Alwyn Young

Structural transformation, the mismeasurement of productivity growth and the cost disease of services

Working paper, etc.

Structural Transformation,
the Mismeasurement of Productivity Growth,
and the Cost Disease of Services

Abstract

If workers self-select into sectors based upon their relative productivity in different tasks, and comparative advantage is aligned with absolute advantage, then as a sector's employment share increases (decreases) the average efficacy of its workforce will fall (rise). This provides a potential explanation for the differential in the measured productivity growth of contracting goods and expanding services. Using changes in defense expenditures as an exogenous shifter of employment shares, I estimate that the elasticity of worker efficacy with respect to employment shares is substantially negative. While conventional estimates indicate that productivity growth in goods is .8% and 1.4% faster than in services in the US and the OECD, respectively, regression point estimates suggest that the true difference might lie between a .5 percent advantage for goods and a .4 percent advantage for services. Taking the middle of this range, the view that goods and services have similar productivity growth rates provides a plausible alternative characterization of growth in developed economies.

Alwyn Young*
Department of Economics
London School of Economics

This Draft: October 2013

*I am greatly indebted to Martin Eichenbaum, Ho Veng-Si and anonymous referees for many helpful comments.

I. Introduction

One of the strongest and seemingly most accurate characterizations of the process and problems of growth in advanced economies is William Baumol's "Cost Disease of Services." Baumol's argument, begun in papers as early as 1965 and continuing to this very day (e.g. Baumol 1965, 1967, 1985 and 2012), starts from the premise that productivity growth is inherently more difficult to achieve in the production of services than in the production of goods. With the two industries competing for factors of production in the same factor markets, the relative cost of producing service output inevitably rises. If the demand for services were income inelastic and price elastic, these trends would not pose a problem, as the share of services in nominal GDP would decline. Alas, precisely the opposite is true, and services garner an increasing share of nominal output. Aggregate productivity growth, equal to the nominal output share weighted average of sectoral productivity growths, must steadily decline.[1]

Decades of data on productivity growth in goods and services have confirmed Baumol's thesis turning it, for all intents and purposes, into a stylized fact of economic growth. Productivity statistics, however, are based on the fundamental assumption that each new worker is qualitatively the same as every old worker.[2] If workers self-select into industries based upon unobservables, this assumption may create a systematic bias, as the type of workers present when an industry is small may not be the same as when the industry becomes large, and vice versa.

In his "Thoughts on the Distribution of Earnings", Roy (1951) identified the mechanism central to this paper. Workers select the industry in which they have the highest relative

[1]Although not mentioned in the papers cited above, implicit in Baumol's argument is the notion that service output is relatively non-tradeable. Otherwise, low productivity growth in services could be met, at least at the individual country level, by exporting more manufactures for services.

[2]To be sure, more sophisticated analyses divide workers into categories based upon observable determinants of human capital such as age and education, but within each category the assumption is ultimately made that all workers are identical.

productivity, i.e. a comparative advantage. If individual productivity in different tasks is uncorrelated or at worst weakly correlated, then individuals having a comparative advantage in an industry will on average also have an absolute advantage in that sector. As a sector expands by offering higher wages to prospective workers elsewhere in the economy, it will draw in individuals with both a lower comparative advantage and a lower absolute advantage in the sector, while leaving individuals with the highest comparative and absolute advantage in competing sectors. Consequently, productivity in expanding sectors will appear to decline and productivity in contracting sectors will appear to rise. In sum, in a Roy world the apparent disparity in the productivity growth of goods and services may come about because services expand by drawing in people who are, as examples, less adept at finance, law and medicine, while goods sectors contract by shedding the least able farmers, manufacturers and miners, all of which is not taken into account in measures of productivity growth. Underlying true levels of productivity growth, i.e. taking into account the average efficacy of the workers present in the two sectors, might not be all that different.

Figure I, which graphs the relative supply and demand for services, summarizes the argument made in this paper. Baumol's supply curve is essentially a horizontal line, determined by the relative productivity of the two sectors.[3] As goods experience more rapid productivity growth, this supply curve shifts up, from S_0^{Baumol} to S_1^{Baumol}, exemplifying the cost disease of services. At the same time, as a consequence of the relatively higher income elasticity of demand for services, the relative demand curve shifts out from D_0 to D_1. The equilibrium moves from E_0 to E_1, with a higher relative output and price of services, which consequently has a growing

[3]If the capital income shares (i.e. factor intensities) of the two sectors differ, the supply curve will be upward sloping even without the effects Roy describes. However, as discussed in Appendix A, empirically the capital income shares of goods and services in the US economy are almost identical and the upward slope in the supply curve attributable to this effect is negligible, i.e. an increase in relative prices of .4 of one percent as relative output goes from 0 to ∞. In the sources cited above Baumol and his co-authors never emphasize a relative price effect emanating from relative factor intensities and, in this regard, appear to be completely correct.

Figure I: Alternative Views of Relative Supply

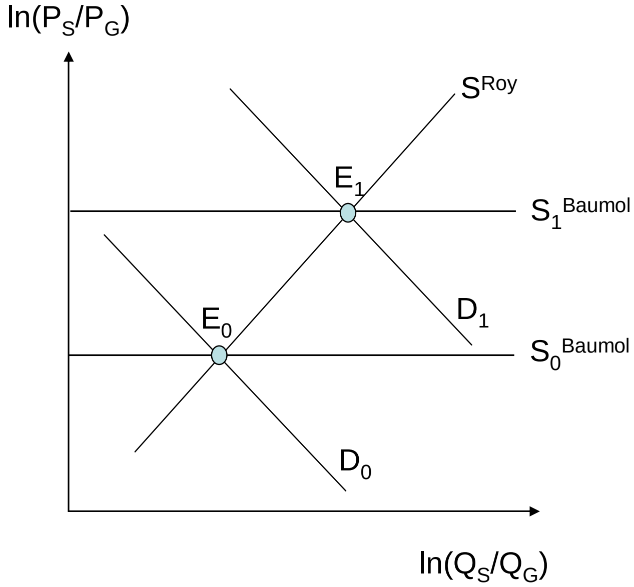

Figure II: Welfare Implications with the Same Equilibrium Prices and Quantities (TFP growth = ΔV/V)

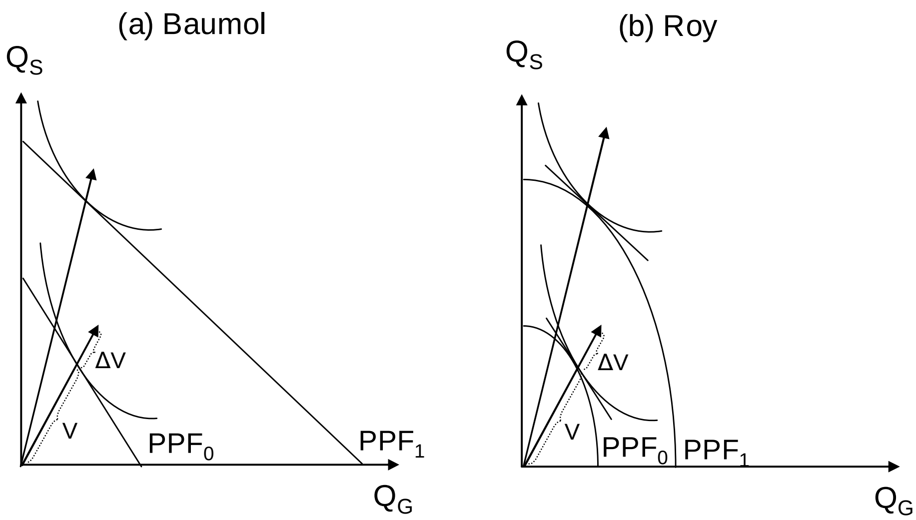

nominal share of the economy. An alternative hypothesis, however, is that the supply curve is substantially upward sloping because of the correlation between comparative and absolute advantage Roy describes. As drawn in the figure, the Roy supply curve S^{Roy} intersects both E_0 and E_1. This describes a situation in which productivity growth is the same in both sectors, so the supply curve does not shift, but the relative demand curve shifts out as incomes rise. Here the rise in the relative price of services is driven purely by the changing efficacy of the average worker in each sector.[4]

Figure I makes clear that the Roy model does not deny the rise in the relative price of services, it merely explains it with a different mechanism. Figure II illustrates why this matters. Panel (a) draws the linear production possibilities frontier implied by the Baumol model, which rotates out as goods experience more rapid productivity growth.[5] Panel (b) draws the Roy production possibilities frontier, which shifts out uniformly when productivity growth is identical in both sectors. This panel shows that the same equilibrium price and quantity relations can be explained with equiproportional shifts of the intercepts of the production possibility frontier and a movement along its concave surface. For the purposes of heuristically illustrating welfare implications, the diagrams also include social indifference curves which, under the assumption of competitive markets, are tangent to the production possibilities frontiers. Aggregate total factor productivity growth is the proportional increase in the length of the ray from the origin to the tangent line on the production frontier ($\Delta V/V$ in the figure).[6] In the Baumol model, as the share

[4]As Figure I makes clear, for Baumol's argument it does not matter whether or not the relative real output of services is rising (only that its nominal share is increasing), but for the Roy argument it does. Baumol et al (1985) argue that there is no change in the relative output of goods and services. This is actually not true. As discussed in Section III, US and OECD data clearly indicate a large rise in the relative real output of services in the post-war era.

[5]For the purposes of this expositional diagram, I assume that factor supplies are constant.

[6]To see this, note that if inputs are constant (as assumed in the diagram), we can describe the problem of maximizing GDP as one of maximizing $P_xX + P_YY$ s.t. $0 \geq F(X,Y,t)$. Differentiating the binding production possibilities constraint, we have (a) $F_XdX + F_YdY + F_tdt = 0$. Rearranging and making use of the first order conditions from the maximization problem ($\lambda F_X = P_X$, etc), one finds that:

of services in total expenditure grows, the growth rate of this vector slows. In the Roy model, the proportional growth rate remains constant. Over time there is a growing discrepancy in the instantaneous rate of welfare growth predicted by the two models.

This paper draws its inspiration from recent interest in the macro implications of Roy's model. Lagakos and Waugh (2011) argue that selection effects of the type described in this paper can explain the greater relative productivity of agricultural workers to non-agricultural workers in countries with larger non-agricultural sectors. Hsieh, Hurst, Jones and Klenow (2012) calculate the inefficiency associated with the historical concentration of women and African-Americans in particular occupations using a Roy model and argue that the gradual elimination of barriers to the participation of these groups in other occupations can explain as much as 1/5th of post-war US aggregate wage growth. Kuralbayeva and Stefanski (2013), independent of the early draft of this paper, argue that the decrease of manufacturing output brought about by the appreciation of the real exchange rate associated with resource windfalls generates a spurious rise in manufacturing productivity as the contraction of the sector leaves only the most productive workers behind. This paper extends these Roy-related analyses to the general consideration of the relative productivity of goods and services. Along the way, I establish the theoretical bias in conventional measures of sectoral and aggregate productivity and clarify the mathematical conditions necessary for Roy effects to be present (i.e. for average worker efficacy to be declining in a sector's employment share). While the papers above calibrate their models, this paper estimates the size of Roy effects using regression techniques.

$$g_{TFP} = -\lambda F_t / GDP = \theta_X g_x + \theta_Y g_Y$$

where λ is the value of relaxing the PPF constraint, Θ_i the GDP share of product i, and g_i the growth of the output of i. Thus, total factor productivity growth, the proportionate value of the time trend relaxation of the PPF constraint, equals the GDP share weighted increase of the output of each product. (a) above, however, holds whether the dX & dY are the observed values or imposed values such that dX = g*X & dY = g*Y. Thus, regardless of the bias of TFP growth, one can equally say that $g_{TFP} = \Theta_X g + \Theta_Y g = g$. g is the proportionate increase in the length of the vector (with slope determined by the current production bundle) from the origin to the production possibilities frontier.

With regards to empirically estimating the elasticity of average worker efficacy with respect to the sectoral employment share, the key parameter in the macro implementation of the Roy model, there has been little prior research. A partial exception is provided by McLaughlin and Bils (2001), who use PSID data to show that the wages of entrants or leavers are lower than those of continuing workers. However, as shown further below, the PSID data used in that paper mostly concern simultaneous entry and exit (a form of employment churning) and are uncorrelated with changes in sectoral employment shares. This paper focuses directly on the impact of changes in sectoral employment, using private sector employment changes driven by changes in military spending to identify the elasticity of average worker efficacy with respect to sectoral employment.

Zvi Griliches, in his presidential address (1994) and earlier (1992), brought to the profession's attention the shortcomings of US measures of service sector output, such as those which extrapolated inputs, eliminating productivity growth by construction. Since his time, however, there have been vast improvements in the national income accounts measures of service sector activity, particularly in regards to the recent time period (1987-2010) which is the focus of this paper's analysis. Triplett and Bosworth (2004) provide a review of these developments and the problems which remain. This paper takes as given the official measures of sectoral output, focusing on the systematic bias brought about by the failure to consider the relation between employment shares and average worker efficacy.

The paper proceeds as follows: I begin in Section II by presenting a simple Roy model, showing how the bias in sectoral measures of total factor productivity growth and the slope of the relative supply curve depend upon a key parameter: the elasticity of average worker efficacy within a sector with respect to that sector's share of total employment. Section II also shows how correlation between an individual's productivity in different activities can eliminate the positive association between comparative advantage and absolute advantage, overturning Roy's prediction that average worker efficacy is inversely related to a sector's employment share. Thus, the relation between worker efficacy and sectoral employment depends upon the process generating individual productivity draws, i.e. it is ultimately something that needs to be estimated empirically rather than identified theoretically.

Section III presents industry level evidence that the elasticity of worker efficacy with respect to sectoral employment is, indeed, substantially negative. Projecting the Bureau of Labour Statistics KLEMS[7] measures for the United States private sector divided into 60 sectors, and the University of Groningen's KLEMS measures for private sector activity in 18 OECD countries divided into 29 sectors, on a variety of instruments, I find that defense spending is the only instrument that robustly satisfies the dual requirements of 1st stage significance and 2nd stage exogeneity (the exclusion restriction) necessary for two stage least squares (2SLS). Estimates of the long run elasticity of worker efficacy with respect to the sectoral employment share range from -.5 to -1, with most observations concentrated in the more negative half of this range. I also find that an elasticity of -.75 equalizes goods and services productivity growth in the US and the OECD at large. It produces a stable Roy supply curve which matches the historical US and OECD data on relative goods and services price and quantity growth, as heuristically illustrated in Figure I above.

Section V concludes the body of the paper. Appendix A provides mathematical proofs of all of the claims made in Section II. While the BLS adjusts its aggregate economy-wide measures of labour input growth for compositional effects, it does not do this in the sectoral KLEMS data base. Appendix B describes how I develop detailed sectoral measures of labour composition which I use to adjust the BLS measures of total factor productivity growth and the

[7] Capital (K), labor (L), energy (E), materials (M) and purchased service inputs (S).

sectoral measures of changing employment shares. Finally, Appendix C provides a review of the PSID data used in the McLaughlin and Bils paper mentioned above, showing that it concerns simultaneous entry and exit, rather than the overall expansion and contraction of sectoral employment, which is the focus of Roy's model and this paper.

II. Structural Transformation and the Mismeasurement of Productivity

In this section I present the main theoretical results of the paper. Throughout the analysis I focus on supply relations alone, leaving the general equilibrium closure of the model with preferences and demand unspecified. This is both because I do not want to take a stand on the nature of preferences and demand (including trade), and because it is unnecessary to do so. All of the implications of the Roy model can be understood in terms of the supply curve and all of the theoretical analysis can be understood in terms of movements along that curve, movements whose causes, while obviously related to demand, do not need to be specified. To focus on intuition, I confine the mathematical proofs of the claims made in this section to Appendix A.

(a) A Simple Model

Consider an economy with two perfectly competitive industries, goods (G) and services (S). Value added in industry i (= G or S) is produced with capital and labour:

$$(1)\quad Q_i = A_i F_i\left(K_i,\ \int_{u\in Set_i} z_i(u)\right)$$

where Set_i is the set of workers u labouring in industry i and $z_i(u)$ is the efficacy or productivity of individual u when working in industry i. Each worker is endowed with a pair of industry productivities (z_G, z_S) which is drawn from some joint cumulative distribution function $G(z_G, z_S)$.

Workers move to the industry providing the highest financial reward. Thus, with w_i

denoting the wage per unit of effective labour offered in industry i, the set of individuals choosing to work in that sector is given by:

$$(2)\quad Set_i = \left\{ u \mid w_i z_i(u) > w_j z_j(u) \right\}$$

where j is the sectoral complement of i. Define π_i as the probability a worker selects industry i or, equivalently, the share of the labour force in industry i. With L denoting the total labour force, L_i, the number of workers in industry i, equals $\pi_i L$. For a given distribution of (z_G, z_S) draws, π_i is determined in a general equilibrium that includes a specification of demand, with $d\pi_i/d\omega > 0$, where $\omega = w_i/w_j$.

