴ The formal definition of TFP that we adopt is the following: $TFP_t = Y_t A_{k,t}^{\psi\alpha(\theta-1)} / K_t^\alpha L_t^{1-\alpha}$, where ψ is the fraction of the variety externality in the creation of new capital not captured by the linear aggregation methodology used by the BEA to compute the capital stock. We calibrate ψ to 0.56 to match the long-run growth rate of TFP implied by the model to the postwar data. In case the variety externality was completely mismeasured, annual TFP growth would be 1.72 percent instead of the actual 1.4 percent. In addition, this definition assumes that the Bureau of Labor Statistics measures of TFP do not correct properly for capacity utilization.

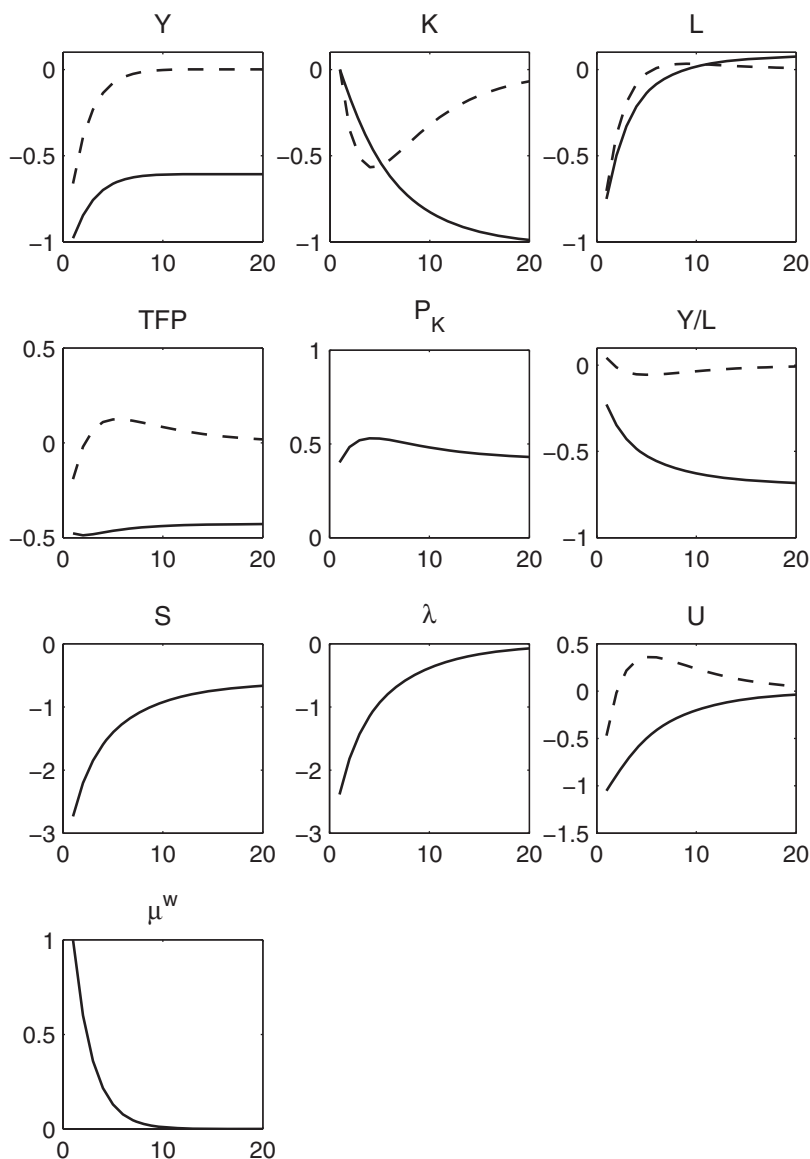


FIGURE 4. IMPULSE RESPONSE FUNCTIONS FOR A UNIT SHOCK TO WAGE MARKUP

tion.²⁵ The initial decline in R&D slows the rate of creation of new intermediate products in both the goods and capital sectors, ultimately leading

to a permanent drop relative to trend in TFP and labor productivity. The slowdown in the adoption rate leads to a near-term decline in disembodied productivity growth. Absent this endogenous slowdown in technology adoption, the decline in productivity following the drop in R&D would take much longer, assuming realistic diffusion lags.

