Volume Title: NBER Macroeconomics Annual 1998, volume 13

Volume Author/Editor: Ben S. Bernanke and Julio Rotemberg, editors

Volume Publisher: MIT Press

Volume ISBN: 0-262-52271-3

Volume URL: http://www.nber.org/books/bern99-1

Publication Date: January 1999

Chapter Title: What Do Technology Shocks Do?

Chapter Author: John Shea

Chapter URL: http://www.nber.org/chapters/c11249

Chapter pages in book: (p. 275 - 322)

John Shea
UNIVERSITY OF MARYLAND AND NBER

What Do Technology Shocks Do?

1. *Introduction*

The real business cycle (RBC) approach to short-run fluctuations, pioneered by Kydland and Prescott (KP) (1982) and Long and Plosser (LP) (1983), has dominated the academic business-cycle literature over the last decade and a half. KP and LP were seminal in several respects. First, they reintroduced the Schumpeterian idea that stochastic technological progress could generate business cycles. Second, they argued that one could explain fluctuations using a frictionless neoclassical framework in which business cycles are optimal and therefore require no smoothing by policymakers. Third, they argued that business cycles could and should be explained using dynamic stochastic general equilibrium models in which preferences and production are explicitly spelled out in a way consistent with microeconomic first principles, such as optimizing behavior.

The RBC literature has broadened considerably since KP and LP. Recent research has introduced frictions such as imperfect competition (e.g., Rotemberg and Woodford, 1995), increasing returns to scale (e.g., Farmer and Guo, 1994), and price stickiness (e.g., Kimball, 1995), as well as alternative sources of shocks, such as government spending (e.g., Christiano and Eichenbaum, 1992), monetary policy (e.g., Christiano and Eichenbaum, 1995), and animal spirits (e.g., Schmitt-Grohe, 1997). The idea that business cycles should be analyzed using explicit dynamic stochastic general equilibrium models seems destined to be the main lasting contribution of KP and LP's work.

Meanwhile, the profession has largely ignored the empirical question of what role technology shocks actually play in business cycles. I believe that this is unfortunate, for four reasons. First, the idea that new prod-

I thank Kortum and Susanto Basu for providing their data, and the editors and participants, as well as seminar participants at Brown, for helpful comments.

ucts and processes are introduced at a time-varying rate is inherently plausible, at least at the disaggregated industry level. Second, much recent research exploring the effects of frictions on business-cycle propagation still assumes that cycles are driven by technology shocks (e.g., Cogley and Nason, 1995; Horvath, 1997; Carlstrom and Fuerst, 1997). It would be useful to know if this modeling strategy has any empirical foundation. Third, while few would argue any more that technology shocks are the *only* source of business cycles, it would still be useful to know if technology shocks can explain *some* part of fluctuations, particularly given that monetary, oil price, and other observable shocks seem unable to account for a large fraction of observed cyclical variation in output (Cochrane, 1994). Finally, even if technology shocks are not responsible for a large share of volatility, the response of the economy to technology could help distinguish between competing views of the economy's propagation mechanisms. In the baseline one-sector flexible-price RBC model, technology shocks shift out the production possibilities frontier, inducing short-run increases in investment, labor, and materials. In multisector models, industry technology shocks reduce input prices to downstream sectors, inducing increases in downstream input and output. Meanwhile, Galí (1996) and Basu, Fernald, and Kimball (1997) demonstrate that favorable technology shocks may reduce input use in the short run if prices are sticky; intuitively, if prices do not fall, output will be unchanged and inputs must fall to accommodate improved total factor productivity (TFP). Thus, one can potentially distinguish between sticky and flexible price models by examining whether technology shocks increase or decrease input use.

To date, the empirical case for technology has largely been made indirectly, by showing that plausibly calibrated models driven by technology shocks can produce realistic patterns of volatility and comovement. Of course, these quantitative exercises, while informative, do not tell us what technology shocks actually do. Two pieces of more direct evidence are that measured TFP is procyclical and that aggregate output potentially has a unit root, suggesting that at least some output shocks are permanent. However, it is now well known that neither of these facts proves that technology is important to business cycles. Observable nontechnology shocks cause procyclical movements in TFP, consistent with imperfect competition, increasing returns to scale, procyclical factor utilization, or procyclical reallocation of factors to high productivity sectors (e.g., Hall, 1988; Evans, 1992; Burnside, Eichenbaum, and Rebelo, 1995; Basu and Fernald, 1997). Meanwhile, demand shocks can have permanent effects on output in endogenous growth models (e.g., Stadler, 1990); and in any case, a unit root is consistent with transitory shocks driving an arbitrarily large fraction of short-run variation (Quah, 1989).

This paper takes a more direct approach to assessing what technology shocks do, an approach inspired by the large literature estimating the impact of monetary policy shocks on the economy (e.g., Christiano, Eichenbaum, and Evans, 1998). Using annual panel data for 19 U.S. manufacturing industries from 1959 to 1991, I employ vector autoregressions (VARs) to document the dynamic impact of shocks to two observable indicators of technological change: research and development (R&D) spending, and patent applications. R&D measures the amount of input devoted to innovative activity, while patent applications measure inventive output. Previous studies (e.g., Griliches and Lichtenberg, 1984; Lichtenberg and Siegel, 1991; Scherer, 1993), as well as results reported below, suggest that variation in R&D and patenting is related to long-run variation in productivity growth across firms and industries. Moreover, industry-level R&D and patents display nontrivial short-run fluctuations, as can be seen in Figure 1, which plots log real R&D and log patent applications by industry of manufacture and use for the U.S. aerospace industry. If technological progress is truly stochastic, then fluctuations in R&D should in part reflect variation in the perceived marginal product of knowledge, while fluctuations in patents should in part reflect shocks to the success of research activity. I use these fluctuations to estimate how a typical industry's inputs and TFP respond over time to technology shocks, and to quantify the share of industry volatility due to technology shocks. I estimate the impact of both own technology shocks and technology shocks in upstream input-supplying industries.

To be sure, fluctuations in R&D and patent applications may not be due to technology shocks alone. Griliches (1989), for instance, argues that patenting fluctuations in the U.S. are in part responses to factors such as changes in patent law and changes in the efficiency and resources of the U.S. Patent Office. Both R&D and patent applications, meanwhile, are a type of investment, and as such they may respond endogenously to output shocks, either because of financial-market constraints or because current shocks are positively correlated with the future marginal product of capital. My preferred VAR specifications address these concerns by including time dummies in the regressions and by placing the technological indicators last in the Choleski ordering used to decompose the VAR innovations into orthogonal components. The time dummies remove fluctuations in R&D and patent applications due to aggregate factors unrelated to true technological progress, such as changes in the number of patent examiners, provided that these factors affect all industries equally. My impulse responses therefore measure the

Figure 1 TECHNOLOGY INDICATORS IN THE AEROSPACE INDUSTRY

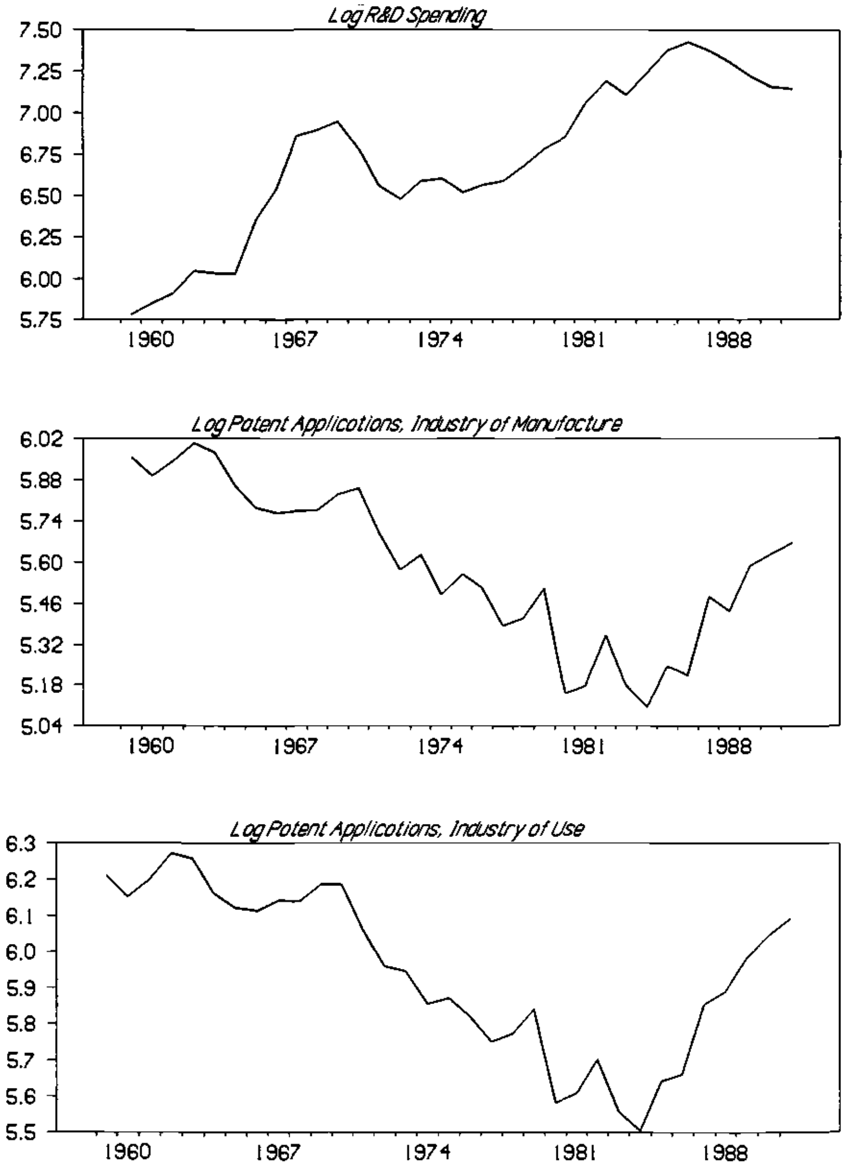

impact of industry-specific technology shocks on industry-specific fluctuations in inputs and TFP, while my variance decompositions estimate the contribution of technology shocks to idiosyncratic industry fluctuations. Placing technology last in the ordering, meanwhile, defines technology shocks as the component of R&D or patenting orthogonal to both lagged technology and lagged and contemporaneous inputs and TFP. Empirically, innovations to industry output are positively correlated with innovations to both R&D and patent applications; placing technology last assumes that this contemporaneous comovement reflects an accelerator mechanism running from industry activity to technology, rather than an instantaneous impact of technology shocks on output. This assumption seems inherently plausible given the likely lags between R&D spending, invention, and diffusion of a new technology (Gort and Klepper, 1982).

My main empirical findings are as follows. First, favorable technology shocks—increases in the orthogonal components of R&D and patents—tend to increase input use, especially labor, in the short run, but to reduce inputs in the long run. Second, technology improvements tend to encourage substitution towards capital relative to materials and labor, as well as substitution towards nonproduction labor relative to production labor. These results are consistent with recent cross-sectional studies establishing a complementary long-run relationship between technological change and equipment (Delong and Summers, 1991) and skilled labor (Berman, Bound, and Griliches, 1994). Third, favorable technology shocks do not significantly increase measured TFP at any horizon, and indeed in some cases reduce TFP. This suggests that procyclical movements in TFP have little to do with the introduction of new products and processes. Fourth, technology shocks explain only a small share of idiosyncratic industry volatility of inputs and TFP at business-cycle horizons. This result is bad news for technology-shock-driven models, particularly given that industry-specific technology shocks are likely to explain industry-specific volatility better than aggregate volatility (Horvath, 1997). However, my results could be consistent with models in which technology contributes to low-frequency fluctuations (e.g., Jovanovic and Lach, 1997); or with models in which the important "real" shocks come from strikes, weather, cartel behavior, and so on; or with models in which "technology shocks" are not due to stochastic scientific and engineering developments, but to stochastic movements in management techniques or industrial organization that cause a given set of inputs to be more or less efficient. Finally, I find that technology improvements are more likely to raise TFP and reduce prices in industries characterized by process innovations than in industries dominated by product innovations. This suggests that my failure to find strong effects of technology on TFP may be due in part to the failure of available price data to capture productivity gains caused by quality improvements and new product introductions.

Two other recent papers (Galí, 1996; Basu, Fernald, and Kimball, 1997) also investigate the short-run impact of technology shocks, in both cases using aggregate data. Galí estimates a structural vector autoregression for labor productivity and labor input in the United States, identifying technology shocks by assuming that only technology affects long-run productivity. Basu, Fernald, and Kimball correct industry-level TFP for variations due to increasing returns to scale, imperfect competition, and cyclical factor utilization, and then measure aggregate technology as an appropriately weighted average of sectoral technology. Interestingly, Galí (1996) and Basu, Fernald, and Kimball (1997) both find that favorable technology shocks reduce input use in the short run, consistent with sticky prices but contrary to my results.

These two papers represent a distinct advance over existing literature. Nevertheless, one might disagree with their methodologies for measuring technological change. Galí's approach rests heavily on the assumption that demand shocks cannot affect productivity in the long run. This assumption is inconsistent both with endogenous growth models and with models in which recessions cleanse the economy by wiping out low-productivity firms (e.g., Caballero and Hammour, 1994, 1996). Cleansing models, in particular, predict that favorable demand shocks will reduce long-run productivity, and Galí himself has in the past argued for such an interpretation of the data (Galí and Hammour, 1992). Interestingly, my impulse response functions suggest that input innovations lead to short-run increases in TFP, consistent with increasing returns or procyclical utilization, but long-run decreases in TFP, consistent with cleansing models.

Basu, Fernald, and Kimball's approach does not rely on long-run restrictions. It does, however, rely on the idea that TFP fluctuations are valid measures of stochastic technological progress at the two-digit industry level, once one corrects for increasing returns, imperfect competition, and cyclical factor utilization. This idea seems plausible, but it is not necessarily true, given that fluctuations in "corrected" sectoral TFP could still be due to nontechnology sources such as measurement error, within-sector factor reallocations, or inadequate corrections for increasing returns or cyclical utilization. Basu, Fernald, and Kimball's methodology would be more convincing if their corrected measure of technology could be linked to some sort of outside measure of technological progress, such as anecdotal evidence on the timing of particular technical changes in particular industries.

The remainder of the paper proceeds as follows. Section 2 describes the data. Section 3 examines long-run and contemporaneous relationships between technological progress and my measures of innovative activity, largely to connect my work to previous studies. Section 4 presents evidence from VARs, and Section 5 concludes.

