

Sector-Specific Technical Change

Susanto Basu, John Fernald, Jonas Fisher, and Miles Kimball¹

November 2013

Abstract: Theory implies that the economy responds differently to technology shocks that affect the production of consumption versus investment goods. We estimate industry-level technology innovations and use the input-output tables to relax the typical assumptions in the investment-specific technical change literature—assumptions that, we find, do not hold in the data. We find that investment-technology improvements are sharply contractionary for hours, investment, consumption, and output. Consumption-technology improvements, on the contrary, are generally expansionary. Thus, disaggregating technology shocks into consumption and investment-specific changes yields two shocks that both produce business-cycle comovement, and also explain a large fraction of annual changes in GDP and its components. Most of the responses we find are consistent with the predictions of simple two-sector models with sticky prices.

JEL Codes: E32, O41, O47

Keywords: Investment-specific technical change, multisector growth models, business cycles

¹ Boston College and NBER, Federal Reserve Bank of San Francisco, Federal Reserve Bank of Chicago, and University of Michigan and NBER, respectively. We thank Faisal Ahmed, Titan Alon, Laurie Glapa, David Thippavong, Kuni Natsuki, and, especially, Kyle Matoba and Stephanie Wang for superb research assistance. We thank seminar participants at several institutions and conferences. The views expressed in this paper are those of the authors, and do not necessarily reflect the views of anyone else affiliated with the Federal Reserve System.

I. Introduction

What shocks drive business cycles? At a minimum, such shocks must move output, consumption, investment and hours worked in the same direction, as this positive comovement is a defining characteristic of business cycles. Technology shocks are attractive candidates, because they can reproduce this comovement in simple, general-equilibrium models of fluctuations. But in the data, technology shocks identified using restrictions from one-sector models do not produce positive comovement. In particular, Gali (1999) and Basu, Fernald and Kimball (2006) find that technology shocks often raise output and consumption, but lower hours worked. Reviewing this evidence, Francis and Ramey (2005) conclude “the original technology-driven ... business cycle hypothesis does appear to be dead.”

We revisit the issue of technology shocks and business-cycle comovement because there is no consensus in the literature on other shocks that could plausibly account for the bulk of the fluctuations that we observe. (For example, monetary shocks produce the right comovements, but are generally estimated to account for a very small fraction of economic volatility.) Our goal is to see whether the conclusions about technology shocks reached earlier might be due to imposing the restrictions of a one-sector model of growth and fluctuations in a world where they do not apply. If there are multiple technologies, each affecting a different type of good, might it be the case that each type of shock individually generates the proper comovements even though their average does not?

We first show that our hypothesis has a firm basis in economic theory. Theory tells us that the composition of technological change, in terms of the type of final good whose production is affected, matters for the dynamics of the economy’s response to technology shocks. For example, we show that a shock to the technology for producing consumption goods has no interesting

business-cycle effects in a standard RBC model. All of the dynamics of hours and investment stressed by RBC theory come from shocks to the technology for producing investment goods.² In an open economy, terms-of-trade shocks affect the economy's ability to provide consumption and investment goods for final use, even if no domestic producer has had a change in technology.

We then present a new and more robust method for estimating sector-specific technical change. Existing methods for investigating the importance of sector-specific technology are based primarily on relative movements in the price deflators for investment and consumption goods, an approach pioneered in an important paper by Greenwood, Hercowitz and Krusell (1997; henceforth GHK). Our evidence, in contrast, is based on an augmented growth-accounting approach, where we estimate technology change at a disaggregated industry level and then use the input-output tables to aggregate these changes.

Our approach offers several advantages. First, we can relax and, indeed, test many of the assumptions necessary for relative-price movements to correctly measure relative changes in domestic technological change. For example, suppose different producers have different factor shares or face different input prices; or suppose markups change over time. Then relative prices do not properly measure relative technical change.³ Second, we discuss extensions to the open economy, where the ability to import and export means that relative investment prices need not measure relative technologies in terms of *domestic* production.

We view our ability to allow for differences in sectoral factor shares as the most important advantage of our method. All methods that use relative price changes to identify technology shocks (with either short- or long-run restrictions) require that all final-use sectors have the same

² Greenwood, Hercowitz and Krusell (2000) provide an insightful approach to business-cycle modeling with sector-specific technical change. However, they used a normalization for sector-specific technology that did not have a pure consumption-augmenting shock, and thus did not uncover this neutrality result.

³ We can also test the orthogonality restrictions used in the structural VAR approach of Fisher (2006), who also uses relative price data but with a long-run restriction.

factor shares. We are able to calculate these final-use shares using our method, and find that they differ significantly across sectors. Thus, one important conclusion of our paper is that relative prices cannot and should not be used to estimate investment-specific technical change.

Applying our method to sectoral US data, we find that consumption-specific and investment-specific technology shocks have very different economic effects. In particular, consumption-specific technology improvements are generally expansionary for all important business-cycle variables (although the rise in hours worked is not statistically significant, the point estimates are positive and of the right economic magnitude). By contrast, investment-specific technology improvements are contractionary for all important business-cycle variables, with statistically significant declines in GDP, consumption, hours worked, and even investment. Importantly, both shocks induce positive comovement among the key business-cycle variables, including output, consumption, investment and hours. Thus, the first major finding of the paper is to present a set of shocks that explain a significant fraction of GDP growth (more than 30 percent), while explaining business-cycle fluctuations in key macro variables.

As we show in the next section, however, neither set of responses is consistent with RBC theory. Clearly RBC theory would not predict a decline in GDP, consumption, hours and, especially, investment following an investment technology *improvement*. Perhaps more surprisingly, the fact that hours and investment rise after a positive consumption technology investment is also evidence against the RBC model. However, Basu, Fernald and Liu (2012) show that these responses are generally consistent with the predictions of simple two-sector sticky-price models.

In conjunction, the theory and empirical work in our paper suggests that we need to revisit several influential empirical literatures using multi-sector economic models. For example, papers in the large literature on the macroeconomic effects of technology shocks typically examine the

responses of macroeconomic variables to aggregate measures of technical change.⁴ But if consumption- and investment-specific shocks are expected to have different economic effects, aggregating the two into a single measure can introduce significant measurement error in an explanatory variable.

To take another example, many papers follow the suggestion of Cochrane (1994) and take the consumption-output ratio to be a measure of permanent income relative to current income. In a general-equilibrium model, the gap is usually due to (expected) changes in productivity.

Rotemberg and Woodford (1996) use this interpretation to fashion an empirical critique of RBC models by comparing the expected transitory changes in macroeconomic variables in the data to those predicted by a one-sector RBC model. But the correlation between expected changes in consumption, output and other variables can be quite different depending on whether the shock affects consumption technology, investment technology, or both. Thus, it may be necessary to revisit Rotemberg and Woodford's influential critique by comparing the data to the predictions of multi-sector RBC models.

Our approach to the data also forces us to think about the role of the open economy in economic fluctuations. That is, our approach requires that we take a stand on the "technology" for producing net exports. Thus, we need to decide whether to take a narrow or broad view of technology. The narrow view is to say that "technology" is just something that shifts a domestic production function. We choose to take a broader view, and define technology from the consumption possibilities frontier: At least in a neoclassical model where there are stable preferences and no distortions, technology is anything that changes consumption for given levels of capital, investment and hours worked. Technology thus defined therefore incorporates changes in

⁴ An exception is Fisher (2006).

the terms of trade. It is clear, of course, that terms of trade changes do not shift domestic production functions, and thus are not technology shocks in the narrow sense.⁵ But it is also true that domestic technology shocks are not the only sources of changes in consumption and welfare, even in a fully neoclassical model. Understanding the behavior of consumption in both the short and long run requires us to take an open-economy, multi-sector approach.

The paper is structured as follows. After this introduction, we provide a simple theoretical example that illustrates that consumption- and investment-specific shocks should have very different economic effects. This example motivates us to write down a macroeconomic model to guide our empirical work. Specifically, the model shows how to map from “messy” microeconomic production and international trade data into simple macroeconomic aggregates. We then discuss the data, present results, and draw conclusions.

II. Consumption-Technology Neutrality

The character of both growth and business cycles depends on the sectoral distribution of technical change. In the neoclassical growth model, capital accumulation arises only if technical change expands the possibilities for producing capital goods. Indeed, as shown by Kimball (1994), with balanced growth, technology change that affects the consumption-producing sector alone has no impact on employment or capital accumulation at all. Hence, the nature of growth is tightly connected to the sectoral distribution of technical change.

The response of a real business cycle model economy to an exogenous technology shock also depends on the sectors of the economy it affects. Hours and investment responses to a pervasive, sector-neutral, positive technology shock are well understood. They follow from the intertemporal substitution of current leisure and consumption for future consumption. The

⁵ See Kehoe and Ruhl (2006).

household is willing to do this because of the high returns to working and saving. These effects are amplified when technical change affects the investment sector alone, because current consumption is even more expensive relative to future consumption in this case (e.g., Fisher, 2007).

Consumption-sector shocks have smaller effects on intertemporal substitution and so have much less effect on decision rules. As Fisher (1994, 1997) and Kimball (1994) discuss, if preferences are logarithmic, then the decision rules for investment and hours, as well as the allocation of capital and labor across producing consumption and investment goods, are invariant to the stochastic process of consumption-specific technology. More generally, for preferences that* are consistent with balanced growth, hours and investment and factor allocations do not respond to a technology shock, if that shock is permanent, unanticipated, and affects only the consumption sector. So the nature of business cycles is also tied to the sectoral distribution of technical change.

We now present a simple model to illustrate these points. We then discuss other recent macroeconomic work that has focused on the final-goods sector in which technical change occurs. The empirical evidence has been drawn almost exclusively from aggregate data, with heavy reliance on relative prices.

Consider the two-sector closed-economy neoclassical growth model. Suppose one sector produces consumption goods C , the other sector produces investment goods, J . Both sectors produce output by combining capital K and labor L with the same function F but separate Hicks-neutral technology parameters, Z_C and Z_J . In particular, consider the social planner's problem for the following problem, where $v(L_t)$ is a convex and δ is the depreciation rate on capital:

$$\max E_0 \sum_{t=0}^{\infty} \beta^t [\ln(C_t) - v(L_t)]$$

Subject to:

$$\begin{aligned} C_t &= Z_{C_t} \cdot F(K_{C_t}, L_{C_t}); & J_t &= Z_{J_t} \cdot G(K_{J_t}, L_{J_t}) \\ K_t &= K_{C_t} + K_{J_t}; & L_t &= L_{C_t} + L_{J_t}; \\ K_{t+1} &= J_t + (1 - \delta)K_t, & K_0 &\text{ given} \end{aligned}$$

This setup appears in the recent literature in various places. For example, this is a two-sector version of the model in Greenwood, Hercowitz, and Krusell (GHK, 1997); Whelan (2000) discusses the mapping to GHK in greater detail. There are various ways to normalize the technology shocks. For example, GHK define $Z_J = Z_C q$, where $q = Z_J/Z_C$. GHK label Z_C as “neutral” technology and q as “investment specific,” a labeling that has been widely followed since. A shock that raises Z_C but leaves q unchanged is neutral in that both Z_C and Z_J increase equally.

For the purposes of discussing consumption-technology neutrality, a different normalization is more natural. In particular, suppose we define $A = Z_C/Z_J$ as “consumption-specific” technology. Then the problem above can be expressed as a special case of the following problem, where we have made use of the assumption of equal factor shares in the two sectors so there is a single aggregate resource constraint:

$$\max E_0 \sum_{t=0}^{\infty} \beta^t [\ln(C_t) - v(L_t)]$$

Subject to:

$$C_t = A_t X_t; \quad X_t + J_t = Z_t F(K_t, L_t)$$

$$K_{t+1} = J_t + (1 - \delta)K_t; \quad K_0 \text{ given}$$

X is an index of inputs devoted to consumption (in the previous example, it would correspond to the production function $Z_J F(K_C, L_C)$). Since $X=C/A$, and A corresponds (in the decentralized equilibrium) to the relative price of investment to consumption, we can interpret X as consumption in investment-goods units (i.e., nominal consumption deflated by the investment deflator). This is essentially the same approach taken by GHK, except that given their normalization they express everything (investment and output, in particular) in consumption units. Given the logarithmic preferences for consumption this last problem can be expressed as:

$$\max E_0 \sum_{t=0}^{\infty} \beta^t [\ln(A_t) + \ln(X_t) - v(L_t)]$$

Subject to:

$$X_t + J_t = Z_t F(K_t, L_t)$$

$$K_{t+1} = J_t + (1 - \delta)K_t; \quad K_0 \text{ given}$$

Consumption-technology neutrality follows directly from this expression of the problem. In particular, because $\ln(A)$ is an additively separable term, *any* stochastic process for A has no effect on the optimal decision rules for L , X , or J . They do not induce capital-deepening or any of the “expected” RBC effects; e.g., employment doesn’t change. Consumption jumps up, which immediately moves the economy to its new steady state.⁶ Consumption shocks affect only real consumption and the relative price of consumption goods. In contrast, investment-sector technology shocks would have much more interesting dynamics. They induce long-run capital-deepening, and even in the short-run affect labor supply and investment dynamics.