²⁵ We are assuming that that R&D involves the same factor intensities as goods production. It is thus not as sensitive to wage behavior as in the Schumpeterian models of Philippe Aghion and Peter Howitt (1992) and Aghion and Gilles Saint-Paul (1991). Barlevy (2004) shows that if productivity shocks in the final goods sector are sufficiently persistent, R&D can be procyclical even in a Schumpeterian model.

Similarly, variable price markup behavior accounts for the initial rise in the relative price of capital, while the increase over time is due to a

TABLE 5—MODEL EVALUATION: STANDARD DEVIATIONS

Frequency Variable	Medium-term cycle 0–50		High-frequency component 0–8		Medium-frequency component 8–50	
	Data	Model	Data	Model	Data	Model
Output	3.87 (3.06–4.67)	4.18	2.20 (1.91–2.48)	2.22	3.15 (2.55–3.75)	3.46
TFP	2.39 (2.07–2.7)	2.42	1.25 (1.15–1.36)	1.03	2.01 (1.76–2.25)	2.15
Labor productivity	2.37 (2.04–2.7)	2.22	1.01 (0.96–1.07)	0.46	2.11 (1.83–2.4)	2.17
Hours	2.84 (2.29–3.09)	3.13	1.59 (1.42–1.74)	1.87	2.33 (1.88–2.48)	2.49
Relative price of capital	4.34 (3.12–5.39)	2.55	1.54 (1.27–1.49)	0.84	4.12 (3.05–5.14)	2.39
R&D	7.65 (4–8.99)	11.69	3.40 (2.65–4.13)	6.4	5.95 (3.64–8.26)	9.62
Investment	9.59 (4.5–14.68)	13.13	4.50 (3.34–5.66)	6.38	8.29 (3.64–12.94)	11.35
Consumption	2.20 (1.83–2.57)	2.45	0.80 (0.73–0.87)	1.07	2.01 (1.72–2.3)	2.17
Total markup	4.35 (3.27–5.43)	4.35	2.38 (2.01–2.74)	2.7	3.65 (2.74–3.56)	3.11
Capacity utilization	4.26 (3.28–5.24)	4.86	3.11 (2.38–3.84)	2.45	2.82 (2.3–3.32)	4.11

Note: For Tables 5–7, model moments are averages over 1,000 simulations of a sample size corresponding to the data.

decline in productivity in the capital goods sector relative to the consumption goods sector. Because investment initially falls proportionately more than output, the capital goods markup rises relative to the consumption goods markup, leading to a jump in the relative price of capital. The disproportionate drop in investment also implies that expected profits in the capital goods sector fall relative to the consumption goods sector. As a consequence, there is a relatively greater drop in both R&D and technology adoption in the capital goods sector, leading to a slowdown in the creation and adoption of specialized intermediate goods, relative to the consumption goods sector. This leads to a permanent increase in the relative price of capital.

We next explore how well the model produces medium-term fluctuations. We do so by comparing moments of artificial data generated by the model economy (driven by shocks to μ_t^w) with the unconditional moments of the actual data. Table 5 reports the standard deviations of the variables for the medium-term cycle and also for the respective high- and medium-frequency components. As we noted earlier, we

normalize the exogenous shock to μ_t^w so that the unconditional variance of the total markup over the medium term generated by the model exactly matches the data (4.35).

Overall, the model does reasonably well. For seven of the nine variables (other than the markup), the model standard deviations for the medium term are within the 95-percent confidence intervals for the data. For output and hours, the model closely matches both the overall variation in the data and the breakdown of the variation between the high and medium frequencies. Even though the exogenous driving force is not technological, the model captures reasonably well the variation in TFP over the medium term (2.42 versus 2.39), and does particularly well at the medium frequencies (2.15 versus 2.01). The model estimate is below the 95-percent confidence interval for TFP at the high frequency, but still is about 82 percent of the mean in the data. The model is within the confidence interval for the variation in labor productivity over the medium-term cycle and over the medium frequencies, in particular. It is too low at the high frequency. However, allowing for unobserved labor effort (which is likely

more important at the high than the medium frequencies) would improve performance on this dimension.