2. *Data Description*

My goal is to examine the time-series interactions between measures of technological change, such as patents and R&D, and measures of economic activity. Ideally, I would estimate these interactions using aggregate data for a single country, following the empirical literature on monetary policy. However, this approach is not feasible in my case. The only readily available data for patents and R&D are annual rather than quarterly or monthly, implying short aggregate time series. Even if higher-frequency data could be constructed, it is not clear that they would be useful, since the impact of technological change on the economy is likely to operate at a somewhat lower frequency than the impact of monetary shocks. To obtain sufficient degrees of freedom to estimate the impact of technology shocks with reasonable precision, I use panel data for 19 manufacturing industries covering 1959–1991, exploiting the fact that technological developments are not perfectly synchronized across industries. An alternative, worth pursuing in future work, would be to use annual aggregate data for a panel of countries, or for panels of both countries and industries.

Data on R&D by industry are taken from the National Science Foundation's annual survey of U.S. firms. I examine only company-financed R&D. Previous research using cross sections of industries and firms (e.g., Terleckyj, 1975; Lichtenberg and Siegel, 1991) has shown that long-run productivity growth is related to company-financed R&D, but not to federally financed R&D, suggesting that public R&D dollars are spent inefficiently or that they are spent in areas, such as defense or space exploration, where productivity measurement is difficult. I convert nominal R&D to 1991 dollars using the GDP deflator, then convert real R&D flows to an R&D capital stock, following Griliches (1973) and most other subsequent research. I employ a linear capital accumulation equation, assuming a 15% annual depreciation rate and setting the 1959 stock equal to the 1959 flow divided by 0.15 plus the industry's average R&D growth over the sample period; these assumptions are standard in recent literature (e.g., Lach, 1995; Keller, 1997). The empirical results are similar if I use real R&D flows instead of R&D stocks. As a timing convention, I include R&D spending in year t in the R&D stock for year

Table 1 SAMPLE MEANS

Industry	*R&D*	*R&D Growth*	*Manuf. Patents*	*Manuf. Patent Growth*	*Use Patents*	*Use Patent Growth*
Food (SIC 20)	879.1	4.62	311.2	0.63	1085	0.06
Textiles (SIC 22–23)	185.5	3.57	620.9	1.17	994	−0.81
Lumber (SIC 24–25)	152.0	4.81	605.9	0.08	597	−0.05
Paper (SIC 26)	611.7	5.53	482.4	0.04	490	0.08
Industrial chemicals (SIC 281–282, 286)	3211.7	2.80	3758.8	0.75	2518	0.47
Drugs (SIC 283)	2629.4	7.75	825.5	5.90	1100	2.73
Other chemicals (other SIC 28)	1053.0	5.27	2517.4	0.54	1261	0.70
Petroleum (SIC 29)	1812.4	2.82	1745.5	−1.30	1659	0.22
Rubber (SIC 30)	693.2	4.01	1586.0	1.03	1348	1.45
Stone (SIC 32)	574.5	3.05	506.8	1.60	557	0.76
Metals (SIC 33)	835.9	1.51	373.1	0.30	795	0.06
Metal prods. (SIC 34)	684.1	2.06	3737.6	0.18	1979	0.24
Computers (SIC 357)	5172.6	6.66	1114.3	2.70	1333	3.09
Other nonelec. equip. (other SIC 35)	2102.3	5.08	10966.1	−0.15	4084	−0.33
Electronics & commun. equip. (SIC 366–367)	5018.4	6.65	5629.4	1.51	4456	1.76
Other electric equip. (Other SIC 36)	2043.7	0.99	4154.1	0.41	2779	0.43
Aerospace (SIC 372, 376)	4022.4	4.81	276.9	−1.22	392	−0.77
Autos & other transp. equip. (SIC 37)	5701.8	4.37	1972.1	−0.28	2787	−0.09
Instruments (SIC 38)	3100.3	7.42	3626.7	2.33	1268	1.59

t, so that I can interpret the correlations between R&D and other variables as reflecting a contemporaneous response of R&D to industry activity. I use data for 19 manufacturing industries; these are listed in Table 1 along with sample means of real R&D flows in millions of dollars and the growth rate of the R&D stock. The largest flows of company R&D are found in automobiles, electronics, and computers; the fastest-growing R&D stocks are in drugs, electronics, computers, and instruments. Note that my baseline sample omits nonmanufacturing industries as well as some manufacturing industries (tobacco, printing and publishing, leather, and miscellaneous manufacturing) whose R&D data are lumped together by the NSF. The share of overall R&D accounted for by these sectors is trivial for most of my sample period.

I must mention two problems with these data. First, to avoid disclosure of individual firms' operations, the NSF suppresses some industry-year observations. In virtually all such cases, the NSF suppresses either company-financed or total (including federally financed) R&D, but not both, so that I can interpolate gaps in company R&D using growth of total R&D. Second, the NSF data are collected at the company level. All R&D spending performed by a company is assigned to the industry in which the company had the most sales, even if part of the R&D was conducted in establishments belonging to another industry. Given that R&D is typically performed in large conglomerated firms, the assignment of R&D to particular industries is presumably subject to error. Particularly troubling is the fact that a given firm's industry classification can change over time as its pattern of sales changes, creating the possibility of large movements in measured industry-level R&D spending unrelated to actual changes in spending at the establishment level. Griliches and Lichtenberg (1984) attempt to overcome this problem by using R&D data grouped by applied product field rather than by industry of origin. Unfortunately, the reporting requirements of the NSF's product field survey were burdensome on participating firms, leading to spotty coverage. The survey was reduced from annual to biannual beginning in 1978, and was discontinued in 1986.

Patent data for U.S. industries are not routinely available. The reason is that the U.S. Patent Office assigns new patents to technological fields, but not to industries. Estimating patents by industry for the U.S. thus requires a mapping from technological fields into industries. The most satisfactory mapping available is the Yale Technology Concordance (YTC), described by Kortum and Putnam (1997). This concordance uses the fact that the Canadian patent office assigns patents to technological fields, to industries of manufacture, and to industries of use; for instance, a new farm tractor invented in an aerospace establishment would be assigned to the agricultural machinery sector (industry of manufacture) and to agriculture itself (industry of use). The YTC estimates mappings between technological field and industries of manufacture and use using the Canadian data, then applies the Canadian mapping to U.S. patents by technological field. For this study, I use annual data on U.S. patent applications grouped both by industry of manufacture and by industry of use, generously provided by Sam Kortum. I convert the annual flows of patents to stocks using the same method as for R&D; the empirical results again are similar if I use flows instead of stocks. Note that patents grouped by date of application are superior to patents grouped by date of grant, both because application presumably coincides with the economic viability of an innovation, and because historically there have been long and variable lags between application and granting in the United States, caused in part by changes in the resources of the U.S. Patent Office (Griliches, 1989).

I must again acknowledge potential problems with these data. First, the assignment of U.S. patents to industries is presumably not perfect, as the mapping between technological fields and industries probably varies between the U.S. and Canada as well as over time. Kortum and Putnam (1997) show that the estimated Canadian mapping forecasts Canadian industry patents out of sample reasonably well, alleviating these concerns somewhat but not entirely. Second, the distinction in the data between industry of manufacture and industry of use is not as sharp as one might hope. Ideally, I would like to interpret manufacture patents as "product innovations" and use patents as "process innovations." However, conversations with Sam Kortum suggested that this interpretation is not entirely correct; for instance, process innovations often wind up being assigned the same industries of manufacture and use even if no new product is created, while new products with broad applicability often wind up being assigned no industry of use. My sense is that we can at least safely assume that manufacture patents contain a higher fraction of product innovations than do use patents, and that use patents contain a higher fraction of process innovations than do manufacture patents.

I present sample means for patent flows and patent stock growth in Table 1. The flows of both manufacture and use patents are highest in nonelectrical machinery and electronics, while patent stocks grow most rapidly in drugs and computers. Notice that manufacture patent flows exceed use patent flows in most industries, and for my sample as a whole; this reflects the fact that many product innovations originating in manufacturing are used in nonmanufacturing, while few innovations originate in nonmanufacturing. The table also documents the fact, discussed in Griliches (1989) and Kortum (1993), that patent stocks have grown more slowly in the postwar United States than R&D stocks, or equivalently that the amount of real R&D per patent has been steadily rising. Some observers assert that this trend is evidence of vanishing technological opportunities; others argue that the cost of patenting has risen secularly and that patenting has become more concentrated in high-value innovations. Recall that my VARs include time dummies, which will control for any economy-wide changes in the cost or benefits of patenting that have affected the ratio of inventive activity to patents.

In addition to examining the impact of own R&D and patents, I examine the impact of innovations in upstream industries. I construct these measures using data from the 1977 U.S. input–output study, following the methods used by Terleckyj (1975), Keller (1997), and others. I begin by constructing a 19-by-20 matrix whose (i, j) element shows the total flow of goods in 1977 from sample industry i to sample industry j, including both intermediate and capital flows; I describe the construction of total flow matrices from raw input–output data in Shea (1991, 1993). The 20th column combines flows to omitted manufacturing industries, nonmanufacturing industries, private consumption, and government. I set diagonal elements to zero, then divide by row sums to obtain the shares of external demand for each sample industry accounted for by each other sample industry. I then multiply these demand shares by each industry's R&D and patent flows, to obtain the implicit "flow" of R&D and patents to and from each sample industry. Taking column sums gives me an estimate of the flows of upstream R&D and upstream patent applications to each manufacturing industry in any year. I cumulate these flows into stocks using the methods described above. These measures exclude R&D or patents coming from omitted industries; as mentioned earlier, however, those industries account for little innovative activity for most of my sample period.

My measures of TFP and inputs for manufacturing industries come from the NBER productivity database, described in Bartlesman and Gray (1996). The NBER data include annual measures of gross output and capital, labor, and materials inputs for 450 four-digit manufacturing industries. I measure labor as total employment multiplied by hours worked per production worker, assuming that production and nonproduction hours per worker are perfectly correlated. I define total input growth as a Divisia index of capital, labor, and materials growth, weighting with factor shares in gross output and measuring the capital share as a residual. TFP growth is defined as output growth minus input growth. I measure input and TFP growth at the four-digit level, aggregate up to the 19 industries listed in Table 1 using shares in nominal gross output, then convert growth rates to level indices. Below, I examine the dynamic impact of technology shocks both on total input and on capital, labor, and materials separately, premultiplying log capital, labor, and materials by their shares in nominal gross output in order to avoid having to impose the condition that factor shares in production are identical across industries.

3. *Preliminary Evidence*

This section examines the univariate time-series properties of my data, and replicates previous work examining cross-section and contemporaneous time-series relationships between technology and TFP. My baseline data are annual observations for 19 manufacturing industries from 1959 to 1991 on TFP; total input and its share-weighted capital, labor, and materials components; stocks of own R&D, manufacture patents, and use patents; and stocks of upstream R&D, manufacture patents, and use patents.

Table 2 PANEL UNIT-ROOT TESTS

X	Other Deterministic Terms: Time Dummies	Other Deterministic Terms: Time Dummies and Sectoral Trends
TFP	−0.003	−0.174
	(0.007)	**(0.022)
Total input	−0.024	−0.193
	(0.011)	**(0.026)
Capital	−0.014	−0.244
	(0.008)	*(0.034)
Labor	−0.107	−0.201
	(0.018)	**(0.028)
Materials	−0.038	−0.218
	(0.013)	**(0.029)
R&D	−0.017	−0.096
	(0.004)	**(0.011)
Manufacture patents	−0.009	−0.092
	(0.003)	**(0.011)
Use patents	−0.014	−0.058
	(0.004)	(0.010)

$\Delta \log X_{it} = \gamma_i +$ (other deterministic terms) $+ \beta \log(X_{it-1}) + \sum_{k=1}^{3} \alpha_k \Delta \log X_{it-k}) + \epsilon_{it}$. This table presents estimates of β from Augmented Dickey–Fuller tests of the null hypothesis that log X contains a unit root, using annual panel data for 19 industries from 1959 to 1991. All regressions include sector-specific intercepts and time dummies; regressions in the right column also contain sector-specific linear time trends. Standard errors are in parentheses. * indicates that β is significant at 10%, while ** indicates significance at 5%. The critical values of 6.816 (10%) and 7.093 (5%) are taken from the asymptotic formula provided in Levin and Lin (1992).

Table 2 presents univariate time-series evidence. For each series, I perform an Augmented Dickey–Fuller test of the null hypothesis that the series has a unit root in log levels, including three lagged growth rates to correct for serially correlated errors. I include sector-specific intercepts and time dummies in each specification, and experiment with including sector-specific time trends; all other coefficients are constrained to be equal across sectors. Since these are panel data, I cannot apply the usual Dickey–Fuller critical values; I instead use the formula provided in Levin and Lin (1992), which with 19 industries implies a 5% critical value of −7.093 and a 10% critical value of −6.816. According to Table 2, I can never reject the null of a unit root when I include only sectoral intercepts and time dummies. However, I can reject a unit root in seven of eight cases when I include sector-specific trends. I conclude that my data are stationary around trends that differ across sectors.

Table 3 looks at the long-run relationship between technology and TFP growth. I estimate cross-section OLS regressions of mean TFP growth on a constant and the mean growth rates of my technology indicators, taken one at a time; the sample size for each regression is 19. The coefficients on own R&D and own manufacture patent growth are positive and significant at 10%, while the coefficients on own use patent growth as well as upstream use and manufacture patent growth are positive and significant at 5%. My sample is too small to allow for multivariate analysis, and the results are fragile; omitting computers, for instance, reduces the coefficient on technology in all cases, and renders the own R&D results insignificant. Still, these results suggest that my technology indicators capture something about technological progress. My findings are consistent with Griliches and Lichtenberg (1984), Scherer (1984, 1993), and Lichtenberg and Siegel (1991), who find that R&D and productivity growth are positively related across firms and industries, and with Terleckyj (1975), who reports a significant positive relationship across industries between TFP growth and upstream R&D. Notice that use patents are more strongly related to TFP growth than manufacture patents, suggesting that process innovations may be better captured by available TFP data than product innovations, a theme to which I return below.

Table 4 estimates contemporaneous time-series relationships between TFP and technology indicators in log levels, while Table 5 does the same in growth rates. I include sectoral intercepts and time dummies in the levels regressions, and experiment with sector-specific trends; I include a constant and time dummies in the growth-rate regressions, and experi-

Table 3 LONG-RUN EVIDENCE

X	γ *Estimate*	β *Estimate*
R&D	−0.005	0.328
	(0.009)	*(0.196)
Manufacture patents	0.006	0.426
	(0.004)	*(0.233)
Use patents	0.003	1.082
	(0.003)	**(0.275)
Upstream R&D	−0.023	0.752
	(0.021)	(0.483)
Upstream manuf. patents	−0.005	1.880
	(0.004)	**(0.768)
Upstream use patents	−0.006	2.345
	(0.006)	**(0.837)

$\Delta \log(\text{TFP}_i) = \gamma + \beta \Delta \log X_i + \epsilon$. This table presents estimates of cross-section relationships between long-run total factor productivity growth and long-run growth in technology indicators. Each variable is entered as a mean industry-level growth rate over 1960–1991; the sample size is 19. Standard errors are in parentheses. * denotes significance at 10%, while ** denotes significance at 5%.