Thus, consumption-sector technology shocks are “neutral.” This consumption-technology neutrality proposition has been implicit in two-sector formulations for a long time; the first explicit references we are aware of are Kimball (1994) and Fisher (1994). Nevertheless, it appears to be a little-known result. One reason is that the seminal work by Greenwood, Hercowitz, and Krusell (1997) used a different normalization, as noted above; they focused on the response of the economy to “investment specific” shocks. A shock to consumption technology alone (leaving investment technology unchanged) would then involve a positive neutral shock, combined with a negative investment-specific shock.⁷

⁶ Kimball (1994) discusses this case further, as well as the extension to King-Plosser-Rebelo (1988) preferences.

⁷ Equivalently, a positive shock to investment-specific technology (with neutral technology unchanged) or a positive shock to neutral technology (leaving investment-specific technology unchanged) would *both* raise investment-technology in the two-sector formulation, and hence would lead to dynamic responses in the frictionless model.

The theoretical motivation for studying sectoral technical change is bolstered by a nascent empirical literature. Most of this literature has followed the GHK normalization of “neutral” and “investment specific” shocks. GHK used data on real equipment prices and argued that investment-specific, not sector-neutral, technical change, is the primary source of economic growth, accounting for as much as 60% of per capita income growth. Cummins and Violante (2000) also find that investment-specific technical change is a major part of growth using more recent data. Several papers also highlight the potential role sector-specific technology shocks in the business cycle. Greenwood, Hercowitz and Huffman (1988) were the first to consider investment-specific shocks in a real business cycle model. Other papers studying investment-specific shocks within the context of fully-specified models are Campbell (1998), Christiano and Fisher (1998), Fisher (1997), and Greenwood, Hercowitz and Krusell (2000). These authors attribute 30-70% of business cycle variation to permanent investment-specific shocks. In the structural VAR literature, Fisher (2006), extending the framework used by Gali (1999) to the case of investment-specific shocks, finds that investment-specific shocks explain 40-60% of the short-run variation in hours and output.⁸

This prior work is based on a “top-down” measurement strategy that relies on investment and consumption deflators (either directly from the BEA, or augmented with equipment deflators from Gordon (1990), and Cummins-Violante (2002)). Our paper contributes to the literature by providing new measures of sector-specific technical change based on industry-level data that, in principle, are more robust. These new measures of technical change can be used to assess the veracity of the existing literature’s findings. The next section describes the theoretical framework underlying our measurement.

⁸ Fisher’s impulse response functions look similar to the subperiod results in Gali, Lopez-Salido, and Valles (2002), with technology shocks reducing hours worked in the pre-1979 period and raising them thereafter.

III. Theoretical Framework

We now describe the theoretical framework we use to identify sector-specific technology shocks. This framework shows how to map from complicated industry-level data on production and trade to key macroeconomic aggregates. Most importantly it embodies the industry/commodity input-output structure of US production, including the activities of exporting and importing, and the fact that the same commodity can be produced domestically and be imported. We assume only three domestically produced final goods and homogeneous factor inputs, but the model is easily extended to accommodate the greater number of final goods and commodity-specific factor inputs that we use in our empirical work. In addition, here we assume perfect competition with constant-returns-to-scale production, but we consider the implications of extensions to this baseline in our empirical work.⁹

The remainder of this section begins by describing the intermediate input economy which maps directly to our data. Next, we discuss the final good economy that corresponds to the intermediate input economy. Third we show how to express the technologies and factor shares of the final good economy in terms of the technologies, factor shares and shares of intermediate goods in gross output of the intermediate input economy.

Intermediate Input Economy

The economy comprises three final goods, consumption, investment and exports, and Φ commodities which are produced domestically and can also be imported. Domestically produced and imported commodities must be combined to produce a *usable commodity* before they can be used as inputs to the production of the final goods or as an intermediate input to the domestic

⁹ Not in this draft, though.

production of commodities. Domestic production of commodities involves capital, labor and intermediate inputs. Final goods are produced as constant elasticity of substitution aggregates of the usable commodities. There is a representative agent who consumes the consumption good, supplies labor, owns the capital stock and engages in borrowing and lending on international capital markets. Capital and labor are homogenous. Net exports, the terms-of-trade and the world interest rate are exogenously given in the model. We relax this assumption in our empirical work.

Competitive equilibrium allocations are obtained as the solution to the following planning problem:

$$\max E_t \sum_{s=t}^{\infty} \beta^{s-t} U(C_s, L_s)$$

Subject to:

$$C_t = \Omega^C(Y_{1t}^C, \dots, Y_{\Phi t}^C); \quad J_t = \Omega^J(Y_{1t}^J, \dots, Y_{\Phi t}^J)$$

$$X_t = \Omega^X(Y_{1t}^X, Y_{2t}^X, \dots, Y_{\Phi t}^X); \quad N_t = X_t - \sum_i (Y_{it}^f / ToT_{it})$$

$$Y_{it}^d = Z_{it} F^i(K_{it}, L_{it}, M_{i1t}, M_{i2t}, \dots, M_{i\Phi t}) \quad , i = 1, \dots, \Phi$$

$$Y_{it} = \Psi(Y_{it}^d, Y_{it}^f) \quad , i = 1, \dots, \Phi; \quad Y_{it} = \sum_{j=1}^{\Phi} M_{ijt} + Y_{it}^C + Y_{it}^J + Y_{it}^X, i = 1, \dots, \Phi$$

$$K_t = \sum K_{it}; \quad L_t = \sum L_{it}$$

$$K_{t+1} = (1 - \delta)K_t + J_t; \quad A_{t+1} = (1 + \bar{r})A_t + X_t; \quad A_0, K_0 \text{ given}$$

Here C, J, X and N denote the final goods consumption, investment, gross exports and net exports; Ω^Λ , $\Lambda = C, J, X, N$ F^i , $i = 1, \dots, \Phi$ and Ψ are constant-returns-to-scale production functions; ToT_{it} denotes the terms-of trade for importing commodity i ; Y_{it}^d and Y_{it}^f denote the domestic production and imports of foreign produced commodity i ; Z_{it} is the exogenous

technology for producing commodity i ; Y_{it} is the production of usable commodity i ; M_{ijt} denotes intermediate input of commodity j used in the production of commodity i ; K_{it} and L_{it} are capital and labor inputs to the production of domestic commodities; K_t and L_t are the aggregate stocks of capital and labor; and \bar{r} is the world interest and A_t is net foreign lending (which evolves exogenously because net exports N is exogenous.)

The one unusual thing about net exports N as a final good is that it can be negative. When net exports are negative, the perspective taken here is that imports can be obtained either by means of the export commodity *or* by means of an electronic credit, A , stated in terms of units of the export commodity. Thus, negative net exports are in units of the export commodity and represent the amount of electronic credits treated as equivalent to units of the export commodity that domestic buyers of imports have been lent. This perspective is valuable because there is only one aggregate export commodity, but many types of imports that the model keeps track of.¹⁰ Because the model conditions on the time path of net exports, it is not essential to keep track of the foreign debt.¹¹ Since keeping track of the foreign debt is not central to the model here, for convenience, one can think of the foreign debt as being stated in terms of units of the export commodity.

Final Good Economy

Our aggregation converts the commodity TFPs Z_{it} into final good TFP. The aggregation is exact for the case where all the production functions are Cobb-Douglas and yields growth rates for TFP that are correct up to second-order approximations of the production functions. To be clear on

¹⁰ Everything would be equivalent if net exports were expressed in terms of any particular imported commodity, but it would be necessary to pick one, since differential changes in the prices of different imported commodities makes them non-equivalent, financially.

¹¹ That is, the main reason one would need to keep track of foreign debt is to determine the time path of net exports. It is part of the relevant intertemporal budget constraint in that determination.

what our aggregation achieves, we now state the social planners problem in terms of final goods only. This problem yields identical allocations for the final goods and factor inputs if the production functions are all Cobb-Douglas. The reduced form planning problem in terms of final goods only is:

$$\max E_t \sum_{s=t}^{\infty} \beta^{s-t} U(C_s, L_s)$$

Subject to:

$$C_t = Z_t^C G^C(K_t^C, L_t^C); \quad J_t = Z_t^J G^J(K_t^J, L_t^J)$$

$$N_t = Z_t^N G^N(K_t^N, L_t^N); \quad K_t = K_t^C + K_t^J + K_t^N$$

$$L_t = L_t^C + L_t^J + L_t^N; \quad K_{t+1} = (1-\delta)K_t + J_t$$

$$A_{t+1} = (1+\bar{r})A_t + N_t; \quad A_0, K_0 \text{ given}$$

As before the terms-of-trade, net exports and the world interest rate are assumed in the model to be exogenous. Note that if we eliminate the international sector then this model would reduce to the two-sector model discussed in Section II except that here the production functions are allowed to differ across sectors.

The third constraint shows the production function for producing net exports. Our treatment of net exports is new and so deserves some discussion. When net exports are negative, the capital and labor inputs are also negative.¹² This arises because consumption and investment goods are produced using imported commodities and when there is a negative trade balance some of these commodities are obtained with the electronic credits associated with international borrowing mentioned above, in addition to the capital and labor used to produce the exported commodities. In

¹² Under constant returns the minus ones multiplying each factor input convert to a minus one in front of the production function. This avoids raising negative numbers to non-integer powers and having to deal with imaginary numbers.

this case, the capital and labor allocated to producing consumption and investment includes the capital and labor that is equivalent in value to the electronic credits. Consequently this capital and labor must be subtracted from the corresponding factor input resource constraints. With balanced trade the quantity of capital and labor used to produce the exported commodities is sufficient to acquire all the imported commodities that end up in the final goods so that the capital and labor inputs to the net export technology are equal to zero. With a positive trade balance the capital and labor employed in the economy exceed the amount necessary to produce consumption and investment and so capital and labor allocated to net exports are positive, reflecting the exports produced to acquire net foreign assets.

Aggregating the Final Good Economy from the Intermediate Input Economy

The final-good sector-specific technologies Z_i^C , Z_i^J and Z_i^N are the focus of our empirical work. These variables can be expressed as weighted averages of domestic commodity TFP and terms-of-trade where the weights depend on the shares of intermediate inputs and trade in gross output. Factor inputs and factor shares in the production of the final goods are similarly defined share-weighted averages of factor shares and commodity shares associated with final goods production. GDP in this economy is measured using the standard chain aggregation with final goods prices derived from the Lagrange multipliers associated with the final good constraints in the planning problem. The appendix describes the aggregation from the model with intermediate goods and final goods to the model with final goods only in detail; here we focus on the intuition.

To build intuition into the nature of the aggregation it is helpful to simplify the underlying intermediate input model to exclude the foreign trade sector. In this closed economy model, consider the production of a final commodity i that is useable for consumption or investment (for

example, a car) Y_i^{final} . That commodity is a function of all the capital and labor that, directly or indirectly, went into producing it, and satisfies

$$Y_i^{final} = \text{constant} \cdot Z_i^{final} K_i^{\sum_j \theta_{ij} a_{jK}} N_i^{1 - \sum_j \theta_{ij} a_{jK}}, \quad (1)$$

where each a_{iK} is capital's share in the gross output production function for commodity i and the θ_{ij} are the elements of the matrix Θ given by

$$\Theta = (I - B)^{-1}.$$

In this last expression I denotes the identity matrix and B is a matrix comprising elements b_{ij} which equal the intermediate input share of commodity j in the gross output production function for commodity i . In equation (1), K_i and L_i includes the capital and labor that was directly used to produce a car, as well as the capital and labor that was used to produce the steel for the car, and the capital and labor used to extract iron ore to produce that steel. The implicit factor shares—the exponents in the equation—account for the factor shares in auto production, steel, and iron mining.

The term Z_i^{final} aggregates the technologies corresponding to gross output of commodity j to yield the technology for producing usable commodity i . Specifically

$$Z_i^{final} = \prod_j Z_j^{\theta_{ij}}$$

The technology for a given final good is then obtained as the weighted average of the usable commodity technologies where the weights are given by the share of each usable commodity in value added of the final good. This follows from the fact that the final goods are constant-returns-to-scale aggregators of the usable commodities. For example, in the Cobb-Douglas case the consumption good technology is given by

$$\ln Z^C = b_C^T \Theta \ln Z. \quad (2)$$

The technology for producing consumption, Z^C , depends on the column vector of underlying commodity technology shocks Z as weighted by the input-output table Θ to get the usable commodity technologies and then weighted again by the column vector of usable commodity shares of consumption, b_C . These usable commodity shares correspond to the shares of usable commodities in the consumption good aggregator Ω^C . The capital and labor shares in G^C , α_C and $1 - \alpha_C$ are similarly obtained as

$$\alpha^C = b_C^T \Theta \alpha_K \quad \text{and} \quad 1 - \alpha^C = b_C^T \Theta \alpha_N,$$

where α_K and α_N are vectors of capital and labor shares in commodity production.