The two variables where the model does less well are the relative price of capital and R&D. The model generates roughly 60 percent of the variability in the relative price of capital observed in the data (2.55 versus 4.34) and about 50-percent greater volatility of R&D relative to the data (11.69 versus 7.65). In each case, however, measurement error in the data could be a factor. First, error in the quality adjustment for the Gordon relative price of capital series could be magnifying the measure of volatility. As we noted in Section I, the BEA measure of the relative price of capital, which employs a simpler method for quality adjustment, has a comovement with output over the medium term that is very similar to the Gordon series. However, the standard deviation of the BEA series over the medium term is only 2.59, which is close to the model's prediction. Since the quality adjustment in the Gordon series is the subject of some controversy, a case can be made that the true volatility lies somewhere in between the Gordon and BEA estimates, implying that the model may not be as far off on the relative price of capital as the Gordon numbers might suggest. Second, the discrepancy between the model and the data regarding R&D volatility may be in part due to the very imprecise measure of private R&D that is available. This measure, based on the NSF definition of R&D that basically includes expenditures to create new goods, is just a subset of total private R&D.²⁶ It is also likely to contain a significant overhead component, implying that variable R&D input is likely to be more volatile than measured total input. Another possibility is that the returns to R&D investments diminish more rapidly than our baseline calibration implies. Interestingly, when we reduce ρ , the elasticity of new intermediate products to R&D, from 0.9 to 0.7 (a number within the lower bound of Griliches's estimates), the volatility of R&D falls to the upper bound for the confidence interval of the data, with no significant effect on the other moments of the other variables.

Figure 5 portrays how well the model captures the cyclical comovements of variables in the data. The top-left panel plots the autocorrelation of output over the medium-term cycle, at lags and leads of ten years. The other panels plot the cross correlation of a particular variable with output at time t . In each case, the dashed line represents the data and the solid line the model. The two dotted lines are 95-percent confidence intervals around the data. For most, the model matches the comovements in the data: for eight of the ten variables the model cross correlations lie mainly within the respective confidence intervals. The notable exceptions again are the relative price of capital and R&D. The model does capture both the contemporaneous and lagged comovement of each of these two variables with output. In each case, however, it implies a smaller lead over output than the data suggest. In the conclusion we discuss some ways to improve the model's ability to account for both the timing and volatility of each of these variables.

It is also worth emphasizing that the model captures the variation of output at low frequencies relative to the high and medium frequencies. The ratio of the standard deviation of output growth at frequencies below 50 years to the standard deviation of output growth at frequencies between 0 and 50 years that the model generates is 0.08, exactly what the data suggest, as we reported in Section I.

Finally, though we do not report the results here, in the working paper version of this paper we show that all the key features of the model, R&D, endogenous technology adoption, and entry and exit are important for explaining the data. Eliminating any one of these features leads to a significant deterioration in model performance.

An Exogenous Technology Shock Model.—We now consider how well a conventional exogenous productivity framework can account for the facts. We thus eliminate endogenous productivity and endogenous markup movements. What is left is an RBC model with variable factor utilization. We keep this latter feature because it improves the overall model performance, as noted by King and Rebelo (1999). Because we want the model to be consistent with the evidence on embodied as well as disembodied technological change, we also intro-

²⁶ See, e.g., Comin (2004) and Comin and Sunil Mulani (2005).