Table 4 CONTEMPORANEOUS EVIDENCE: LOG LEVELS

	Other Deterministic Terms	
X	*Time Dummies*	*Time Dummies and Sectoral Trends*
R&D	0.334	0.056
	**(0.041)	*(0.035)
Manufacture patents	0.396	−0.300
	**(0.055)	**(0.063)
Use patents	1.102	0.121
	**(0.073)	(0.078)
Upstream R&D	0.773	−0.460
	**(0.095)	**(0.126)
Upstream manuf. patents	2.121	−0.441
	**(0.217)	**(0.149)
Upstream use patents	3.289	−0.784
	**(0.250)	**(0.263)

$\log(\mathrm{TFP}_{it}) = \gamma_i +$ other deterministic terms $+ \beta \log X_{it} + \epsilon_{it}$. This table presents estimates of contemporaneous relationships between log levels of total factor productivity and technology indicators, using annual panel data on 19 industries from 1959 to 1991. All regressions include sector-specific intercepts and time dummies; regressions in the right column also contain sector-specific linear time trends. Standard errors are in parentheses. * indicates significance at 10%, ** indicates significance at 5%.

Table 5 CONTEMPORANEOUS EVIDENCE: GROWTH RATES

	Other Deterministic Terms	
X	*Time Dummies*	*Time Dummies and Sectoral Trends*
R&D	0.132	0.047
	**(0.049)	(0.050)
Manufacture patents	0.195	−0.070
	**(0.074)	(0.101)
Use patents	0.512	0.052
	**(0.099)	(0.124)
Upstream R&D	0.235	−0.058
	**(0.123)	(0.142)
Upstream manuf. patents	0.530	−0.291
	**(0.217)	(0.254)
Upstream use patents	1.140	−0.057
	**(0.280)	(0.367)

$\Delta \log(\mathrm{TFP}_{it}) = \gamma +$ other deterministic terms $+ \beta\, \Delta \log X_{it} + \epsilon_{it}$. This table presents estimates of contemporaneous relationships between growth rates of total factor productivity and technology indicators, using annual panel data on 19 industries from 1960 to 1991. All regressions include a constant and time dummies; regressions in the right column also include sector-specific intercepts. Standard errors are in parentheses. * indicates significance at 10%, while ** indicates significance at 5%.

ment with sectoral intercepts. Results omitting sectoral trends in Table 4 suggest strong, positive contemporaneous relationships between TFP and all six technology indicators. However, including sectoral trends weakens the relationship substantially for own R&D and use patents, and reverses the sign in the other four cases. Similarly, in Table 5 there is a strong positive relationship between TFP growth and technology growth when I control only for time dummies, but this relationship vanishes when I add sectoral intercepts. I conclude that cross-industry differences in trend productivity growth are positively related to cross-industry differences in trend technology growth, but that once I control for these differences there is little correlation between TFP and technology. My results contradict Lach (1995), who reports a contemporaneous positive relationship between patent stock growth and TFP growth in a sample similar to mine, as well as Griliches and Lichtenberg (1984), who find no time-series relationship between TFP and R&D even when omitting sectoral trends.

Tables 4 and 5 suggest that there is no contemporaneous *within-industry* relationship between TFP and technology in annual data. Table 6 asks whether such a relationship exists over a longer horizon, by regressing *medium-run* TFP growth on technology growth measured over the sixteen-year intervals 1960–1975 and 1976–1991. There are two

Table 6 MEDIUM-HORIZON EVIDENCE: 16-YEAR GROWTH RATES

	Other Deterministic Terms	
X	*Time Dummy*	*Time Dummy and Fixed Effect*
R&D	0.285	0.140
	*(0.154)	(0.288)
Manufacture patents	0.297	−0.347
	(0.166)	(0.416)
Use patents	0.960	0.614
	**(0.233)	(0.451)
Upstream R&D	0.462	−1.206
	(0.406)	(0.872)
Upstream manuf. patents	1.209	0.149
	**(0.572)	(0.835)
Upstream use patents	2.348	2.371
	**(0.729)	(2.078)

$\Delta \log (\mathrm{TFP}_{it}) = \gamma +$ other deterministic terms $+ \beta \Delta\log X_{it} + \epsilon_{it}$. This table presents estimates of the relationship between medium-horizon growth rates of total factor productivity and technology indicators, using data on 19 industries for two 16-year periods, 1960–1975 and 1976–1991. All regressions include a constant and a dummy for the second period; regressions in the right column also include a sector-specific fixed effect. Standard errors are in parentheses. * indicates significance at 10%, while ** indicates significance at 5%.

observations per industry, implying a sample size of 38. In the first column, I control only for a constant and a dummy for the second period; these results suggest a positive and significant relationship between TFP and technology growth in the medium run. However, these results rely on both cross-industry and within-industry variation. Adding a sectoral fixed effect in the second column makes the relationship between TFP and technology insignificant, as standard errors rise substantially in all cases and point estimates fall substantially in five of six cases. These results suggest that most of the variation in medium-run technology growth is cross-industry rather than within-industry, and that within-industry medium-run variation in technology and TFP are only weakly related. I obtain similar results when I experiment with different starting and ending dates, as well as with four- and eight-year horizons.

One might wonder if Tables 4 through 6 obviate the need for any further investigation of time-series relationships between TFP and technology. The answer is no. Had I found a robust positive contemporaneous relationship between (say) R&D and TFP, I could not have concluded that R&D shocks cause TFP to rise, because of potential omitted-variables bias; shocks to industry output could raise measured TFP due to (say) cyclical utilization, while at the same time increasing R&D for accelerator reasons. Similarly, the absence of a contemporaneous relationship does not prove that R&D has no impact on TFP, since such an impact is likely to emerge only with a lag. Both of these problems can be addressed by using vector autoregressions.

4. *VAR* Evidence

In this section I present results from a series of vector autoregressions using annual industry panel data on inputs, TFP, and technological indicators from 1959 to 1991. All variables are in log levels, following the panel unit-root tests presented in Table 2, although the impulse response functions in log levels are broadly similar if I estimate using growth rates. All specifications include sector-specific intercepts, sector-specific time trends, and time dummies. The time dummies are intended to control for aggregate shocks that affect R&D and patenting intensity, but are unrelated to true technological progress, such as the changes in the U.S. Patent Office discussed in Griliches (1989). Of course, time dummies will also remove any variation due to aggregate technology shocks, which may bias my results against technology-shock models; fortunately, my results are broadly similar if I omit time dummies. I use four lags; experiments with other lag lengths yielded similar results.

Figures 2 through 4 present the complete set of estimated impulse response functions, along with 1.65 Monte Carlo standard error bands, for three-variable VARs estimated on the manufacturing sample. The VARs are ordered as total input, TFP, and either own R&D (Figure 2), own manufacture patents (Figure 3), or own use patents (Figure 4). I enter technology indicators one at a time for presentational simplicity; results are similar if I enter multiple indicators simultaneously. Placing technology last reflects my belief that shocks to R&D or patenting are likely to affect industry activity only with lags. Placing technology first would generate a significant but small expansionary initial impact of technology on inputs and TFP, but with otherwise similar impulse responses and variance decompositions. Figure 5 breaks the input responses to own technology into disaggregated components, taken from five-variable VARs ordered as capital, labor, materials, TFP, and technology. Figures 6 and 7 present the responses of input, TFP, capital, labor, and materials to upstream R&D and patents; these estimates are similar if I control for input–output weighted measures of upstream input demand, suggesting that upstream technology is not merely proxying for upstream activity. I summarize these impulse responses in Table 7, which lists the sign and horizon of every significant effect of technology shocks on nontechnology variables. Table 8 presents Granger causality evidence, and Table 9 presents variance decompositions.

While the results vary somewhat across specifications, some robust patterns emerge. First, the impulse responses of TFP to technology are not significantly positive at any horizon, and indeed are significantly negative in the long run for all three upstream technology measures. This result is reinforced in Table 8, which indicate that TFP is Granger-caused only by upstream R&D and upstream use patents (and in these cases with negative coefficients).

Second, the impulse responses of total input to technology tend, if anything, to be positive in the short run but negative in the long run. I find that technology shocks significantly raise input use in the short run for three of six technology indicators, while significantly decreasing long-run input use in four of six cases.

Third, the impulse responses of individual inputs (particularly labor and materials) to technology are stronger than the responses for total input, as the disaggregated impacts tend to wash out due to both staggering and conflicting signs. This pattern is mirrored in the Granger causality results, which show that technology shocks forecast total input for only three of six technology indicators, but forecast labor and materials in five of six cases. Favorable technology shocks increase short-run labor use significantly in five of six cases, while decreasing long-run materials