Define the expected efficacy of a worker in sector i, i.e. their productivity conditional on working in that sector, as

$$(3)\quad \bar{z}_i = E\left(z_i(u) \mid u \in Set_i\right) = \frac{\int_{u \in Set_i} z_i(u)\,\mathrm{du}}{\int_{u \in Set_i} \mathrm{du}} = \frac{\int_{u \in Set_i} z_i(u)\,\mathrm{du}}{L * \pi_i}$$

As proven in Appendix A, regardless of the distribution function generating the paired draws (z_G, z_S) , the elasticity of expected worker efficacy with respect to the sectoral employment share is greater than -1:

$$(4)\quad \xi = \frac{d\bar{z}_i}{d\pi_i}\frac{\pi_i}{\bar{z}_i} > -1$$

From (3), we see that if we ignore the numerator the elasticity of $\bar{z}_i$ with respect to π_i is -1. The numerator, however, is increasing in π_i, as anything that increases the total number of workers will increase the cumulative sum of their productivities. Consequently, the overall elasticity of $\bar{z}_i$ with respect to π_i will be greater than -1 (examples for particular functional forms are provided in Appendix A). None of the empirical estimates presented later in Section III rejects this prediction. While ξ may be positive or negative, Roy (1951), as explained in the Introduction, argued that it should be negative, i.e. average worker efficacy declines as a sector expands and draws in less productive workers. For the moment, I will assume this to be true.

Aggregate labour input in an industry is a product of the number of workers times the average efficacy per worker, so the production function is usefully reexpressed as:

(5) $Q_i = A_i F_i\left(K_i, L_i \bar{z}_i\right)$

From this, we see that total factor productivity growth, properly calculated, is given by[8]

(6) $\hat{A}_i(true) = \hat{Q}_i - \Theta_{Ki}\hat{K}_i - \Theta_{Li}(\hat{L}_i + \hat{\bar{z}}_i)$

where a ^ denotes a proportional change and Θ_{Ki} and Θ_{Li} are the factor income shares of capital and labour in sector i, respectively. Unfortunately, in estimating total factor productivity growth growth accountants treat each new worker as the equivalent of existing workers,[9] estimating total factor productivity growth to be

(7) $\hat{A}_i(est) = \hat{Q}_i - \Theta_{Ki}\hat{K}_i - \Theta_{Li}\hat{L}_i \;=\; \hat{A}_i(true) + \Theta_{Li}\hat{\bar{z}}_i \;=\; \hat{A}_i(true) + \xi\,\Theta_{Li}\hat{\pi}_i$

If average worker efficacy depends inversely on a sector's share of the labour force ($\xi < 0$), growth accountants will systematically overestimate productivity growth in sectors whose labour share is contracting, such as goods industries, and systematically underestimate it in sectors whose labour share is expanding, such as services.

In a world in which heterogeneous workers choose their occupation based upon their unobserved productivity in different tasks there may also be biases in conventional measures of aggregate economy-wide total factor productivity growth. Aggregate productivity growth equals

[8]The derivation is the usual one for total factor productivity calculations. With perfect competition the capital rental and wage per unit of effective labour equal the value marginal product of each factor, so the elasticity of the production function with respect to each factor equals the factor income share.

[9]A more refined practice is to differentiate workers into types based upon observable characteristics such as age and education. Within each type, however, marginal workers are treated as identical to average workers, producing the same problem, as I show when I extend the model further below.

the sum of sectoral productivity growths weighted by their shares of aggregate value added:

$$(8)\quad \hat{A}(true) = \sum_i \Omega_i \hat{A}_i(true) \quad \text{where} \quad \Omega_i = {P_i Q_i}\Big/_{GDP} \quad \text{and} \quad GDP = \sum_i P_i Q_i$$

As proven in Appendix A, regardless of the characteristics of the joint distribution from which the paired individual productivities (z_G, z_S) are drawn or the elasticity of average worker efficacy with respect to the employment share (ξ), the growth accountant's estimate of aggregate productivity is related to actual productivity growth by the formula[10]

$$(9)\quad \hat{A}(est) = \sum_i \Omega_i \hat{A}_i(est) = \hat{A}(true) + \sum_i \frac{L}{GDP} W_i \bar{z}_i d\pi_i = \hat{A}(true) + \sum_i \Omega_i \Theta_{Li} \hat{\pi}_i$$

where $d\pi_i$ is the change in sector i's share of total employment and $W_i = w_i \bar{z}_i$ equals the average earnings of a worker in sector i.

While the last equality in (9) expresses the bias in terms of the traditional proportional changes of growth accounting, the next to last equality, with its emphasis on average earnings, is more illuminating. Within the standard growth accounting framework, which treats all workers within a sector as identical, any difference in average sectoral wages is an indication of inefficiency, an inequality in the value marginal product of labour. Consequently, growth accountants implicitly treat the movement of workers from sectors with low average wages to those with high average wages as an improvement in aggregate efficiency, a "gain from reallocation". In a Roy model the marginal worker *u* finds that his value marginal product is equalized across sectors, $w_i z_i(u) = w_j z_j(u)$, but there is nothing forcing the average value marginal product of workers in a sector ($W_i = w_i \bar{z}_i$) to be equalized across industries. Average marginal products reflect characteristics of the infra-marginal distribution of heterogenous talent within

[10]From (7) and (8) one can see that the aggregate bias involves some weighted sum of changes in sectoral average worker efficacies. The proof in the appendix shows how these can be transformed into a wage weighted sum of changes in employment shares.

industries. They may equalize across sectors, as happens to be the case when the z_i are draws from fréchet distributions, or they may not, as is the case for draws from exponential distributions (see Appendix A), but gaps in these average products neither reflect the degree of inefficiency nor potential gains from reallocation. As it so happens, in US and OECD data of the past few decades the average wages of goods and service workers are quite similar (see Section III), so I find little aggregate bias in estimates of total factor productivity growth for these countries, but these issues could be quite relevant in the study of productivity growth and structural change in developing countries where, for example, there are often large differences between agricultural and non-agricultural wages.

Finally, I turn to the derivation of the shape of the relative supply curve. I make two empirical assumptions which, although not universal characteristics of the model, approximately characterize the US and OECD economies (see end of Section III below): (1) average wages per worker are proportional across sectors; and (2) factor income shares are the same in the two sectors. Mathematically, these amount to:

$$(10)\quad W_G = w_G \bar{z}_G \propto w_S \bar{z}_S = W_S \quad \text{and} \quad \frac{rK_G}{W_G L_G} = \frac{rK_S}{W_S L_S}$$

$$\text{so} \quad \hat{w}_G - \hat{w}_S = \hat{\bar{z}}_S - \hat{\bar{z}}_G \quad \text{and} \quad \hat{K}_G - \hat{L}_G = \hat{K}_S - \hat{L}_S$$

where r is the common rental per unit of capital. As $Q_i = A_i F_i(K_i, L_i \bar{z}_i)$ and $L_i = L\pi_i$, we have

$$(11)\quad \hat{Q}_G - \hat{Q}_S = \hat{A}_G - \hat{A}_S + \Theta_K(\hat{K}_G - \hat{K}_S) + \Theta_L(\hat{L}_G + \hat{\bar{z}}_G - \hat{L}_S - \hat{\bar{z}}_S)$$

$$= \hat{A}_G - \hat{A}_S + (\hat{\pi}_G - \hat{\pi}_S) + \Theta_L(\hat{\bar{z}}_G - \hat{\bar{z}}_S) = \hat{A}_G - \hat{A}_S + (1/\xi + \Theta_L)(\hat{\bar{z}}_G - \hat{\bar{z}}_S)$$

From the dual measure of total factor productivity growth $\hat{A}_i = \Theta_K \hat{r} + \Theta_L \hat{w}_i - \hat{P}_i$,[11] so

$$(12)\quad \hat{P}_S - \hat{P}_G = \Theta_L(\hat{w}_S - \hat{w}_G) + (\hat{A}_G - \hat{A}_S) = \Theta_L(\hat{\bar{z}}_G - \hat{\bar{z}}_S) + (\hat{A}_G - \hat{A}_S)$$

[11]Totally differentiating $P_i Q_i = rK_i + w_i L_i \bar{z}_i$: $\hat{P}_i + \hat{Q}_i = \Theta_K(\hat{r} + \hat{K}_i) + \Theta_L(\hat{w}_i + \hat{L}_i + \hat{\bar{z}}_i)$. Substituting for $\hat{Q}_i$ gives the equation in the text.

Finally, substituting for $\hat{\hat{z}}_G - \hat{\hat{z}}_S$ using (11), we derive the shape of the Roy supply curve:

$$(13)\quad \hat{P}_S - \hat{P}_G = \left(\frac{-\Theta_L \xi}{1+\Theta_L \xi}\right)\left(\hat{Q}_S - \hat{Q}_G\right) + \left(\frac{1}{1+\Theta_L \xi}\right)\left(\hat{A}_G - \hat{A}_S\right) \quad [\text{Roy}]$$

The first term on the right-hand side of (13) gives the slope of the supply curve; the second term gives the vertical shift associated with a change in relative total factor productivities. For $0 > \xi > -1$, the supply curve is upward sloping, as drawn in Figure I of the Introduction. In the special case where $\xi = 0$ and average worker productivity does not vary with the sectoral employment share, labour is, for all intents and purposes, homogenous and the supply curve reduces to:

$$(14)\quad \hat{P}_S - \hat{P}_G = \hat{A}_G - \hat{A}_S \quad [\text{Baumol}]$$

With P_S/P_G independent of Q_S/Q_G, this is, of course, Baumol's horizontal relative supply curve.

Equation (13) highlights the fact that, in the absence of differences in productivity growth rates, there is a limit to the relative price growth that can be explained by Roy's model of self selection. With the labour share of 2/3 observed in the US and OECD economies, as ξ goes from 0 to -1 the slope parameter $-\Theta_L\xi/(1-\Theta_L\xi)$ goes from 0 to 2. Thus, the Roy supply curve can be no steeper than 2, i.e. the historical growth of the relative output of services to goods has to be at least ½ the historical growth of the relative price if one wants to eliminate Baumol type effects from the story. As it so happens, the historical growth rates of relative goods and services outputs and prices in the US and the OECD at large appear to be about equal (see Section III), which can be explained, in the absence of any differences in productivity growth, with a ξ of -.75. This value is comfortably within the range of long run estimates using defense spending as an instrument reported later in Section III.

(b) Comparative and Absolute Advantage and the Sign of ξ

In Appendix A I prove that sufficient conditions for ξ, the elasticity of average worker efficacy with respect to a sector's share of total employment, to be less than zero are that (a) the sectoral productivity draws z_i are independent of each other; and (b) the elasticity of the cumulative distribution function for each of the draws, $(dG/dz)*(z/G)$, is decreasing in the productivity of the draw. The latter characteristic is true of all of the popular distribution functions defined on non-negative numbers, i.e. the chi-squared, exponential, F, fréchet, gamma, lognormal, pareto, rayleigh, uniform and weibull distributions,[12] so I relegate a discussion of its role to the appendix. The assumption of independence is more problematic, so I explore its role here with a simple example and diagram.

Consider a two sector example where the draw for sector i is deterministically related to that of sector j by the equation $z_i = z_j^{\eta}$, with z_j drawn from any distribution function. Workers will select sector i if $w_i z_i > w_j z_j$ or, equivalently, $w_i/w_j > z_j^{1-\eta}$. Figure III illustrates how the characteristics of the resulting equilibrium vary with η. Panel (a) considers the case where $\eta<0$, i.e. the productivity draws are negatively correlated. The upper quadrant of the diagram shows that there exists a marginal draw z_j^* such that all workers with draws greater than z_j^* work in sector j and all workers with draws less than z_j^* work in sector i. The productivity of workers in sector i is illustrated in the lower quadrant, where the axis, despite its location below the horizontal line, should be read as representing positive numbers. With $\eta<0$, the productivity of workers in sector i is negatively related to the z_j draws. $\bar{z}_j$ is given by the average of the workers to the right of z_j^*, while $\bar{z}_i$ is given by the average of the workers below (i.e. south of) $z_i^* = z_j^{*\eta}$.

[12]While this condition may be true for all of the well known distributions, I should note that it isn't hard to think of distribution functions where it is not. Thus, the distribution function G(z) = (exp(z)-1)/(exp(1)-1) defined on [0,1] violates the condition and, in a simple two sector example, produces regions where the average productivity of workers in a sector is rising in the sector's share of total employment. I should also note that for the uniform distribution defined on [a,b], for a > 0 the elasticity of the cumulative distribution is strictly decreasing in z but for a = 0 it is constant and a weaker form of the theorem applies (ξ is non-positive).

Figure III: Correlated draws, $z_i=z_j^{\eta}$

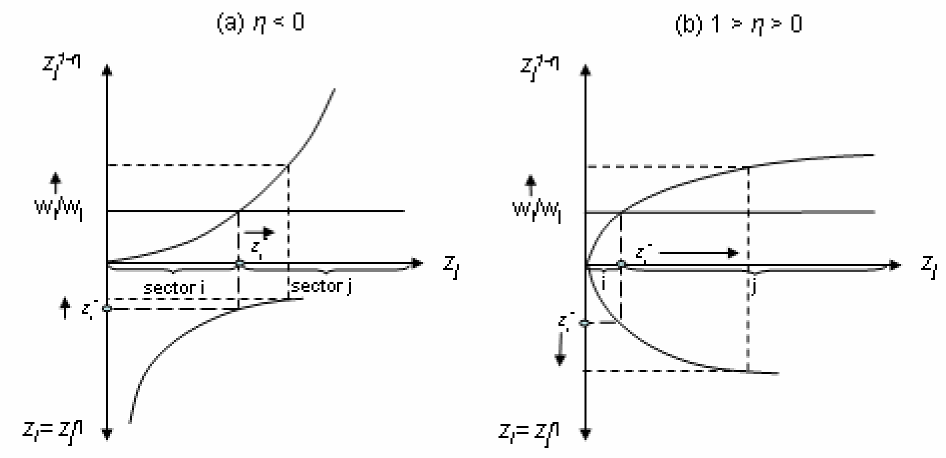

As w_i/w_j increases, sector j sheds workers with less than the average productivity in that industry, while sector i gains workers with less than the average productivity in that sector. Average productivity rises in the sector losing workers and falls in the sector gaining workers, so $\xi < 0$.

Turning to panel (b) of the figure, we consider the case where the draws are positively correlated, so $1 > \eta > 0$.[13] With a positive relationship between z_i and z_j, $\bar{z}_j$ is once again given by the average of the workers to the right of z_j^*, but $\bar{z}_i$ now equals the average of the workers *above* z_i^*. As w_i/w_j rises, industry j sheds workers with less than the average productivity in that sector, but industry i gains workers with more than its average sectoral productivity. ξ is still negative for sector j, but it is now positive for sector i.

Returning to the trade terminology used in the Introduction of this paper, if there is a positive correlation between comparative advantage and absolute advantage, then marginal workers entering or exiting an industry will have less than the average sectoral productivity. If, however, the correlation between comparative and absolute advantage is negative, marginal workers will have more than the average productivity. In panel (a) of Figure III, workers who

[13]For the case $\eta > 1$, rearrange $z_i = z_j^{\eta}$ as $z_j = z_i^{1/\eta}$, rename i as j and j as i and proceed with diagram (b).

choose to work in industry i or j (a consequence of comparative advantage) are absolutely more productive in that sector than workers who choose to work in the other sector, so comparative advantage is positively correlated with absolute advantage. In panel (b), this is true for sector j, but it is no longer true for sector i. In the case of sector i, workers who choose to work in the industry (those with z_i lying north of z_i^* on the vertical axis) are absolutely less productive in that sector than those who choose to work elsewhere (those with z_i lying south of z_i^* on the vertical axis), so comparative advantage is negatively correlated with absolute advantage.

Roy argued that if a worker's productivities in different sectors are independent of each other, then the marginal worker entering or exiting an industry will be less efficient than the average worker in that sector. The theorem described above and proven in Appendix A shows that, modulo a technical density condition, Roy's conjecture is true. Figure III shows that positive correlation between an individual's productivity in different sectors undermines the association between comparative and absolute advantage, producing an indeterminate association between average and marginal productivities. In constructing total factor productivity growth estimates, as discussed shortly below, the growth accountant typically adjusts for observables such as age and education that create positive correlations in individual productivity across industries and tasks. These adjustments are, however, by no means exhaustive and it remains an empirical question whether or not comparative advantage is positively or negatively correlated with absolute advantage. The empirical results of the next section, interpreted in the light of the Roy model, provide some evidence in favour of the view that the elasticity of average worker efficacy with respect to a sector's employment share is negative, i.e. that by and large comparative and absolute advantage are indeed positively correlated.