Since it is crucial to our measurement of final good technologies and factor shares, it is helpful to discuss the matrix Θ further. Consider the effects of an innovation to the vector of commodity technologies, Z . First, the vector Z gives the direct effect of commodity-level innovations on the Φ different commodity outputs. Second, the extra output of each commodity increases its use proportionately as an intermediate input into other commodities; this second-round effect is BZ . These second round effects on the quantity of each commodity yield a third-round effect of B^2Z and so on. In the end, as the technological improvements in the production of each commodity repeatedly work their way through the input-output matrix, there is an “input-output” multiplier of

$$I + B + B^2 + B^3 + \dots,$$

which is equivalent to $(I - B)^{-1}$ and therefore also Θ . A similar logic explains the factor shares in the production of the usable commodities. For example, the vector of capital shares for the vertically integrated production of each usable commodity is given by the vector

$$\alpha = \Theta a_K = [I + B + B^2 + B^3 + \dots] a_K,$$

where a_k is the vector of capital shares in the direct production of commodity i . The capital share for the vertically-integrated production of a commodity is the capital share for the direct production of that commodity, plus the contribution from the capital shares of the materials needed to produce that commodity (the term in the expansion involving B), plus the contribution from the materials needed to make those materials (the term in the expansion involving B^2 , and so on.

Adding back the trade sector requires that we adjust the matrix B to address the addition of exports and the fact that final goods comprise both domestically and foreign produced commodities. These changes complicate the derivation of Θ but do not change the logic of the aggregation. The main change to our measurement is to introduce the terms-of-trade into the measurement of the final good technologies and factor shares. The terms-of-trade enter because we have modeled imports as intermediate goods used to produce the final goods and the terms-of-trade are the technology for producing imports from exports.¹³ Details are in the appendix.

Having derived final good technologies it is useful to see how they compare to other measures in the literature. Greenwood, Hercowitz and Krusell (2003) propose measuring the technology for producing investment goods relative to that for producing consumption goods using the price of investment goods relative to consumption goods and this practice is common in the New Keynesian DSGE literature. Similarly, Fisher (2006) proposes identifying the dynamic response to the relative technology using the assumption that only a shock to the relative technology has a long run impact on the relative price. In the appendix we derive a detailed

¹³ Of course, while terms of trade shocks function theoretically just like technology shocks, their econometric properties will be much different. Terms of trade are likely to be much closer to stationary, and are also likely to be more radically endogenous than technology shocks proper.

expression showing that differences in sales taxes and share-weighted factor prices drive a wedge between relative technologies and relative final-goods prices.

To build intuition into the underlying sources of that wedge it is helpful to consider the first order conditions for labor and capital in the consumption and investment sectors under the additional assumptions that factor inputs are heterogeneous and final goods are subject to excise taxes. At any given date, these first order conditions can be written

$$\frac{P_C}{P_J} = \frac{1 + \tau_C}{1 + \tau_J} \frac{Z_J}{Z_C} \frac{R_C^{\alpha_C} W_C^{1-\alpha_C}}{R_J^{\alpha_J} W_J^{1-\alpha_J}} \quad (3)$$

where P_C and P_J are market prices, τ_C and τ_J are excise taxes, R_C and R_J are rental prices of capital, W_C and W_J are wage rates, and α_C and α_J are capital shares in production. In log first differences we have

$$dp_C - dp_J = dz_J - dz_C + [\alpha_C dr_C + (1 - \alpha_C) dw_C] - [\alpha_J dr_J + (1 - \alpha_J) dw_J] + [d\tau_C - d\tau_J] \quad (4)$$

Clearly, at any point in time, unless factor shares and growth rates of taxes and factor prices are all identical then relative prices and relative technologies will not be equal. If we add imperfect competition to the model then there is an additional wedge due to differences in the growth rates of mark-ups. This makes it clear that the convention of equating relative prices with relative technologies at each point in time can be problematic. Our measurement is not subject to this drawback. Note however, that as long as there are no permanent shocks to the wedges, that

In the context of our model, one can show that if we aggregate according to equation () and use standard commodity-level TFP residuals (without controls for utilization or non-constant returns to scale), then relative final-goods prices and relative final-use TFP are related by the following identity:

$$\begin{aligned}
d\text{tfp}_j - d\text{tfp}_c &= (dp_c - dp_j) + (b_j - b_c)(I - B)^{-1} [s_k dr + s_L dw] + (b_j - b_c)(I - B)^{-1} [dp^{\text{Dom}} - dp^{\text{Dom,Producer}}] \\
&= \text{Relative price} + \text{Factor-Price Wedge} + \text{Tax Wedge}
\end{aligned}$$

In the equation, $(b_j - b_c)$ is a row vector for a given year of commodity-specific shares in final investment versus consumption; $[a_k dr + a_L dw]$ is a column vector of commodity-specific share-weighted input price growth (where only domestic commodities have non-zero entries); similarly, $[dp^{\text{Dom}} - dp^{\text{Dom,Producer}}]$ is a column vector of the difference between price paid by a purchaser and the price received by a producer.¹⁴

IV. Data

We use the KLEM dataset produced by Dale Jorgenson and his collaborators, along with the underlying input-output data.¹⁵ The main vintage of data runs from 1960-2005, but we have merged them with an earlier vintage of Jorgenson data back to 1950. The dataset provides consistent industry output and inputs as well as commodity final-use data for 35 industries/commodities. (The dataset has the full input-output use matrix only back to 1960; the so-called “make” table is available only back to 1977.)¹⁶

The Jorgenson data offer several key advantages. First, they provide a unified dataset that allows for productivity analysis by providing annual data on gross output and inputs of capital, labor, and intermediates. Second, the productivity data are integrated with annual input-output data that are measured on a consistent basis. The I-O data include not just intermediate-input flows

¹⁴ For this equation to hold as an identity in the data, we need to allow final-use shares, the input-output matrix B , and factor shares to vary period-by-period. Otherwise, this is only an approximation. In the results, we show this decomposition using period-by-period shares.

¹⁵ The dataset updates Jorgenson, Gollop, and Fraumeni (1987). We thank Dale Jorgenson as well as Barbara Fraumeni, Mun Ho, and Kevin Stiroh for helpful conversations about the data. Jon Samuels helped tremendously with data availability. The main data are available at <http://dvn.iq.harvard.edu/dvn/dv/jorgenson> (accessed April 26, 2010).

¹⁶ The original data source for input-output and industry gross output data is generally the Bureau of Labor Statistics which has, for a long time, produced annual versions of both nominal and real (chained) I-O tables. (The current vintage of data from the BLS Employment Projections Program is available at <http://www.bls.gov/emp/#data> (accessed April 26, 2010).

of the different commodities, but also commodity-specific final expenditures on consumer non-durables and services, consumer durables, investment, government, and exports and imports. Although our model is in terms of just consumption and investment, the logic extends naturally to this larger set of final expenditure categories. Third, the input data incorporate adjustments for composition; e.g., services from different types of labor are taken to be proportional to wages, as standard first-order-conditions imply, not simply to total hours worked.

Our data do not comprise all of GDP. First, they exclude the service flow from owner-occupied housing. Second, they exclude services produced directly as part of government administration (they do include government enterprises; and they include direct purchases by the government of goods and services produced by non-government businesses). Thus, our output measure corresponds to the output of businesses plus nonprofit institutions serving households.

We have separate prices for domestic commodity production and for commodity imports. For each commodity, we calculate a composite supply price as a Tornquist index, which we apply to all uses of the commodity. In essence, we assume that all users of a commodity consume the same composite of domestic and imported supplies of the commodity.¹⁷

Much of the macroeconomics literature on investment-specific technical change has used alternative quality-adjusted deflators for investment goods, following Gordon (1983) and its extension by Cummins and Violante (2002). We have incorporated these deflators in our work. First, since the Cummins-Violante measures end in 2000, we updated their empirical estimates, allowing us to obtain adjusted deflators for detailed equipment investment categories through 2005. Second, we re-estimate the capital-input measures for each industry in the Jorgenson data. To do

¹⁷ Relatively recently, the BEA has begun producing an “import matrix” that shows, for each commodity, the estimated industry uses of imports. But they create this table using the “assumption that each industry uses imports of a commodity in the same proportion as imports-to-domestic supply of the same commodity,” (http://www.bea.gov/industry/xls/1997import_matrix.xls, ‘readme’ tab) rather than using separate data on where imports flow. So the BEA makes the same assumption we do.

so, we use the alternative Gordon-Cummins-Violante (GCV) deflators to obtain modified measures of real equipment investment from the BEA's detailed investment-by-industry dataset. For each type of capital asset, we then estimate perpetual inventory measures of the detailed asset-by-industry capital stocks; we then do Tornquist aggregation to industry capital-input measures using estimated user-cost weights. Third, since these investment goods are produced by an industry, we also need to modify commodity and industry output measures to account for their production. For example, if the extended Gordon deflators suggest an adjustment to, say, special industrial machinery, we modify the deflators for the Jorgenson commodity/industry "non-electrical equipment" to take account of this modification. The necessary modifications are fairly complex, and involve digging into the detailed BLS input-output tables that underlie the Jorgenson data (e.g., to obtain appropriate weights for, say, special industrial machinery in total non-electrical equipment); details are available from the authors. In general, these modifications affect the means of the data (for example, equipment technology growth is faster, but consumption technology growth (equipment users) is slower), but regression coefficients discussed below are not much affected.

From our disaggregated data, we obtain measures of industry/commodity technology to feed through the input-output tables. One easy measure would be the sectoral Solow residuals, i.e., $dy_i - dx_i$, where dx_i is share-weighted inputs of capital, labor, and intermediates. However, a sizeable literature emphasizes that, in the short run particularly, measures of the Solow residual even at a disaggregated level contain non-technological components such as unobserved variations in factor utilization and the effects of non-constant returns to scale.

Hence, we update the estimated industry technology residuals from Basu, Fernald, and Kimball (BFK, 2006). For this version of the paper, we focus on controlling for factor utilization and maintain the assumption of constant returns. In particular, we estimate the following

regressions at the level of each industry, using oil price increases, government defense spending, and a measure of monetary innovations (from a VAR) as instruments:

$$dy_i - dx_i = c_i + \beta dh_i + \varepsilon_i,$$

where dy is output growth, dx is share-weighted input growth, and dh is hours-per-worker growth (as a proxy for utilization). These estimates control for variations in factor utilization, which can drive a wedge between technology and measured TFP. For each industry, technology is measured as $c_i + \varepsilon_i$. See BFK (2006) or Appendix B of this paper for further details.¹⁸

BFK (2006) implement their adjustments for non-mining private industries. For other industries, including mining, government enterprises, and the international trade industry, we simply use standard TFP. This implies that the technology measures may not be fully exogenous. Where this matters for regression analysis, however, there is a natural instrument: The measures of final-use technology, where we zero out the endogenous elements. Of greatest importance there, we zero out the terms of trade (our measure of technology for the artificial international trade industry), which is clearly endogenous.

V. Results

We now consider various issues of relevance for macroeconomists. Except where indicated, our measures of relative prices, TFP growth (with industry Solow residuals), and technology change (measured using utilization-adjusted TFP) are based on the Jorgenson data adjusted with Gordon-Cummins-Violante deflators. We will often refer to utilization-adjusted TFP simply as “technology,” since under our maintained hypothesis of constant returns to scale, utilization adjustments are sufficient to control for mismeasurement in the Solow residual and identify technology change.

¹⁸ These residuals correspond to domestic industry production. We convert them to commodity-supply residuals with the input-output make tables and by rescaling by the domestic supply share.

We first examine the long-run properties of the data from 1950 to 2005, which we also break into sub-periods of 1950-1982 and 1983-2005. Table 1 shows the sample-average growth rates of final-use TFP (the top panel) and technology (the bottom panel). Since utilization should be stationary in the long run, both TFP and utilization-adjusted TFP should have much the same means. This prediction holds in the data, with differences between the top and bottom panels showing up in the second or third decimal place. The data show the pattern we have come to expect, with the technology for producing equipment and durable consumption advancing at a rapid rate, on the order of 3 percent per year over the last two decades. The technology for producing non-durables and services (NDS) consumption improves very slowly over the whole period, and is essentially flat over the post-1983 period. The technology for producing structures shows a definite and puzzling decline over the whole period, which is particularly pronounced over the later period. We ascribe this decline to a lack of adequate adjustment for the increasing quality of structures in the official data. Similar mismeasurement probably also explains the lack of growth in NDS technology over the post-1983 period.¹⁹

Our mapping of industry production residuals to final-use categories also provides novel estimates of technology for two categories of output, government purchases and exports. Both have increased strongly over our sample period, particularly export technology, which grows at 2.1 percent from 1950 to 2005.