Figure 2 INPUTS, TFP, AND R&D

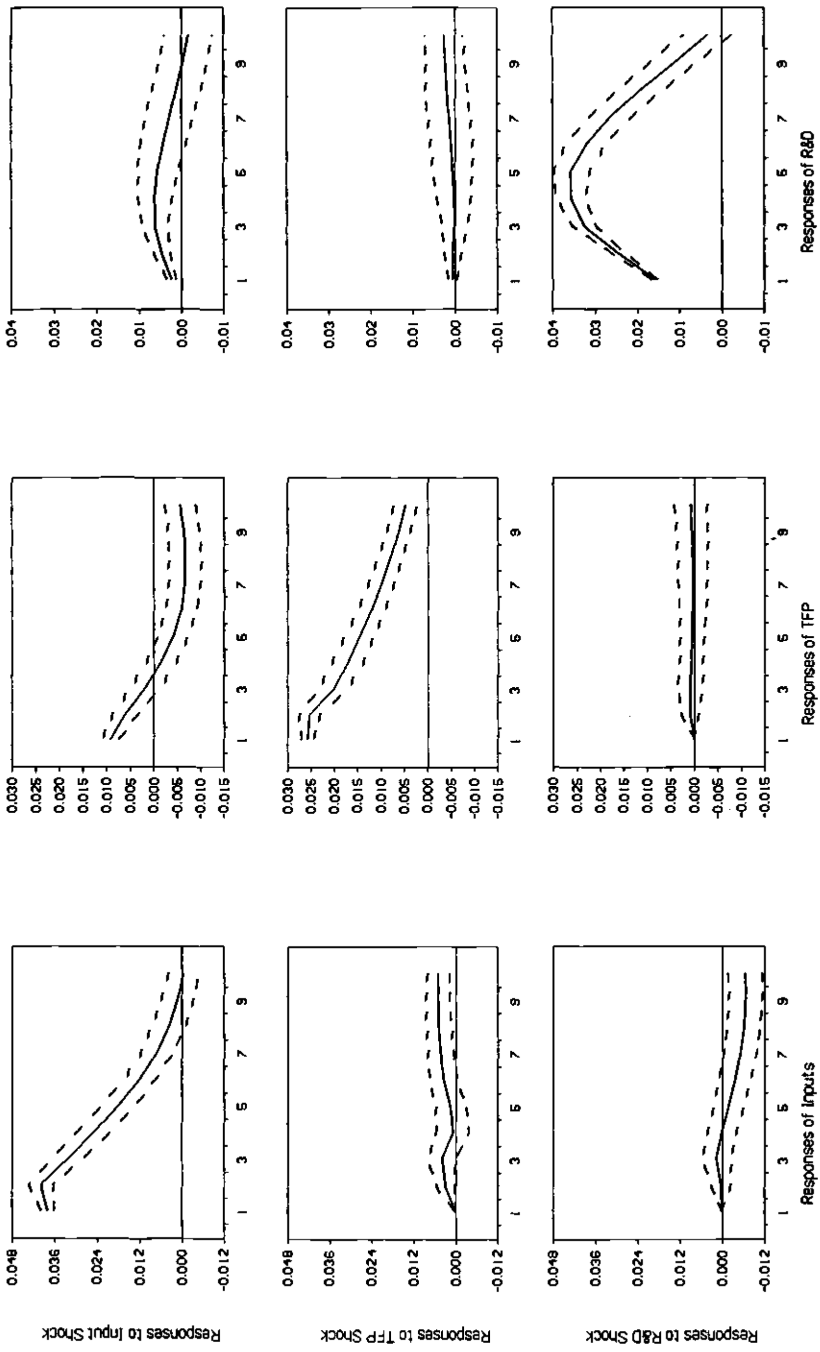

Figure 3 INPUTS, TFP, AND MANUFACTURE PATENTS

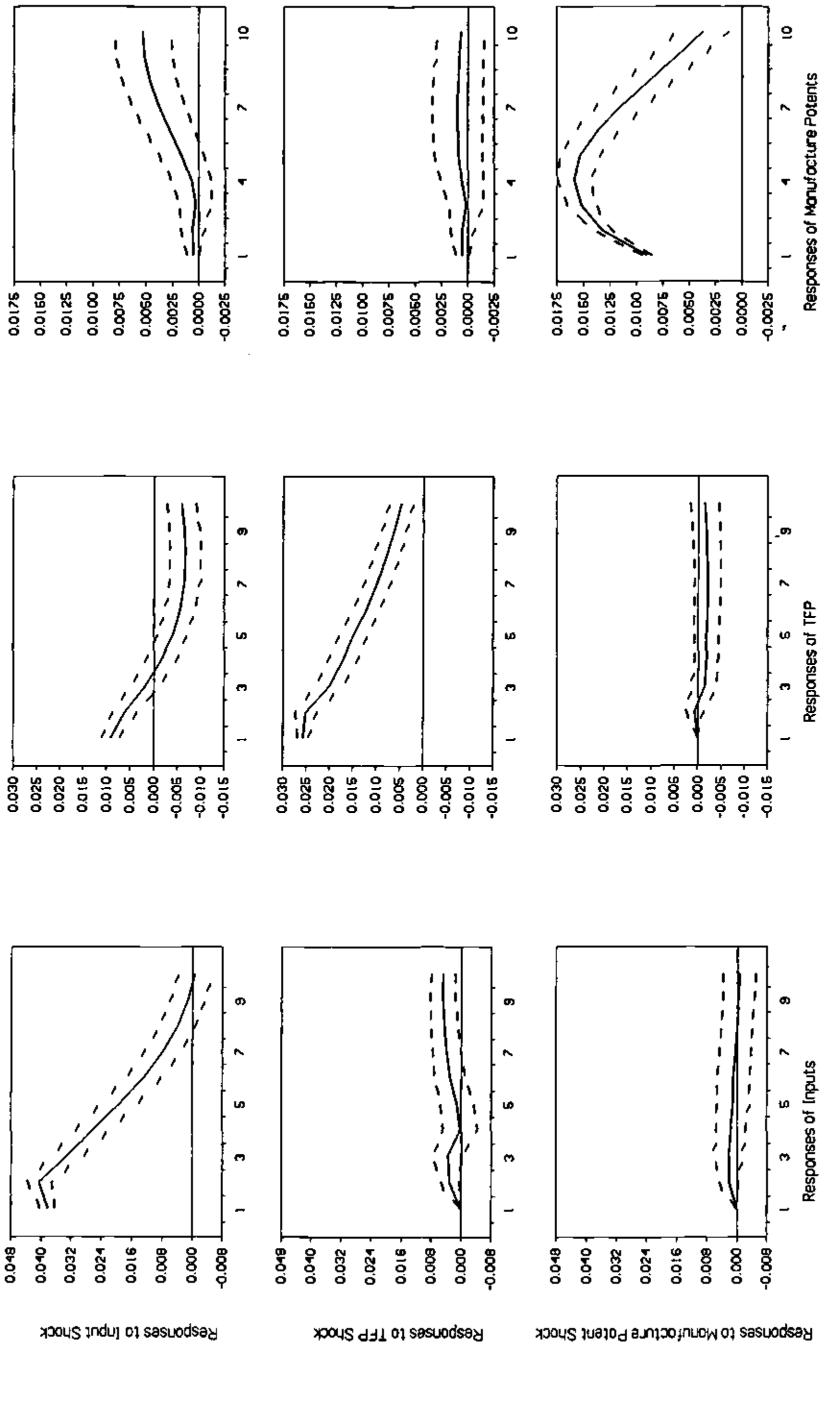

Figure 4 INPUTS, TFP, AND USE PATENTS

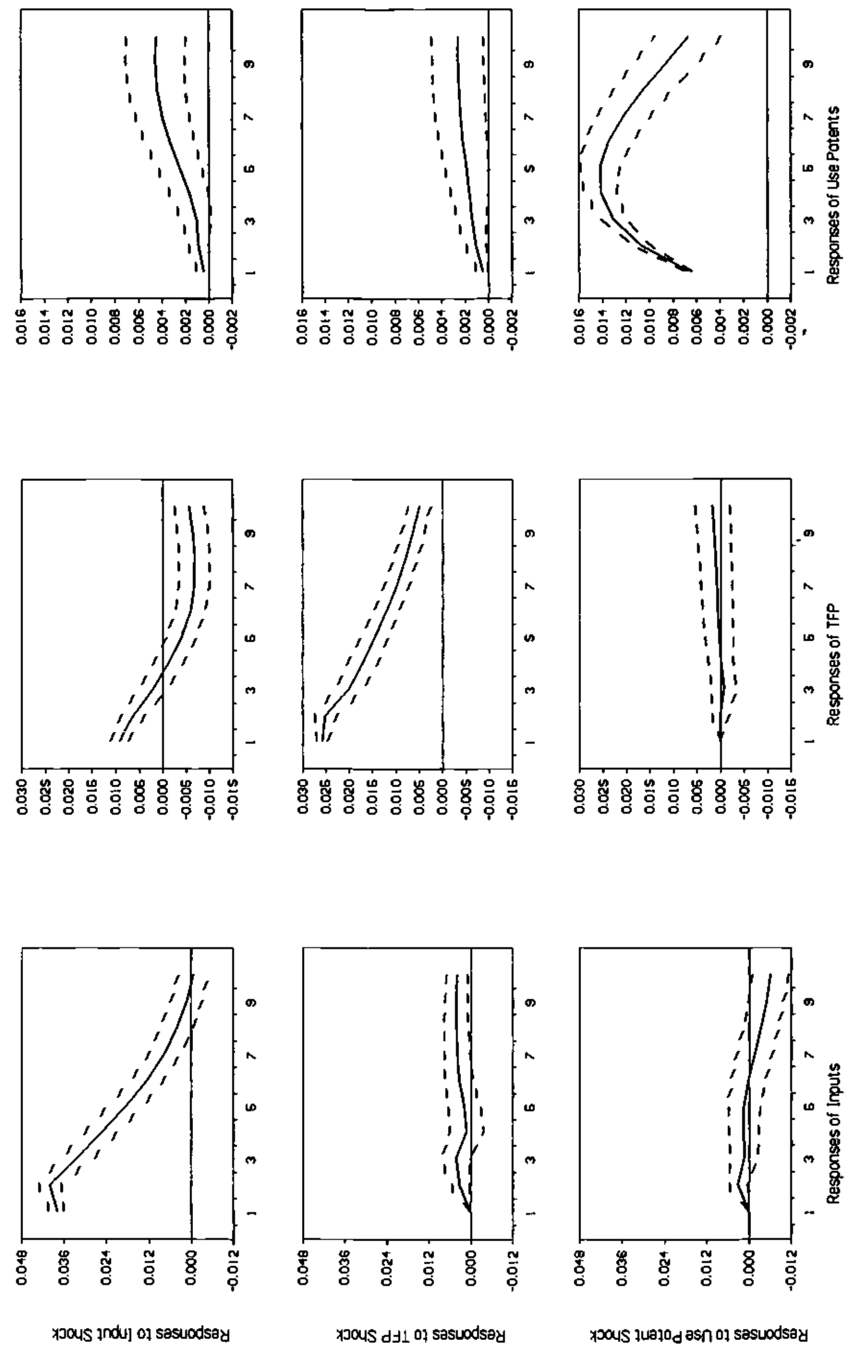

Figure 5 CAPITAL, LABOR, MATERIALS, AND OWN TECHNOLOGY SHOCKS

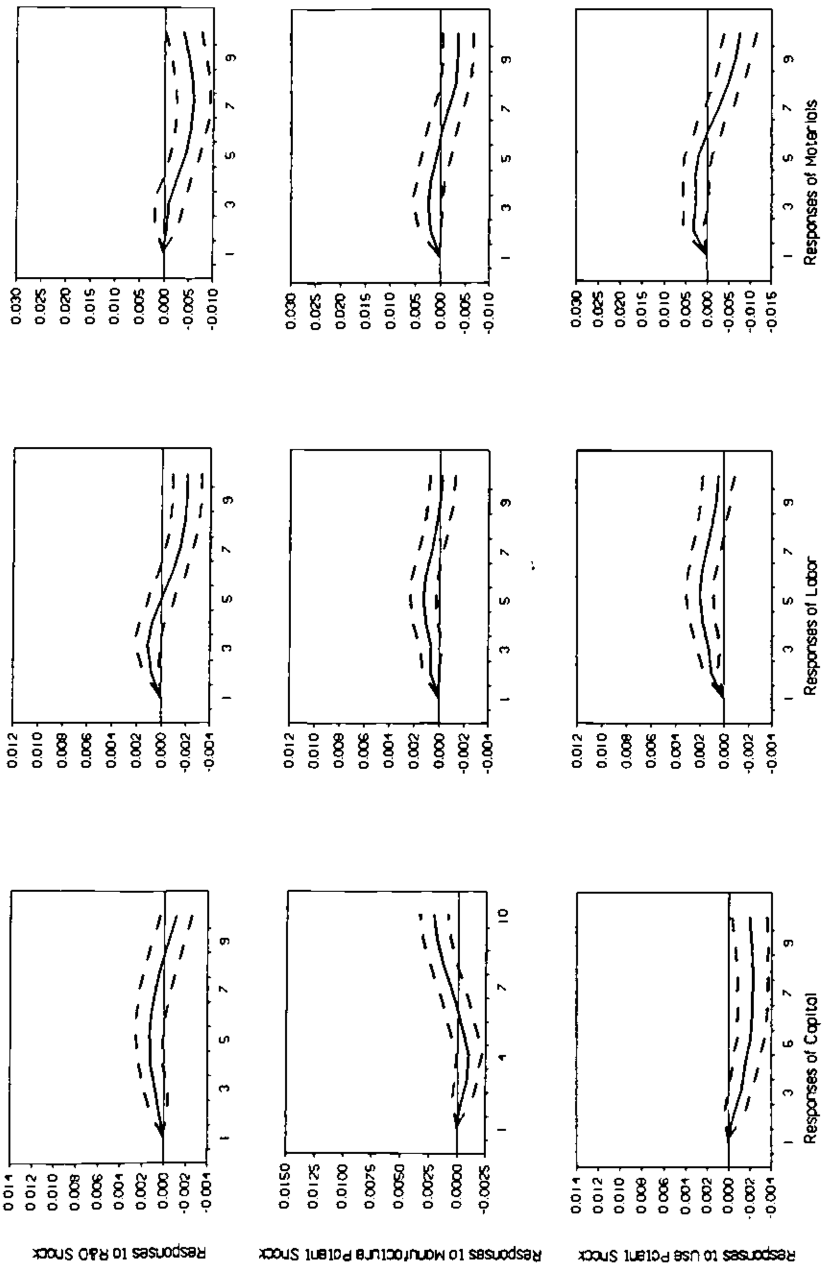

Figure 6 INPUTS, TFP, AND UPSTREAM TECHNOLOGY SHOCKS

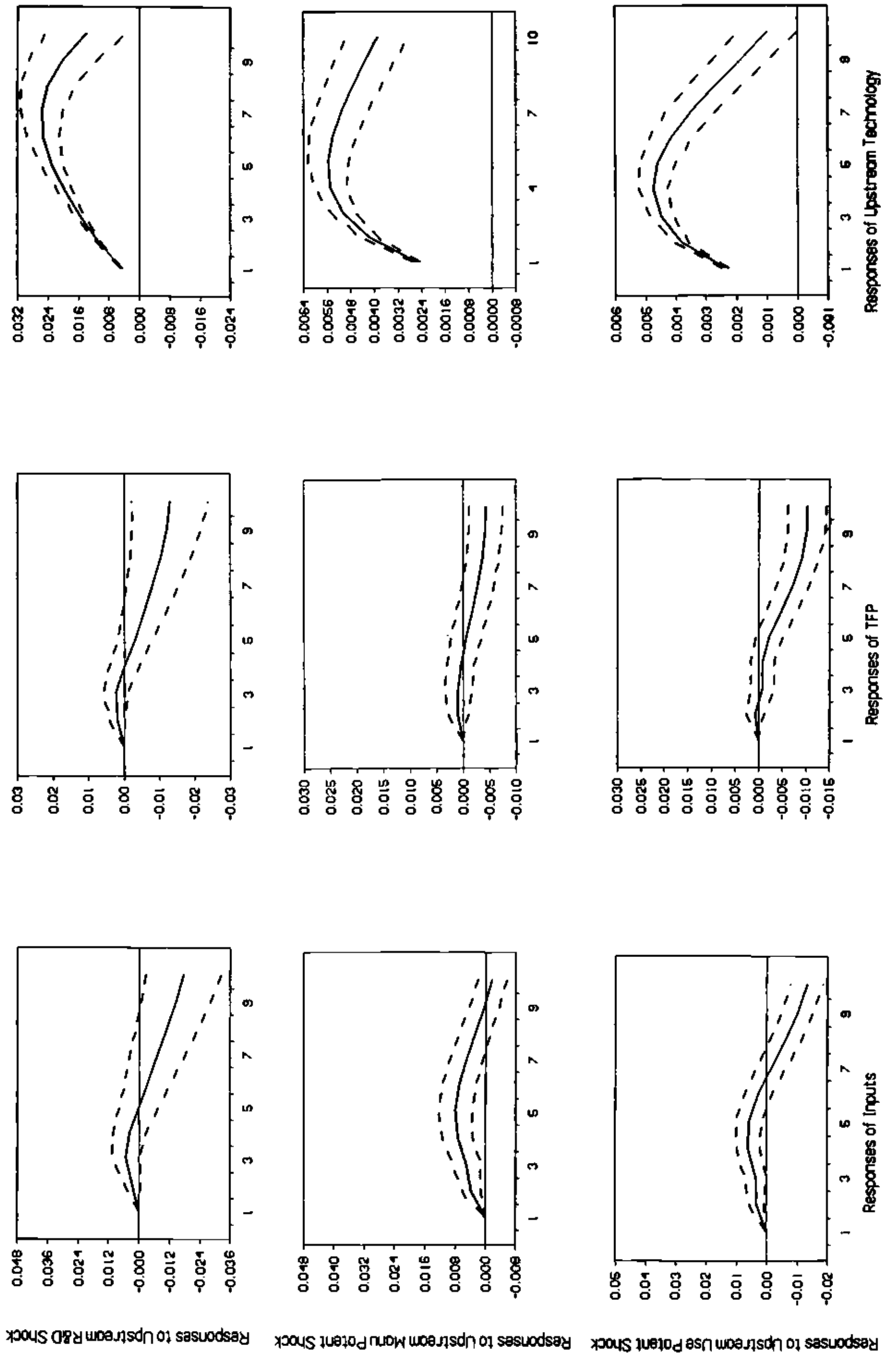

Figure 7 CAPITAL, LABOR, MATERIALS, AND UPSTREAM TECHNOLOGY SHOCKS

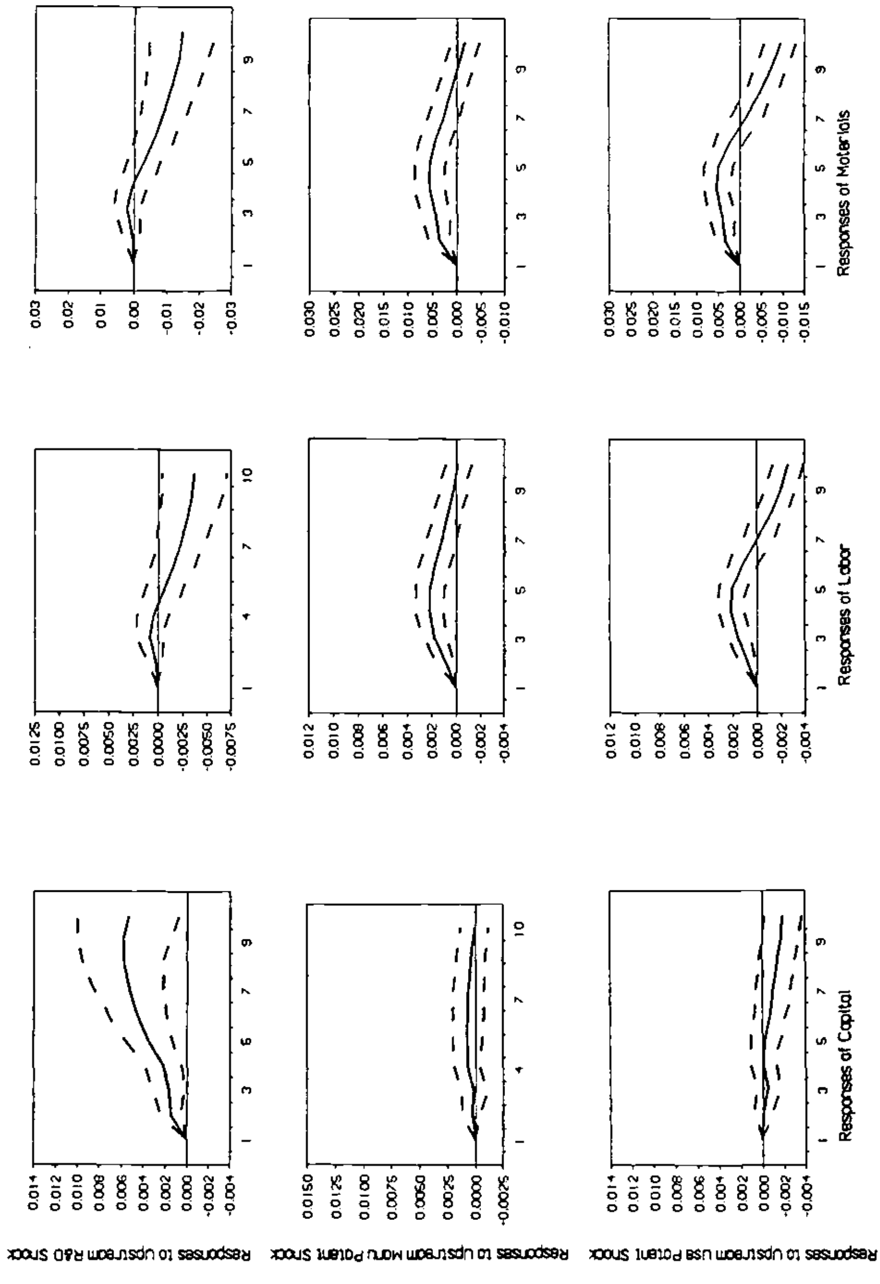

Table 7 IMPULSE RESPONSE FUNCTIONS: SUMMARY

	Three-Variable VAR		Five-Variable VAR		
Tech. Indicator	*Total Input*	*TFP*	*Capital*	*Labor*	*Materials*
R&D	↓ 7–10	—	↑ 4	↑ 2–3 ↓ 7–10	↓ 4–10
Manufacture patents	—	—	↓ 4 ↑ 8–10	↑ 4–5	↓ 8–10
Use patents	↑ 2 ↓ 9–10	—	↓ 3–10	↑ 2–7	↑ 2 ↓ 8–10
Upstream R&D	↓ 9–10	↓ 7–10	↑ 2–10	↓ 8–10	↓ 6–10
Upstream manuf. patents	↑ 2–7	↓ 7–10	—	↑ 2–6	↑ 2–6
Upstream use patents	↑ 2,4–5 ↓ 8–10	↓ 6–10	↓ 9–10	↑ 2–5 ↓ 8–10	↑ 2–5 ↓ 8–10

This table summarizes the VAR impulse functions by reporting all cases of a significant (10%) impact of technology on industry variables, along with the relevant horizons in years. The impulse responses are calculated from VARs estimated using annual panel data for 19 industries from 1959 to 1991. The results in the first two columns are based on three-variable VARs ordered as total input, TFP, and technology, while the results in the last three columns are based on five-variable VARs ordered as capital, labor, materials, TFP, and technology. All VARs are estimated in log levels and include sector-specific intercepts and trends as well as time dummies. The standard errors are computed using Monte Carlo integration.

Table 8 GRANGER-CAUSALITY TESTS: *P*-VALUES

	Three-Variable VAR		Five-Variable VAR		
Technology Indicator	*Total Input*	*TFP*	*Capital*	*Labor*	*Materials*
Panel A: Does Technology Granger-Cause Inputs or TFP?					
R&D	0.21	0.95	0.42	0.01	0.02
Manufacture patents	0.85	0.51	0.04	0.05	0.13
Use patents	0.05	0.92	0.01	0.01	0.00
Upstream R&D	0.12	0.07	0.02	0.12	0.01
Upstream manuf. patents	0.01	0.46	0.47	0.03	0.01
Upstream use patents	0.00	0.01	0.15	0.00	0.00
Panel B: Do Inputs and TFP Granger-Cause Technology?					
R&D	0.38	0.62	0.34	0.95	0.63
Manufacture patents	0.05	0.41	0.58	0.14	0.21
Use patents	0.11	0.56	0.72	0.25	0.01
Upstream R&D	0.00	0.35	0.04	0.06	0.00
Upstream manuf. patents	0.06	0.48	0.08	0.03	0.51
Upstream use patents	0.00	0.32	0.01	0.02	0.16

This table presents *P*-values from Granger causality tests from technology to industry activity and vice versa. The tests are based on VARs estimated using annual panel data for 19 industries from 1959 to 1991.

Table 9 VARIANCE DECOMPOSITIONS

		Percentage of Variance Due to Technology				
		Three-Variable VAR		*Five-Variable VAR*		
Tech. Indicator	*Years*	*Total Input*	*TFP*	*Capital*	*Labor*	*Materials*
R&D	3	0.05	0.07	0.23	0.66	0.04
	6	0.38	0.06	1.08	0.95	2.06
	9	2.42	0.07	1.20	3.89	4.06
Manufacture patents	3	0.19	0.16	0.31	0.27	0.37
	6	0.22	0.60	0.53	1.22	0.42
	9	0.23	0.97	1.60	1.25	1.32
Use patents	3	0.26	0.04	0.54	0.87	0.70
	6	0.28	0.05	2.81	3.49	0.93
	9	0.98	0.19	5.22	4.13	3.49
Upstream R&D	3	0.76	0.44	1.47	0.22	0.16
	6	1.15	2.17	8.04	0.99	2.23
	9	8.21	12.51	20.33	8.01	14.10
Upstream manuf. patents	3	1.00	0.08	0.03	1.38	1.34
	6	3.98	0.25	0.28	4.65	3.94
	9	4.49	1.82	0.39	4.86	3.95
Upstream use patents	3	0.60	0.08	0.09	1.09	1.19
	6	2.11	1.48	0.23	3.69	2.97
	9	5.35	10.13	1.38	5.27	5.57

This table summarizes the VAR variance decompositions by reporting the share of variance of industry activity variables accounted for by shocks to technology at 3-, 6-, and 9-year horizons. The variance decompositions are calculated from VARs estimated using annual panel data for 19 industries from 1959 to 1991. The results in the first two columns are based on three-variable VARs ordered as total input, TFP, and technology, while the results in the last three columns are based on five-variable VARs ordered as capital, labor, materials, TFP, and technology. All VARs are estimated in log levels and include sector-specific intercepts and trends as well as time dummies.

use significantly in five of six cases. Note that shocks to both own and upstream R&D significantly increase capital accumulation in the short run. This result is consistent with Lach and Rob (1996), who find that R&D Granger-causes physical investment in industry panel data, and with Lach and Schankerman (1989), who find the same result in firm-level data. In my data, R&D does not Granger-cause capital, but it does Granger-cause investment.

Fourth, technology shocks explain only a small fraction of input and TFP variation at business-cycle horizons. Technology explains less than 2% of three-year volatility in all cases, and less than 5% of six-year volatility in all but one case. Technology has somewhat more explanatory power at longer horizons, particularly for upstream R&D and upstream use patents; recall, however, that the impulse responses for these cases suggest significant long-run *contractions* of inputs and TFP following favorable technology shocks. The fact that technology explains a larger share of variance at longer horizons is consistent with Jovanovic and Lach (1997), who model technology shocks as having long diffusion lags and find that technology shocks underexplain short-run volatility but overexplain long-run volatility.