(c) Practical Extensions

A modest amount of notational and algebraic complexity must be added to the model to bring it to the data. To this end, imagine that there are N sectors with gross output in each sector i a function of J types of labour input and M types of other inputs:

$$(15)\quad Q_i = A_i F_i\left(\int\limits_{u\in Set_i^1} z_i^1(u)\,,\ \int\limits_{u\in Set_i^2} z_i^2(u)\,,\ \ldots\ ,\ \int\limits_{u\in Set_i^J} z_i^J(u)\,,\ M_i^1,\ M_i^2,\ \ldots\ ,M_i^M \right)$$

where I now use superscripts to denote the type of input and subscripts the industry. The switch from value added to gross output reflects the fact that my data sources, the BLS and Groningen KLEMS, measure total factor productivity growth at the sectoral level, using the gross output concept, so the list of M additional inputs moves beyond capital and includes intermediate inputs such as materials, services and energy. Good estimates of total factor productivity growth typically adjust for "labour quality" by decomposing labour into mutually exclusive categories based upon observable determinants of human capital such as sex, age and education. This decomposition not only produces more accurate measures of total factor productivity growth, it also implicitly controls for factors that produce a positive correlation in individual productivity across tasks, as noted above.

While the Groningen KLEMS adjust for labour quality, the BLS KLEMS measures do not adjust for labour quality, using only total labour hours as the measure of labour input. Using Current Population Survey data, I have constructed measures of labour input for each of the 60 KLEMS sectors cross-classified by sex, age (6 categories) and education (5 categories). I follow a methodology very similar to that used by the BLS in producing its measures of labour quality for the aggregate economy, using the CPS data to determine the distribution of workers by characteristic, but benchmarking the sectoral totals of hours and workers using the BLS Current Employment Statistics data. Details are provided in Appendix B. I use these estimates to adjust the BLS TFP growth measures for the changing composition of the workforce and to calculate the changing shares of workers by characteristic, as in (17) below. The main results, however, can just as easily be found with the unadjusted BLS data, as reported in footnotes later.

To extend the model to this environment, let each worker of type j be endowed with a set of N industry productivities $(z_1^j, z_2^j, \dots, z_N^j)$ drawn from some joint distribution function and let w_i^j denote the wage per unit of effective labour of type j in industry i. A worker chooses to work in sector i if $w_i^j z_i^j(u) > w_k^j z_k^j(u) \;\; \forall \;\; k \neq i$. Total factor productivity growth in each sector is given by

$$(16) \quad \hat{A}_i(true) = \hat{Q}_i - \sum_j \Theta_{Li}^j (\hat{L}_i^j + \hat{\bar{z}}_i^j) - \sum_m \Theta_{Mi}^m \hat{M}_i^m$$

where L_i^j is the number of workers of type j employed in sector i, $\bar{z}_i^j$ is their average efficacy, and the Θ_{Li}^j and Θ_{Mi}^m represent the factor income shares of workers of type j and other inputs of type m in sector i, respectively. Conventional measures of total factor productivity growth, by ignoring changes in the average efficacy of workers, have a bias equal to:

$$(17) \quad \hat{A}_i(est) = \hat{Q}_i - \sum_j \Theta_{Li}^j \hat{L}_i^j - \sum_m \Theta_{Mi}^m \hat{M}_i^m = \hat{A}_i(true) + \sum_j \Theta_{Li}^j \hat{\bar{z}}_i^j$$

$$= \hat{A}_i(true) + \xi \sum_j \Theta_{Li}^j \hat{\pi}_i^j$$

Growth accounting calculations intrinsically assume that all workers of a given type are the same. Unless the list of observable worker characteristics completely exhausts the determinants of individual productivity, the productivity of the marginal worker entering or exiting an industry will generally be different than that of the sectoral average for that type of worker. If the elasticity of average worker efficacy with respect to the employment share is negative (ξ <0),[14] conventional growth accounting will under or overstate productivity growth in sectors with

[14](17) assumes that ξ is the same for all sectors and types at all times. This is precisely true for some distribution functions (e.g. independent draws from Fréchet distributions with the same dispersion parameter). Otherwise, one must take ξ as an average of the differing elasticities.

expanding or contracting employment shares, respectively.

Finally, I note that for productivity measures based upon gross output the appropriate measure of economy-wide productivity growth and the bias in conventional estimates are once again given by equations (8) and (9) earlier, except that, with Q_i now denoting gross-output rather than value added, the sectoral weights $\Omega_i = P_iQ_i/GDP$ are the "Domar weights" (Domar 1961, Hulten 1978) and sum to more than one in aggregate. Because sectoral output is also used as an input in other industries, rather than simply a component of final demand, each sector's productivity growth has a multiplied effect on aggregate productivity. We can define aggregates for goods and services, separately, by using the GDP originating in each sector as the denominator:

$$(18)\quad \hat{A}_j = \sum_{i\in I(j)} \Omega_i \hat{A}_i \quad \text{where} \quad \Omega_i = {P_iQ_i}/{GDP_j}, \ \text{and} \quad \hat{A} = \hat{A}_G \frac{GDP_G}{GDP} + \hat{A}_S \frac{GDP_S}{GDP}$$

and where *I(j)* is the set of industries in sector j = G or S and, as shown, aggregate productivity growth is the GDP share weighted average of goods and services productivity growth. I use this formula in my calculation of total goods and services productivity growth in Section III.[15]

III. Industry Evidence on the Elasticity of Worker Efficacy with Respect to Employment Shares

(a) Empirical Specification

I use the following two stage least squares specification to explore the bias in sectoral measures of total factor productivity growth brought about by changing labour allocations:

[15]Since TFP growth measured using the value added approach equals the ratio of sectoral gross output to sectoral value added times TFP growth measured using the gross output approach, (18) transforms the gross output TFP measures into a value added share weighted sum of sectoral value added TFP measures. For the aggregate economy, Domar's method aggregates sectoral gross output TFP measures into the implied aggregate value added TFP measure. (18) is a simple extension of his method to allow the calculation of sub-total value added measures.

$$(19)\quad \hat{Y}_{ict} = \alpha_{ic} + \delta_{ct} + \gamma_{ic}\hat{U}_{ct} + \xi\hat{X}_{ict} + \varepsilon_{ict}$$

$$\hat{X}_{ict} = \alpha_{ic}^{X} + \delta_{ct}^{X} + \gamma_{ic}^{X}\hat{U}_{ct} + \beta_{ic}\hat{Z}_{ct} + \eta_{ict} \qquad E(\varepsilon_{ict}\eta_{ict}) \neq 0$$

where $\hat{Y}_{ict}$ is total factor productivity growth in industry i of country c in period t, the α_{ic} are industry x country dummies capturing mean productivity growth by sector and the δ_{ct} are country x year dummies capturing economy-wide fluctuations in average productivity growth. There is a well known association between the business cycle and measured productivity growth, driven perhaps by mismeasurement due to changes in capacity utilization and the role real technology shocks play in producing the business cycle. While the country x year dummies account for mean economy-wide changes, the ln change in the national unemployment rate $\hat{U}_{ct}$, entered separately by industry x country (γ_{ic} is an industry x country effect), corrects for the cyclical variation in relative industry productivity growth that might otherwise appear as correlation with other variables. Finally, $\hat{X}_{ict}$ equals the labour income share weighted sum of the change in national employment shares by worker type, as shown in the right-hand side of (17) earlier. The coefficient ξ, by the theory described earlier above, is the elasticity of worker efficacy with respect to employment shares, the principal object of interest in the regression.[16]

The OLS relation between productivity and employment shares potentially has both exogenous and endogenous components. On the one hand, movements in relative industry demand, due to the growth of aggregate income and non-homothetic preferences, will lead to exogenous changes in relative employment shares. On the other hand, the response of relative demand to relative price movements brought about by productivity growth may lead to an endogenous response of employment shares to productivity growth. There are special cases where these effects disappear, such as with homothetic utility and unitary income elasticities of

[16] This specification estimates a single ξ, but it is compatible with a world in which ξ varies, say, by industry (ξ_i). As long as the variation is independent of the right hand side variables, there is no bias and the estimated value of ξ can be interpreted as the mean value of the ξ_i Estimating ξ by industry is not sensible, as the resulting sample sizes are tiny (e.g. 20+ observations per industry in the US), while the properties of 2SLS rely on asymptotics.

demand (no exogenous variation of relative demand) or with Hicks-Neutral technical change and unitary price elasticities of demand (no endogenous variation of factor allocations with sectoral productivity growth), but it seems reasonable to allow for the existence of both in the data.[17] As shown in the second line of (19), to correct for potential endogeneity I run a first stage regression in which the labour income weighted changes in sectoral employment shares are regressed on the exogenous variables of the total factor productivity equation plus an excluded instrument. The relation of the excluded instrument with $\hat{X}_{ict}$ is allowed to vary across industries and countries (β_{ic} varies by industry x country). Variation by industry is necessary, as for an instrument to influence employment shares it must raise employment in some industries at the expense of others, and variation by countries allows for differences in the composition of otherwise nominally "identical" sectoral aggregates. Because the instrument is interacted by industry x country, i.e. appears multiple times in the regression, it is possible to perform a valid overidentification test of the exclusion restriction, even though only "one" instrument appears in the regression.[18]

I draw on two datasets which provide comprehensive measures of private sector total factor productivity broken down by sector ($\hat{Y}_{ict}$ above). First, I use data on total factor productivity growth by sector drawn from the Bureau of Labor Statistics' KLEMS (capital, labour, energy, materials and business services) database, which provides estimates of US private

[17]Ngai-Pissarides (2007) provide an analysis of the case with homothetic utility, Hicks-Neutral technical change and inelastic demand, where all of the relation between labour allocations and productivity is endogenous. Homothetic utility, however, provides a poor characterization of demand, as it implies that relative quantities fall with relative prices whereas, as discussed below, the overwhelming trend in the OECD is for relative quantity to rise with relative price (reflecting non-unitary income elasticities). Hicks-Neutral technical change misses interesting interactions between factor biased technical change and the elasticity of substitution. For example, Bustos, Caprettini & Ponticelli (2013) show that despite an infinite elasticity of demand (free trade), labour augmenting technical change in the presence of a low elasticity of factor substitution can actually lead to a reduction in sectoral employment.

[18]Lest the reader think there is an error here, I confirm the exactness of the overidentification test using simulated data that satisfy the exclusion restriction, as discussed further below.

sector productivity growth disaggregated into 60 comprehensive industries from 1987 to 2010. As noted earlier, these data do not adjust for the changing composition of the labour force, so I use Current Population Survey data to develop industry level measures of the distribution of workers by sex x age x education and use these to adjust the total factor productivity growth and calculate a compositionally adjusted measure of changing labour shares, as described in Appendix B.[19] Second, I use the EU KLEMS database, developed by the University of Groningen with a consortium of diverse partners, which divides private sector productivity growth in a variety of advanced economies into 29 comprehensive sectors.[20] After removing transition economies, where productivity growth and factor allocations are likely to be driven by considerations outside this paper, the sample consists of 18 countries, namely: Australia, Austria, Belgium, Canada, Denmark, Finland, France, Ireland, Italy, Japan, Korea, Luxembourg, Netherlands, Portugal, Spain, Sweden, the United Kingdom, and the United States. The productivity estimates run from 1970 to 2005, with the available years varying by country. I shall refer to these data as the OECD or OECD 18 sample, notwithstanding their development in the European Union. Measures of annual unemployment for the US and the OECD countries are drawn from the Federal Reserve St. Louis FRED database.

Turning to potential instruments, I consider simple measures of my own alongside the

[19]The adjusted and unadjusted industry measures of total factor productivity growth are available on my website. My calculations indicate that adjustments for the changing sex x age x education composition of the labour force lower economy-wide private sector total factor productivity growth between 1987 and 2010 from an average of .0125 per annum to .0097 per annum.

[20]There are actually 31 private sectors, but two ("private households with employed persons" and "extra-territorial organizations") are relatively minor and do not appear in all instances. Employment shares are always calculated relative to national totals (including the public sector). Although the EU KLEMS TFP calculations adjust for the composition of the workforce, the data provided only allow for the calculation of the distribution of total workers by sector (not workers by type), so I use the labour income share times the change in the total employment share as the X variable, as in eqn (7) above. (7) and (17) are identical if the distribution of workers by type is proportional to the industry share of total employment, i.e. $L_i^{\,j} = L^j(L_i/L)$. For the US KLEMS, I find that substituting the changing shares of total employment for the changing shares of employment by worker type yields virtually identical results, as reported in a footnote below.

more sophisticated constructions of others. Using FRED, Stockholm International Peace Research Institute (SIPRI) and World Bank data, the instruments I prepare are: (1) the ln change in country defense expenditures over GDP; (2) the average ln change in metal prices (aluminum, copper, iron ore, lead, nickel, platinum, tin and zinc); and (3) the average ln change in oil prices (Dubai and West Texas Intermediate). Changes in defense expenditures, driven by events such as the collapse of the Soviet Union and 9/11, are arguably exogenous to sectoral productivity growth. There is less reason to feel confident in the exogeneity of metals and oil prices. Productivity change in key producing or using industries in the US and the OECD countries, which are large actors in the global markets for these materials, might produce endogenous responses in prices. While US defense spending and materials and oil prices are available for all years of my TFP data, because of changes in concepts and coverage, the SIPRI data on OECD country military expenditures only extend back to 1988.[21]

I expand the list of potential instruments by adding all 15 of the non-technology shock instruments considered by Stock and Watson (2012) in their dynamic factor model analysis of the US economy. Covering oil prices, monetary policy, uncertainty, liquidity and fiscal policy, these are:[22] (1) Hamilton's (2003) measure of the increase of the oil price PPI relative to the max of the previous 3 years, available for 1962-2010; (2) Kilian's (2008) measure of OPEC production shortfall from wars and civil strife, available for 1971-2004; (3) the residuals of Ramey & Vine's (2010) measure of full gasoline prices regressed on lagged macroeconomic variables, based on

[21]The SIPRI website notes that SIPRI has not been able to construct a consistent series extending back to earlier dates, and the SIPRI data has now become the standard, reproduced in other on-line sources (such as the World Bank) to the exclusion of any other information. I tried to construct an alternative series of my own using historical paper issues of The Military Balance, but ultimately concluded that SIPRI's concerns about coverage and data quality are correct.

[22]In most cases I use the data provided on-line by Stock and Watson and follow their procedures (e.g. AR(2)s, regressions on lagged macro variables, etc) to construct the instruments. The dataset, however, contains a major misreporting of the Ramey-Vine figures (formulas rather than values were copied into the Stock & Watson spreadsheet), so I use the updated data from Valerie Ramey's website.

their updated spreadsheet (available 1959-2011); (4) Romer and Romer's (2004) residual of Fed monetary intentions regressed on internal Fed forecasts (1969-1996); (5) Smets and Wouters' (2007), updated by King and Watson (2012), measure of the shock to the monetary policy reaction function in a dynamic stochastic general equilibrium model (1959-2004); (6) Sims and Zha's (2006) monetary policy shock estimated in a structural VAR (1960-2002); (7) Gürkaynak, Sack and Swanson's (2005) measure of surprise changes in the federal funds rate (1990-2004); (8) innovations in an AR(2) of the VIX, as suggested by Bloom (2009) (1962-2011); (9) innovations in an AR(2) of Baker, Bloom and Davis's (2012) policy uncertainty index calculated from media references to economic policy (1985-2011); (10) innovations in an AR(2) of the TED spread, as provided by Stock & Watson (1971-2011); (11) innovations in an AR(2) of Gilchrist-Zakrajšek's (2012) bond premium (1973-2010); (12) Bassett et al's (2011) measure of unpredictable changes in bank-level lending standards (1992-2010); (13) Ramey's (2011) measure of news of changes in the net present value of military spending divided by nominal GDP (1959-2010); (14) Fisher and Peters' (2010) measure of excess returns on stocks of military contractors (1959-2008); and (15) Romer and Romer's (2010) measure of tax changes relative to GDP (1959-2007). I average quarterly or monthly shocks to annual levels.

With the exception of Kilian's oil production shortfall, the Stock & Watson instruments listed above are US-centered and not appropriate for an OECD analysis. However, as shown in the pages below, none of these instruments performs at all well in the analysis of the US KLEMS. Hence, undertaking the monumental task of developing similar instruments country by country is not likely to be profitable. In fact, the only instrument that consistently satisfies the first stage requirement of significance and the second stage exclusion restriction is defense spending. Thus, my main point in using Stock & Watson's extensive list is to highlight the difficulty of finding alternative instruments for sectoral labor allocations.