Figure 1 shows the ordering of growth rates visually by plotting the estimated levels of final-use technology (with utilization-adjusted residuals) that correspond to the bottom panel of Table 1. Note that our “government technology” is *not* the technology in the production of services by government workers, which are not included in our measure of output. Government technology

¹⁹ For example, the official statistics do not attempt, even in principle, to capture the gains from increased product variety.

growth captures improvement in the production technology for the goods that the government buys from private industries. The government purchases a mix of equipment, non-durables and services, and structures. Similarly, the correct interpretation of our “trade technology” is that it is the technology for producing exports. Since exports are primarily goods (and about half durable goods), trade technology grows at a rate comparable to equipment technology and consumer durables.

GHK and the literature building on their work use relative prices as a proxy for relative technology. Of course, under some assumptions this procedure is rigorously correct, but as discussed above these assumptions are restrictive. Now that we have an independent measure of sectoral technical change, derived under more general assumptions than GHK, we can ask an important question: To what extent do relative final-use output prices provide a good approximation to relative final-use TFP or technology? Table 2 shows the comparison with relative TFP (without utilization adjustments, so the identity derived in equation (4) holds exactly). As expected, the price of non-durables and services consumption has risen on average relative to the price of the equipment investment. This trend is more pronounced in the post-1982 period. Even so, relative prices do not change as fast as relative TFP. Table 2 shows that on average relative TFP has risen about 0.5 percent per year faster than the relative price. As discussed in the theory section, this can happen if there are differences in the growth rates of share-weighted primary input prices or long-run changes in the relative tax wedge. In fact, about two-thirds of the gap between relative prices and relative TFP is due to long-run changes in relative primary input prices, with the remaining gap due to changes in relative sales taxes. (The important factor for the tax wedge turns out to be rising sales taxes on consumer goods, which worked to raise consumer prices; the effect of tax changes on investment prices alone is relatively small.)

How can the growth rate of primary input prices diverge across sectors over a period of decades? One would think that prices for identical factors should equalize across industries over such long periods, and the composition corrections in the Jorgenson data should ensure that we are indeed comparing constant-quality inputs across sectors. (Thus, for example, if computer manufacturing requires a higher fraction of college-educated workers than auto manufacturing, we would not misinterpret this fact as implying that the price of homogeneous “labor” is higher in computers than in autos.) Recall, however, that the primary input prices are weighted by their cost shares. We know that over time the price of labor rises faster than capital rentals, so industries with high labor shares will, *ceteris paribus*, experience faster price growth. In Table 3 we report capital shares by final-use sectors, and show that this pattern does indeed hold in the data. That is, equipment investment does indeed have a lower capital share (and thus a higher labor share) than does NDS consumption. In addition to verifying the reason for the long-run gap between relative prices and relative TFP, the result that capital shares differ noticeably across sectors invalidates both short- and long-run identification based on uncovering relative technology change from the behavior of relative prices. Thus, the conclusions of a large literature are immediately called into question.

Tables 4A and 4B show selected correlations for standard TFP and BFK technology change, respectively.. The sectoral shocks themselves are very highly correlated, as shown in the first rows of each panel. This is not surprising: In addition to any truly “common” shocks, the final-sector shocks are different weighted averages of the same vector of underlying commodity shocks. However, the correlations are lower using utilization-adjusted TFP, as one expects: This control eliminates the common mismeasurement from business-cycle variations in factor utilization.

Our procedure allows us to test the orthogonality assumptions typically used in the preceding empirical literature. In particular, GHK and Fisher (2006) explicitly assume orthogonality of what they label “neutral” and (equipment) “investment-specific” productivity innovations. In our two-sector framework, these labels correspond to the consumption-sector productivity shock and the *relative* equipment-investment-consumption shock. Line 3 of the two panels shows the correlation between the “neutral” and “equipment specific” shocks.

A priori, there seems little reason to expect orthogonality, given that these shocks are different mixtures of the same underlying productivity residuals. The orthogonality assumption appears shaky in the data, for both the standard and utilization-adjusted TFP series, where the correlations are about (plus or minus) 0.10. Note that Fisher (2006) does not assume that observed, standard TFP measures technology, so the utilization-adjusted result is probably more relevant. Note also that orthogonality with consumption TFP is a much better approximation than is orthogonality with investment TFP, as shown in line 4.²⁰ However, whether or not the orthogonality assumptions are verified, GHK’s and Fisher’s identification schemes cannot uncover relative technology from relative prices, since factor shares differ across sectors.

Finally, what are the dynamic effects of exogenous shocks to different types of technology? Table 5 shows dynamic responses of business-cycle variables to changes in sectoral technology. The first set of variables we study consists of GDP and its components, with all the dependent variables entered in growth rates.²¹ The table also shows the same regressions for other interesting macroeconomic variables: hours, wages, a variety of prices, and interest rates, with variables other

²⁰ The correlations for the entire period are not the averages of the correlations for the subperiods, because the standard deviations over the entire period exceed the average of the standard deviations for the subperiods. An interesting question that we have not yet pursued is what orthogonality of relative TFP and consumption TFP implies about the input-output structure of the economy.

²¹ The GDP and investment data are taken directly from published BEA tables. They have not been adjusted using the Gordon investment deflators.

than interest rates entered in log differences. The R^2 shown are for bandpass filtered data, i.e., where the dependent variable and the fitted values are first bandpass filtered to focus on variation between 2 and 8 years.

The challenge is the desire to be comprehensive in including all final-use shocks with the cost of using greater disaggregation (e.g., including structures separately), when multicollinearity becomes a serious problem. We aggregate final-use sectors in the way that corresponds to typical assumptions in two-sector models: One sector produces all investment goods, including equipment for both the private sector and the government, consumer durables, structures and exports; the other produces non-durables and services consumption for the household and government sectors. We group exports with investment goods, for two reasons. First, exports are mostly durable. Second, exports are a way of adding to a country's durable wealth, just like equipment investment, for exports can be used to purchase capital held in the rest of the world. For all variables, we include the current value and two lags.

As the data section noted, we do face a problem in interpreting the terms of trade as exogenous. First, the terms of trade should change with changes in domestic consumption and investment technology, if these goods are exported. Second and more importantly, the terms of trade may change for reasons that are clearly not exogenous. For example, monetary policy can change the terms of trade if prices are sticky. Monetary policy may also react to changes in output or hours worked due to non-technology shocks, thus exacerbating the endogeneity problem.

We address this problem by constructing a second set of measures of technology that we *do* think are exogenous, using only industry technology shocks based on the utilization-adjusted TFP residuals. Operationally, we set to zero changes in the terms of trade as well as changes in the TFP of the natural resource industries and government enterprises. We then use this group of rigorously exogenous measures as instruments for the various types of technology change in Table 5.

Focusing first on the effects of investment-specific shocks, we find results similar to those documented by Gali (1999) and BFK (2006), which are puzzling from the standpoint of a simple one- or two-sector RBC model with flexible prices. Aggregate GDP, non-durables consumption, purchases of consumer durables, residential investment and imports all fall significantly after a positive shock to equipment technology. Most strikingly, equipment and software investment falls for two years after a positive shock to the technology for producing investment goods, and the decline is highly significant. This last finding is especially striking if one considers the likelihood of measurement error in our measures of investment output. Since we are regressing equipment growth on the technology for producing equipment, measurement error in output would create a *positive* bias in the coefficient. Adjusting for measurement error, the true coefficients are likely to be even more negative than our estimates. (There is little autocorrelation in the investment technology growth series, so this result doesn't seem to reflect predictability that equipment technology will be even better in the future. In some models, news of a future technology improvement will cause demand to contract until the improvement is actually realized.)

Table 5 shows that hours worked fall significantly when investment technology improves, and do not recover their pre-shock value even after three years. Note that this is a decline in *total* hours worked in response to a *positive* technology shock to a sector that (even with the inclusion of consumer durables) produces less than 20 percent of GDP! As shown in Table 6, the updated BFK results show a distinct and significant negative contemporaneous correlation between hours growth and aggregated BFK technology. (The negative effect is more pronounced with the original BFK industries, i.e., excluding agriculture, mining, and government enterprises).

As expected, an equipment-technology improvement does lower the relative price of equipment. Somewhat surprisingly, consumption prices also fall, as does the GDP deflator.

Next we examine the effects of consumption-technology shocks. Here again we find results at odds with simple RBC models which, as we noted, generally predict consumption technology neutrality. We see that consumption-technology shocks increase GDP and non-durables consumption significantly, but also increase equipment investment, consumer durables consumption and (on impact but not cumulatively) residential housing, with at least one statistically significant impact in each category. Hours worked are basically unchanged over the three years. As expected, the relative price of consumption to equipment falls, but the decline is not statistically significant.

The final column of Table 5 shows the R^2 of the regression, after bandpass-filtering both the dependent variable and the fitted value in order to focus on business-cycle frequencies. For quantity variables, especially, the explanatory power is quite high. Figure 3 shows the fitted values of hours and equipment to technology shocks. The explanatory power of the regression is clear. Comparing the fitted values in the top and bottom panels, it is also clear that, conditional on technology, hours and equipment investment move comove positively. But it is also clear that, according to the regressions, technology innovations did not cause the deep recessions of the 1970s and 1980s.

Our simple model of Section II predicts that business investment should not change in response to a consumption technology improvement. The model did not include residential investment and consumer durables. We conjecture that if the model is extended to allow for durable goods to provide services that directly enter the utility function, and incorporates non-separability between non-durables consumption and durables services, it would predict roughly the results that we find, at least qualitatively,

Taken together, these results suggest that the earlier findings of Gali (1999) and BFK (2006) about the contractionary effects of technology improvements are driven primarily by the

effects of investment technology shocks. They also raise an interesting question: Does any fairly standard model predict that consumption technology improvements should be expansionary, but investment technology improvements should not? We believe that a simple two-sector model with sticky prices in both sectors and a standard Taylor rule for monetary policy can go some way towards matching the results we find. Preliminary results obtained by Basu, Fernald and Liu (2012) suggest that we are right.

A reason the sticky price model may be appropriate is the extremely slow pass-through of technology to relative prices. The pass-through was almost imperceptible in the regression results. However, other models of slow pass-through (for example, in flexible-price models with time-varying markups), might also explain the results we find.

Considering the results altogether, we find strong evidence that we cannot summarize technology shocks using a single aggregate index. Aggregation would hold either if (1) all final-use technology shocks were perfectly correlated or (2) if the different shocks had identical economic effects. Unfortunately, neither condition is satisfied, since Table 4B shows that the relevant correlation is about 0.50, and we can statistically reject the hypothesis that all the coefficients of investment- and consumption-technology shock in Tables 5 are identical at all lags. Hence, it appears important to use multi-sector models to investigate economic fluctuations due to technology shocks.

Table 5 has some interesting findings that bear on the sticky-price interpretation of our results, and suggest some new thinking about optimal monetary policy. Note that the Fed Funds rate falls significantly and sharply in the first two years following a consumption technology improvement, but barely changes following an investment technology improvement. Correspondingly, the GDP deflator falls for all three years following an investment shock, but only on impact and insignificantly following a consumption shock. If the Fed targets consumption

prices, but not investment prices, then that would go a long way towards explaining why the effects of an investment shock are so contractionary but the effects of the consumption shock are expansionary. These results also suggest that optimal monetary policy should target investment as well as consumption goods prices.

VI. Conclusions

Theory suggests that the final-use sector in which technology shocks occur matters for the dynamic response of those shocks. We make this point with the example of consumption-technology neutrality in the real-business-cycle model. Consistent with the predictions of the model, our empirical results show that consumption- and investment-specific shocks do indeed have significantly different economic effects, and the differences persist for at least several years.

The details of the results present several puzzles. The results for investment-sector shocks show that hours decline when investment technology improves, GDP declines sharply and, most surprisingly, investment itself declines strongly. These results are consistent with the findings of Gali (1999) and BFK (2006), who interpret them as evidence in favor of nominal price rigidities. Our findings for consumption shocks are quite different, however: they show significant increases in GDP and in the consumption of both non-durables and durable goods, including housing. Whether these results taken together can be explained by richer or differently-parameterized flexible-price or sticky-price models with multiple sectors is at this point an open question, although we find two-sector sticky price models promising in this regard.