Along with the results for technology shocks, two other features of my estimates are worth noting. First, own R&D and own patents respond positively and significantly to input shocks; moreover, R&D increases immediately, whereas patents increase only after five years. One possible interpretation is that industry expansions generate increased R&D immediately, and that this investment eventually leads to an increased flow of patentable inventions. Second, input shocks lead to short-run increases but long-run decreases in TFP. A possible interpretation is that industry expansions raise measured TFP in the short-run due to increasing returns or cyclical utilization, but reduce long-run productivity. An interesting question is how these two features of the data can coexist—if expansions increase R&D and patents, then why don't they increase productivity? One possibility is that expansions raise inventive activity but also allow lower-productivity firms to enter and survive, generating a net decline in productivity.

4.1 RESULTS FOR INPUT MIX, WORKER MIX, AND PRICES

While this paper is primarily concerned with the impact of technology shocks on total input and TFP, technology shocks are likely to affect other variables as well. Conventional models predict that favorable technology shocks reduce the relative price of industry output. Meanwhile, if technology shocks are permanent and the supply of capital is more elastic in the long run than the supply of labor or materials (as in the baseline RBC model), then favorable technology shocks increase the long-run ratio of capital to other inputs. Finally, some have hypothesized that technological advances have been biased towards skilled workers during the postwar period, either because skilled workers have an advantage in learning new technologies (Greenwood and Yorukoglu, 1997) or because technology shocks are investment-specific and capital is complementary with skill (Krusell *et al.* 1996).

In the working version of this paper (Shea, 1998), I present impulse response functions to own and upstream technology shocks for three variables: the *worker mix,* defined as the ratio of nonproduction to total employment; the *input mix*, defined as the log ratio of capital's product (capital raised to the power of capital's share of revenue) to labor and materials' product; and the industry's relative price, defined as the implicit gross output deflator divided by the GDP deflator. I assume that increases in nonproduction employment are positively correlated with changes in the ratio of skilled to unskilled employees, following Berman, Bound, and Griliches (1994). My input-mix variable is one of several ways I could quantify changes in capital relative to other variables; results are similar when I use more familiar measures such as the capital–labor ratio. The impulse responses are taken from four-variable VARs in which the new variables are ordered after inputs and TFP but before technology. Data for nonproduction employment, total employment and prices are taken from the NBER productivity database.

The results conform to prior intuition in two out of three cases. Favorable technology shocks cause significant long-run substitution towards capital for five of six technology indicators; technology improvements also significantly increase the ratio of nonproduction to total employment in five of six cases, although these increases often occur in the short run rather than the long run. However, the estimated impact of technology shocks on price is not robust; own R&D and own use patent shocks significantly reduce price in the long run, but manufacture patent shocks raise price in the medium run, while upstream use patent shocks raise price in the long run. The fact that own use patents (which should reflect process innovations) reduce prices while own manufacture patents (which should reflect product innovations) raise prices suggests that available price data might not accurately reflect product innovations, an idea to which I return below.

4.2 RESULTS FOR NONMANUFACTURING

My empirical results to this point have relied exclusively on manufacturing industries. However, technology shocks originating in manufacturing, such as the introduction of the jet engine in the late 1950s, often have important downstream impacts in nonmanufacturing. While disaggregated data on own R&D and own patenting in nonmanufacturing industries is not readily available—in part because virtually all R&D and patenting occurred in manufacturing until very recently—I can construct measures of upstream technology for nonmanufacturing using the techniques described in Section 2. Figure 8 presents impulse responses of inputs and TFP to upstream technology shocks for a panel of 10 nonmanufacturing industries: agriculture; mining; construction; transportation; communications; electric utilities; gas utilities; trade; finance, insurance and real estate (FIRE); and services. Data on inputs and TFP

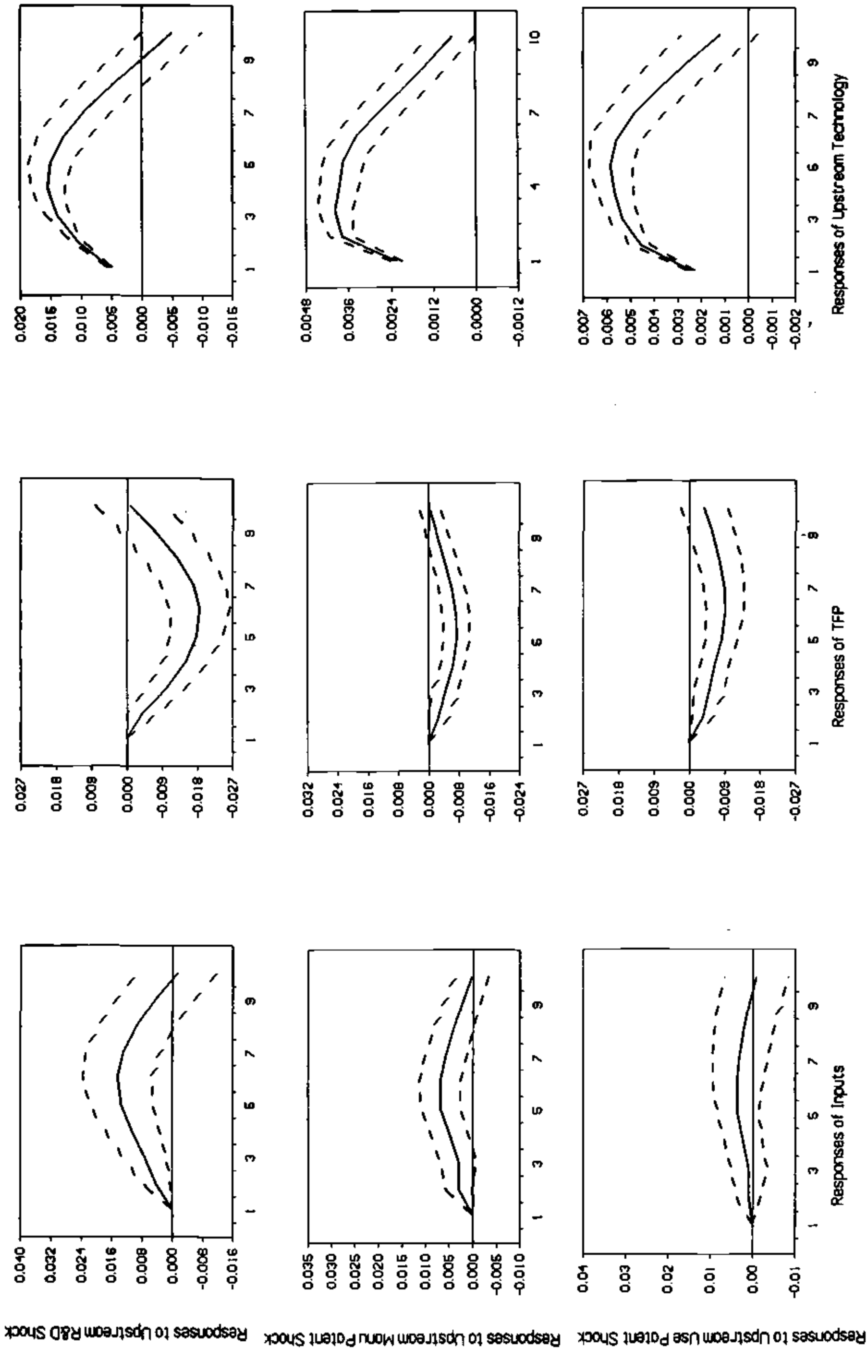

Figure 8 INPUTS, TFP, AND UPSTREAM TECHNOLOGY IN NONMANUFACTURING

come from an updated version of the KLEM database described by Jorgenson, Gollop, and Fraumeni (1987), generously provided by Susanto Basu. Although their preferred measure of labor input corrects for variations in labor-force composition, I use man-hours to be consistent with the manufacturing data. The impulse responses for nonmanufacturing are striking and robust: favorable upstream technology shocks significantly increase total input in the short run, but reduce measured TFP in the short run; total input and TFP return to trend in the long run. In the working version of the paper, I show that results for capital, labor, and materials are similar to those for total input. The variance decompositions (available from the author) assign technology a substantial share of TFP and input volatility at six years, particularly for upstream R&D; however, technology has a much smaller role for output volatility, as the input and TFP effects cancel each other out.