(b) Results

I begin by evaluating the suitability of the various instruments to the problem at hand. In Table I below I run the 1st stage regression of the specification of equation (19) using one instrument at a time, reporting the p-value of the F test on the instrument[23] and the total number of observations. In the case of the OECD, I only use my instruments and Kilian's oil production shortfall, which can be considered part of global trends. There are two notable aspects of Table I. First, virtually all of the factors considered by Stock and Watson (instruments d through r) are not meaningful determinants of labour allocations. Only the oil price max measure and Federal Funds surprises are significant at the 5% level, and these results are suspect as other measures of oil prices and monetary policy are quite insignificant. Second, in the OECD sample none of the instruments are even close to being significant.

Table I's results are perhaps not terribly surprising. To generate a significant reallocation of labour across sectors, an instrument must not merely shift macroeconomic supply and demand, it must substantially alter relative industry supply or demand away from the norm. Many shocks which have strong aggregate macroeconomic consequences and serve as good instruments for the analysis of macro aggregates might not have sufficiently strong relative effects for the objective of this paper. In this regard it is noteworthy that Ramey's (2011) measure of news of changes in the NPV of military spending is insignificant. Ramey argues that, in explaining changes in macroeconomic aggregates in the United States, her news variable dominates actual defense spending changes. The macroeconomic influence of Ramey's news variable, however, most likely represents the response of private economic actors to the foreseen aggregate consequences (e.g. on demand and tax burdens) of that spending. Continuity of private demand suggests that

[23] Although in each case there is only one instrument, its coefficient is allowed to vary by industry x country, hence an F-test rather than a t-statistic.

Table I: 1st Stage P-value in Regression of Weighted Employment Share Changes on Instruments (instruments evaluated one at a time using specification of eqn 19)				
	United States 60 sectors 1987-2010		OECD 18 29 sectors 1970-2005	
	F p-v.	N	F p-v.	N
(a) Δ ln Country Defense Expenditures/GDP	.000	1380	.192	8049
(b) Δ ln Metals Prices	.000	1380	.374	12109
(c) Δ ln Oil Prices	.833	1380	.367	12109
(d) Oil Price Increase Over Prior Maximum (Hamilton 2003)	.005	1380		
(e) OPEC Oil Production Shortfall (Kilian 2008)	.253	1020	.762	11617
(f) Residual of US Gasoline Prices (Ramey & Vine 2010)	.965	1380		
(g) Monetary Policy Shock (Romer & Romer 2004)	.866	540		
(h) Monetary Policy Reaction Shock (Smets & Wouters 2007)	.084	1020		
(i) Monetary Policy Shock (Sims & Zha 2006)	.884	900		
(j) Fed. Funds Surprises (Gürkaynak et al 2005)	.000	900		
(k) VIX Innovation (Bloom 2009)	.863	1380		
(l) Policy Uncertainty Index Innovation (Baker et al 2012)	.092	1380		
(m) TED Spread Innovation (Stock & Watson 2012)	1.00	1380		
(n) Bond Premium Innovation (Gilchrist & Kayrajšek 2012)	1.00	1380		
(o) Bank Lending Shocks (Basett et al 2011)	.992	1140		
(p) NPV Defense Spending News/GDP (Ramey 2011)	.104	1380		
(q) Excess Returns on Defense Stocks (Fisher & Peters 2010)	.432	1260		
(r) Tax Changes/GDP (Romer & Romer 2010)	.108	1200		

Notes: F p-v. = F-test p-value on the industry x country coefficients associated with the instrument. N = observations, sample changes with the availability of the instrument. Instruments (d) – (r) calculated using data from Stock and Watson 2012; instruments (a)-(c) based upon FRED, SIPRI and World Bank data, as described in the text. Each regression follows the 1st stage specification given in (19), with industry x country and country x year fixed effects and the national unemployment rate change and instruments entered separately for each industry x country. The dependent variable is the labour share weighted change in the share of employment by worker type. Each row represents a separate analysis with the indicated instrument alone.

these are unlikely to have large effects on the distribution of economic activity, even if they affect levels. In contrast, actual defense spending shifts the pattern of demand away from the private norm, resulting in more significant changes in sectoral employment shares. Thus, Ramey's finding for macroeconomic aggregates need not extend to my analysis of labour allocations. When entered jointly with actual defense spending changes in the 1st stage regression for the US, I find the p–value on the F-test of Ramey's news variable to be .313, while that on actual defense spending changes remains .000. The insignificance of defense spending in the OECD regressions stems from the fact that for 3358 of the 8049 observations defense spending changes are zero. Defense spending as a share of GDP is extremely stable in most OECD countries and, with low values and one decimal precision in the SIPRI data, the sudden changes that do occur are mostly likely reflective of rounding error (e.g. moving from .9 percent of GDP previously to 1.0 percent ever after in one year in Japan).

Table II presents 2nd stage results using each of the four instruments which are significant at the 5% level in the 1st stage regressions for the United States in Table I (EU KLEMS results are presented later). Aside from the estimate of ξ, the elasticity of worker efficacy with respect to the sectoral employment share, I also report the p-value of the 1st stage F-test (which will vary across specifications) and the 2nd stage χ^2 overidentification test.[24] In the top panel, which follows the baseline specification of equation (19), three of the four instrumental variables estimates of ξ are substantially negative, although the only statistically significant estimate is that found using defense expenditures. Defense spending, however, is the only instrument which

[24] As noted earlier, an overidentification test is possible with one instrument because, since it is entered separately for each industry, there are technically actually I (equal to the number of industries) instruments. The overidentification test is whether these instruments have any predictive value in the regression beyond their association with changes in employment shares.

As the reader might worry that this is somehow econometrically wrong, I have used Monte Carlo simulations to confirm the accuracy of the test statistic. Using the covariance of the residuals from the first and second stage regressions of the baseline specification with defense expenditures, I produce 500,000 simulated draws of the data under the assumption that (modulo their influence on employment shares) defense expenditures are exogenous in the second stage regression. The resulting test statistic is nearly exact, i.e. the nominal rejection values are very close to the actual rejection probabilities (see the next paragraph).

I should also note that Basmann (1960) argues that the standard (Sargan 1958) χ^2 overidentification test is too conservative (i.e. rejects the null too frequently) in finite samples and proposes a small sample adjustment to the test statistic. I have confirmed his argument, for my case, using the Monte Carlo simulations described above. I find Sargan's χ^2 test to be grossly conservative (rejecting, as examples, 15.5% of the time at the 5% level and 4.6% of the time at the 1% level), while Basmann's small sample correction is only slightly conservative (rejecting 5.8% of the time at the 5% level and 1.3% of the time at the 1% level) but very close to being exact. Consequently, throughout this paper I use Basmann's statistic as the overidentification test.

Table II: Annual TFP Growth on Changes in Employment Shares (United States: 60 sectors x 1987-2010)					
	OLS	2SLS by type of instrument			
		Δ Defense Spending	Δ Metals Prices	Oil Price Maximum	Fed Funds Surprises
(1) Baseline specification (equation 19)					
ξ (s.e.)	-.218 (.108)	-.922 (.266)	-.546 (.318)	.372 (.384)	-.468 (.318)
F & χ^2 p-v.		.000 & .148	.000 & .004	.005 & .000	.000 & .000
N/K/L	1380	1380/199/59	1380/199/59	1380/199/59	900/191/59
(2) Dropping unemployment controls by industry (business cycle adjustment)					
ξ (s.e.)	-.167 (.100)	-.359 (.226)	-.245 (.452)	.359 (.396)	-.742 (.412)
F & χ^2 p-v.		.000 & .031	.440 & .000	.033 & .000	.371 & .002
(3) Substituting ln changes in capacity utilization for unemployment controls					
ξ (s.e.)	-.240 (.100)	-.689 (.222)	-.465 (.346)	.363 (.375)	-.654 (.343)
F & χ^2 p-v.		.000 & .009	.003 & .478	.029 & .000	.044 & .950
(4) Adding ln changes in capacity utilization to unemployment controls					
ξ (s.e.)	-.207 (.109)	-.771 (.254)	-.457 (.332)	.372 (.364)	-.596 (.319)
F & χ^2 p-v.		.000 & .260	.000 & .427	.003 & .000	.000 & .663
(5) Dropping country x year dummies (common component of TFP growth)					
ξ (s.e.)	-.257 (.107)	-1.03 (.263)	-.738 (.318)	.372 (.390)	-.541 (.317)
F & χ^2 p-v.		.000 & .146	.000 & .001	.007 & .000	.000 & .000
(6) Dropping one industry at a time					
Max ξ (s.e.)	-.119 (.107)	-.812 (.264)	-.300 (.325)	.636 (.441)	-.045 (.315)
Min ξ (s.e.)	-.328 (.113)	-1.13 (.312)	-.915 (.318)	-.007 (.386)	-.872 (.363)
Max F p-v.		.000	.015	.048	.003
Min F p-v.		.000	.000	.001	.000
Max χ^2 p-v.		.582	.067	.004	.000
Min χ^2 p-v.		.075	.001	.000	.000
(7) Adding 4 lags of employment share changes					
∑ ξ (s.e.)	-.685 (.209)	-.750 (.283)	-.547 (.338)	-.233 (.348)	-.621 (.359)
F & χ^2 p-v.		.000 & .068	.048 & .002	.083 & .009	.000 & .000

Notes: ξ (s.e) = coefficient (standard error) on labour share weighted changes of employment shares by worker type. F & χ^2 p-v. = p-value on 1st stage significance and 2nd stage overidentification tests. N/K/L = number of observations/number of regressors in 1st stage/excluded instruments in 2nd stage. Because of the joint year and industry dummies, one of the industry coefficients for each of the variables entered by industry (i.e. unemployment and capacity changes and instruments) is co-linear with other variables and is dropped in all specifications other than those without year dummies. Thus, there are only 59 excluded instruments in the baseline specification. ∑ ξ = sum of the coefficients on current & four lags of weighted employment share changes.

does not strongly reject the 2^{nd} stage exclusion restriction. I confirm the likely endogeneity of the oil price instrument by correlating its 1^{st} stage industry coefficients with the average energy share of gross output in those industries. If this instrument represents exogenous shifts in prices, then its effect should be substantially negatively correlated with the energy intensity of production, i.e. industries which are more energy intensive should see their relative employment share fall with exogenous increases in oil prices, as their supply curves shift up. In practice, I find a correlation coefficient of .232. While not significant (p-value = .077), the correlation is of the wrong sign. This might occur if some of the increases in the price of oil represent an endogenous positive response to rising energy demand in using industries. In sum, of 18 potential instruments, only 1 (defense expenditures) satisfies the dual requirements of 1^{st} stage significance and 2^{nd} stage exogeneity, and that instrument produces a strongly negative (-.922) estimate of ξ.

The lower panels of Table II examine the sensitivity of the results to the specification. In panel (2) I remove the unemployment rate entered by industry. This has a very large impact on the estimates, dramatically reducing the estimate of ξ for both defense expenditures and metal prices, raising it for Federal Funds surprises, and rendering both metals prices and Fed surprises completely insignificant in the 1^{st} stage regression. In panel (3) I substitute the Federal Reserve's estimate of aggregate mining, manufacturing and utilities capacity utilization for the unemployment rate, interacting it by industry as was done for unemployment. As shown, this moves ξ back to the estimates of panel (1), although the value using defense expenditures (-.689) is less extreme than in the baseline specification (-.922). The Fed's measure of capacity utilization, however, does not exhaust the association of industry productivity and labour allocations with the business cycle. Adding the measure of aggregate capacity utilization to the baseline specification with unemployment and defense spending, I find that the industry coefficients on the unemployment rate in both the 1^{st} and 2^{nd} stage regressions remain highly significant (F p-values of .000 & .003, respectively), suggesting that the business cycle characteristics of relative industry productivity and employment may go beyond capacity utilization and mismeasurement to something real. The estimate of ξ from defense spending in this specification is -.711 (panel 4). In general, controlling for the association between the business cycle and relative labour allocations and productivity seems appropriate[25] and this matters in the regression because the correlation between defense spending changes and changes in the unemployment rate in this time period is quite strong (.652 with a p-value of .001). Nevertheless, the reader looking to see whether the defense spending results can be rendered insignificant need look no further than panel (2). Panel (5) of Table II shows that removing the year dummies, but retaining the unemployment controls, generally increases the magnitude of ξ, with the negative estimate using metals prices now appearing significant.

Panel (6) of Table II explores whether identification and significance come from one particular industry by rerunning the baseline specification 60 times, removing one industry each time, and reporting the maximum-minimum range of the estimates of ξ and the F & χ^2 p-values. As shown, the estimates of ξ based upon the non-defense instruments vary enormously, but the range for defense expenditures is much more limited. Also of note is the stability of the 1st and 2nd stage tests for defense expenditures. Regardless of which industry is removed, defense spending is always found to be highly significant in the 1st stage regression and exogenous in the 2nd stage overidentification test. In fact, removing all possible combinations of two and even three industries, the 1st stage p-value on defense spending never rises above 2.3×10^{-8}, the p-value on its 2nd stage overidentification test never falls below .011, and the coefficient never becomes less negative than -.590 (274). Thus, the correlations between defense expenditures, employment

[25] To see this, the reader might introspect and consider their reaction if I had informed them that the estimate of ξ was substantially negative, but only when measures of the business cycle are *excluded* from the regression.

and productivity that lie behind the significant coefficients reported in the top panel of Table II go far beyond one, two or even three key industries.

The estimates using annual data in Table II might not provide an accurate representation of long run effects. On the one hand, it is possible that short run coefficients overstate the negative influence of the employment share on sectoral productivity as workers entering a sector are likely to be less productive initially than they will be in the long run, once they acquire sector specific human capital. On the other hand, it is possible that short run coefficients actually understate the negative effect of the employment share on sectoral productivity. Worker reallocations come about through changes in equilibrium output, either due to a shift of supply or demand. A sudden increase in output will lead to an influx of workers and, typically, a transitory rise in capacity utilization, producing a transitory overstatement of productivity.[26] Thus, this mismeasurement of productivity will be positively correlated with the movement of workers into a sector, understating the negative influence this otherwise has on measured productivity.[27]

Panel (7) in Table II addresses the issue of long run effects by adding four lagged values of the labour share weighted change in employment shares as pre-determined exogenous right-hand side variables to the baseline specification, with current employment reallocations instrumented with the instrument specified in each column. The cumulative effect on long run measured productivity is given by the sum of the current and lagged coefficients, which is presented in the table. Comparing these with the baseline results at the top of the table, one sees

[26]This applies even for instruments that shift the supply curve, provided they satisfy the exclusion restriction, i.e. are not directly correlated with total factor productivity growth. If something shifts the supply curve down without changing fundamental productive capacity, it will lead to an expansion of output which, along with the rise in the employment share, should produce a transitory increase in capacity utilization.

[27]As there are now a variety of mismeasurements, I should clarify. The object of interest in this paper is the mismeasurement of productivity due to the failure to account for the changing efficacy of workers as a sector's employment share expands. The transitory mismeasurement due to capacity utilization, however, works in the opposite direction and may temporarily conceal the effect I'm studying.

that ξ is now somewhat smaller in magnitude in the defense expenditures analysis (-.750 vs. -.922 earlier), while the oil price maximum, which earlier reported an insignificant positive coefficient, is no longer 1^{st} stage significant and now produces a negative point estimate of ξ. The metals prices coefficient is unchanged, while that for Fed surprises is more negative.

Table II also reports OLS results, running each specification without instruments. Although the baseline OLS relation between employment share changes and productivity (-.218) is small, the long run cumulative association, as evidenced by panel (7)'s regression with lags of past employment changes, is much more negative (-.685). It is difficult to explain how past employment changes relate negatively to current productivity growth within a framework where employment shares reflect the endogenous response of demand to shifts of the supply curve brought about by productivity change. The result is easier to comprehend, however, if one moves to a framework where changes in employment shares reflect exogenous shifts of the demand curve brought about by non-unitary income elasticities of demand and other shocks to relative demand. When demand shifts out in an industry, it produces a transitory rise in capacity utilization and a spurious rise in productivity, minimizing the negative effect of employment shifts on measured productivity. Over time, however, capacity adjusts and the full impact is revealed. Evidence in favour of this argument can be found by regressing total factor productivity growth on industry output growth, with industry, year and unemployment x industry controls as in the baseline specification. With only current output growth in the regression, the OLS coefficient (s.e.) is .219 (.025). With four lags of past output in the regression, the cumulative OLS coefficient is .076 (.058). Thus, past output increases, like past employment increases, lead to lower current productivity growth, which is consistent with the capacity utilization story outlined above.