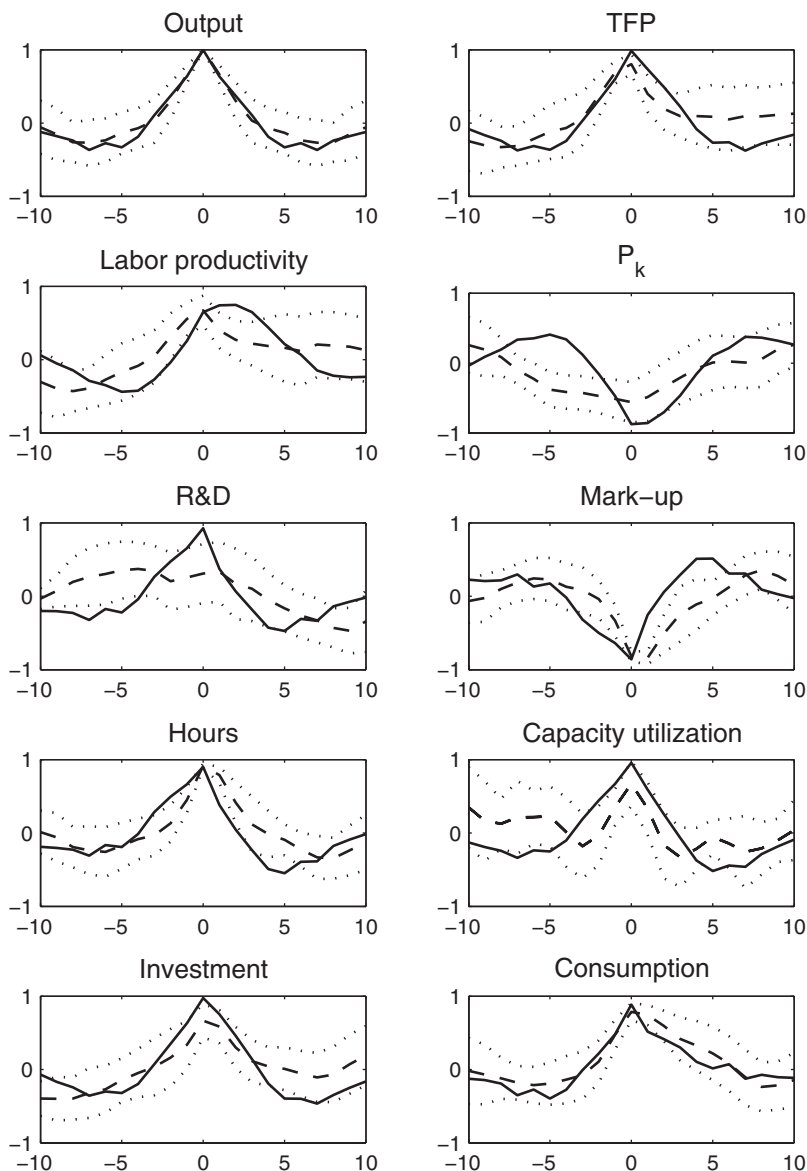


FIGURE 5. CROSS CORRELATION WITH OUTPUT: MODEL (SOLID), DATA (DASHED), AND 95-PERCENT CONFIDENCE BANDS (DOTTED)

duce a shock to the production of new capital goods. In particular, we make the creation of new capital goods a multiple of the level of investment, where the multiple is an exogenous random variable. The inverse of this multiple is then the relative price of capital in units of consumption goods. Where relevant, we use the

same parameter values as in our baseline model. One exception is that to enhance the ability of the model to account for hours volatility, we raise the Frisch labor supply elasticity from unity to two, a value recommended by Thomas F. Cooley and Edward C. Prescott (1995).

We consider three types of experiments: (a) a

TABLE 6—TWO-SECTOR RBC MODEL: STANDARD DEVIATIONS (0–50)

Variable	Data	TFP shock	$\rho(P_k) = 0.99$		$\rho(P_k) = 0.90$	
			P_k shock	Both shocks	P_k shock	Both shocks
Output	3.87 (3.06–4.67)	3.20	0.53	2.64	2.77	3.86
TFP	2.39 (2.07–2.7)	2.39	1.42	2.39	1.76	2.67
Labor productivity	2.37 (2.04–2.7)	2.22	1.52	2.35	1.45	2.35
Hours	2.84 (2.29–3.09)	1.22	1.12	1.5	1.58	1.88
Relative price of capital	4.34 (3.12–5.39)	—	4.34	4.34	4.34	4.34
Investment	9.59 (4.5–14.68)	10.00	2.63	8.48	14.76	16.94
Consumption	2.20 (1.83–2.57)	1.84	2.06	2.55	1.07	1.87
Capacity utilization	4.26 (3.28–5.24)	2.83	4.26	4.85	5.27	5.74

shock to TFP; (b) a shock to the relative price of capital; and (c) simultaneous shocks to the relative price of capital. We begin with an exogenous shock to TFP. We set the auto-correlation of the shock at 0.88, which is consistent with value that King and Rebelo (1999) use for quarterly data. This results in a serial correlation of TFP over the medium-term cycle generated by the model of 0.63, which is slightly below the actual number of 0.7, though well within the 95-percent confidence interval. Because the model performance would deteriorate if we were to use a more persistent process, we stick with the value proposed by King and Rebelo (1999). We then adjust the standard deviation of the shock to match the unconditional standard deviation of TFP over the medium-term cycle. Since the measure of TFP in the data is not corrected for utilization, we adjust the exogenous process for TFP to match the serial correlation and variance of the model Solow residual (which is similarly not corrected for utilization).