4.3 MEASUREMENT ERROR IN PRICE INDICES

The fact that favorable technology shocks do not significantly increase measured TFP raises suspicions about the quality of the TFP data. Much recent research has criticized BLS price data for not registering implicit price changes due to quality improvements or new product introductions (e.g., Gordon, 1990), or, for sectors outside of manufacturing, for not registering price changes at all (e.g., Baily and Gordon 1988). If product innovations do not reduce measured prices, they are less likely to increase measured output or TFP; this is especially troublesome given that roughly 80% of U.S. R&D is devoted to product rather than process innovation (Scherer, 1984). Similarly, if nonmanufacturing prices are measured poorly, then upstream innovations that reduce true prices and increase true activity may not raise measured output; if measured inputs rise (perhaps because inputs are easier to measure than output), then measured TFP is likely to fall.

To examine whether measurement errors in prices are important for my results, I divide the 19 sample manufacturing industries into process-innovating vs. product-innovating sectors. Table 10 presents the average percentage of R&D spending in gross output over the period 1959–1991, as well as the percentage of process R&D in total R&D spending in 1974, as reported in Scherer (1984). The table indicates that there is a fairly sharp break between process- and product-innovating sectors, and that the most R&D-intensive industries are typically product-intensive. I assign food, textiles, lumber, paper, industrial chemicals, petroleum, rubber, stone, and primary metals to the process-innovating group, and the other ten industries to the product-innovating group. Figures 9 and 10 present the responses of total input, TFP, and price to

Table 10 PROCESS VS. PRODUCT R&D INTENSITY

Industry	*R&D Intensity*	*Percentage of Process R&D*
Food (SIC 20)	0.2	56.1
Textiles (SIC 22–23)	0.1	61.3
Lumber (SIC 24–25)	1.5	59.0
Paper (SIC 26)	0.7	33.1
Industrial chemicals (SIC 281–282, 286)	3.4	47.6
Drugs (SIC 28)	8.2	12.0
Other chemicals (other SIC 28)	1.4	14.6
Petroleum (SIC 29)	1.4	64.0
Rubber (SIC 30)	1.1	48.0
Stone (SIC 32)	0.9	52.5
Metals (SIC 33)	0.5	75.2
Metal prods. (SIC 34)	0.4	14.6
Computers (SIC 357)	12.5	5.5
Other nonelec. equip. (other SIC 35)	1.2	4.1
Electronics & commun. equip. (SIC 366–387)	5.4	21.2
Other electric equip. (other SIC 36)	2.3	11.5
Aerospace (SIC 372, 6)	4.4	21.7
Autos & other transp. equip. (SIC 37)	2.6	4.8
Instruments (SIC 38)	5.0	8.0

The first column reports the average value of R&D spending as a percentage of nominal gross output over the sample period 1959–1991. The second column reports the fraction of 1974 R&D spending devoted to process R&D, as reported in Scherer (1984).

technology shocks for the two groups. The results indicate a sharp distinction between process- and product-innovating industries: for process-innovating sectors, favorable technology shocks induce a significant long-run increase in TFP, a significant long-run decline in price, and a significant long-run decline in inputs; for product-innovating sectors, favorable technology shocks do not raise TFP in any instance, and reduce long-run inputs and prices in only one case. These results suggest that the failure of technology to increase TFP in my full sample may be due to the failure of available data to reflect price declines and productivity gains due to quality improvements and new-product introductions. Another interesting result is that process-industry TFP declines significantly in the short run in two of three cases. This result is consistent with models in which technological advances cause a short-run productivity decline as workers move down the new technology's learning curve (e.g., Greenwood and Yorukoglu, 1997; Hornstein and Krusell, 1996).

Figure 9 OWN TECHNOLOGY SHOCKS IN PROCESS-INNOVATING INDUSTRIES

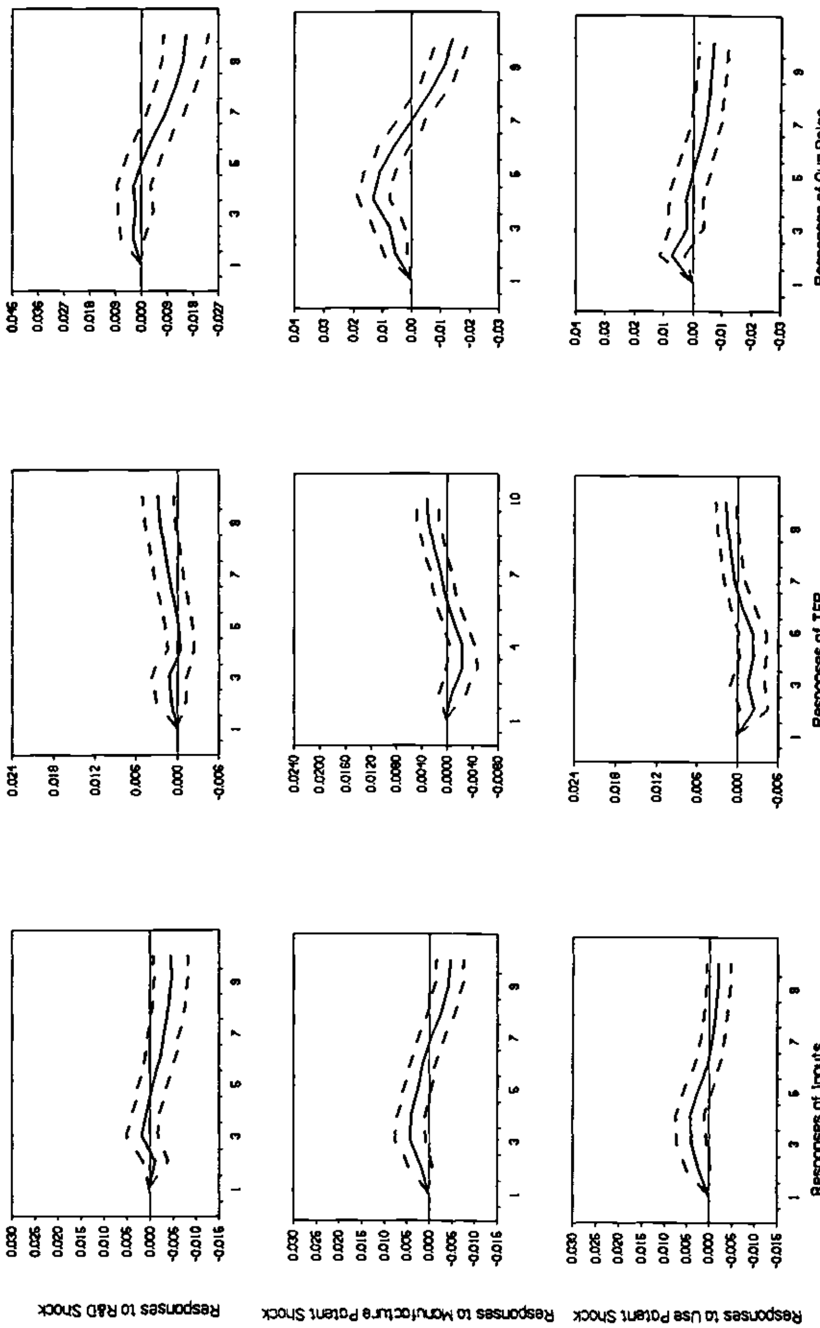

Figure 10 OWN TECHNOLOGY SHOCKS IN PRODUCT-INNOVATING INDUSTRIES

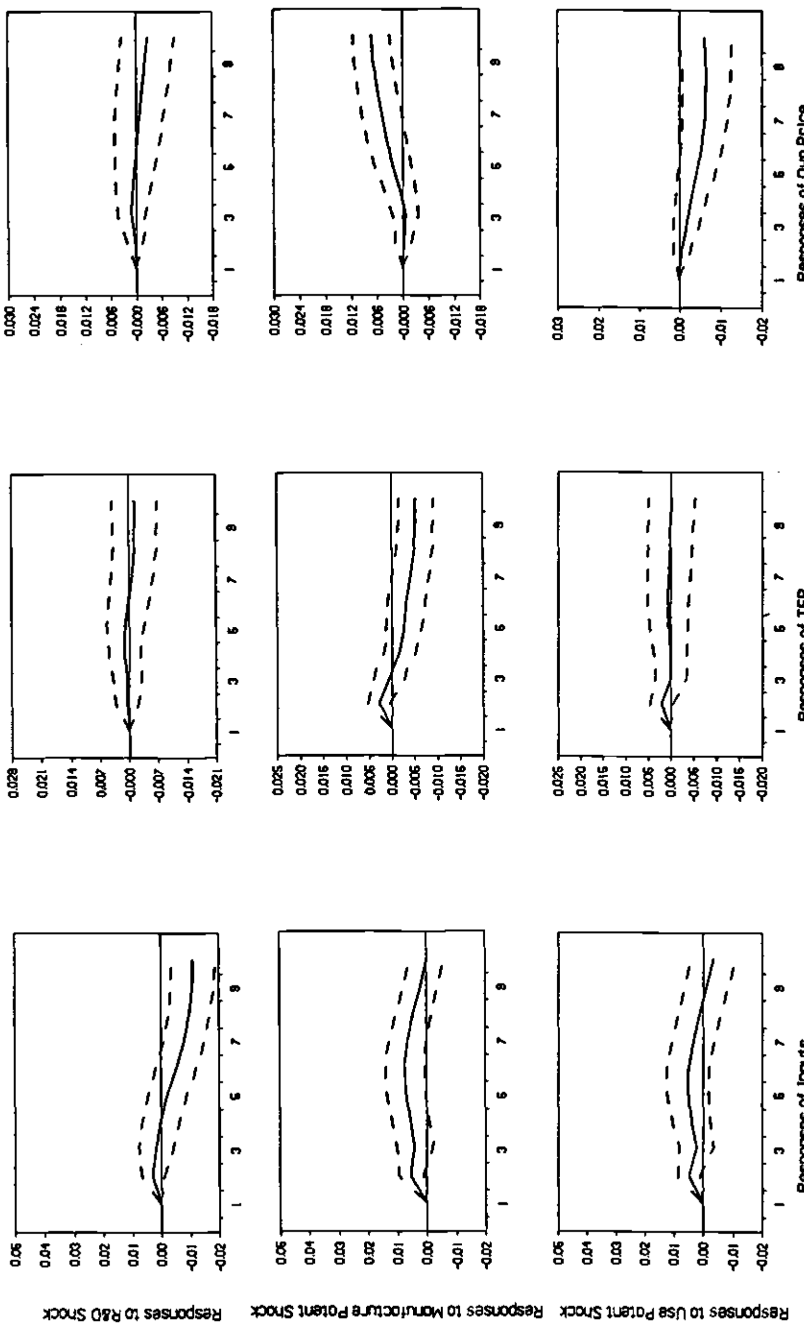

5. *Conclusion*

This paper's contribution is to estimate the impact of technology shocks on the economy using R&D spending and patent applications rather than observed total factor productivity to measure technology. The most surprising finding is that favorable technology shocks do not raise measured TFP at any horizon. Taken at face value, this suggests that observed procyclical variation in TFP is entirely due to factors such as increasing returns, cyclical utilization, and factor reallocation, and not at all due to procyclical technology. It also suggests that efforts to measure short-run changes in true technology by purging measured TFP of movements due to cyclical utilization and so on (e.g., Basu, Fernald, and Kimball, 1997; Burnside, Eichenbaum, and Rebelo, 1996) may be doomed from the start.

Of course, another interpretation of my results is that my R&D and patent data are riddled with measurement error that biases me against finding a significant impact of technology on TFP. While the R&D and patent data are certainly vulnerable to criticism, my results cannot be so easily dismissed. Measurement error should bias me against finding a significant impact of technology on anything. Yet I find that favorable technology shocks have a significant short-run expansionary impact on labor, a significant long-run contractionary impact on total input, and a significant positive impact on capital and nonproduction worker intensities. I also find that technology shocks raise long-run TFP and reduce long-run prices in a subsample of industries dominated by process R&D. These results suggest that the important measurement error is not in R&D or patents, but in output prices. Most R&D in the United States is devoted to product innovations, yet many observers believe that available price data systematically ignore real price declines due to quality improvements and new-product introductions. Similarly, a good deal of the impact of industrial R&D is felt in downstream nonmanufacturing sectors, yet many observers argue that price changes of all kinds in nonmanufacturing are poorly measured.

It is quite possible, then, that technology shocks are more important to actual output and TFP fluctuations than they are to observed fluctuations. To paraphrase Ed Prescott (1986), theory may be ahead of business-cycle measurement. If real-business-cycle enthusiasts want to convince the profession that technology shocks are genuinely important to business cycles, their first order of business should be to construct historical price series for manufacturing and nonmanufacturing sectors that correct for quality improvements and new product introductions, following the painstaking work of Gordon (1990) on durable goods.

Such a project will surely require many hours of research into the history of product innovations in particular sectors, but imagine how different the profession would be today had Friedman and Schwartz (1963) not devoted many hours of research to the history of monetary institutions and monetary shocks.

REFERENCES

Baily, M., and R. Gordon (1988). The productivity slowdown, measurement issues and the explosion of computer power. *Brookings Papers on Economic Activity* 2:347–431.

Bartlesman, E., and W. Gray. (1996). The NBER manufacturing productivity database. Cambridge, MA: National Bureau of Economic Research. NBER Technical Working Paper 205.

Basu, S., and J. Fernald. (1997). Returns to scale in US production: Estimates and implications. *Journal of Political Economy* 105:249–283.

———, ———, and M. Kimball. (1997). Are technology improvements contractionary? University of Michigan. Mimeo.

Berman, E., J. Bound, and Z. Griliches. (1994). Changes in the demand for skilled labor within U.S. manufacturing: Evidence from the Annual Survey of Manufactures. *Quarterly Journal of Economics* 109:367–397.