Table III: Response of Capacity Utilization and Productivity to Output and Employment Share Changes (22 industries, 1987-2010)				
X variable	Δ Output		Δ Employment Share	
Y variable	Δ Cap U	Δ TFP	Δ Cap U	Δ TFP
(1) OLS: Current value of X				
Coef (s.e.)	.548 (.028)	.186 (.036)	1.70 (204)	.137 (.215)
(2) OLS: Adding four lagged values of X				
∑ Coef (s.e.)	.035 (.050)	-.096 (.084)	-.364 (.343)	-1.11 (.420)
(3) 2SLS: Current value of X instrumented with Δ defense expenditures				
Coef (s.e.)	.070 (.086)	-.192 (.097)	-.030 (.487)	-1.57 (509)
F & χ^2 p-v.	.000 & .891	.000 & .103	.000 & .929	.000 & .236
N/K/L	506/85/21	506/85/21	506/85/21	506/85/21
(4) 2SLS: Adding four pre-determined lagged values of X				
∑ Coef (s.e.)	-.041 (.060)	-.248 (.100)	-.234 (.399)	-1.78 (.509)
F & χ^2 p-v.	.000 & .568	.000 & .087	.005 & .771	.005 &.343
N/K/L	418/85/21	418/85/21	418/85/21	418/85/21

Notes: Δ Cap U = ln change in Federal Reserve Board's measure of industry capacity utilization; Δ TFP = ln change in TFP index, adjusted for labour quality (dependent variable in Tables I and II); Coef (s.e) = coefficient (standard error) on the current X variable; ∑ Coef = sum of the coefficients on current & four lags of the X variable. F & χ^2 p-v. and N/K/L as in Table II.

Table III supports the preceding argument using the Federal Reserve Board's industry level measures of capacity utilization, defined as current output over maximum sustainable output.[28] These measures are only available for the 22 mining, manufacturing and utilities industries in the 60 sector KLEMS disaggregation of private sector activity. In this table I run regressions with either the change in capacity utilization or total factor productivity growth as the "Y" variables, and either the growth of output or the labour income share weighted change in the share of economywide employment by type, the right-hand side variable of interest in the

[28]These measures are based upon the Survey of Plant Capacity and are defined as "the greatest level of output the plant can maintain within the framework of a realistic work schedule after factoring in normal downtime and assuming sufficient availability of inputs to operate the capital in place" (Gilbert, Morin & Raddock 2000, p. 194). The survey measures are then regressed on a time trend, ln capital and dummies which correct for outliers. This suggests that the reported series is basically a smoothed version of the original data, allowing outliers that the Fed believes represents real changes.

regressions reported above, as the "X" variable. Each regression includes a complete set of industry and time dummies and the change in the unemployment rate entered separately by industry, as in the baseline specification of equation (19). Aside from results with the current value of "X" alone, I also report the cumulative sum of the coefficients in a specification with the current value and four predetermined lags of "X."

I begin by taking both X variables as exogenous, running OLS specifications in the top two panels of the table. In the first two columns we see that an increase in current output raises both capacity utilization and measured total factor productivity growth, but that the cumulative long run effect, once lags are allowed, is insignificantly different from zero in both cases. In the third column of the table, we see that a 1% increase in a sector's labour income share weighted employment share is associated with a large 1.7% short run rise in capacity utilization, but has no long run effects. Regarding measured TFP, in the fourth column, an increase in the sectoral employment share has no significant short run impact on productivity, but a very large (-1.1) long run effect. These results are completely consistent with a view of exogenous demand fluctuations producing transitory movements in capacity utilization which obscure the true effect of labour allocations on measured productivity.[29]

The preceding is intended to be heuristic, and should not be taken completely literally. In particular, one cannot interpret the results as necessarily indicating that *all* changes in equilibrium quantity demanded (and labour allocations) are exogenous to productivity. To proceed more

[29]One can try to use the Fed's industry capacity utilization measures to directly adjust productivity, but this raises additional issues. First, an OLS regression approach is unsuitable, because industry capacity utilization is endogenous to industry productivity, but instruments for industry level capacity utilization are hard to find, as defense spending is uncorrelated with capacity utilization (see below). Second, one can use the utilization estimates to mechanically adjust productivity, but this requires some assumptions about what is being over and underutilized (capital, capital and labour, or capital, labour and some material inputs like energy) and what would have to be changed to reach sustainable output. For my purposes, however, it is sufficient to simply show that as capacity utilization effects disappear in the long run, the OLS relation between employment shares and productivity becomes decidedly negative.

carefully, the third and four panels of Table III instrument each X with defense expenditures, the instrument which I have previously found to be consistently 1st stage significant and 2nd stage exogenous. As before, I enter the instrument separately for each industry, and as before the 1st and 2nd stage test statistics satisfy the requirements of 2SLS in an admirably robust and consistent fashion.

Turning to coefficient estimates, the first notable result is that the first and third columns of panels (3) and (4) indicate that defense expenditures, while moving around output and labour allocations, have absolutely no effect on industry level capacity utilization. This is consistent with Ramey's (2011) argument that defense spending changes are well anticipated by public news announcements. While Ramey's news variable is completely insignificant in the 1st stage regressions for this sample, as it was before, this merely confirms that the timing of news is different than the timing of actual expenditures.[30] Nevertheless, actual expenditures, when they arrive, may be well anticipated, so that capacity utilization expands evenly with production needs. Because defense spending has no observable impact on capacity utilization, the long and short term coefficients for productivity growth, in the second and fourth columns, are virtually identical. The elasticity of observed productivity with respect to output is estimated to be around -.2 (s.e. of about .1). The coefficient on labour share weighted changes in employment shares, which following the theory above is interpretable as the elasticity of average worker efficacy with respect to the employment share, is found to be about -1.5 in this sample of only 22 industries

[30]Using Ramey's variable as the instrument in the 1st stage regressions for output and employment share changes, I get p-values on the F-tests of .161 and .483 respectively. Running Ramey's instrument jointly with current expenditures in these regressions, I get p-values of .837 and .992 on her news variable and .000 and .000 on actual expenditures. As emphasized earlier above, none of this invalidates Ramey's point that her news variable does a better job of explaining changes in macroeconomic aggregates, which will be influenced by the reaction of private economic actors to the anticipated future consequences of those expenditures. This is a distinct, however, from moving actual patterns of production away from the private norm, in which actual expenditures have a more significant effect.

(s.e. of about .5). This is greater in absolute magnitude than the maximum of -1 allowable by theory, but not (statistically) significantly so.

To summarize the results for the US KLEMS, out of 18 potential instruments, defense spending is the only one that consistently and strongly satisfies the dual requirements of 1st stage significance and 2nd stage exogeneity. Long and short term effects for defense spending are quite similar, as defense spending does not have much of an influence on capacity utilization. The long term OLS association between changing labour allocations and measured productivity is much more negative than the short term relation, and this appears to reflect transitory capacity utilization changes consistent with exogenous shifts in demand. The long term OLS estimate of the elasticity of worker efficacy with respect to employment shares in the total US KLEMS sample (-.685 in Table II) is not significantly different from that arrived at using defense expenditures as an instrument (-.750). Thus, while there may be some endogeneity of labour allocations, it probably accounts for a relatively small share of the total variation (exogenous plus endogenous) in this variable.[31]

Turning now to the EU KLEMS OECD data, as I do not have any instrument that is 1st stage significant in the analysis of the entire data set, I focus on country specific results. Since defense spending is a robustly significant and exogenous instrument in the US KLEMS data, I begin by running country by country 1st stage regressions using defense spending as an

[31] As noted earlier, while the preceding analysis is based upon my labour composition adjustment of BLS TFP growth and my estimates of changing sectoral employment shares by type, results are quite similar if I use the original BLS data on productivity and labour allocations without differentiation by worker type. For example, using defense spending as an instrument, I get the following estimates (s.e.) of ξ for the panels in Table II: Panel I: -1.06 (.275); II: -.373 (.218); III: -.722 (.218); IV: -1.03 (.275); V: -1.17 (.269); and VII: -.769 (.292). These follow the patterns presented in the table. The corresponding short term and long term OLS results (panels I & VII) are -.377 (.122) and -.809 (.218). In Table III, looking at columns (4)-(6) in panels III and IV, where employment share changes are instrumented with defense expenditures, I get insignificant short and long term coefficients for capacity utilization of -.062 (.475) and -.390 (.376), short and long term coefficients for BLS measured TFP growth of -1.82 (.502) and -1.84 (.514), and short and long term coefficients for capacity utilization adjusted BLS TFP growth of -1.75 (.540) and -1.45 (.580). Again, these results parallel those reported above.

Table IV: Country level analysis using EU KLEMS Data (29 sectors, 1970-2005)				
	Australia	Finland	Netherlands	United Kingdom
(1) Baseline specification (equation 19) with Δ defense expenditures/GDP as instrument				
ξ (s.e.)	-1.09 (.185)	-.310 (.359)	-.264 (.335)	-.886 (.153)
F & χ^2 p-v.	.004 & .901	.000 & .841	.005 & .290	.000 & .631
N/K/L	493/100/28	493/100/28	493/100/28	493/100/28
(2) Dropping one industry at a time				
Max ξ (s.e.)	-1.03 (.176)	-.079 (.395)	.152 (.406)	-.727 (.165)
Min ξ (s.e.)	-1.18 (.193)	-.682 (.406)	-.472 (.366)	-1.10 (.198)
Max F p-v.	.012	.004	.038	.001
Min F p-v.	.002	.000	.002	.000
Max χ^2 p-v.	.960	.941	.691	.809
Min χ^2 p-v.	.653	.712	.026	.534
(3) Adding 4 lags of employment share changes				
$\sum \xi$ (s.e.)	-1.15 (.307)	-.801 (.424)	-.560 (.407)	-.985 (.206)
F & χ^2 p-v.	.014 & .873	.000 & .728	.001 & .272	.000 & .928
Notes: As in Table II.				

instrument. I then proceed to the 2nd stage analysis for the four non-US countries where I find defense spending to be 1st stage significant at the 5% level (namely Australia, Finland, the Netherlands and the United Kingdom). As shown in Table IV, in each of these countries defense spending satisfies the 2nd stage exclusion requirement and produces negative estimates of ξ, although only the large point estimates of Australia and the United Kingdom are statistically significant. Removing one industry at a time, I find that defense spending robustly satisfies the 1st and 2nd stage significance and exclusion requirements. The point estimates of ξ vary greatly for Finland and the Netherlands and much less so for Australia and the United Kingdom, in keeping with their relative standard errors in the baseline specification. Adding lags of employment share changes to the regression produces a much larger estimate of the cumulative negative effect of reallocation on productivity, particularly for Finland and the Netherlands.

The EU KLEMS data base has two sets of estimates for the United States, one covering 1977-2005 based upon the current NAICS (North American Industry Classification System) used in the US KLEMS, and another covering 1970-2005 based upon the historical SIC (Standard Industrial Classification). The industrial sectors in both series share the same nominal titles and have TFP estimates grouped into the same 29 private sector division that I use in the general analysis of (SIC-based) EU KLEMS data for other countries.[32] Both of these series provide a longer time series than the BLS' US KLEMS (covering 1987-2010) and appear to be developed independently of that source. As in the case of the US KLEMS, I run 1st stage regressions for each of the 18 instruments in Table I and then proceed to a 2nd stage analysis with those instruments which are significant at the 5% level.

Table V reports second stage results for the six instruments which are 1st stage significant at the 5% level in the EU KLEMS US NAICS data. Defense expenditures operate much as in the analysis of the US KLEMS, producing an extremely large negative estimate of ξ in the baseline specification, 1st and 2nd stage significance and exclusion test statistics that are quite robust to the removal of one industry at a time, and (once lags are accounted for) a somewhat smaller estimate of the cumulative effect of employment changes. The oil price maximum, which produced a positive point estimate of ξ earlier in Table I, generates a ξ of -1.1 in this case. However, notwithstanding its statistical significance in the baseline specification, with the removal of one industry this coefficient is easily made positive. The remaining four instruments produce a cornucopia of insignificant results in the baseline regression, are often quite sensitive to the removal of one industry at a time and, when employment change lags are added, produce big cumulative negative estimates of ξ and are found to be utterly insignificant in the 1st stage

[32] Whenever I refer to results using all of the EU KLEMS data, as in Table I's 1st stage regressions, I use the SIC version of the US data, in keeping with the SIC definitions used for other countries.

Table V: US analysis using NAICS based US data in EU KLEMS (by instrument) (29 sectors, 1977-2005)						
	Δ Defense Spending	Δ Oil Prices	Oil Price Maximum	Smets/Wouters M Shock	Sims/Zha M Shock	TED Spread Innovation
(1) Baseline specification (equation 19)						
ξ (s.e.)	-1.06 (.425)	-.838 (.424)	-1.13 (.331)	-.740 (.413)	.932 (.476)	.047 (.328)
F & χ^2 p-v.	.003 & .524	.003 & .000	.000 & .000	.002 & .307	.007 & .869	.000 & .001
N/K/L	812/111/28	812/111/28	812/111/28	783/110/28	725/108/28	812/111/28
(2) Dropping one industry at a time						
Max ξ (s.e.)	-.873 (.427)	.667 (.554)	.282 (.404)	-.520 (.426)	1.08 (.510)	.241 (.361)
Min ξ (s.e.)	-1.20 (.455)	-1.43 (.435)	-1.72 (.363)	-1.10 (.463)	.763 (.530)	-.380 (.333)
Max F p-v.	.012	.155	.000	.079	.048	.000
Min F p-v.	.001	.002	.000	.000	.003	.000
Max χ^2 p-v.	.697	.001	.011	.440	.942	.009
Min χ^2 p-v.	.362	.000	.000	.205	.693	.000
(3) Adding 4 lags of employment share changes						
$\sum \xi$ (s.e.)	-.851 (.394)	-2.16 (.500)	-1.86 (.404)	-.944 (.456)	-.626 (.607)	-.809 (.457)
F & χ^2 p-v.	.045 & .528	.339 & .020	.018 & .082	.271 & .066	.865 & .988	.396 & .000
Notes: As in Table II.						

regression. In sum, as in the analysis of the BLS US data, only defense spending consistently satisfies the 1st and 2nd stage tests, and that instrument produces an estimate of ξ close to -1.

Table VI reports second stage results for the six instruments which are 1st stage significant at the 5% level in the EU KLEMS US SIC data. Three of these instruments (defense spending, oil price maximum, and the TED spread) overlap with the list for the EU KLEMS US NAICS data. While the oil price maximum and TED spread produce results that are similar to those in the preceding table, those with defense expenditures are dramatically different. Although defense spending is 1st stage significant and 2nd stage exogenous in the baseline specification, it produces a small and statistically insignificant estimate of ξ. With lags, however, the coefficient becomes considerably more negative, albeit not statistically significant. With regards to the remaining

Table VI: US analysis using SIC based US data in EU KLEMS (by instrument) (29 sectors, 1970-2005)						
	Δ Defense Spending	Oil Price Maximum	Residual Gas Prices	Romer/Romer M Shock	Fed Funds Surprises	TED Spread Innovation
(1) Baseline specification (equation 19)						
ξ (s.e.)	-.118 (.251)	-.821 (.350)	-.329 (.413)	.162 (.373)	-.234 (.485)	.086 (.326)
F & χ^2 p-v.	.000 & .151	.000 & .000	.031 & .000	.036 & .385	.001 & .000	.000 & .644
N/K/L	1015/118/28	1015/118/28	1015/118/28	754/109/28	435/98/28	1015/118/28
(2) Dropping one industry at a time						
Max ξ (s.e.)	-.048 (.240)	-.173 (.369)	-.035 (.400)	.358 (.472)	.412 (.557)	.251 (.320)
Min ξ (s.e.)	-.191 (.468)	-1.14 (.399)	-2.49 (.981)	-.006 (.384)	-.599 (.483)	-.072 (.446)
Max F p-v.	.000	.004	.926	.058	.013	.000
Min F p-v.	.000	.000	.019	.003	.000	.000
Max χ^2 p-v.	.723	.000	.008	.676	.004	.876
Min χ^2 p-v.	.009	.000	.000	.191	.000	.363
(3) Adding 4 lags of employment share changes						
$\sum \xi$ (s.e.)	-.550 (.449)	-1.55 (.395)	-1.42 (.459)	-1.14 (.532)	-.351 (.640)	-.760 (.384)
F & χ^2 p-v.	.034 & .540	.000 & .000	.048 & .000	.206 & .748	.002 & .000	.000 & .879
Notes: As in Table II.						

instruments, the point estimates are generally quite sensitive to the removal of one industry or the 1st stage regression is rendered insignificant once lags are introduced. With the introduction of lags the cumulative effect of employment changes becomes much more negative, although the TED spread is the only instrument in this specification that is strongly significant and exogenous. Its estimate of ξ is both substantially negative (-.760) and statistically significant.