A novel contribution of this paper is a new method to measure sector-specific technologies that does not rely simply on relative price changes, as frequently used in the macro literature. In special cases, the relative-price measures are appropriate; but in general, they are not. We show that the conditions necessary for relative prices to measure relative technologies do not hold in the

data. Since our method allows us to test hypotheses that are identifying assumptions in other works, it is a robust, albeit data-intensive, way of checking whether these assumptions hold in any situation where their validity is in doubt.

Appendix A: The Optimal Mode of Production

Closed Economy: Stretching terminology a bit, define “final” or “net” output of a *commodity* as the quantity produced that is used in the *direct* production of consumption, investment, government purchases—that is, the amount *not* used in the production of intermediate goods. Gross output of a commodity is the *total* amount of the commodity produced for all purposes.

Let R be the rental rate for capital K , W the wage for labor N , and Z a vector of multiplicative technology shifters. Consider the cost minimizing plan for making one unit of final or net output of commodity i available for use in producing C, J, G :

$$\begin{aligned}
 P_i(R, W, Z) &= \min_{\hat{K}_{ij}, \hat{N}_{ij}, \hat{X}_{ij}, \hat{M}_{ij\ell}} R \sum_j \hat{K}_{ij} + W \sum_j \hat{N}_{ij} \\
 \text{s.t. } \hat{Y}_{ii} &= 1 + \sum_{\chi} \hat{M}_{i\chi i} \\
 \hat{Y}_{ij} &= \sum_{\chi} \hat{M}_{i\chi j} \quad (\forall j \neq i) \\
 \ln(\hat{Y}_{ij}) &= \ln(Z_j) + a_{jK} \ln(\hat{K}_{ij}) + a_{jN} \ln(\hat{N}_{ij}) + \sum_{\ell} b_{j\ell} \ln(\hat{M}_{ij\ell}) \quad (\forall j),
 \end{aligned}
 \tag{1.5}$$

The i subscript refers throughout to the fact that the program is intended to produce *net* output of commodity i . The hats on key variables are a reminder that the program is for *one* net unit of commodity i . $P_i(R, W, Z)$ is the minimum average (and marginal) cost of producing commodity i . \hat{K}_{ij} is the amount of capital used directly to produce commodity j in order to produce one net unit of commodity i . \hat{N}_{ij} is the amount of labor used directly to produce commodity j in order to produce one net unit of commodity i . \hat{Y}_{ij} is the *total gross amount* of commodity j used to produce one net unit of commodity i . $\hat{M}_{ij\ell}$ is the amount of commodity ℓ used directly to produce commodity j in order to produce one net unit of commodity i . Z_j is the multiplicative technology shifter for producing commodity j . a_{jK} is capital’s share in the direct inputs to commodity j . a_{jN} is labor’s share in the direct inputs to commodity j . Finally, $b_{j\ell}$ is the share of commodity ℓ in the direct inputs to commodity j . Our objective in this appendix is to find the boiled down production function for net output of commodity i as a function of the technology vector and the total capital and labor inputs used directly or indirectly to produce net output of commodity i . We will use Shephard’s Lemma (also called the derivative principle) to find capital’s share α_i and labor’s share $1 - \alpha_i$, and the corresponding method to find the elasticity ξ_{ij} of net output with respect to each component Z_j of the technology vector Z . To be specific, by Shephard’s Lemma, capital’s share in the boiled-down net output production function is given by

$$\alpha_i = \frac{R}{P_i(R, W, Z)} \frac{\partial P_i(R, W, Z)}{\partial R},
 \tag{1.6}$$

labor's share in the net output production function is given by

$$(1.7) \quad 1 - \alpha_i = \frac{W}{P_i(R, W, Z)} \frac{\partial P_i(R, W, Z)}{\partial W},$$

and the elasticity of the net output production function with respect to Z_j is equal in magnitude but opposite in sign to the elasticity of the unit cost function with respect to Z_j

$$(1.8) \quad \xi_{ij} = - \frac{Z_j}{P_i(R, W, Z)} \frac{\partial P_i(R, W, Z)}{\partial Z_j}.$$

These elasticities of the unit cost function can be found by using the Envelope Theorem. First, write the Lagrangian for (1.5) as

$$(1.9) \quad \begin{aligned} L_i = & R \sum_j \hat{K}_{ij} + W \sum_j \hat{N}_{ij} \\ & + \sum_j \lambda_{ij} \left[\ln(\hat{Y}_{ij}) - \ln(Z_j) - a_{jK} \ln(\hat{K}_{ij}) - a_{jN} \ln(\hat{N}_{ij}) - \sum_{\ell} b_{j\ell} \ln(\hat{M}_{ij\ell}) \right] \\ & + \zeta_{ii} + \sum_j \zeta_{ij} \left[\left(\sum_{\chi} \hat{M}_{i\chi j} \right) - \hat{Y}_{ij} \right]. \end{aligned}$$

The first order conditions with respect to \hat{K}_{ij} and \hat{N}_{ij} yield this pair of equations:

$$(1.10) \quad R \hat{K}_{ij}(R, W, Z) = a_{jK} \lambda_{ij}(R, W, Z)$$

$$(1.11) \quad W \hat{N}_{ij}(R, W, Z) = a_{jN} \lambda_{ij}(R, W, Z).$$

where we have made explicit the dependence of quantities and Lagrange multipliers coming out of the optimal program on R , W , and Z . By the definition of a_{jK} as the share of capital in direct production,

$$(1.12) \quad R \hat{K}_{ij}(R, W, Z) = a_{jK} P_j(R, W, Z) \hat{Y}_{ij}(R, W, Z).$$

Together, Equations (1.10) and (1.12) imply that

$$(1.13) \quad \lambda_{ij}(R, W, Z) = P_j(R, W, Z) \hat{Y}_{ij}(R, W, Z).$$

To get more insight into λ_{ij} , define θ_{ij} as the total gross nominal output of commodity j used to produce $\frac{1}{P_i(R,W)}$ net units of commodity i . That is,

$$(1.14) \quad \theta_{ij}(R,W,Z) = \frac{P_j(R,W,Z)}{P_i(R,W,Z)} \hat{Y}_{ij}(R,W,Z),$$

Using Equation (1.14), Equation (1.13) can be written as

$$(1.15) \quad \lambda_{ij}(R,W,Z) = \theta_{ij}(R,W,Z)P_i(R,W,Z).$$

To find θ_{ij} , consider the relationship between the net output production functions and the direct production functions. Because of the properties of Cobb-Douglas production functions, the value of the commodity j used *directly* to produce 1 net unit of commodity i is b_{ij} . In turn, obtaining b_{ij} net units of commodity j for direct use in the production of commodity i requires, in turn, the overall production of a value $b_{ij}\theta_{j\ell}$ of commodity ℓ . Overall, the total value of commodity ℓ produced must be $\theta_{i\ell} = \sum_j b_{ij}\theta_{j\ell}$, except that when $\ell = i$, $\theta_{ii} = 1 + \sum_j b_{ij}\theta_{ji}$, since one must account for the original unit value of commodity i that is to be produced. In matrix form,

$$(1.16) \quad \Theta = I + B\Theta,$$

where $\Theta = [\theta_{ij}]$, $B = [b_{ij}]$, and I is the identity matrix. Solving for Θ ,

$$(1.17) \quad \Theta = (I - B)^{-1}.$$

Note that for the Cobb-Douglas case, the matrix Θ is a constant.

Now return to the problem of finding the parameters of the boiled-down, net output production function. By the Envelope Theorem,

$$(1.18) \quad \alpha_i = \frac{R}{P_i} \frac{\partial P_i}{\partial R} = \frac{R}{P_i} \frac{\partial L_i}{\partial R} = \frac{R \sum_j \hat{K}_{ij}}{P_i} = \frac{\sum_j \lambda_{ij} a_{jk}}{P_i} = \frac{\sum_j \theta_{ij} P_i a_{jk}}{P_i} = \sum_j \theta_{ij} a_{jk}.$$

Note that P_i , \hat{K}_{ij} , λ_{ij} , and a_{jk} are all functions of R , W , and Z , since the optimal program for producing one unit of net output of commodity i depends on R , W , and Z . (L_i is a function of almost everything in sight.) Similarly,

$$(1.19) \quad 1 - \alpha_i = \sum_j \theta_{ij} a_{jN}.$$

In vector form,

$$(1.20) \quad \alpha = \Theta a_K$$

and

$$(1.21) \quad 1 - \alpha = \Theta a_N.$$

where 1 in Equation (1.21) represents a vector of 1's. These are all constants in the Cobb-Douglas case. Finally, the elasticity of output with respect to component j of the technology vector is

$$(1.22) \quad \xi_{ij} = - \frac{Z_j}{P_i} \frac{\partial P_i}{\partial Z_j} = - \frac{Z_j}{P_i} \frac{\partial L_i}{\partial Z_j} = \frac{\lambda_{ij}}{P_i} = \theta_{ij}.$$

Since all the elasticities of the boiled-down net output production function are constant, it must be given by

$$(1.23) \quad Y_i^{final} = \text{constant} \cdot K_i^{\alpha_i} N_i^{1-\alpha_i} \prod_j Z_j^{\xi_{ij}} = \text{constant} \cdot K_i^{\sum_j \theta_{ij} a_{jN}} N_i^{1-\sum_j \theta_{ij} a_{jN}} \prod_j Z_j^{\theta_{ij}},$$

where Y_i^{final} represents *final* or net output of commodity i .

The last step in obtaining the technologies for producing final goods, for example consumption, involves specifying how those goods are aggregated from the net outputs. Continuing the example, suppose consumption is produced according to

$$C = \Omega^C(Y_1^C, \dots, Y_\Phi^C),$$

where Ω^C is a constant-returns-to-scale Cobb-Douglas function. Then the consumption technology is

$$C = Z^C G^C(K^C, L^C)$$

with

$$\begin{aligned} \ln Z^C &= b_C^T \ln Z^{final} \\ \ln K^C &= \end{aligned}$$

Departures from Cobb-Douglas: As long as the elasticities of factor substitution are strictly between zero and infinity, one can handle departures from Cobb-Douglas in a straightforward way. Define the logarithmic price function p_i for final net output as follows:

$$p_i(r, w, z) = \ln(P_i(e^r, e^w, e^z)).$$

That is, $p_i = \ln(P_i)$, $r = \ln(R)$, $w = \ln(W)$ and $z_j = \ln(Z_j)$ for each element of the vectors z and Z so that the function p_i is a log-log version of P_i . Similarly, define the logarithmic price function q_i for gross output with the arguments

$$q_i(r, w, z_i, p_1, p_2, \dots)$$

Then since the price of a commodity is the same whether used as final output or as an intermediate good, the functions p_i and q_i obey the functional equation

$$p_i(r, w, z) = q_i(r, w, z_i, p_1(r, w, z), p_2(r, w, z), \dots)$$

Differentiating with respect to r , w , and z_j ,

$$\frac{\partial p_i}{\partial r} = \frac{\partial q_i}{\partial r} + \sum_j \frac{\partial q_i}{\partial p_j} \frac{\partial p_j}{\partial r}$$

$$\frac{\partial p_i}{\partial w} = \frac{\partial q_i}{\partial w} + \sum_j \frac{\partial q_i}{\partial p_j} \frac{\partial p_j}{\partial w}$$

$$\frac{\partial p_i}{\partial z_\ell} = \frac{\partial q_i}{\partial z_\ell} + \sum_j \frac{\partial q_i}{\partial p_j} \frac{\partial p_j}{\partial z_\ell} = -\delta_{i\ell} + \sum_j \frac{\partial q_i}{\partial p_j} \frac{\partial p_j}{\partial z_\ell}.$$

Where $\delta_{i\ell}$ is the Dirac delta that is equal to 1 when its two subscripts are equal and zero otherwise. The last equality holds because technology is multiplicative and defined so that the a one percent improvement in a commodity's own technology increases its gross output by one percent when all the other inputs are fixed, implying that the direct effect of this technology improvement on cost is a one percent reduction.