Table 6 reports the results. Overall, the simple TFP shock model does about as well at capturing volatility over the medium-term cycle as it does for the conventional high-frequency cycle. The model accounts for roughly 80 percent of the volatility of output. The volatilities of the other variables relative to output are reasonably in line with the data. The main exception involves hours. As in the typical high-

frequency analysis, the model generates only about half the relative volatility of hours that appears in the data. From our vantage, however, a significant shortcoming of the TFP shock model is that, by construction, it cannot account for the cyclical pattern of embodied technological change, which we have argued is a key feature of the medium-term cycle. We thus next consider a shock to the relative price of capital.

As with the TFP shock, we attempt to feed in an exogenous price that matches the autocorrelation and standard deviation of the filtered data for the medium-term cycle. Because the relative price of capital is very persistent, to come close to the autocorrelation in the data it is necessary to feed in a near unit root process.²⁷ (Our numerical simulation procedure requires stationary forcing processes, thus ruling out a pure unit root shock.) In particular, we set the autocorrelation of the shock, $\rho(P_k)$, at 0.99. Having a highly persistent process for the relative price of

²⁷ Greenwood et al. (2000) also consider a model of how exogenous shocks to the relative price of capital might generate business fluctuations. They differ by considering only equipment prices and also by adding investment adjustment costs, where the latter is chosen so that the model variance matches the data. Also, they just focus on the high frequency. Finally, they use a forcing process for the relative price of capital that is much less persistent than our data for the medium term would suggest is reasonable.

TABLE 7—FIRST-ORDER AUTOCORRELATIONS (0–50)

Variable	Data	Baseline model	RBC, TFP shock	RBC, $\rho(P_k) = 0.99$		RBC, $\rho(P_k) = 0.90$	
				P_k shock	Both shocks	P_k shock	Both shocks
TFP	0.70 (0.51–0.89)	0.75	0.65	0.82	0.71	0.65	0.65
Relative price of capital	0.91 (0.78–1.04)	0.83	—	0.74	0.74	0.70	0.70

capital is also plausible on a priori theoretical grounds: that is, it seems reasonable to presume that embodied technological changes are permanent, implying that a (near) unit root process with drift is appropriate.²⁸

As the third row of Table 6 shows, the model with a highly persistent shock to the relative price of capital does not fare well. It has only a minimal impact on output volatility for two reasons. First, in contrast to a disembodied technology shock which directly affects the production of all output, this shock has a direct effect only on the production of new capital goods. Second, because the shock is highly persistent, it creates little incentive for intertemporal substitution of hours.

We next allow for simultaneous shocks to the two technology variables. We keep the autocorrelations of each shock unchanged, but adjust the variance of each shock innovation to match the corresponding unconditional standard deviations of the filtered data. As the fourth column of Table 6 shows, there is a clear improvement in model performance, relative to the case with just shocks to the relative price of capital. Nonetheless, the two-shock model explains less of the overall output volatility than the simple TFP shock case, slightly less than 70 percent, as compared to roughly 80 percent.²⁹ Otherwise, the relative volatilities of the other variables are largely similar to the pure TFP shock case. One exception is that utilization is significantly more volatile, due to a direct impact of movements on the relative price of capital on the incentive to adjust capital utilization.

The reason the two-shock model does not explain more the 70 percent of the volatility of output involves the high persistence of the shock to the relative price of capital, which greatly minimizes intertemporal substitution.

To confirm, we redid the experiment, reducing $\rho(P_k)$ from 0.99 to 0.90. As the last column of Table 6 shows, in this case the exogenous technology model exactly matches the volatility of output. There are some shortcomings, however. The enhanced intertemporal substitution leads to investment volatility that is far greater than in the data (16.94 versus 9.59).³⁰ On the other hand, hours volatility remains too low relative to the data (1.88 versus 2.84).