As Tables V and VI suggest, there are peculiar differences between the SIC-based and NAICS-based EU KLEMS data for the United States. The correlation between the annual industry x year total factor productivity growth in one dataset and the other, for the 29 nominally identical[33] large private sector industry groupings and the 28 years that the two datasets overlap,

[33] E.g. "mining & quarrying", "education", "rubber and plastics", etc.

is only .502 (i.e. an R^2 of .25), despite the fact that they ostensibly measure exactly the same thing. The labour income share weighted labour reallocation measures, however, are much more similar, with a correlation of .860 (R^2 = .74). Not surprisingly, this produces radically different regression results. There are also some disturbing anomalies in the EU KLEMS SIC based US data and in the EU KLEMS data set as a whole.[34] Such concerns are, however, somewhat beside the point, as it cannot be taken as altogether surprising that a single instrument, such as defense spending, will in some specifications or some data sets produce weaker results.

Before concluding, I present the OLS results for the EU KLEMS data. As shown in Table VII, the results here closely parallel those for the United States. Whether in the four European countries examined in the tables above, either of the SIC and NAICS versions of EU KLEMS US data, or the EU KLEMS database as a whole, the association between employment share changes and productivity growth is negative, but becomes much more so when past employment share changes are added to the regression. As in the case of the US data, the difference between the current and cumulative coefficients lends itself to the interpretation that exogenous movements in demand produce transitory changes in capacity utilization which obscure the strongly negative long term association between employment shares and measured productivity. The cumulative

[34]For example, between 1970 and 1981, according to the EU KLEMS SIC data, the relative value added price of private sector services to goods in the United States fell 27%, while the relative quantity rose by 25%, for a -2% change in relative nominal value added. According to the current official US National Income and Product Accounts, however, during this same period the relative value added price of private sector services to goods actually declined only 5% (reflecting rising energy prices), while the relative quantity rose by 14%, for a +9% change in relative nominal value added (Chain47on.xls and VA74on.xls available at www.bea.gov). The historical SIC series on the BEA website (GDPbyInd_VA_SIC.xls) does not provide real indices back to 1970, but in the 1977-1981 period it shows an 8% increase in relative real service quantity and 5% decrease in relative price (similar to the current series 7% and 3% figures for the same period), while the EU KLEMS SIC data show a 15% increase in relative quantity and 13% decline in relative price. In making these comparisons, I follow the BEA's definition of goods (agriculture, mining, manufacturing and construction) and services (all other private sector).

As another example, in the EU KLEMS data one finds that in 3% of the observations with total factor productivity estimates capital income is negative, averaging -.14 of value added and -.05 of gross output and ranging as far as -5.7 times value added or -.33 of gross output. In these observations, appearing in 16 countries, one gets very close (R2 = .985) to the EU KLEMS estimate of total factor productivity growth by dropping capital growth while using the gross output shares of intermediate inputs and labour (i.e. weights which combined now exceed 1) to calculate the contribution of these inputs to output growth.

Table VII: OLS Regressions using EU KLEMS (29 sectors, 1971-2005)							
	Australia	Finland	Netherlands	UK	US NAICS	US SIC	OECD 18
(1) Baseline specification with employment share changes (equation 19)							
ξ (s.e.)	-.875 (.061)	-.266 (.095)	-.485 (.092)	-.777 (.068)	-.518 (.119)	-.344 (.094)	-.422 (.023)
N	667	1015	754	986	812	1015	12109
(2) Adding 4 lags of employment share changes							
$\sum \xi$ (s.e.)	-.941 (.202)	-.726 (.148)	-.756 (.181)	-1.05 (.126)	-.946 (.197)	-.929 (.196)	-.615 (.048)
N	551	889	638	870	696	899	10025
Notes: As in Table II.							

OLS coefficients are in most cases quite close to the corresponding cumulative coefficients using 2SLS, suggesting that much of the variation in labour shares is exogenous. I recognize of course that this interpretation, taking employment shares as being exogenous and OLS coefficients as accurate representations of causal relations, is awfully convenient in a paper which struggles to find more than one robust instrument.

The EU KLEMS results, by and large, confirm the analysis using the US KLEMS. Defense spending is the only instrument which is consistently 1st stage significant, 2nd stage exogenous and robust, both in terms of test statistics and coefficient point estimates, to the selective removal of industries. Long term OLS elasticities are more negative than short term relations. The cumulative estimate of ξ, both OLS and 2SLS with defense spending, is always more negative than -.5 and often much closer to the theoretical limit of -1. Standard errors, however, are very large and coefficient estimates in particular specifications and samples are not significantly different from zero. Thus, while the preponderance of evidence suggests that average worker efficacy does indeed fall with a sector's employment share, there is substantial uncertainty regarding the precise magnitude of the elasticity.

I conclude by simply considering how different values of ξ change our assessment of

Table VIII: Average Total Factor Productivity Growth with Roy Adjustments						
	United States (1987-2010, based on US KLEMS)			18 OECD countries (1970-2005, based on EU KLEMS)		
ξ	Goods	Services	Aggregate	Goods	Services	Aggregate
0.00	.0157	.0073	.0097	.0157	.0017	.0070
-0.25	.0134	.0078	.0094	.0131	.0036	.0072
-0.50	.0110	.0084	.0091	.0104	.0055	.0074
-0.75	.0087	.0090	.0088	.0077	.0074	.0076
-1.00	.0064	.0095	.0085	.0051	.0094	.0079

Notes: Goods, services and aggregate calculated from BLS KLEMS and EU KLEMS 60 sector and 29 sector, respectively, gross output TFP measures using equation (18), with adjustments for bias as indicated by equation (17) earlier above.

relative goods and services productivity growth. In Table VIII I combine the 60 sector US KLEMS and 29 sector EU KLEMS sectoral estimates of gross output private sector productivity growth into goods & services value added aggregates. With an ξ of 0, i.e. no adjustment for Roy effects, the US and EU KLEMS data indicate that productivity growth is .8 % faster per annum in goods than services in the United States and 1.4 % faster per annum in the OECD 18 as a whole. Moving down, as ξ becomes more negative the gap between goods and services productivity growth narrows until, at a value of -.750, it disappears altogether in both samples.[35] Aggregate productivity growth is relatively insensitive to ξ, i.e. there is not much aggregate bias, reflecting the fact that average wages per worker are relatively equal across the two sectors.[36]

In the US National Income and Product Accounts, between 1947 and 2011 the ln relative price of services to goods increases at an average annual rate of .0083, while the ln relative quantity increases by .0090. According to the EU KLEMS data, between 1970 and 2005 the ln relative price of services to goods in the OECD 18 increases at an average annual rate of .0114,

[35]The Domar weighted sum of sectoral reallocations is larger in the OECD 18 than in the US alone, and hence eliminates a larger productivity gap with, interestingly, the same value of ξ.

[36]In the US KLEMS the ln average annual wage per hour is .059 higher in goods, with a time trend of -.0014 (.0005). In the EU KLEMS, across 471 country x year observations ln relative goods wages are -.084 lower than in services and, with country dummies, show an annual trend of .0052 (.0004).

while the ln relative quantity rises by .0106. Thus, the long run rate of increase of the relative price of services to goods is roughly equal to the long run rate of increase of their relative quantity. Section II earlier showed that, under the assumptions of equal sectoral factor income shares and proportionate wages, assumptions which are tolerably satisfied in the data,[37] the slope of the Roy supply curve equals $-\Theta_L\xi/(1+\Theta_L\xi)$. Setting Θ_L equal to 2/3 and ξ to -.75, one gets a slope of 1. As the Roy supply curve shows, there are bounds on the explanatory power of the Roy model. If ξ is to lie within its theoretical limit of -1, there must be a sufficient movement of relative quantity and, more precisely, labour allocations relative to the observed sectoral relative price and (measured) productivity movements. Both this simple back-of-the-envelope calculation and the more careful computations of Table VIII show that these movements exist.

The reported difference in goods and services productivity growth in the US and the OECD is .8 and 1.4 percent per annum, respectively. Examining the values in Table VIII for ξ from -.5 to -1, the range of defense spending based long run elasticities found earlier, the adjusted difference ranges from +.5 percent in favour of goods to +.4 percent in favour of services. Thus, while it provides indications that the productivity growth gap between the two sectors is grossly overstated, this paper does not have a definitive point estimate to deliver to the reader. A value of ξ equal to -.75, however, lies in the middle of the point estimates, and allows for the reinterpretation of historical productivity, price and quantity data as representing a world in which true productivity growth in goods and services is roughly equal but Roy worker efficacy effects give rise to relative cost changes and the appearance of productivity growth differences. Thus, the "Roy supply curve" is a plausible, albeit not proven, explanation of the cost disease of services. This is the main point of this paper.

[37]Relative wages are addressed in the preceding footnote. Regarding factor shares, in the US KLEMS the average annual labour share in goods is .65, while in services it is .68, and their ln difference has an annual trend of -.0039 (.0005). Across the EU KLEMS, the average annual labour share in goods is .68 and in services is .64 and the ln difference, with country dummies, has an annual trend of .0002 (.0003)

V. Conclusion

William Baumol's cost disease of services has become part of the intellectual landscape of the profession, a truism taught, at least by this author, to generations of students. The profession, however, is also mindful of the fact that total factor productivity growth is a residual, Abramovitz's (1956) famous "measure of our ignorance", and has constantly sought new ways of explaining it. This paper follows a growing literature showing the role Roy's model of self-selection amongst heterogeneous workers can play in explaining macroeconomic phenomena. It finds evidence in the relation between employment shares and measured productivity that average worker efficacy declines as a sector's employment share increases, systematically biasing standard measures of productivity growth. While there is considerable uncertainty about the precise magnitude of these effects, the depiction of the relative supply of goods and services as being based upon equal goods and services productivity growth, with a rising relative cost brought about by an association between average worker efficacy and sectoral employment shares, is a plausible alternative characterization of developments in the US and the OECD.

As noted by Jones (2002), barring the Great Depression and World War II, the growth of income per capita in the United States has been a remarkably steady 2% per annum for more than 130 years, despite enormous structural changes in the US economy. Theoretically, it is difficult to think about this historical record in a framework in which aggregate economic growth is asymptotically drawn down to that of the slowest, most stagnant, sector. Practically, it is hard to sustain a fear of prospective stagnation in the face of such a lengthy retrospective history of constant growth. The alternative view that, by and large, a rising tide of technology raises all boats (industries), while changes in relative prices simply reflect movements along a standard classroom concave production possibilities frontier, provides an easier way to think about the past history and future prospects of the US economy.

VI. Bibliography

Abramovitz, Moses. "Resource and Output trends in the U.S. since 1870." American Economic Review Vol. 46, No. 2 (May 1956 Papers and Proceedings): 5-23.

Agresti, Alan. Categorical Data Analysis. New York: John Wiley and Sons, 1990.

Baker, Scott R., Nicholas Bloom, and Steven J. Davis. "Measuring Economic Policy Uncertainty." Manuscript, 2012.

Bassett, William F., Mary Beth Chosak, John C. Driscoll, and Egon Kakrajšek. "Changes in Bank Lending Standards and the Macroeconomy." Washington D.C.: Division of Monetary Affairs, Board of Governors of the Federal Reserve System, 2011.

Basmann, R. L. "On Finite Sample Distributions of Generalized Classical Linear Identifiability Test Statistics." Journal of the American Statistical Association Vol. 55, No. 292 (Dec. 1960): 650-659.

Baumol, William J. and W.G. Bowen. "On the Performing Arts: The Anatomy of Their Economic Problem." American Economic Review Vol. 55, No. 1/2 (March 1965): 495-502.

Baumol, William J. "Macroeconomics of Unbalanced Growth: The Anatomy of Urban Crisis." American Economic Review Vol. 57, No. 3 (June 1967): 415-426.

Baumol, William J., Sue Anne Batey Blackman and Edward N. Wolff. "Unbalanced Growth Revisited: Asymptotic Stagnancy and New Evidence." American Economic Review Vol. 75, No. 4 (September 1985): 806-817.

Baumol William J. The Cost Disease. New Haven: Yale University Press, 2012.

Bloom, Nicholas. "The Impact of Uncertainty Shocks." Econometrica Vol. 77, No. 3 (May 2009): 623-85.

Bustos, Paula, Bruno Caprettini and Jacopo Ponticelli. "Agricultural Productivity and Structural Transformation: Evidence from Brazil." Manuscript, August 2013.

Domar, Evsey D. "On the Measurement of Technological Change." The Economic Journal Vol. 71, No. 284 (December 1961): 709-729.

Fisher, Jonas D.M., and Ryan Peters. "Using Stock Returns to Identify Government Spending Shocks." The Economic Journal Vol. 120, Issue 544 (May 2010): 414-36.

Gilbert, Charles, Norman Morin and Richard Raddock. "Industrial Production and Capacity Utilization: Recent Developments and the 1999 Revision." Federal Reserve Bulletin (March 2000): 188-205.

Gilchrist, Simon and Egon Zakrajšek. "Credit Spreads and Business Cycle Fluctuations." American Economic Review Vol. 102, No. 4 (June 2012): 1692-1720.

Griliches, Zvi, ed. Output Measurement in the Service Sectors. NBER Studies in Income and Wealth, Vol. 56. Chicago: University of Chicago Press, 1992.

Griliches, Zvi. “Productivity, R&D, and the Data Constraint.” American Economic Review Vol. 84, No. 1 (1994): 1-23.

Gürkaynak, Refet S., Brian Sack, and Eric T. Swanson. “Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements.” International Journal of Central Banking Vol. 1, No. 1 (2005): 55-93.

Hamilton, James D. “What Is an Oil Shock?” Journal of Econometrics Vol. 113 (2003): 363-398.

Hsieh, Chang-Tai, Erik Hurst, Charles I. Jones, and Peter J. Klenow. “The Allocation of Talent and U.S. Economic Growth.” Manuscript, May 2012.

Hulten, Charles R. "Growth Accounting with Intermediate Inputs." The Review of Economic Studies Vol. 45, No. 3 (Oct. 1978): 511-518.

Jones, Charles I. "Sources of US Economic Growth in a World of Ideas." American Economic Review Vol. 92, No. 1 (March 2002): 220-239.

Kilian, Lutz. “Exogenous Oil Supply Shocks: How Big Are They and How Much Do They Matter for the U.S. Economy?” Review of Economics and Statistics Vol. 90, No. 2 (May 2008): 216-40.

Kuralbayeva, Karlygash and Radoslaw Stefanski. “Windfalls, Structural Transformation and Specialization.” Journal of International Economics Vol. 90, No. 2: 273-301.

Lagakos, David and Michael E. Waugh. "Selection, Agriculture, and Cross-Country Productivity Differences." Manuscript, Nov. 2011.

McLaughlin, Kenneth J. and Mark Bils. “Interindustry Mobility and the Cyclical Upgrading of Labor.” Journal of Labor Economics Vol. 19, No. 1 (January 2001): 94-135.

Ngai, L. Rachel and Christopher A. Pissarides. “Structural Change in a Multisector Model of Growth.” American Economic Review Vol. 97, No. 1 (March 2007): 429-443.

Ramey, Valerie A. “Identifying Government Spending Shocks: It’s All in the Timing.” Quarterly Journal of Economics Vol. 126, No.1 (Feb. 2011): 1-50.

Ramey, Valerie A. and Daniel J. Vine. “Oil, Automobiles, and the U.S. Economy: How Much Have Things Really Changed?” NBER Macroeconomics Annual Vol. 25 (2010): 333-367.

Romer, Christina D. and David H. Romer. “A New Measure of Monetary Shocks: Derivation and Implications.” American Economic Review Vol. 94, No. 4 (Sept. 2004): 1055-84.

Romer, Christina D., and David H. Romer. “The Macroeconomic Effects of Tax Changes: Estimates Based on a New Measure of Fiscal Shocks.” American Economic Review Vol. 100, No. 3 (June 2010): 763-801.

Roy, A.D. "Some Thoughts on the Distribution of Earnings." Oxford Economic Papers, New Series, Vol. 3, No. 2 (June 1951): 135-146.

Sargan, J. D. "The Estimation of Economic Relationships using Instrumental Variables." Econometrica Vol. 26, No. 3 (July 1958): 393-415.

Sims, Christopher A. and Tao Zha. "Were There Regime Switches in U.S. Monetary Policy?" American Economic Review Vol. 96, No. 1 (March 2006): 54-81.

Smets, Frank and Raph Wouters. "Shocks and Frictions in U.S. Business Cycles: A Bayesian DSGE Approach." American Economic Review Vol. 97, No. 3 (June 2007): 586-606.

Stock, James H. and Mark W. Watson. "Disentangling the Channels of the 2007-09 Recession." Brookings Papers on Economic Activity (Spring 2012): 81-135.

Triplett, Jack E. and Barry P. Bosworth. Productivity in the U.S. Services Sector: New Sources of Economic Growth. Washington, D.C.: Brookings Institution Press, 2004.