By the principles of duality with unit cost functions, the factor shares for final net output can be written $\alpha_i = \frac{\partial p_i}{\partial r}$, $1 - \alpha_i = \frac{\partial p_i}{\partial w}$, while the factor shares for gross output can be written

$$a_{iK} = \frac{\partial q_i}{\partial r}, a_{iN} = \frac{\partial q_i}{\partial w} \text{ and } b_{ij} = \frac{\partial q_i}{\partial p_j}. \text{ Also, define } \psi_{i\ell} = \frac{\partial p_i}{\partial z_\ell}. \text{ Form column vectors } \alpha = [\alpha_i],$$

$a_K = [a_{iK}]$, $a_N = [a_{iN}]$, $\mathbf{1} = [1]$ and the matrices $B = [b_{ij}]$ and $\Psi = [\psi_{i\ell}]$. Then, the equations above can be cast in matrix form as

$$\alpha = a_K + B\alpha$$

$$\mathbf{1} - \alpha = a_N + B(\mathbf{1} - \alpha)$$

$$\Psi = -I + B\Psi.$$

Solving for α , $1 - \alpha$ and Ψ ,

$$\alpha = (I - B)^{-1} a_K$$

$$\mathbf{1} - \alpha = (I - B)^{-1} a_N$$

$$\text{(implying } \mathbf{1} = (I - B)^{-1}(a_K + a_N) \text{ or } a_K + a_N = (I - B)\mathbf{1})$$

and

$$\Psi = -(I - B)^{-1}.$$

Open Economy: The strategy for extending the results to the open economy is as follows. First, model exports as an additional commodity, produced by a CRTS CES production function that aggregates the other commodities shown as exported in the input-output table. Second, model each of the Φ types of imported commodities as distinct commodities (superscripted *foreign*) from both the home-produced commodities (superscripted *domestic*) and the aggregate of the home-produced and imported commodity of each type that is used to produce other commodities and for final use (no superscript). Thus, the open economy version of our model has $3\Phi + 1$ commodities to the closed-economy model's Φ commodities. The aggregate of home-produced and imported commodities that is used to produce other commodities and for final use is sometimes called "supply," but for clarity, we will call it the "usable" commodity. All commodities, *including exports* are produced from the usable commodity. (This is an assumption that could be questioned. We also considered a model in which exports are produced from domestic commodities alone, but decided to emphasize the simpler and more readily interpretable model in which all commodities are produced from usable commodities.)

Third, model imported commodities as produced by a CRTS production function with one input: the aggregate export commodity. For each imported commodity, the ratio of the international price of the export commodity to the international price of each imported commodity acts as the technology parameter for that imported commodity.

Third, since the emphasis in this appendix is on the production side of the economy, it focuses on equilibrium conditional on the time paths of the levels of aggregate final output: that is, consumption C , investment J , government purchases G and net exports X . (Note the nonstandard use of X for *net exports*.) Alternatively, one can easily consider an equilibrium conditioning on the time paths of government purchases G and net exports X , in which C and J arise from optimizing behavior. Modeling the determination of the time path of net exports X is beyond the scope of this paper, just as modeling the determination of the time path of government purchases G is beyond

the scope here. Fourth, the model also conditions on the path of international prices. (That is, it is solved conditional on the terms of trade for each imported commodity.)

The one unusual thing about net exports X as a final good is that it can be negative. When net exports are negative, the perspective taken here is that imports can be obtained either by means of the export commodity *or* by means of an electronic credit stated in terms of units of the export commodity. Thus, negative net exports are in units of the export commodity and represent the amount of electronic credits treated as equivalent to units of the export commodity that domestic buyers of imports have been lent. This perspective is valuable because there is only one aggregate export commodity, but many types of imports that the model keeps track of. (Everything would be equivalent if net exports were expressed in terms of any particular imported commodity, but it would be necessary to pick one, since differential changes in the prices of different imported commodities makes them non-equivalent, financially.)

Because the model conditions on the time path of net exports, it is not essential to keep track of the foreign debt. That is, the main reason one would need to keep track of foreign debt is to determine the time path of net exports. (It is part of the relevant intertemporal budget constraint in that determination.) Since keeping track of the foreign debt is not central to the model here, for convenience, one can think of the foreign debt as being stated in terms of units of the export commodity.

Computationally, the key to extending results to the open economy is to lay out the open economy version of the matrix B , which we will call B^* . First, define a row vector $b_x' = [b_{x_i}]$, of the shares of each domestic commodity in aggregate exports, where the prime is a reminder that this is a row vector. Next, define the diagonal matrix

$\Delta = [\delta_{ij} \nu_i]$ of import shares for each commodity, where δ_{ij} is again the Dirac delta equal to 1 when its subscripts are equal and zero otherwise, and ν_i is the share of commodity i that is imported. Then with rows representing outputs, columns representing inputs and the commodities in the order *exports, foreign commodities, domestic commodities, usable commodities*,

$$B^* = \begin{bmatrix} 0 & 0 & 0 & b_x' \\ \iota & 0 & 0 & 0 \\ 0 & 0 & 0 & B \\ 0 & \Delta & I - \Delta & 0 \end{bmatrix}$$

where ι is a column vector of ones. That is, exports are produced from domestic commodities with the shares in the row vector $b_x' = [b_{x_i}]$. Imported commodities are “produced” solely from the export commodity, which therefore has a share of 1 for each imported commodity. Domestic commodities are produced from usable commodities according to the matrix of input shares B discussed above for the closed economy. Finally, usable commodity i is produced from the corresponding foreign commodity with share ν_i and the corresponding domestic commodity with share $1 - \nu_i$. These shares show up in the diagonal matrices Δ and $I - \Delta$. All other input shares for this extended matrix are zero.

Computationally, it is also important to specify the counterparts to the capital and labor share vectors and the counterpart to the technology shock vector. Since capital and labor are only used in the production of domestic commodities, one can write:

$$a_K^* = \begin{bmatrix} 0 \\ 0 \\ a_K \\ 0 \end{bmatrix}, \quad a_N^* = \begin{bmatrix} 0 \\ 0 \\ a_N \\ 0 \end{bmatrix}.$$

That is, capital and labor are only *directly* used in the production of domestic commodities. Finally, let $\tau_i = p_x - p_i^{\text{foreign}}$ be the logarithm of the terms of trade for the foreign produced version of commodity i and $\tau = [\tau_i]$ be the vector of terms of trade. Then

$$dz^* = \begin{bmatrix} 0 \\ d\tau \\ dz \\ 0 \end{bmatrix}.$$

That is, changes in the terms of trade for each foreign-produced commodity acts as the technology shock for that foreign-produced commodity. Domestically-produced commodities have the technology shocks specified above in the discussion of the closed-economy model. By construction, there are no technology shocks to the aggregation procedures for aggregating the different commodities that are exported or for aggregating foreign and domestic versions of a commodity into usable commodities.

It is possible to algebraically compute $(I - B^*)^{-1}$:

$$I - B^* = \begin{bmatrix} 1 & 0 & -b_x' & 0 \\ -t & I & 0 & 0 \\ 0 & 0 & I & -B \\ 0 & -\Delta & \Delta - I & I \end{bmatrix}$$

and

$$(I - B^*)^{-1} = \begin{bmatrix} 1 + b_x' \Omega^{-1} \Delta t & b_x' \Omega^{-1} \Delta & b_x' \Omega^{-1} (I - \Delta) & b_x' \Omega^{-1} \\ t + t b_x' \Omega^{-1} \Delta t & I + t b_x' \Omega^{-1} \Delta & t b_x' \Omega^{-1} (I - \Delta) & t b_x' \Omega^{-1} \\ B \Omega^{-1} \Delta t & B \Omega^{-1} \Delta & I + B \Omega^{-1} (I - \Delta) & B \Omega^{-1} \\ \Omega^{-1} \Delta t & \Omega^{-1} \Delta & \Omega^{-1} (I - \Delta) & \Omega^{-1} \end{bmatrix}$$

where

$$\Omega = I - B + \Delta B - \Delta t b_x'.$$

This can be verified by multiplying to confirm that

$$(I - B^*)(I - B^*)^{-1} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & I & 0 & 0 \\ 0 & 0 & I & 0 \\ 0 & 0 & 0 & I \end{bmatrix}.$$

Substituting in to the formulas for the objects of interest:

$$\alpha^* = (I - B^*)^{-1} a_K^* = \begin{bmatrix} b_x' \Omega^{-1} (I - \Delta) a_K \\ t b_x' \Omega^{-1} (I - \Delta) a_K \\ [I + B \Omega^{-1} (I - \Delta)] a_K \\ \Omega^{-1} (I - \Delta) a_K \end{bmatrix},$$

$$1 - \alpha^* = (I - B^*)^{-1} a_N^* = \begin{bmatrix} b_x' \Omega^{-1} (I - \Delta) a_N \\ t b_x' \Omega^{-1} (I - \Delta) a_N \\ [I + B \Omega^{-1} (I - \Delta)] a_N \\ \Omega^{-1} (I - \Delta) a_N \end{bmatrix},$$

and

$$\Psi^* dz^* = -(I - B^*)^{-1} \begin{bmatrix} 0 \\ d\tau \\ dz \\ 0 \end{bmatrix} = \begin{bmatrix} -b_x' \Omega^{-1} \Delta d\tau - b_x' \Omega^{-1} (I - \Delta) dz \\ -(I + t b_x' \Omega^{-1} \Delta) d\tau - t b_x' \Omega^{-1} (I - \Delta) dz \\ -B \Omega^{-1} \Delta d\tau - [I + B \Omega^{-1} (I - \Delta)] dz \\ -\Omega^{-1} \Delta d\tau - \Omega^{-1} (I - \Delta) dz \end{bmatrix}.$$

The fourth (bottom) component of each of these is the component relevant to *usable* commodities, and appropriately aggregated, to consumption, investment and government purchases. The first (top) component is relevant to the export commodity. In other words, for final goods usable commodities, the reduced-form change in effective value-added productivity (including both technology and terms of trade effects) is

$$\Omega^{-1} (I - \Delta) dz + \Omega^{-1} \Delta d\tau$$

(Note the sign change going from change in price to change in productivity.) For exports, the reduced-form change in productivity is

$$b_x' \Omega^{-1} (I - \Delta) dz + b_x' \Omega^{-1} \Delta d\tau$$

Thus, the reduced form change in effective value-added productivity for exports is just the export share vector times the reduced form change in effective value-added productivity for commodities.

Note that if Δ , the import shares, is zero, then the fourth component reduces to the closed-economy formula, since

$$\begin{aligned} B^{-1}[(I - B)^{-1} - I] &= B^{-1}[-I + I + B + B^2 + \dots] = B^{-1}[B + B^2 + B^3 + \dots] = [I + B + B^2 + \dots] \\ &= (I - B)^{-1} \end{aligned}$$

(or multiply both the far left and far right of this equation by $I - B$ to verify).

Appendix B: Updating Basu, Fernald, and Kimball (2004) Industry Residuals

We used our new dataset to update the estimated industry technology series along the lines of Basu, Fernald, and Kimball (2006). BFK follow Hall (1990), but add industry hours per worker to the regression to control for unobserved variations in labor effort and capital's workweek. The basic idea behind the utilization proxy is that a cost-minimizing firm operates on all margins simultaneously, both observed and unobserved. As a result, changes in observed margins can proxy for unobserved utilization changes. If labor is particularly valuable, for example, firms will work existing employees both longer (observed hours per worker rise) and harder (unobserved effort rises).

More specifically, let industries have the following gross-output production function:

$$(24) \quad Y_i = F^i(A_i K_i, E_i H_i N_i, M_i, Z_i).$$

The industry produces gross output, Y_i , using the capital stock K_i , employees N_i , and intermediate inputs of energy and materials M_i . The capital stock and number of employees are quasi-fixed, so their levels cannot be changed costlessly. But industries may vary the intensity with which they use these quasi-fixed inputs: H_i is hours worked per employee; E_i is the effort of each worker; and A_i is the capital utilization rate (that is, capital's workweek). Total labor input, L_i , is the product $E_i H_i N_i$. The production function F^i is (locally) homogeneous of γ_i in total inputs. γ_i exceeding one implies increasing returns to scale, reflecting overhead costs, decreasing marginal cost, or both. Z_i indexes technology.

With cost minimization, the standard first-order conditions give the output elasticities, which we use to weight growth of each input. Let dx_i be observed share-weighted input growth, and du_i be unobserved share-weighted growth in utilization, i.e., the workweek of capital and labor effort. (For any variable J , we define dj as its logarithmic growth rate $\ln(J_t / J_{t-1})$.) This yields:

$$(25) \quad dy_i = \gamma_i(dx_i + du_i) + dz_i,$$

where

$$(26) \quad \begin{aligned} dx_i &= s_{K_i} dk_i + s_{L_i} (dn_i + dh_i) + s_{M_i} dm_i, \\ du_i &= s_{K_i} da_i + s_{L_i} de_i, \end{aligned}$$

and s_{j_i} is the ratio of payments to input J in total cost. (We have normalized the elasticity with respect to technology, $(F_Z Z / F)$, to unity, since it isn't separately identified.)

The BFK proxy for for unobserved variations in utilization, du_i , is derived from a dynamic cost-minimizing firm problem.. Specifically, suppose firms seek to minimize the present discounted value of costs for any given path of output. There is a convex cost of adjusting the quasi-fixed factors (K and N). Total compensation per worker is $WG(H, E)V(A)$, where W is the base wage; the function G specifies how the hourly wage depends on effort, E , and the length of the workday, H ; and $V(A)$ captures a shift premium associated with varying capital's workweek. (This true shadow wage might be implicit, and not observed in the data at high frequency.) Thus, firms must pay a higher wage to get workers to work longer (on a given shift), or to work harder, or to work at night. In this environment, BFK show that, as a first-order approximation, the following expression holds:

$$dy_i = \gamma_i dx_i + \beta_i dh_i + dz_i.$$

The coefficient β_i , which can be estimated, relates observed hours growth to unobserved variations in labor effort and capital's workweek. That coefficient incorporates various elasticities including, in particular, the elasticity of unobserved effort with respect to hours, from the implicit function relating them (which came out of optimization).