Burnside, C., M. Eichenbaum, and S. Rebelo. (1995). Capital utilization and returns to scale. In *NBER Macroeconomics Annual 1995*, B. Bernanke and J. Rotemberg (eds.). Cambridge, MA: The MIT Press, pp. 67–110.

———, ———, and ———. (1996). Sectoral Solow residuals. *European Economic Review* 40:861–869.

Caballero, R., and M. Hammour. (1994). The cleansing effect of recessions. *American Economic Review* 84:1350–1368.

———, and ———. (1996). On the timing and efficiency of creative destruction. *Quarterly Journal of Economics* 111:805–852.

Carlstrom, C., and T. Fuerst. (1997). Agency costs, net worth and business fluctuations: A computable general equilibrium analysis. *American Economic Review* 87:893–910.

Christiano, L., and M. Eichenbaum. (1992). Current real business cycle theory and aggregate labor market fluctuations. *American Economic Review* 82:430–450.

———, and ———. (1995). Liquidity effects, monetary policy and the business cycle. *Journal of Money, Credit and Banking* 27:1113–1136.

———, ———, and C. Evans. (1998). Monetary policy shocks: What have we learned and to what end? Cambridge, MA: National Bureau of Economic Research. NBER Working Paper 6400.

Cochrane, J. (1994). Shocks. *Carnegie-Rochester Conference Series on Public Policy* 41:296–364.

Cogley, T., and J. Nason (1995). Output dynamics in real business cycle models. *American Economic Review* 85:492–511.

DeLong, B., and L. Summers. (1991). Equipment investment and economic growth. *Quarterly Journal of Economics* 106:445–502.

Evans, C. (1992). Productivity shocks and real business cycles. *Journal of Monetary Economics* 29:191–208.

Farmer, R., and J.-T. Guo. (1994). Real business cycles and the animal spirits hypothesis. *Journal of Economic Theory* 63:42–72.

Friedman, M., and A. Schwartz. (1963). *A Monetary History of the United States, 1867–1960.* Princeton: Princeton University Press.

Galí, J. (1996). Technology, employment and the business cycle: Do technology shocks explain aggregate fluctuations? Cambridge, MA: National Bureau of Economic Research. NBER Working Paper 5721. *American Economic Review,* forthcoming.

———, and M. Hammour. (1992). Long run effects of business cycles. Mimeo.

Gordon, R. (1990). *The Measurement of Durable Goods Prices.* Chicago: University of Chicago Press.

Gort, M., and S. Klepper. (1982). Time paths in the diffusion of product innovations. *Economic Journal* 92:630–653.

Greenwood, J., and M. Yorukoglu. (1997). 1974. *Carnegie-Rochester Conference Series on Public Policy* 46:49–95.

Griliches, Z. (1973). Research expenditures and growth accounting. In *Science and Technology in Economic Growth,* B. R. Williams (ed.). New York: John Wiley and Sons, pp. 59–83.

———. (1989). Patents: Recent trends and puzzles. *Brookings Papers: Microeconomics,* 291–319.

———, and F. Lichtenberg. (1984). R&D and productivity growth at the industry level: Is there still a relationship? In *R&D, Patents and Productivity,* Zvi Griliches (ed.). Chicago: University of Chicago Press, pp. 465–501.

Hall, R. (1988). The relation between price and marginal cost in US industry. *Journal of Political Economy* 96:921–947.

Hornstein, A., and P. Krusell. (1996). Can technology improvements cause productivity slowdowns? In *NBER Macroeconomics Annual 1996,* B. Bernanke and J. Rotemberg (eds.). Cambridge: The MIT Press, pp. 209–259.

Horvath, M. (1997). Cyclicality and sectoral linkages: Aggregate fluctuations from independent sectoral shocks. Stanford University. Mimeo.

Jorgenson, D., F. Gollop, and B. Fraumeni. (1987). *Productivity and U.S. Economic Growth,* Cambridge: Harvard University Press.

Jovanovic, B., and S. Lach. (1997). Product innovation and the business cycle. *International Economic Review* 38:3–22.

Keller, W. (1997). Trade and the transmission of technology. Cambridge, MA: National Bureau of Economic Research. NBER Working Paper 6113.

Kimball, M. (1995). The quantitative analytics of the basic neomonetarist model. *Journal of Money, Credit and Banking* 27:1241–1289.

Kortum, S. (1993). Equilibrium R&D and the patent–R&D ratio: US evidence. *American Economic Review* 83(May):450–457.

———, and J. Putnam. (1997). Assigning patents to industries: Tests of the Yale Technology Concordance. *Economic Systems Research* 9:161–175.

Krusell, P., L. Ohanian, J. V. Rios-Rull, and G. Violante. (1996). Capital–skill complementarity and inequality. University of Rochester. Mimeo.

Kydland, F., and E. Prescott. (1982). Time to build and aggregate fluctuations. *Econometrica* 50:1345–1370.

Lach, S. (1995). Patents and productivity growth at the industry level: A first look. *Economics Letters* 49:101–108.

———, and R. Rob. (1996). R&D, investment and industry dynamics. *Journal of Economics and Management Strategy* 5:217–249.

———, and M. Schankerman. (1989). Dynamics of R&D and investment in the scientific sector. *Journal of Political Economy* 97:880–904.

Levin, A., and C.-F. Lin. (1992). Unit root tests in panel data: Asymptotic and finite-sample properties. University of California, San Diego. Mimeo.

Lichtenberg, F., and D. Siegel. (1991). The impact of R&D investment on productivity—new evidence using linked R&D–LRD data. *Economic Inquiry* 29:203–228.

Long, J., and C. Plosser. (1983). Real business cycles. *Journal of Political Economy* 91:39–69.

Prescott, E. (1986). Theory ahead of business cycle measurement. *Federal Reserve Bank of Minneapolis Quarterly Review* 10:9–22.

Quah, D. (1989). Permanent and transitory movements in labor income: An explanation for "excess smoothness" in consumption. *Journal of Political Economy* 98:449–475.

Rotemberg, J., and M. Woodford. (1995). Dynamic general equilibrium models with imperfectly competitive product markets. In *Frontiers of Business Cycle Research,* Thomas Cooley (ed.). Princeton, NJ: Princeton University Press, pp. 243–293.

Scherer, F. (1984). Using linked patent and R&D data to measure interindustry technology flows. In *R&D, Patents and Productivity,* Z. Griliches (ed.). Chicago: University of Chicago Press, pp. 417–464.

———. (1993). Lagging productivity growth: Measurement, technology and shock effects. *Empirica* 20:5–24.

Schmitt-Grohe, S. (1997). Comparing four models of aggregate fluctuations due to self-fulfilling expectations. *Journal of Economic Theory* 72:96–147.

Shea, J. (1991). The input–output approach to instrument selection: Technical appendix. University of Maryland. Mimeo.

———. (1993). The input–output approach to instrument selection. *Journal of Business and Economic Statistics* 11:145–154.

———. (1998). What do technology shocks do? Cambridge, MA: National Bureau of Economic Research. NBER Working Paper no. 6632, July.

Stadler, G. (1990). Business cycle models with endogenous technology. *American Economic Review* 80:763–778.

Terleckyj, N. (1975). Direct and indirect effects of industrial research and development on the productivity growth of industries. In *New Developments in Productivity Measurement,* J. Kendrick and B. Vaccara (eds.). Chicago: University of Chicago Press, pp. 359–386.

Comment

JORDI GALÍ
New York University and NBER

1. Technology and Business Cycles, in the Theory and in the Data

Under the world view advocated by real business cycle (RBC) economists, observed economic fluctuations can be interpreted, to a first approximation, as the result of agents' optimal responses to changes in aggregate technology, in an environment with perfect competition, market clearing, and flexible prices. The empirical basis for that claim lies in the ability of calibrated RBC models to match patterns of unconditional second moments of a number of macroeconomic time series.

Though that ability is largely acknowledged, some recent research has undertaken "more direct" assessments of RBC models—and of alternative business-cycle frameworks—by identifying the dynamic effects of variations in technology on different macroeconomic variables, and by evaluating quantitatively their role as a source of short-run economic fluctuations. Many interesting questions that may shed light on the nature of business cycles are addressed in that literature: What are the effects of technology shocks in actual economies? How do they differ from the predictions of standard RBC models? What is their contribution to business-cycle fluctuations? Shea's present paper fits squarely into that line of research.

The key challenge facing such an inquiry lies in the empirical identification of exogenous technology shocks, since it is generally accepted that conventional Solow residuals cannot be taken as reliable measures of "true" total factor productivity. Different approaches to identification have been pursued in the literature. Thus, Basu, Fernald, and Kimball (1997) identify technology shocks as the innovation in an "adjusted" Solow residual series, where the adjustment attempts to correct for the bias associated with the potential presence of increasing returns, imperfect competition, variable input utilization, and sectoral reallocation of inputs across heterogeneous sectors. In my own work (Galí, 1996) identification is achieved by restricting permanent technology shocks to be the only source of the unit root in labor productivity, i.e., the only shocks that may have a permanent effect on the level of that variable—a restriction that can be shown to hold under assumptions typically made in standard models.

Despite the different methodologies and data used, a number of common findings emerge in those papers, including a result that appears to be very robust: exogenous improvements in technology tend to reduce employment, at least in the short run.[1] That result is in stark contrast to the prediction of the standard RBC models, where a positive response of employment to an exogenous increase in TFP is at the center of the mechanism underlying business cycles. On the other hand, and as argued by the above-mentioned authors, the estimated effects of technol-

1. Basu, Fernald, and Kimball (1997) use postwar U.S. annual data. Galí (1996) uses quarterly data for the U.S. as well as the remaining G7 countries. Galí's methodology has been applied by Kiley (1997) to two-digit industry-level U.S. data.

ogy shocks are consistent with the prediction of New Keynesian models characterized by imperfect competition and sticky prices. The intuition underlying that prediction is straightforward: if nominal aggregate demand is predetermined and prices respond only sluggishly to shocks, real aggregate demand (and thus, the position of the demand schedule facing each firm) will change little in response to an increase in TFP. Accordingly, the quantity of goods each firm will wish to produce and sell will also remain largely unchanged (since, by assumption, adjusting prices downward is either unfeasible or too costly). Since firms now have access to a more efficient technology, that level of output can now be produced with less inputs, thus leading to a decline in employment (and, presumably, a lower capital utilization as well). Needless to say, such effects and the mechanism through which they are transmitted have little in common with those underlying RBC models.

2. *Shea's Empirical Framework*

Shea's "What Do Technology Shocks Do?" has a similar motivation to the papers by Basu, Fernald, and Kimball (1997) and Galí (1996), but uses a different strategy in order to identify exogenous technology shocks. Specifically, Shea's approach exploits the availability of data on "tangible" activities associated with technological innovation, and the fact that indicators of such activities display non-negligible short-run fluctuations.

Of course, the study of the empirical links between innovative activities and productivity measures has been the subject of a time-honored empirical literature.[2] Work in that tradition generally takes variations in technology (over time and/or across firms or industries) as the phenomenon to be explained, trying to detect and quantify the relationship between measures of technological change—typically, estimates of total factor productivity (TFP)—and indicators of innovative activity (R&D expenditures or patent applications). From the point of view of that literature, Shea's main contribution lies in the use of a structural VAR to model the connection between technological innovation and TFP growth, an approach which allows for largely unrestricted dynamics, including the possibility of an endogenous response of technological innovation to fluctuations in each industry's level of economic activity.

Shea uses two different variables as indicators of technological innovation: R&D (i.e., a measure of the input in the innovation process) and the number of patent applications (a measure of the output of that process). For each industry he estimates the responses of inputs and TFP to or-

2. See, e.g., the Griliches (1984) NBER volume devoted to the subject.

thogonalized shocks (1) to the industry's own technology indicators, and (2) to the technology indicators of upstream industries (constructed using input–output data).

It is important to distinguish between the two types of shocks, since the predicted responses are likely to be very different. In particular, technological innovation in upstream industries should not affect an industry's own TFP (unless technological spillovers are present), but only its marginal cost (to the extent that it is reflected in suppliers' prices). Thus, there is no reason why, in the presence of sticky prices, input use in an industry should decline in the face of a positive shock in upstream industries, as opposed to an analogous shock in the own industry. Hence, and for the sake of brevity and comparability with the existing literature, I will concentrate the remainder of my comments on the results based on own technological innovations.

Shea assumes that the level of R&D (or, alternatively, the number of patents) in industry i is determined by the equation

$$z_t^i = \sum_{j\geq 1} \alpha_j z_{t-j}^i + \sum_{j\geq 0} \beta_j y_{t-j}^i + \epsilon_t^i,$$

where z_t is the value taken by the industry technology indicator, y_t is a vector including industry TFP and inputs measures, and ϵ_t^i represents the "exogenous" technology shock. The key identifying assumption is that ϵ_t^i is orthogonal to y_t^i (as well as its lags), i.e., that a technology shock does *not* have a *contemporaneous* (within the year) impact on industry aggregates such as TFP or inputs.

The dynamic effects of technology on any industry variable y^i are then given by the sequence of coefficients $\{\phi_j\}$ of the regression equation

$$y_t^i = \sum_{j\geq 1} \phi_j \epsilon_{t-j}^i + u_t^i.$$

The structural VAR approach adopted by Shea seems, in principle, clearly suitable for the issue at hand, for it allows for an endogenous component in innovative activities, as well as rich, largely unrestricted dynamics. Furthermore, the recursive restriction used (namely, that shocks to R&D or patents cannot have a contemporaneous impact on industry aggregates such as TFP or inputs) seems reasonable when R&D data are used, given the likely lags between R&D expenditures and actual implementation of the resulting innovation (it may be more questionable when patent data are used).

Given the previous setup, Shea's question can be formulated as follows: what are the effects of exogenous variations in industry R&D or patents on the industry's TFP and its level of activity?

3. *Shea's Main Results and their Interpretation*

Let me focus on two of the results emphasized by Shea among those that are claimed to be reasonably robust across specifications:

RESULT 1 A positive own technology shock (i.e., one associated with an increase in R&D or the number of patents in the same industry) has no significant effect on industry TFP at any horizon.

RESULT 2 A positive own technology shock tends to increase inputs (especially labor), in the short run, but decrease them in the long run.