Appendix A: Characteristics of the Model's Equilibrium

This appendix provides the mathematical details behind the assertions made in Section II. All of the proofs are couched in terms of a two sector economy (goods and services). Their extension to the more complex N-sector case is straightforward.

(a): Regardless of the cumulative distribution function describing the paired draw (z_G, z_S), $\xi = (d\bar{z}_i / d\pi_i)(\pi_i / \bar{z}_i) > -1$.

Let $G_{j|i}(y|z)$ describe the conditional probability $z_j \leq y$ given that $z_i = z$, i.e. the cumulative distribution function of z_j given z_i, and let $g_{j|i}$ describe the corresponding conditional density and g_i the marginal density of z_i. Then, with $\omega = w_i/w_j$

$$\text{(a1)} \quad \bar{z}_i = \frac{N(\omega)}{\pi_i(\omega)}, \text{ with } N = \int_0^\infty z_i g_i(z_i) G_{j|i}(\omega z_i|z_i)\,\mathrm{dz_i}\ ,\ \pi_i = \int_0^\infty g_i(z_i) G_{j|i}(\omega z_i|z_i)\,\mathrm{dz_i}$$

$$dN/d\omega = \int_0^\infty z_i^{\,2} g_i(z_i) g_{j|i}(\omega z_i|z_i)\,\mathrm{dz_i},\ \text{and}\ d\pi_i/d\omega = \int_0^\infty z_i g_i(z_i) g_{j|i}(\omega z_i|z_i)\,\mathrm{dz_i},$$

where, in cases where the domains of g_i and $G_{j|i}$ do not include all positive real numbers I extend them, for the purpose of the integration, by defining g_i and $g_{j|i}$ as equal to zero in the extended region (and similarly for other proofs below). Note that $g_i(z_i)g_{j|i}(\omega z_i,|z_i) = g_{i,j}(z_i,\omega z_i)$, the joint distribution of z_i and z_j at the point $(z_i,\omega z_i)$. Assuming this joint distribution has mass along a positive measure of the ray with slope ω from the origin, we have $dN/d\omega > 0$ and $d\pi_i/d\omega > 0$,[38] and it follows that

[38]If it does not, we are at a value of ω where neither π_i nor $\bar{z}_i$ vary with ω, so the derivative of one with respect to the other is not well defined.

$$\text{(a2)}\quad \frac{d\bar{z}_i}{d\pi_i}\frac{\pi_i}{\bar{z}_i} = -1 + \frac{\frac{dN}{d\omega}\frac{d\omega}{d\pi_i}}{\bar{z}_i} > -1.$$

As noted in the text, this result is fairly obvious.

For particular distributional forms, it is easy to calculate closed form solutions for ξ illustrating the properties imposed by different distributional assumptions. Thus, for the case where the z_i are independent draws from fréchet distributions with cumulative distribution functions $G_i(z_i) = exp(-(z_i/\lambda_i)^{-\theta})$, we have:

$$\text{(a3)}\quad \pi_i = \int_0^\infty \theta\lambda_i^\theta z_i^{-\theta-1} \exp\left[-(z_i/\lambda_i)^{-\theta}\right]\exp\left[-(w_i z_i / w_j\lambda_j)^{-\theta}\right]\mathrm{dz_i} = \frac{(w_i\lambda_i)^\theta}{\sum_i (w_i\lambda_i)^\theta}$$

and

$$\text{(a4)}\quad \bar{z}_i = \frac{\int_0^\infty \theta\lambda_i^\theta z_i^{-\theta} \exp\left[-z_i^{-\theta}\left(\lambda_i^\theta + (w_j\lambda_j/w_i)^\theta\right)\right]\mathrm{dz_i}}{\pi_i}$$

$$= \frac{\lambda_i^\theta}{\pi_i C^{(\theta-1)/\theta}}\int_\infty^0 -x^{-1/\theta}e^{-x}\,\mathrm{dx} \quad \text{where} \quad C = \lambda_i^\theta + (w_j\lambda_j/w_i)^\theta$$

$$= \frac{\lambda_i(w_i\lambda_i)^{\theta-1}}{\pi_i(w_i^\theta C)^{(\theta-1)/\theta}}\int_0^\infty x^{-1/\theta}e^{-x}\,\mathrm{dx} = \lambda_i\Gamma\left(\frac{\theta-1}{\theta}\right)\pi_i^{-1/\theta}$$

where I have used the substitution $x = z_i^{-\theta}C$ in the second line. Consequently:

$$\text{(a5)}\quad \xi = (d\bar{z}_i/d\pi_i)(\pi_i/\bar{z}_i) = -1/\theta$$

Thus, for independent draws from fréchet distributions ξ is a constant, a function of the distribution's dispersion parameter.

It is not difficult to find distributions with different characteristics. Thus, if z_G and z_S are independent draws from exponential distributions with densities $\lambda_i exp[-z_i\lambda_i]$, allowing $\tilde{\omega} = w_i\lambda_j / w_j\lambda_i$ we have:

$$\text{(a6)}\quad \bar{z}_i = \frac{\int_0^\infty z_i\lambda_i \exp(-\lambda_i z_i)[1-\exp(-\lambda_i\tilde{\omega}z_i)]\mathrm{dz_i}}{\int_0^\infty \lambda_i \exp(-\lambda_i z_i)[1-\exp(-\lambda_i\tilde{\omega}z_i)]\mathrm{dz_i}} = \frac{1}{\lambda_i}\frac{\tilde{\omega}+2}{\tilde{\omega}+1}$$

$$\pi_i = \int_0^\infty \lambda_i \exp(-\lambda_i z_i)[1-\exp(-\lambda_i\tilde{\omega}z_i)]\mathrm{dz_i} = \frac{\tilde{\omega}}{1+\tilde{\omega}}$$

$$\frac{d\bar{z}_i}{d\pi_i}\frac{\pi_i}{\bar{z}_i} = \frac{d\bar{z}_i/d\omega}{d\pi_i/d\omega}\frac{\pi_i}{\bar{z}_i} = -\frac{\tilde{\omega}}{\tilde{\omega}+2}$$

so ξ once varies between 0 and -1 as ω goes from 0 to ∞ or, equivalently, π_i goes from 0 to 1. In this case ξ is a function of sectoral size alone. The log-normal distribution provides an interesting third example. With independent productivity draws with ln means μ_i and (for simplicity) common standard deviation σ, we have (the proof involves the convolution and incomplete moment properties of the lognormal):

$$\text{(a7)}\quad \bar{z}_i = \exp\left[\mu_i + .5\sigma^2\right]\frac{N\left(\frac{\tilde{\omega}}{\sqrt{2}\sigma}+\frac{\sigma}{\sqrt{2}}\right)}{N\left(\frac{\tilde{\omega}}{\sqrt{2}\sigma}\right)},\ \pi_i = N\left(\frac{\tilde{\omega}}{\sqrt{2}\sigma}\right)$$

$$\xi = \exp\left[-\frac{\tilde{\omega}}{\sigma}\frac{\sigma}{2}-\frac{\sigma^2}{4}\right]\frac{N\left(\frac{\tilde{\omega}}{\sqrt{2}\sigma}\right)}{N\left(\frac{\tilde{\omega}}{\sqrt{2}\sigma}+\frac{\sigma}{\sqrt{2}}\right)}-1$$

where $N()$ is the cumulative standard normal and $\tilde{\omega} = ln(w_i/w_j)+\mu_i-\mu_j$. Holding constant the size of each sector (i.e. $\tilde{\omega}/\sigma$), as σ goes from 0 to ∞ ξ goes from 0 to -1 in both sectors.

(b): The bias in the growth accountant's measure of aggregate productivity growth for any distribution of paired draws (z_G,z_S).

Aggregate total factor productivity growth is the GDP share weighted sum of sectoral value added productivity growth

$$\text{(b1)}\quad \hat{A}(est) = \sum_i \Omega_i \hat{A}_i(est) = \hat{A}(true) - \sum_i \Omega_i \Theta_{Li}\hat{\bar{z}}_i$$

$$= \hat{A}(true) - \frac{L}{GDP}\sum_i w_i\pi_i d\bar{z}_i,\quad \text{where}\ \Omega_i = \frac{P_iQ_i}{GDP}, \Theta_{Li} = \frac{w_i\bar{z}_iL\pi_i}{P_iQ_i}$$

and where I have made use of (7) and (8) in the text for the second equality and simplified using the definitions of Ω_i and Θ_{Li} in the last equality. Following the notation in (a) above for marginal, conditional and joint densities, with $\omega = w_G/w_S$ we have

$$\text{(b2)}\quad \bar{z}_G = \frac{N_G(\omega)}{\pi_G(\omega)} = \frac{\int_0^\infty z_G g_G(z_G) G_{S|G}(\omega z_G|z_G)\,\mathrm{dz_G}}{\int_0^\infty g_G(z_G) G_{S|G}(\omega z_G|z_G)\,\mathrm{dz_G}}, \quad \bar{z}_S = \frac{N_S(\omega)}{\pi_S(\omega)} = \frac{\int_0^\infty z_S g_S(z_S) G_{G|S}(\omega^{-1} z_S|z_S)\,\mathrm{dz_S}}{\int_0^\infty g_S(z_S) G_{G|S}(\omega^{-1} z_S|z_S)\,\mathrm{dz_S}}$$

and

$$\text{(b3)}\quad \frac{dN_G}{d\omega} = \int_0^\infty z_G{}^2 g_G(z_G) g_{S|G}(\omega z_G|z_G)\,\mathrm{dz_G} = \int_0^\infty z_G{}^2 g_{G,S}(z_G, \omega z_G)\,\mathrm{dz_G}$$

$$\frac{dN_S}{d\omega} = -\int_0^\infty \omega^{-2} z_S{}^2 g_S(z_S) g_{G|S}(\omega^{-1} z_S|z_S)\,\mathrm{dz_S} = -\int_0^\infty \omega^{-2} z_S{}^2 g_{G,S}(\omega^{-1} z_S, z_S)\,\mathrm{dz_S}$$

$$= -\omega \int_0^\infty z_G{}^2 g_{G,S}(z_G, \omega z_G)\,\mathrm{dz_G} = -\omega \frac{dN_G}{d\omega}$$

where I have used the substitution $z_G = \omega^{-1} z_S$ in the integral of the last line. Using

$$\text{(b4)}\quad d\bar{z}_i = \frac{dN_i/d\omega}{\pi_i} d\omega - \bar{z}_i \frac{d\pi_i/d\omega}{\pi_i} d\omega = \frac{dN_i/d\omega}{\pi_i} d\omega - \bar{z}_i \frac{d\pi_i}{\pi_i}$$

we substitute into the last line of (b1) and get

$$\text{(b5)}\quad \hat{A}(est) = \hat{A}(true) - \sum_i \frac{L}{GDP} w_i \pi_i d\bar{z}_i = \hat{A}(true) - \frac{L}{GDP} \sum_i w_i \pi_i \left[\frac{dN_i/d\omega}{\pi_i} d\omega - \bar{z}_i \frac{d\pi_i}{\pi_i} \right]$$

$$= \hat{A}(true) + \sum_i \frac{L}{GDP} w_i \bar{z}_i d\pi_i = \hat{A}(true) + \sum_i \Omega_i \Theta_{Li} \hat{\pi}_i$$

which is equation (9) in the text.

(c): Equality of mean sectoral wages with different distributions.

In the case of independent draws from fréchet distributions, equilibrium wages per worker equalize across sectors. Using (a4) earlier:

$$\text{(c1)}\quad w_i \bar{z}_i = w_i \lambda_i \Gamma\left(\frac{\theta - 1}{\theta}\right) \left(\frac{(w_i \lambda_i)^\theta}{\sum_i (w_i \lambda_i)^\theta} \right)^{-1/\theta} = \Gamma\left(\frac{\theta - 1}{\theta}\right) \left(\frac{1}{\sum_i (w_i \lambda_i)^\theta} \right)^{-1/\theta}$$

which is independent of i. This is not a general characteristic of this type of model. For example, for the case where the productivities are independent draws from the exponential distribution, we use (a6) and see:

$$\text{(c2)}\quad w_i\bar{z}_i = \frac{w_i}{\lambda_i}\left[\frac{(w_i\lambda_j / w_j\lambda_i)+2}{(w_i\lambda_j / w_j\lambda_i)+1}\right] \neq \frac{w_j}{\lambda_j}\left[\frac{(w_j\lambda_i / w_i\lambda_j)+2}{(w_j\lambda_i / w_i\lambda_j)+1}\right] = w_j\bar{z}_j \quad \text{unless} \quad w_i/\lambda_i = w_j/\lambda_j$$

(d): Independence of the paired productivity draws and *η(z) = zg(z)/G(z)*, the elasticity of the distribution function generating the draws, declining in *z* are, together, sufficient conditions for $d\bar{z}_i / d\pi_i < 0$ and $\xi < 0$, i.e. for average labour efficacy to be declining in a sector's share of total employment.

Equation (a1) above gives the formulas for $\bar{z}_i$ and π_i for a general joint distribution function $g_{i,j}(z_i,z_j)$ determining the paired productivity draws (z_i,z_j). (a1) also notes that these are functions of the endogenous variable $\omega = w_i/w_j$. From this we see that:

$$\text{(d1)}\quad \frac{d\bar{z}_i}{d\pi_i} = \frac{d\bar{z}_i / d\omega}{d\pi_i / d\omega} = \frac{1}{\pi_i}\left[\frac{dN/d\omega}{d\pi_i / d\omega} - \bar{z}_i\right]$$

where *N(ω)* is defined earlier in (a1). As *dN/dω* divided by $d\pi_i/d\omega$ equals $dN/d\pi_i$, which is the quality of the marginal worker, we see intuitively that the condition we are looking for is that the quality of the marginal worker entering the industry is less than that of the average worker. Substituting using the formulas in (a1) we have

$$\text{(d2)}\quad \frac{d\bar{z}_i}{d\pi_i} = \frac{1}{\pi_i}\left[\frac{\int_0^\infty z_i^2 g_i(z_i) g_{j|i}(\omega z_i|z_i)\mathrm{dz_i}}{\int_0^\infty z_i g_i(z_i) g_{j|i}(\omega z_i|z_i)\mathrm{dz_i}} - \frac{\int_0^\infty z_i g_i(z_i) G_{j|i}(\omega z_i|z_i)\mathrm{dz_i}}{\int_0^\infty g_i(z_i) G_{j|i}(\omega z_i|z_i)\mathrm{dz_i}}\right] = \frac{1}{\pi_i}\left[E(a) - E(b)\right]$$

$$\text{where}\quad F_\mathrm{a}(x) = \frac{\int_0^x \eta(\omega z_i) g_i(z_i) G_{j|i}(\omega z_i|z_i)\mathrm{dz_i}}{\int_0^\infty \eta(\omega z_i) g_i(z_i) G_{j|i}(\omega z_i|z_i)\mathrm{dz_i}},\quad F_b(x) = \frac{\int_0^x g_i(z_i) G_{j|i}(\omega z_i|z_i)\mathrm{dz_i}}{\int_0^\infty g_i(z_i) G_{j|i}(\omega z_i|z_i)\mathrm{dz_i}},\quad \eta(\omega z_i) = \frac{\omega z_i\, g_{j|i}(\omega z_i|z_i)}{G_{j|i}(\omega z_i|z_i)}$$

and where I have redefined the terms in [] as the difference between the expectation of two random variables with cumulative density functions $F_a(x)$ and $F_b(x)$. As is well known, if $F_a(x) \geq$

$F_b(x)$ for all x, then $E(a) \le E(b)$.[39] Note that $F_a(x)$ is the same as $F_b(x)$ except for the weighting function η. If z_i and z_j are independent, then η becomes

$$\text{(d3)}\quad \eta(z) = \frac{z g_j(z)}{G_j(z)}$$

which is the elasticity of the distribution function. If this is non-increasing in its argument, then $F_a(x) \ge F_b(x)$ for all x[40] and $E(a) \le E(b)$. Strict inequality follows if η is strictly decreasing.[41] Note that this is a sufficient but not necessary condition, as $E(a) < E(b)$ does not imply $F_a(x) > F_b(x)$ for all x.