As we noted earlier, however, deviating significantly from a (near) unit root process for exogenous embodied technological change is not appealing from a theoretical perspective. It is also not appealing from an empirical one. In Table 7 we show the autocorrelations for TFP and the relative price of capital over the medium-term cycle for the different cases of the exogenous technology shock that we considered. For comparison, we also report these autocorrelations for our baseline endogenous productivity model. For all the exogenous technology shock cases, the model autocorrelation for the relative price of capital is low relative to the data. For

²⁸ A similar argument can be made regarding disembodied technology shocks; however, convention in this literature quickly evolved into considering stationary processes that are not close to being unit root. This complicates the interpretation of what these shocks actually are.

²⁹ Shocks to the relative price of capital affect measured TFP by affecting utilization, which in the two-shock case reduces that size of the embodied technology shock required to match TFP. It is for this reason that output volatility is lower in the two-shock case than in the case with just shocks to disembodied productivity.

³⁰ One way to dampen investment would be to add adjustment costs to investment. It is unclear, though, how important these costs may be at the annual frequency. Investment adjustment costs also significantly dampen the response of hours to exogenous productivity shocks (see, e.g., Neville and Ramey, 2002).

the highly persistent forcing process, $\rho(P_k) = 0.99$, it is just below the lower bound for the 95-percent confidence interval (0.74 versus 0.78). For the less persistent forcing process, $\rho(P_k) = 0.90$, it moves farther away from the target, dropping to 0.70.

By contrast, as column 2 of Table 1 shows, our baseline endogenous productivity model succeeds at capturing the high degree of persistence in the relative price of capital at the medium frequencies (0.87 versus 0.91 in the data). It similarly captures the persistence of TFP (0.75 versus 0.70).

IV. Concluding Remarks

We have explored the idea that the postwar business cycle should not be associated with only high-frequency variation in output (as generated by the conventional business cycle filter), but rather should also include the medium-frequency oscillations between periods of robust growth and periods of relative stagnation. We first construct measures of the cycle based on this notion and demonstrate that the resulting fluctuations are considerably more volatile and persistent than are the usual high-frequency measures. These results suggest that postwar business fluctuations may be more important phenomena than conventional analysis suggests.

The medium-term cycle we have identified features significant procyclical movements in both embodied and disembodied technological change. These facts, among others, motivated us to approach modeling the medium-term cycle by modifying a reasonably conventional business cycle framework to allow for R&D, technology adoption, and variation in markups. A virtue of our approach is that we are able to fully endogenize cyclical productivity dynamics, including both embodied and disembodied dynamics. In addition, our endogenous productivity mechanism offers an avenue through which nontechnological disturbances at the high frequency can have sustained effects on productivity over the medium term. As with many recent quantitative macroeconomic frameworks, the volatility of output at the high frequency is associated with countercyclical markup variation. Within our framework, however, these high-frequency fluctuations influence the pace of both R&D and adoption,

in turn generating subsequent movements in productivity. The net effect is a propagation mechanism that generates the kind of persistent business cycles that our filter uncovers.

There are, however, several places where the model performance could be improved. As in the data, the model generates a countercyclical movement in the relative price of capital over the medium term and procyclical movement in R&D. The model, however, generates only about 60 percent of the variability in the relative price of capital that is observed in the data. As we noted, it may be that the conventionally used Gordon series overstates the volatility. The official BEA series has similar cyclical properties over the medium term, but an overall volatility that is in line with our model estimates. Another possible solution might be to introduce a shock to the innovation process for the production of new capital goods, in order to add an independent source of volatility to the relative price of capital. Conversely, the model R&D variability is about 50 percent greater than what we observe in the data. As we discussed earlier, here it may be important to try to gather improved measures of R&D that are not only more comprehensive, but also can separate variable from overhead costs. Another possibility is that the returns to R&D investment may diminish more rapidly than our baseline calibration implies. Although we do not report the results here, by using an elasticity of new products with respect to R&D that is in the lower range of Griliches's estimates, it is possible to bring R&D volatility closer in line with the evidence, without causing the performance of the model on other dimensions to deteriorate.