In this paper, we impose constant returns in order to be consistent with the final-use aggregation theory. Results are, however, generally robust to allowing for industry non-constant returns. Previous work (e.g., Basu and Fernald, 1997, as well as BFK) find that the typical industry has close to constant returns. In addition, an advantage of imposing constant returns is that in some industries it is difficult to distinguish returns to scale from utilization, since they are positively correlated measures of inputs.

We estimate these equations as a system using instrumental variables, in order to control for correlation between technology shocks and input growth. As described in BFK, we use updated versions of two of the Hall-Ramey instruments: oil prices and growth in real government defense spending. For oil, we use increases in the U.S. refiner acquisition price. We also use, as a third instrument, an updated estimate of Burnside's (1996) quarterly Federal Reserve "monetary shocks" from an identified VAR. In all cases, we have quarterly instruments; we sum the four year $t-1$ quarterly shocks as instruments. For the money shock, we lag the instrument an extra year in case the VAR does not fully control for the systematic response to technology shocks.²²

[Specific regression results will be available in a later version of the paper. NEED TO ADD.]

²² BFK find that results are robust to different combinations and lags of the instruments; results are not driven by any one of these instruments. They also discuss the small sample properties of instrumental variables.

We need residuals for all commodities (industries) that are used as intermediate inputs or that use intermediate inputs. BFK did not estimate residuals for agriculture, mining industries, or government enterprises. TFP is highly volatile, and the instruments have no explanatory power.

To conserve parameters, we restrict the utilization coefficient within five groups. We split durables manufacturing into two groups, based on their weights in final equipment demand. The reason is that we care a lot about the coefficient in the “core” group, which comprises fabricated metals, non-electrical and electrical machinery, motor vehicles, and other transportation equipment. Within non-durables, we split out chemicals and petroleum from other industries, because the data clearly preferred to put a different coefficient on these industries. Finally, we include non-manufacturing (8).

For the commodity-specific terms of trade, we use the aggregate export deflator relative to the commodity-specific import deflator. Even if there is no foreign supply of the commodity, there are still imports of non-competing goods. In any cases where there are no imports into an industry, we set the import-supply price to one (unchanging over time). This is innocuous, since it gets multiplied by a weight of zero.

The commodity-specific terms of trade is, of course, endogenous, and may respond to a wide range of other shocks (e.g., demand shocks at home and abroad that affect nominal exchange rates). As a result, in our regression analysis, we instrument by recalculating the final-use shocks where we zero out the terms-of-trade vector as well as the technology shocks in industries where we don't have controls for utilization (agriculture, mining, and government enterprises).

Further details on the production-function estimates will be available in a later draft.

References

- Basu, Susanto and John Fernald (1997). "Returns to Scale in U.S. Production: Estimates and Implications." *Journal of Political Economy*.
- Basu, Susanto, John Fernald, and Miles Kimball (2006). "Are Technology Improvements Contractionary?" *American Economic Review*, December.
- Basu, Susanto and Miles Kimball (2000). "A Two-Sector DGE Model with Nominal Rigidities." Unpublished, University of Michigan.
- Campbell, Jeffrey (1998). "Entry, Exit, Embodied Technology, and Business Cycles," *Review of Economic Dynamics*, 1 (April): 371-408.
- Christiano, L. and J. Fisher, "Stock Market and Investment Good Prices: Implications for Macroeconomics," Chicago Fed Working Paper 98-6, 1998.
- Cochrane, J. (1994). "Shocks."
- Cummins, J. and G. Violante (2002). "Investment-Specific Technical Change in the United States (1947-2000): Measurement and Macroeconomic Consequences," *Review of Economic Dynamics*, 5(2), (April): 243-284.
- Fisher, J, "Relative Prices, Complementarities and Co-movement among Components of Aggregate Expenditures," *Journal of Monetary Economics*, 39 (August 1997), 449-74.
- Fisher, Jonas (2006). "The Dynamic Effects of Neutral and Investment-Specific Technology Shocks." *Journal of Political Economy*.
- Galf, Jordi (1999), "Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?" *American Economic Review*, 89 (March 1999).
- Gordon, Robert J. (1990). *The Measurement of Durable Goods Prices*, Chicago: University of Chicago Press.
- Greenwood, J., Z. Hercowitz and G. Huffman, "Investment, Capacity Utilization and the Business Cycle," *American Economic Review*, 78 (June 1988), 402-17.
- Greenwood, Jeremy, Zvi Hercowitz, and Per Krusell (1997). "Long-Run Implications of Investment-Specific Technological Change," *American Economic Review*, vol. 87(3), pages 342-62.
- Greenwood, J., Z. Hercowitz and P. Krusell (2000). "The Role of Investment-Specific Technological Change in the Business Cycle," *European Economic Review*.
- Herrendorf, Berthold & Valentinyi, Akos (2007). "Measuring Factor Income Shares at the Sector Level - A Primer," CEPR Discussion Papers 6271, C.E.P.R. Discussion Papers.
- Kehoe, T. J. and Ruhl, K. J. (2006). "Are Shocks to the Terms of Trade Shocks to Productivity?" NBER wp 1311.
- Kimball, Miles (1994). "Proof of Consumption Sector Neutrality." University of Michigan manuscript, August.
- Lawson, Ann M., Kurt S. Bersani, Mahnaz Fahim-Nader, and Jiemin Guo (2002). "Benchmark Input-Output Accounts of the United States." *Survey of Current Business*, December.
- Rotemberg, J. J. and Woodford, M. (1996). "Real-Business-Cycle Models and the Forecastable Movements in Output, Hours, and Consumption," *American Economic Review*, vol. 86.

Whelan, Karl (2000). "Balanced Growth Revisited: A Two-Sector Model of Economic Growth." FEDS discussion paper, Board of Governors of the Federal Reserve System.

Table 1. Sample-average growth rates of sector-specific technologies
(Standard TFP and utilization-adjusted technology estimates)

A. Final-goods sector TFP growth (percent change, annual rate)			
	1950-2005	1950-1982	1983-2005
(1) Consumption TFP, dz_C	0.11	0.20	-0.01
(2) Durables Consumption, TFP dz_{CD}	2.04	1.52	2.82
(3) Investment TFP, dz_J	1.45	1.05	2.04
3a. Equipment dz_{JE}	2.99	2.20	4.19
3b. Structures dz_{JS}	-0.35	-0.09	-0.75
(4) Government TFP, dz_G	1.34	1.16	1.61
(5) Trade TFP, dz_{Trade}	2.22	1.63	3.10
B. Final-goods sector technology, using utilization-adjusted industry shocks (percent change, annual rate)			
	1950-2005	1950-1982	1983-2005
(1) Consumption TFP, dz_C	0.10	0.21	-0.05
(2) Durables Consumption, TFP dz_{CD}	2.07	1.57	2.82
(3) Investment TFP, dz_J	1.10	0.74	1.62
3a. Equipment dz_{JE}	3.07	2.39	4.09
3b. Structures dz_{JS}	-1.19	-0.93	-1.58
(4) Government TFP, dz_G	1.05	0.89	1.29
(5) Trade TFP, dz_{Trade}	2.26	1.74	3.04

Notes: In the top panel, TFP is calculated for the 35 domestic industries in the Jorgenson dataset and aggregated to final-use sectors as described in the text. The bottom panel uses industry technology residuals, as described in the text. Those estimates control for variations in factor utilization.

Table 2. Final-Use Capital Shares

Final-good category	Capital share
Consumption (Non-durables and services)	0.36
Consumer Durables	0.31
Equipment Investment	0.30
Structures	0.26
Government	0.30
Government Consumption	0.35
Government Equipment	0.26
Government Structures	0.28
Net Exports	0.35
Sample average 1961-2004	

Notes: For final-use vector b_h , where $h \in \{C, CD, JE, JS, G, GC, GE, GS, NX\}$, capital's share is calculated as $b_h[I - B]^{-1}s_k'$, where s_k' is the vector of commodity capital shares

Table 3. Relative output prices and relative technology
(percent change, annual rate)

A. Prices of nondurable and services consumption relative to price of equipment investment, BEA and Jorgenson datasets

	1960-2005	1960-1982	1982-2005
BEA	2.69	1.58	3.70
Jorgenson-adjusted with Gordon deflators	2.67	1.61	3.67

Notes: Price index for consumption of non-durables and services is calculated as a chained-price index.

B. Jorgenson data (adjusted with Gordon deflators):
Relative prices and relative sectoral TFPs

$$dz_j - dz_c = (dp_c - dp_j) + (b_j - b_c)(I - B)^{-1} [s_K dp_K + s_L dp_L] + (b_j - b_c)(I - B)^{-1} [p^{Dom} - dp^{Dom, Producer}]$$

= Relative price + Factor-Price Wedge + Tax Wedge

	Full Sample	1960-1982	1982-2004	Acceleration
	(1)	(2)	(3)	((3)-(2))
(1) Relative TFP (Equip. invest. relative to NDS consump.)	3.15	2.11	4.14	2.04
(2) Relative price (NDS consump. relative to equip. investment)	2.68	1.60	3.71	2.10
(3) Relative primary input prices (equip relative to NDS cons.)	0.33	0.40	0.27	-0.14
(4) Relative sales tax wedge (equip. relative to NDS cons.)	0.14	0.10	0.17	0.07

Notes: Lines (2) through (4) sum to line (1). NDS is non-durables and services. In the equation, $(b_j - b_c)$ is a row vector for a given year of commodity-specific shares in final investment versus consumption; $[s_K dp_K + s_L dp_L]$ is a column vector of commodity-specific share-weighted input price growth, where only domestic commodities have non-zero entries; similarly, $[p^{Dom} - dp^{Dom, Producer}]$ is a column vector of the difference between price paid by a purchaser and the price received by a producer. Terms are defined and discussed further in the text.

Table 4. Correlations of sector-specific residuals

A. Selected correlations of final-use TFP			
	1950-2005	1950-1982	1983-2005
(1) $\text{Corr}(dz_J, dz_C)$	0.82	0.88	0.73
(2) $\text{Corr}(dz_{JE}, dz_C)$	0.69	0.80	0.63
(3) $\text{Corr}(dz_{JE} - dz_C, dz_C)$	0.16	0.28	0.08
(4) $\text{Corr}(dz_{JE} - dz_C, dz_J)$	0.57	0.52	0.68
B. Selected correlations of final-use technology			
	1950-2005	1950-1982	1983-2005
(1) $\text{Corr}(dz_J, dz_C)$	0.52	0.54	0.58
(2) $\text{Corr}(dz_{JE}, dz_C)$	0.24	0.29	0.29
(3) $\text{Corr}(dz_{JE} - dz_C, dz_C)$	-0.21	-0.18	-0.22
(4) $\text{Corr}(dz_{JE} - dz_C, dz_J)$	0.54	0.52	0.55

Notes: dz_J is TFP or technology for overall investment; $dz_{J,E}$ is for equipment and software; dz_C is for non-durables and services consumption. With 56 observations, a correlation of 0.22 is significant at the 10 percent level and 0.27 is significant at the 5 percent level; with 22 observations, a correlation of 0.35 is significant at the 10 percent level and 0.41 is significant at the 5 percent level. See also notes to Table 1.