My main concern has to do with the economic interpretation of those results, i.e., what we learn from them regarding the merits of alternative business-cycle models. Standard business-cycle frameworks do not explicitly model the process of innovation in the R&D sector (or, for that matter, the process of patent application, grant, and diffusion of associated knowledge). Hence, Shea's technology shocks do not have an exact counterpart in those models, in which changes in TFP are taken to be exogenous (wisely or not). Given that the estimated response of TFP to a "Shea technology shock" is essentially flat, it is not clear what sort of prediction of those models could be refuted by looking at Results 1 and 2. In particular, models with imperfect competition and sticky prices predict a decline in employment in response to a technology shock *only if the latter is associated with an increase in TFP.* Since in Shea's evidence the level of TFP remains essentially unchanged in response to an orthogonalized innovation in own technology shocks (R&D or patents), the absence of a response of employment and other inputs to the same shock would seem to be consistent with the prediction of a conventional sticky price model. Thus, Results 1 and 2 can hardly be interpreted as providing evidence against New Keynesian models. Neither can they be seen as being in contradiction with the results reported in Galí (1996) and Basu, Fernald, and Kimball (1997), for those authors found evidence of a decline in employment after a technology shock that "succeeds" in raising productivity.

Most interestingly, in the only two VAR specifications for which Shea detects a significant short-run change in TFP in response to an (own) technology shock (namely, when the analysis is restricted to process-innovating industries and patents are used as a technology indicator), inputs are shown to respond in the direction opposite to the movement in TFP, i.e., in a way consistent with the predictions of sticky-price models!

In the remaining cases, the only hypothesis supported by the Shea's evidence is that innovations to industry R&D or patents have *no* significant dynamic effects on the same industry's TFP. That is an intriguing result, but not one with obvious implications for business-cycle theory.

4. *Are All Industries Alike?*

In spite of some of the rhetoric found both in the paper and in the present discussion, it is not completely true that the dynamics of the estimated model are "essentially unrestricted": Shea's methodology constrains the dynamic responses to shocks to be the same for all industries. Of course, if that restriction is satisfied in the population, its imposition in the estimation procedure can only increase the precision of the estimates. But a look at some simple statistics gives us a reason to be somewhat suspicious. Table 1 reports cross-correlations of R&D (or patents) and TFP, industry by industry, using Shea's data set. In addition to the contemporaneous correlation, the highest cross-correlation (in absolute

Table 1 INDUSTRY CROSS-CORRELATIONS

	R&D		*Patents*	
Industry	$\rho(A_t, Z_t)$	$\rho(A_{t+k}, Z_t)$	$\rho(A_t, Z_t)$	$\rho(A_{t+k}, Z_t)$
Food	0.61	0.65(−1)	−0.37	−0.49 (+4)
Textiles	0.05	−0.57 (+4)	0.11	0.51 (−2)
Lumber	−0.54	−0.54 (+0)	0.22	0.61 (−3)
Paper	−0.09	0.42 (+3)	−0.17	−0.17 (+0)
Ind. chem.	0.29	0.52 (+3)	0.53	0.58 (−1)
Drugs	−0.92	−0.92 (+0)	−0.85	−0.85 (+0)
Other chem.	−0.40	−0.58 (−4)	0.71	0.71 (+0)
Petroleum	0.29	0.73 (+3)	0.11	0.67 (+4)
Rubber	0.03	−0.38 (+4)	0.35	0.62 (−3)
Stone	0.39	0.62 (+2)	0.53	0.72 (−2)
Metals	−0.26	−0.34 (+2)	0.38	0.48 (−2)
Metal prod.	−0.05	−0.17 (−2)	0.43	0.49 (−2)
Computers	0.97	0.97 (+0)	−0.29	−0.72 (+4)
Other nonelec.	−0.41	−0.50 (+2)	0.74	0.74 (+0)
Elec. & Commun.	−0.20	0.42 (−4)	−0.18	−0.38 (+3)
Oth. elec. eq.	0.07	0.55 (−4)	0.58	0.58 (+0)
Autos & Trans.	0.05	−0.37 (−4)	0.54	0.64 (−1)
Aerospace	−0.51	−0.57 (+1)	−0.06	0.36 (+4)
Instruments	−0.72	−0.72 (+0)	0.32	0.57 (+3)

$\rho(A_t, Z_t)$ is the contemporaneous correlation between TFP and R&D (or patents). $\rho(A_{t+k}, Z_t)$ is the element of the cross-correlogram of the same variables with the highest (absolute) value, with the corresponding lead or lag being shown in brackets. Data are detrended.

value) and the lag or lead at which the latter is found are shown in parentheses. No clear common pattern emerges: values and signs are all over the place, pointing to substantial heterogeneity across industries (even across indicators) in the size and timing of the effects of technology shocks, and/or in the properties of the endogenous component of the innovation indicators. That result may not be surprising, since after all, industries as different in the nature of their production processes as lumber and aerospace are included in Shea's sample. But it clearly raises some doubts about the usefulness (and meaning) of any estimates that fail to take such heterogeneity into account.

5. *Concluding Comments*

John Shea has written a paper that addresses an issue at the center of some of the macroeconomic controversies of the past twenty years: the link between technological change and business cycles. In doing so he makes use of a data set as well as an empirical approach that seem very well suited to the issue at hand. Yet, many readers are likely to find some of the results somewhat disappointing, in the sense that they seem to raise more questions than they answer. The inability, for most specifications, to detect a significant effect of an industry's R&D expenditures or patent applications on its own TFP is worrisome, for it is hard to think of many other factors that may underlie variations in productivity measures at horizons other than the long run (though mismeasurement, cyclical or other, is always a likely candidate).

Most of the results also seem to fall short of yielding any obvious lessons that could further our understanding of business cycles and/or help evaluate the empirical merits of alternative models. Exceptions apply, however: as I have argued above, some of the few significant results seem to lend some additional support to the existing evidence that points to a short-run negative comovement between TFP and inputs in response to technology shocks, which is consistent with sticky-price models.

I am convinced that some of the interesting issues raised by Shea's approach and results will stimulate further work on the subject. We can only hope that, when alternative empirical models or data sets are used to ask the question that gives Shea's paper its title, the data choose to speak somewhat louder.

REFERENCES

Basu, S., J. Fernald, and M. Kimball. (1997). Are technology improvements contractionary? University of Michigan. Mimeo.

Galí, J. (1996). Technology, employment and the business cycle: Do technology shocks explain economic fluctuations? Cambridge, MA: National Bureau of Economic Research. NBER Working Paper 5721. Forthcoming in *American Economic Review.*
Griliches, Z. (1984). *R&D, Patents, and Productivity.* Chicago: University of Chicago Press.
Kiley, M. (1997). Labor productivity in U.S. manufacturing: Does sectoral comovement reflect technology shocks? Federal Reserve Board. Mimeo.

Comment

ADAM B. JAFFE
Brandeis University and NBER

This paper describes a valuable empirical exercise that was carried out with care and is presented with clarity. It provides a comprehensive summary of the empirical relationships, at the level of approximately two-digit SIC industries, among R&D, patents, other inputs, measured output, measured total factor productivity (TFP), and measured output prices. The results can be distilled into the following set of stylized facts:

1. There is a clear cross-sectional relationship across industries between indicators of technological activity and the growth rate of measured TFP.
2. For the set of industries taken as a whole, it is impossible to "find" this effect of technology on measured TFP in the form of significant impulse responses at any time horizon to technology shocks.
3. For a set of industries dominated by process rather than product innovation, there is a significant long-run positive response of measured TFP to technology shocks.
4. For all industries taken together, technology shocks have a significant short-run expansionary effect on labor input.
5. For all industries taken together, technology shocks have a significant long-run contractionary effect on total input, and induce significant long-run substitution of capital and nonproduction labor for production labor.

Shea concludes from these stylized facts that (1) technology shocks cannot be viewed as a significant driver of short-run movements in *measured* output or productivity for the economy as a whole; (2) technology shocks may be driving *actual* output and productivity even in the short run; and (3) macroeconomists should devote significant time and attention to the history of product innovation in particular sectors, in order to develop price series that would permit accurate measurement of output and productivity. While I am certainly not going to argue with the third conclusion, I believe that there is sufficient uncertainty about what is really going on in the data that the first conclusion must be somewhat qualified.

The separation of industries between those dominated by product innovations and those dominated by process innovations is a valuable and important improvement in the final version of the paper. I agree that the results for the process-innovation-dominated industries suggest that the absence of technology shock responses in the overall pooled sample is probably due, at least in part, to the failure of the output series to incorporate appropriately the effects of new products and quality improvements. I also agree that this partition confirms the results in the overall panel in which patents classified by industry of use typically have clearer effects on TFP than patents classified by industry of manufacture. Quite apart from its macro implications, this is a nice contribution to the industry-level literature on interindustry technology flows.

The only cloud over this otherwise sunny picture is that the greater effect of patents by industry of use relative to patents by industry of manufacture continues to hold when one looks at *upstream* patents. In principle, innovations "used" in my upstream industries should lower their costs and hence the prices of my inputs, but should not affect my measured productivity. In contrast, innovations "manufactured" in my upstream industries create precisely the kind of unmeasured product improvement that we are worried about; if my (actual) inputs are improving in quality but my (measured) inputs are not, then I should be having improvements in my (measured) TFP. In other words, the very story that explains the greater effect on "own" productivity of patents by industry of use relative to patents by industry of manufacture implies that this relationship ought to be reversed in the upstream patents. The fact that it is not suggests that some other factor may be at work. Econometrically, the pattern of results could be explained by inherently greater measurement error in the classification by industry of manufacture. But this doesn't seem very plausible; if anything, it is harder to figure out where something will be used than it is to figure out where it is made. This issue probably merits further exploration, looking in more detail at specific industries and examining how differences between the two patent totals drive the results.

The finding that the results differ significantly between product-innovating and process-innovating industries also confirms what one would suspect more broadly, which is that neither the magnitude nor the timing of technology responses is likely to be the same across such diverse industries. While it is perhaps possible to interpret the overall results as mean effects in a random-coefficients framework, the likely fragility of the results with respect to other partitions suggests that we should at least hesitate before concluding that technology shocks have no effect at the industry level.

More fundamentally, it remains quite plausible to me that, even at the level of individual industries (as defined here), one might find no effect of technology shocks at any specific time horizon, even if they did have real effects on measured TFP. The simplest way to view the above stylized facts about the relationship between technology and TFP is that technology does affect (measured) TFP, but it does so with lags that are so variable that it is impossible to pin down the timing of the effects. I am not a macroeconomist and do not know if a model in which technology shocks produce productivity and output impacts with highly variable lags will generate business cycles. If not, then these results may suggest that technology shocks are not a likely explanation for observed business cycles, but this is a little different from saying that they have no impact on measured TFP at the industry level.

Shea does a very good job of discussing the potential problems and limitations with the R&D and patent series as indicators of technological improvement. In addition to the problems that he discusses, I would emphasize the highly aggregated nature of the sectors that are the units of observation here, and the difficulty of R&D and patent series in measuring really big innovations. If one thinks of likely candidates for technology shocks, it is not hard to think of examples of innovations that have had perceptible effects on output or productivity, but would probably not be visible in the R&D or patents series for sectors as broad as those defined here. Consider Viagra relative to all prescription and nonprescription drugs; Aspartame relative to the entire food sector, or the electric-arc steel furnace relative to the entire metals sector. It is in some sense the essence of important shocks that their effects on productivity or output are highly disproportionate to the R&D or patents associated with them. For this reason, it is possible for the industry TFP series to be reflecting real technology shocks even if those shocks were not associated with R&D or patents.

Shea correctly points out that this kind of measurement error in the technology indicators ought to bias the results towards finding no effects of technology shocks, while he does find significant effects at some horizons on other inputs. He concludes from this that his negative conclusion about the effects of technology shocks on TFP and output should be taken at face value. While I agree that finding significant technology effects on inputs puts some qualitative limit on the amount of measurement error, it seems a bit of a jump to rule out its having a material impact on the TFP results.

My queasiness about using the significant effects of technology on inputs to conclude that the negative conclusion about technology and TFP is real is strengthened by the nature of these significant input effects. Specifically, the significantly negative long-run contractionary effect of technology shocks on total inputs seems inconsistent with the rest of the story. If inputs are falling because of process innovation, then there is no reason why this should not be seen in a long-run positive impact on measured TFP. On the other hand, if the story is that in the complete panel technology is largely producing product innovations that are not picked up in measured TFP, how can it be that technology shocks significantly reduce total input in the long run? If TFP is unchanged, than a significant decline in total inputs implies a significant decline in (measured) output. While we know that the *increase* in real output corresponding to improved products may not be reflected in measured output, it is hard for me to understand how product innovation would lead to a significant long-run *decline* in measured output. Overall, then, the significant pattern of effects of the technology indicators on inputs is not really consistent with the interpretation given to those indicators, so I cannot take much comfort from it.

It is the nature of a Comment to focus on criticisms and areas of disagreement. Despite my having done so, it is still true that I learned a lot from this paper, and believe that this kind of empirical analysis is crucial to a better understanding of the linkages between the macro economy and the microeconomic phenomena of innovation and technological change.

Discussion

The finding that the link between innovative activity and short-run output fluctuations is weak generated considerable discussion. Several participants, including David Backus, Susanto Basu, and Russell Cooper, stressed the point that, in the business-cycle literature, a "technology shock" occurs only at the time at which output is affected—not at the time that the inventive process begins or a technological innovation is patented. Various factors, including slow diffusion of knowledge, the need to work out details of implementation, and the time needed to adjust factors of production, can lead to long and varying lags between inventive activity and any effect on output.

The fact that measures of R&D fail even to affect total factor productivity (TFP) is a significant puzzle, as Basu noted. Julio Rotemberg pointed out that this result seems inconsistent with the finding that R&D activity and TFP growth are positively correlated across industries. Shea responded that the cross-sectional finding is not inconsistent with the conclusion that R&D activity has little effect on TFP or output in the short run. Michael Woodford added that the cross-sectional results may be spurious, in that R&D activity and TFP growth might be jointly determined across industries by some third factor. Henning Bohn suggested that as Shea's measures of R&D do not affect TFP, it is possible that there exist unmeasured sources of technical change which determine TFP; and therefore that we shouldn't necessarily conclude that TFP changes are irrelevant for business fluctuations. Backus noted that short-run output and TFP dynamics are more likely to reflect the underlying transmission mechanism than the nature of the shocks themselves, and that these shocks need not be technological to account for observed behavior.

A potentially important but hard-to-measure source of TFP growth is organizational change. As an illustration, Backus stressed the important effects on productivity of organizational changes in the airline industry. Shea and Rotemberg expressed skepticism that changes in organization were likely to have high-frequency effects.

John Cochrane questioned the implicit restriction that only unanticipated changes in research and development have real effects. Unlike anticipated changes in monetary policy, anticipated changes in R&D should have important effects on measured TFP and output. Shea defended his approach as an identification assumption; while anticipated R&D can have real effects, it is difficult to disentangle those effects from the effects of other shocks that change both R&D and output.

Simon Gilchrist pointed out that the effects of technology shocks on input use depends on auxiliary assumptions about the economy. For example, they will depend on the nature of the monetary authority's responses to various shocks. Jordi Galí cited results from his own work supporting the view that the Fed does not accommodate technology shocks sufficiently.