We also show that, under certain circumstances, a pure exogenous productivity framework can explain a good fraction of the variation in the data. To capture the movements in both embodied and disembodied technological change, it is of course necessary to have shocks to both total factor productivity and the relative price of capital. To generate sufficiently large output and employment fluctuations, however, it appears necessary to allow for exogenous shocks to the relative price of capital that are less persistent than the data suggest. Beyond this empirical consideration, as we noted, there are also good theoretical reasons to believe that movements in embodied productivity should

not be transitory. On the other hand, it is surely worth exploring exogenous productivity models more thoroughly than we have done in this paper.

Finally, to minimize complexity, we modeled markup variation at the high frequency in a very simple way. By contrast, as we noted in the introduction, recent quantitative macroeconomic frameworks endogenize markup variation by introducing money together with nominal price and wage rigidities. It would be straightforward to enrich our model along this dimension. Doing so would yield a rather different perspective on what's at stake over the performance of monetary policy. Within these more conventional frameworks, medium-frequency variation is driven mainly by exogenous shocks. With our endogenous productivity mechanism, however, (large) high-frequency fluctuations can have very persistent effects. To the extent monetary policy influences this high frequency variation, it will have ramifications for medium-frequency dynamics.

REFERENCES

- Aghion, Philippe and Howitt, Peter.** "A Model of Growth through Creative Destruction." *Econometrica*, 1992, 60(2), pp. 323–51.
- Aghion, Philippe and Saint-Paul, Gilles.** "Uncovering Some Causal Relationships between Productivity Growth and the Structure of Economic Fluctuations: A Tentative Survey." National Bureau of Economic Research, Inc., NBER Working Papers: No. 4603, 1991.
- Barlevy, Gadi.** "On the Timing of Innovation in Stochastic Schumpeterian Growth Models." National Bureau of Economic Research, Inc., NBER Working Papers: No. 10741, 2004.
- Basu, Susanto.** "Procyclical Productivity: Increasing Returns or Cyclical Utilization?" *Quarterly Journal of Economics*, 1996, 111(3), pp. 719–51.
- Basu, Susanto and Fernald, John G.** "Returns to Scale in U.S. Production: Estimates and Implications." *Journal of Political Economy*, 1997, 105(2), pp. 249–83.
- Baxter, Marianne and King, Robert G.** "Measuring Business Cycles: Approximate Band-Pass Filters for Economic Time Series." *Review of Economics and Statistics*, 1999, 81(4), pp. 575–93.
- Baxter, Marianne and Farr, Dorsey.** "Variable Factor Utilization and International Business Cycles." National Bureau of Economic Research, Inc., NBER Working Papers: No. 8392, 2001.
- Blanchard, Olivier J.** "The Medium Run." *Brookings Papers on Economic Activity*, 1997, 0(2), pp. 89–141.
- Burnside, Craig; Eichenbaum, Martin and Rebelo, Sergio.** "Capital Utilization and Returns to Scale," in Ben S. Bernanke and Julio J. Rotemberg, eds., *NBER macroeconomics annual 1995*. Cambridge, MA: MIT Press, 1995, pp. 67–110.
- Caballero, Ricardo J. and Hammour, Mohamad L.** "Jobless Growth: Appropriability, Factor Substitution, and Unemployment." *Carnegie-Rochester Conference Series on Public Policy*, 1998, 48(0), pp. 51–94.
- Christiano, Lawrence J.; Eichenbaum, Martin and Evans, Charles L.** "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy." *Journal of Political Economy*, 2005, 113(1), pp. 1–45.
- Christiano, Lawrence J. and Fitzgerald, Terry J.** "The Band Pass Filter." *International Economic Review*, 2003, 44(2), pp. 435–65.
- Comin, Diego.** "R&D: A Small Contribution to Productivity Growth." *Journal of Economic Growth*, 2004, 9(4), pp. 391–421.
- Comin, Diego and Gertler, Mark.** "Procyclical Speed of Technological Diffusion." Unpublished Paper, 2004.
- Comin, Diego and Mulani, Sunil.** "A Theory of Growth and Volatility at the Aggregate and Firm Level." National Bureau of Economic Research, Inc., NBER Working Papers: No. 11503, 2005.
- Cooley, Thomas F. and Prescott, Edward C.** "Economic Growth and Business Cycles," in Thomas F. Cooley, *Frontiers of business cycle research*. Princeton: Princeton University Press, 1995, pp. 1–38.
- Cummins, Jason G. and Violante, Giovanni L.** "Investment-Specific Technical Change in the United States (1947–2000): Measurement and Macroeconomic Consequences." *Review of Economic Dynamics*, 2002, 5(2), pp. 243–84.
- Davies, Stephen.** *The diffusion of process innovations*. Cambridge: Cambridge University Press, 1979.
- Evans, George W.; Honkapohja, Seppo and Ro-**