Table 5: Regression results

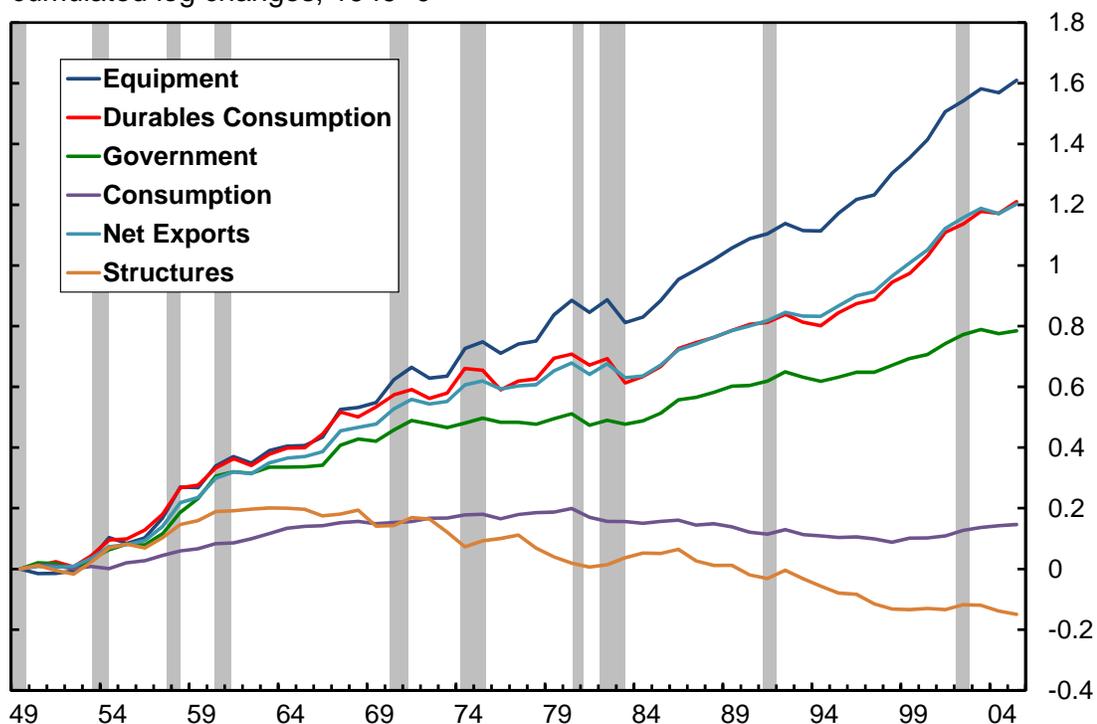
		Technology shocks						
		Investment and Exports			Consumption (nondurables and services)			
		dz_J	$dz_{J,t-1}$	$dz_{J,t-2}$	dz_C	$dz_{C,t-1}$	$dz_{C,t-2}$	R^2
(1)	GDP, GCV adjusted.	-0.33 (0.09)	-0.01 (0.14)	0.12 (0.11)	0.90 (0.25)	0.01 (0.43)	0.00 (0.31)	0.34
(2)	Equip. and Soft. Invest, GCV adj.	-1.03 (0.38)	-0.89 (0.45)	0.41 (0.35)	2.76 (1.15)	0.97 (1.40)	-1.05 (1.01)	0.20
(3)	Consumer durables, GCV adj.	-0.51 (0.37)	0.20 (0.41)	0.27 (0.31)	3.19 (0.80)	-0.75 (1.27)	-1.02 (0.95)	0.26
(4)	Consump. of nondur. and services, GCV adj.	-0.18 (0.05)	0.02 (0.05)	0.02 (0.07)	0.38 (0.15)	0.07 (0.14)	0.01 (0.16)	0.47
(5)	Non-res. structures invest.	-0.32 (0.45)	-0.92 (0.37)	-0.03 (0.51)	-0.75 (1.28)	1.88 (0.87)	1.01 (1.00)	0.23
(6)	Residential structures invest.	-1.59 (0.68)	1.22 (0.62)	0.11 (0.63)	6.16 (1.80)	-0.75 (1.56)	-3.30 (1.71)	0.44
(7)	Exports, GCV adj.	-0.39 (0.34)	-0.25 (0.38)	0.91 (0.42)	0.53 (0.80)	1.17 (0.80)	-1.09 (1.07)	0.23
(8)	Imports, GCV adj.	-0.67 (0.24)	-0.35 (0.32)	0.58 (0.36)	1.99 (0.84)	-0.20 (0.95)	-1.25 (0.92)	0.26
(9)	Hours worked	-0.39 (0.08)	-0.19 (0.12)	0.17 (0.11)	0.29 (0.36)	0.18 (0.39)	-0.32 (0.25)	0.29
(10)	Wage (Jorgenson, composition-adj.)	-0.24 (0.13)	-0.26 (0.13)	-0.26 (0.13)	0.41 (0.32)	0.29 (0.30)	0.04 (0.30)	0.12
(11)	GDP deflator (Jorgenson), GCV adj.	-0.39 (0.11)	-0.27 (0.11)	-0.30 (0.14)	-0.31 (0.34)	0.19 (0.28)	0.18 (0.27)	0.45
(12)	Rel price: CNDS to Equip (Jorgenson), GCV-adj.	0.27 (0.10)	0.09 (0.10)	0.01 (0.11)	-0.28 (0.30)	-0.02 (0.26)	0.01 (0.24)	0.16
(13)	Price of NDS Consumption (Jorgenson), GCV adj.	-0.28 (0.11)	-0.26 (0.13)	-0.32 (0.15)	-0.42 (0.37)	0.21 (0.29)	0.32 (0.33)	0.39
(14)	Price of equipment (Jorgenson), GCV adj.	-0.54 (0.15)	-0.35 (0.13)	-0.33 (0.21)	-0.14 (0.47)	0.23 (0.38)	0.31 (0.33)	0.16
(15)	Fed Funds rate	-0.09 (0.13)	-0.31 (0.10)	-0.15 (0.12)	-1.25 (0.34)	-0.59 (0.31)	-0.21 (0.32)	0.46
(16)	10-year Treasury	-0.17 (0.10)	-0.22 (0.13)	-0.12 (0.11)	-0.55 (0.24)	-0.85 (0.24)	-0.51 (0.29)	0.45

Notes: Instrumental variables regressions from 1952-2005. All variables except interest rates are in annual growth rates; interest rates are entered in levels. Dependent variables are shown in the left column. Regressors are current and lagged values of BFK final-use-sector technology shocks shown, along with a constant term (not shown). Regressions with prices or the wage also include decade dummies for the 1970s and 1980s. Instruments are corresponding variables where the terms of trade and technology in agriculture, mining, and government enterprises are zeroed out before aggregation. Heteroskedastic and auto-correlation robust standard errors in parentheses.

Figure 1: Cumulated Final-Use Sector Technology Series

Final-Use Technology

cumulated log changes, 1949=0



Note: All series are calculated by cumulating log changes of utilization-adjusted technology from 1950-2005. These series “zero out” the terms-of-trade shocks and technology in sectors for which no utilization adjustment is available. The value for 1949 is normalized to zero.

Figure 2: Relative Prices, Relative TFP, and Relative Technology

Relative Prices, Technology, and TFP

cumulated log growth rates

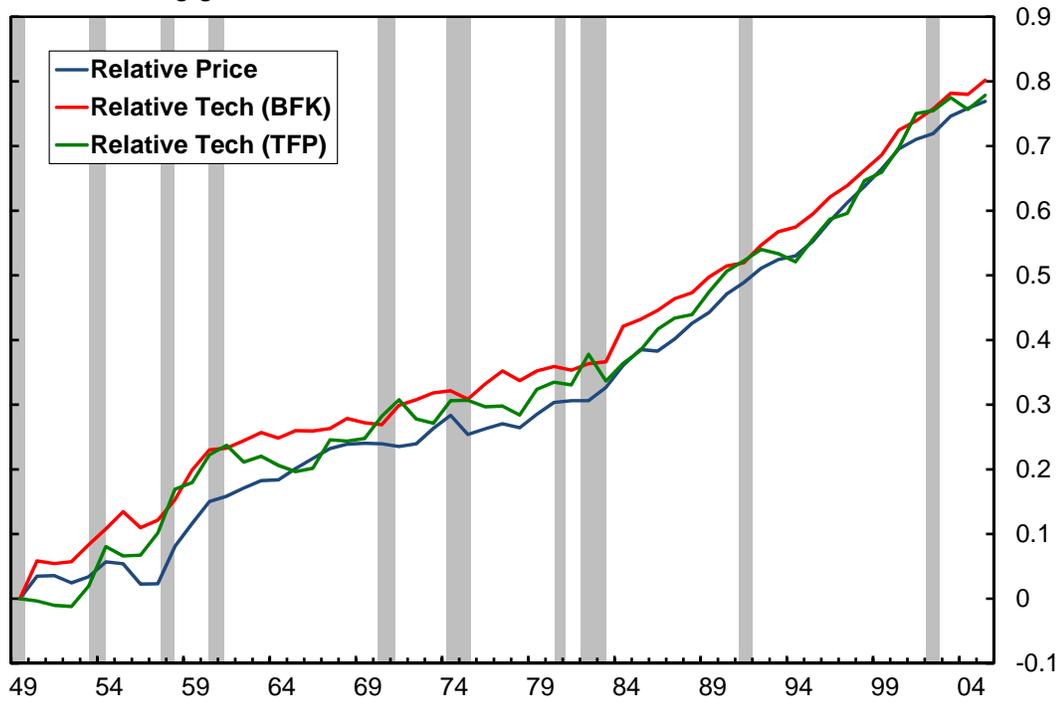
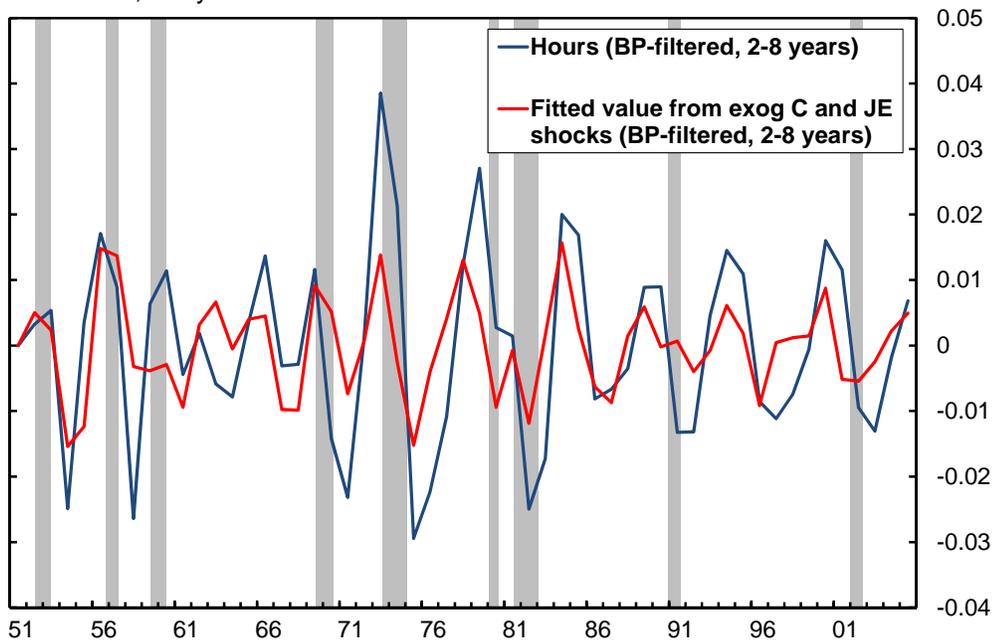


Figure 3: Fitted values of hours and equipment to technology
Business-cycle frequency

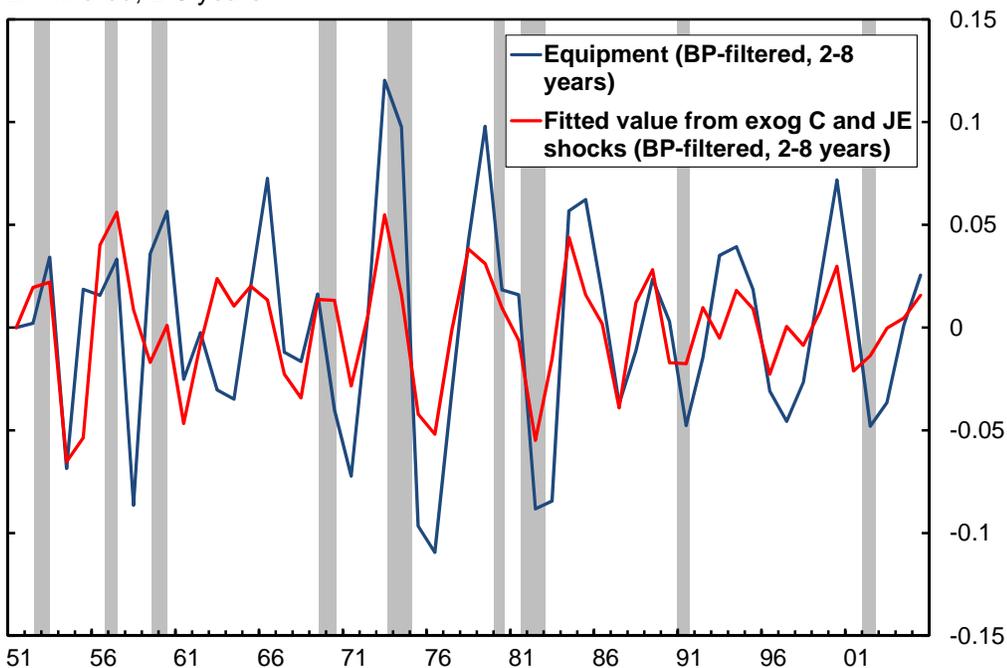
Hours

BP-filtered, 2-8 years



Equipment

BP-filtered, 2-8 years



Fitted values are from regressions shown in Table 5.