- mer, Paul.** "Growth Cycles." *American Economic Review*, 1998, 88(3), pp. 495–515.
- Galí, Jordi; Gertler, Mark and Lopez-Salido, J. David.** "Markups, Gaps and the Welfare Costs of Business Cycles." National Bureau of Economic Research, Inc., NBER Working Paper: No. 8850, 2002.
- Galí, Jordi and Zilibotti, Fabrizio.** "Endogenous Growth and Poverty Traps in a Cournotian Model." *Annales d'économie et de statistique*, 1995, 0(37–38), pp. 197–213.
- Gordon, Robert J.** *The measurement of durable goods prices*. National Bureau of Economic Research Monograph Series. Chicago: University of Chicago Press, 1990.
- Greenwood, Jeremy; Hercowitz, Zvi and Huffman, Gregory W.** "Investment, Capacity Utilization, and the Real Business Cycle." *American Economic Review*, 1988, 78(3), pp. 402–17.
- Greenwood, Jeremy; Hercowitz, Zvi and Krusell, Per.** "The Role of Investment-Specific Technological Change in the Business Cycle." *European Economic Review*, 2000, 44(1), pp. 91–115.
- Griliches, Zvi.** "Patent Statistics as Economic Indicators: A Survey." *Journal of Economic Literature*, 1990, 28(4), pp. 1661–707.
- Hall, Robert E.** "Macroeconomic Fluctuations and the Allocation of Time." *Journal of Labor Economics*, 1997, 15(1), pp. S223–50.
- Hobijn, Bart.** "Is Equipment Price Deflation a Statistical Artifact?" Federal Reserve Bank of New York, Staff Report No. 139, 2001.
- ▶ **Hodrick, Robert J. and Prescott, Edward.** "Post-War U.S. Business Cycles: An Empirical Investigation." *Journal of Money, Credit and Banking*, 1997, 29(1), pp. 1–16.
- King, Robert G. and Rebelo, Sergio T.** "Resuscitating Real Business Cycles," in John B. Taylor and Michael Woodford, eds., *Handbook of macroeconomics*. Vol. 1B. Amsterdam: Elsevier Science, North-Holland, 1999, pp. 927–1007.
- Mansfield, Edward.** "The Diffusion of Industrial Robots in Japan and the United States." *Research Policy*, 1989, 18(4), pp. 183–92.
- Neville, Francis and Ramey, Valerie.** "Is the Technology-driven Real Business Cycle Hypothesis Dead? Shocks and Aggregate Fluctuations Revisited." *Journal of Monetary Economics* (forthcoming).
- Rebelo, Sergio.** "Real Business Cycle Models: Past, Present, and Future." National Bureau of Economic Research, Inc., NBER Working Papers: No. 11401, 2005.
- Romer, Paul M.** "Endogenous Technological Change." *Journal of Political Economy*, 1990, 98(5), pp. S71–102.
- Rotemberg, Julio J.** "Stochastic Technical Progress, Smooth Trends, and Nearly Distinct Business Cycles." *American Economic Review*, 2003, 93(5), pp. 1543–59.
- Rotemberg, Julio J. and Woodford, Michael.** "Dynamic General Equilibrium Models with Imperfectly Competitive Product Markets," in Thomas F. Cooley, ed., *Frontiers of business cycle research*. Princeton: Princeton University Press, 1995, pp. 243–93.
- Smets, Frank and Wouters, Raf.** "Comparing Shocks and Frictions in US and Euro Area Business Cycles: A Bayesian DSGE Approach." *Journal of Applied Econometrics*, 2005, 20(2), pp. 161–83.
- Solow, Robert M.** "Toward a Macroeconomics of the Medium Run." *Journal of Economic Perspectives*, 2000, 14(1), pp. 151–58.