Figure A1: Necessity for Additional Restrictions on the Cumulative Distribution Function

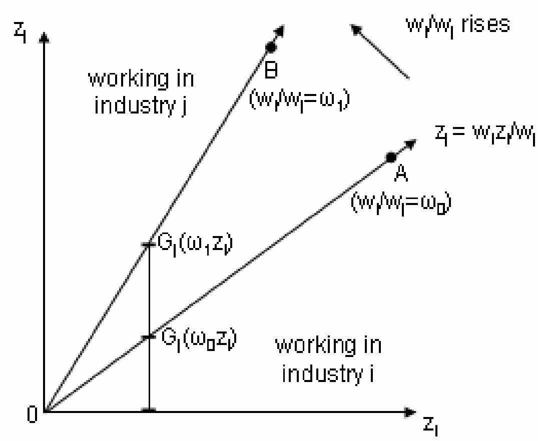

Figure A1 provides some intuition as to why independence alone is not sufficient to guarantee $\xi < 0$ and an additional condition on $zg(z)/G(z)$, such as that specified above, is needed. Individual talent is distributed across the (z_i,z_j) space depicted in the diagram. The ray $z_j = w_i z_i/w_j$ determines the division between sectors, with workers with (z_i,z_j) draws above the ray working in industry j and those with draws below the ray working in sector i. Initially, workers below ray $\overrightarrow{OA}$ work in industry i, but as w_i/w_j rises from ω_0 to ω_1 workers in the region encompassed by the rays $\overrightarrow{OA}$ and $\overrightarrow{OB}$ shift to the sector. The average quality of pre-existing sector i workers depends on the z_i weighted integral of the joint density in the area below $\overrightarrow{OA}$, while the quality of marginal workers depends upon the z_i weighted integral of the joint density in the area between $\overrightarrow{OA}$ and $\overrightarrow{OB}$. Even if z_i and z_j are independent, it is possible for the marginal worker to be of higher quality if the ratio $[G_j(\omega_1 z_i)- G_j(\omega_0 z_i)]/ G_j(\omega_0 z_i)$ (the relative cumulative density for the z_j draws) rises with z_i in some regions, i.e. more relative

[39] As $\int_0^\infty x f(x)\mathrm{d}x = \int_0^\infty f(x)\int_0^x 1\mathrm{d}t\mathrm{d}x = \int_0^\infty \int_t^\infty f(x)\mathrm{d}x\,\mathrm{d}t = \int_0^\infty [1-F(t)]\mathrm{d}t.$

[40] Note that $F_a(x) = A/(A+1)$ and $F_b(x) = B/(B+1)$ where

$$A = \frac{\int_0^x \eta(\omega z_i)h(z_i)\mathrm{dz_i}}{\int_x^\infty \eta(\omega z_i)h(z_i)\mathrm{dz_i}} \ge \frac{\eta(\omega z_i)\int_0^x h(z_i)\mathrm{dz_i}}{\eta(\omega z_i)\int_x^\infty h(z_i)\mathrm{dz_i}} = \frac{\int_0^x h(z_i)\mathrm{dz_i}}{\int_x^\infty h(z_i)\mathrm{dz_i}} = B \text{ and } h(z_i) = g_i(z_i)G_{j|i}(\omega z_i \mid z_i).$$

[41] Thus, for uniform distributions on [a,b], where η is a constant if a = 0, $d\bar{z}_i / d\pi_i = 0$ for some values of ω.

weight is placed on higher values of z_i in the marginal worker integral. Thinking of ω_1/ω_0 as the same relative change applied for each z_i, avoiding this everywhere amounts to an elasticity restriction on G_j. The condition is sufficient, but not necessary, because it is possible for $[G_j(\omega_1 z_i)\text{-} G_j(\omega_0 z_i)]/ G_j(\omega_0 z_i)$ to be rising in some areas and falling elsewhere and yet, depending upon the distribution of z_i, for the average quality of the marginal worker to still be lower than that of pre-existing workers.

(e): The range of prices supported by the supply curve in the standard Cobb-Douglas model with unequal sectoral factor intensities (footnote in the Introduction).

For the standard Cobb-Douglas model with homogenous labour and production functions $Q_i = A_i K_i^{\alpha_i} L_i^{1-\alpha_i}$, the first order conditions for the optimal use of labour and capital imply:

$$\text{(e1)}\quad w = P_i(1-\alpha_i)A_i(K_i/L_i)^{\alpha_i} \quad \text{and} \quad r = P_i\alpha_i A_i(K_i/L_i)^{\alpha_i-1}$$

where w is the wage and r the rental. From this it follows that

$$\text{(e2)}\quad \frac{K_i}{L_i} = \frac{\alpha_i}{1-\alpha_i}\frac{w}{r}\ ,\ \text{so}\quad P_i = \frac{r(K_i/L_i)^{1-\alpha_i}}{\alpha_i A_i} = \frac{r(w/r)^{1-\alpha_i}(\alpha_i/1-\alpha_i)^{1-\alpha_i}}{\alpha_i A_i}$$

Consequently:

$$\text{(e3)}\quad \frac{P_i}{P_j} = \frac{A_j}{A_i}\left(\frac{w}{r}\right)^{\alpha_j-\alpha_i}\frac{\alpha_j(\alpha_i/1-\alpha_i)^{1-\alpha_i}}{\alpha_i(\alpha_j/1-\alpha_j)^{1-\alpha_j}} = \frac{A_j}{A_i}\ \text{if}\ \alpha_i = \alpha_j$$

The last equality simply notes that if the factor income shares are identical, the standard model yields a horizontal Baumol supply curve.

Focusing on the first equality in (e3), we see that, holding constant the productivities A_i and A_j, the equilibrium variation in relative prices depends upon the equilibrium variation in w/r. The question is what range of variation in w/r is possible given constant total factor productivities and a constant endowment of capital and labour. Let sector j be the sector with the higher capital intensity ($\alpha_j > \alpha_i$), and note that in equilibrium it must be the case that $K_j/L_j \geq K/L \geq K_i/L_i$, i.e. the economy-wide capital-labour ratio must lie between the two sectoral capital-labour ratios. From (e2) this implies

$$\text{(e4)}\quad \frac{\alpha_j}{1-\alpha_j}\frac{w}{r} \geq \frac{K}{L} \geq \frac{\alpha_i}{1-\alpha_i}\frac{w}{r} \quad \text{or} \quad \frac{1-\alpha_i}{\alpha_i}\frac{K}{L} \geq \frac{w}{r} \geq \frac{1-\alpha_j}{\alpha_j}\frac{K}{L}$$

Since the economy-wide capital labour ratio is the weighted average of the sectoral capital labour ratios (with weights equal to their employment shares), as *w/r* moves to its lower bound the output of sector i goes to 0, while as it reaches its upper bound the output of sector j goes to zero. Consequently, as *w/r* moves from its lower to its upper bound the relative output Q_i/Q_j goes from 0 to ∞. This traces out the supply curve. Combining (e3) and (e4) we see that the relative price change associated with this movement is:

$$\text{(e5)}\quad \ln\left[\left(\frac{P_i}{P_j}\right)^{\max} \Big/ \left(\frac{P_i}{P_j}\right)^{\min}\right] = (\alpha_j - \alpha_i)\ln\left[\frac{1-\alpha_i}{\alpha_i} \Big/ \frac{1-\alpha_j}{\alpha_j}\right]$$

In the BLS KLEMS 1987 to 2010 database the average annual capital income shares of value added for aggregate goods and services are $\alpha_G = .35$ and $\alpha_S = .32$, respectively. Plugging these numbers into (e5), we get a variation in the ln relative price of goods to services from the bottom to the top of the supply curve of .03*ln(119/104) = .0040, i.e. 4/10ths of one percent. For all intents and purposes, this is a horizontal supply curve. Thus, operating as if goods and services share the same factor income share provides a very close approximation to the actual relative supply curve generated by their differing factor intensities.

Appendix B: Labour Quality Measures for the BLS KLEMS

As noted in the text, the BLS KLEMS total factor productivity estimates do not differentiate by worker type. For its aggregate private business and private non-farm business TFP measures, however, the BLS constructs measures of differentiated labour input using March Supplement Current Population Survey (CPS) data to construct measures of differentiated labour input and then adjusting the hours totals to match Current Employment Statistics (CES). I use a similar methodology to construct differentiated labour measures for the 60 private sectors in the KLEMS and the government sector.

The first difficulty one encounters lies in matching the industrial sector definitions of the CPS and the KLEMS. From 2003 to 2010, the CPS data uses aggregations of the categories in the 2002 NAICS (North American Industry Classification System), which are a close match to the NAICS categories used in the 60 sector KLEMS. The only exceptions are NAICS 523 (securities, commodity contracts and investments) and 525 (funds, trusts and other financial vehicles), which are separate in the KLEMS but combined in the CPS data. I assume that the distribution of workers by type within the two sectors is the same as in the combined CPS sector. Pre-2003 data, however, are based upon the 1972, 1980 and 1987 SIC (Standard Industrial

Classification) codes. While the differences between one SIC and another are minor, and easily reconciled by renumbering and combining a few detailed sub-categories, the differences between the SIC and the NAICS appear more substantial.

The BLS and I address the issue of changing sectoral definitions in labour statistics using 2000-2002 CPS data. In the 2000, 2001 and 2002 iterations of the CPS, industry and occupation data were collected using both the old and new classification systems. In its published labour statistics, the BLS uses the cross-distribution of employment between old industry and new industry in the dual coded data to convert the old data series to the new industrial definitions (http://www.bls.gov/cps/constio198399.htm). I follow a similar methodology, except that I use the cross distribution from old system industry x occupation categories to new industry. However, there are hundreds of industry and occupation categories, so not every industry x occupation cross-classification present in the 1987-2002 data appears in the 2000-2002 sample. For those missing observations, I use higher levels of aggregation, using first the old system industry x detailed (46 categories) occupation cross-classification, then the old system industry x major (14 categories) occupation cross-classification, and, when all else fails, for a handful of observations, simply the old system industry to new system industry distribution.

The second problem that arises is that of zeros. I cross-classify workers by 61 sectors (60 private plus public administration), 2 sexes, 6 age groups, 5 educational categories, and 24 years.[42] Given the limited samples in the CPS, this inevitably creates lots of zeros. Zeros are a serious problem, as total factor productivity calculations involve calculating ln changes. I address this issue by using iterative proportional fitting (Agresti 1990) to estimate the full five-dimensional cross distribution using sub-dimensional totals. Iterative proportional fitting fits a model that assumes independence at higher dimensions. To illustrate with the three dimensional example where X is cross classified by i, j and k, one can use the observed X_{ij}, X_{jk}, and X_{ik} totals to produce estimates $\hat{X}_{ijk}$ which are ln-linearly related to implicit interaction factors λ_{ij}, λ_{jk}, and λ_{ik}, with no interactions at the i x j x k level. By using sub-dimensional totals to estimate the full array, one eliminates the zeros in the detailed cross-classifications. For my estimates of wages per hour, where the samples are particularly sparse as the data are not available for all workers, I use all of the two dimensional cross-classifications to estimate the five dimensional array. I calculate total hours and total income for each two-dimensional sub-array, iteratively proportional fit the entire five dimensional array, and then take ratios of cells to calculate wages

[42]The age categories are 15-24, 25-34, 35-44, 45-54, 55-64, and 65+; the educational categories are less than high school, completed high school, some college, completed college, and more than college; the years are 1987-2010.

per hour. For my worker and hours data, the samples are larger. I begin by defining 12 major sector aggregations (the principal sectors, with manufacturing sub-divided into durables and non-durables) for the 61 detailed sectors. I then iteratively proportionally fit the five dimensional array using every available three dimensional array based upon major industry classification and every two dimensional array based upon detailed industry classification.[43] The use of major industry aggregations allows me to include interactions at higher dimensions without introducing zeros into cells, while the detailed industry two dimensional arrays retain the information on cross-distributions at that level.

To summarize my procedure, I begin by using the 2000-2002 CPS SIC industry x occupation to NAICS industry population distribution to convert 1987-2002 industry data to 2002 NAICS definitions. I then use the CPS March Supplement individual weights and aggregate to the 60 KLEMS sectors plus the government sector. I treat as a "worker" anyone who reports more than zero hrs of work in the previous week. I then adjust the population totals and hours of work totals by year x industry to match the BLS estimates of workers and hours by year x industry[44] and iteratively proportionally fit workers and total hours to calculate workers and hours by industry x sex x age x education x year. For wages per hour, I take all individuals for which the BLS is able to calculate a wage per hour (based upon the direct report or data on "usual hours"), aggregate into 61 sectors using the CPS weights, adjust hours totals by industry using the BLS CES data, and then iteratively proportionally fit total earnings and hours, taking the ratio of the two five dimensional arrays to calculate wages per hour. The combination of hours and wages per hour then allow me to calculate sub-factor income shares by industry (Θ_{Li}^{j} in the paper) and the data on hours per worker allow me to calculate Tornqvist measures of the growth of labour quality by sector which are comparable to those the BLS calculates for the aggregate private sector:

$$\text{(B1)}\quad \sum_j \left(\frac{\Theta_{Lit}^{j}+\Theta_{Lit-1}^{j}}{2}\right)\ln\left(\frac{H_{it}^{j}}{H_{it-1}^{j}}\right)-\left(\frac{\Theta_{Lit}+\Theta_{Lit-1}}{2}\right)\ln\left(\frac{H_{it}}{H_{it-1}}\right)$$

[43] Thus, allowing D to denote detailed industry, M major industry, S sex, A age, E education and Y year, I use the sub-dimensional arrays DS, DA, DE, DY, MSA, MSE, MSY, MAE, MAY, MEY, SAE, SAY, SEY, and AEY. In iterative proportional fitting, one can aggregate a dimension into sub-categories. As long as that sub-category contains additional cross-distributions, it is not redundant (i.e. MS is redundant given DS, but MSA is not) and provides an additional interaction factor.

[44] The KLEMS TFP database only contains indices of hours. I take levels of hours and workers from the Industry Employment and Hours Data Tables of the BLS labour productivity database. These are not strictly consistent with the hours indices of the BLS KLEMS total factor productivity database. However, I do not use these totals to change the measure of the growth of total labour input (hours) in the KLEMS database calculations, but only to calculate distributions of workers by characteristic, as shown shortly in (B1) and (B2).

where H_{it}^{j} denotes total hours of worker type j in industry i at time t and H_{it} denotes total hours in sector i at time t. The measures are added to the growth of labour input and subtracted from the growth of total factor productivity in the BLS data. The data on the distribution of the population by worker characteristic then allow me to calculate weighted and unweighted Tornqvist measures of the changing shares of the labour force:

$$\text{(B2)}\quad \sum_{j}\left(\frac{\Theta_{Lit}^{j}+\Theta_{Lit-1}^{j}}{2}\right)\ln\left(\frac{\pi_{it}^{j}}{\pi_{it-1}^{j}}\right) \quad\text{and}\quad \left(\frac{\Theta_{Lit}+\Theta_{Lit-1}}{2}\right)\ln\left(\frac{\pi_{it}}{\pi_{it-1}}\right)$$

where π_{it}^{j} denotes the share of the aggregate working population of type j in industry i at time t and π_{it} denotes the share of the aggregate working population in sector i at time t. These measures are used as the instrumented dependent variable in Section III. Since everything is benchmarked to the BLS totals, the H_{it} and π_{it} measures are simply the original BLS data and are consistent with the totals of H_{it}^{j} and π_{it}^{j} across j. The two measures in (B2) are different, but highly correlated, with a correlation coefficient of .917.

Appendix C: Existing Micro-Data Estimates (McLaughlin & Bils 2001)

McLaughlin & Bils (2001, tables 4 and 5) using PSID data from 1979 to 1992 report that the average ln wage of industry leavers relative to stayers in industries with contracting employment shares and industry entrants relative to stayers (continuing workers) in industries with expanding employment shares is about -16 or -17 percent without adjustment for worker characteristics and -6 or -7 percent with adjustment for worker characteristics. These estimates might lead one to conclude that comparative advantage is indeed aligned with absolute advantage, but that Roy worker efficacy effects are rather small. In this appendix I show that the data examined by McLaughlin & Bils have little to do with the expansion and contraction of industries and are mostly related to a form of “churning” whereby workers simultaneously exit and enter industries.

I work with the annual 1971-1997[45] records of the PSID, using both the low income sample and the census based representative sample, focusing on the industry of employment of household heads. I use two industrial classifications: (a) 9 aggregate sectors, a measure which should eliminate spurious industry shifts brought about by minor errors and misclassifications; (b) 24 sectors, which is the greatest detail I can achieve while keeping industry definitions

[45]Prior to 1971 industry is not reported; after 1997 the PSID moves to a biannual framework and hence the calculation of movers and stayers is not comparable.

Table C-1: Entry, Exit and Sectoral Growth in the PSID (observations are industry x year)									
		Representative & Low Income Sample				Representative Sample Alone			
		7 Private Sectors		22 Private Sectors		7 Private Sectors		22 Private Sectors	
		Rates	ln Rates	Rates	ln Rates	Rates	ln Rates	Rates	ln Rates
Correlation Between Entry and Exit Rates									
General (p-value)		.675 (.000)	.717 (.000)	.842 (.000)	.759 (.000)	565 (.000)	.577 (.000)	.747 (.000)	.664 (.000)
Partial (p-value)		.282 (.000)	.167 (.025)	.302 (.000)	.128 (.002)	.224 (.002)	.118 (.114)	.229 (.000)	.106 (.012)
N		182	181	572	571	182	180	572	564
Regression on Change in Industry Employment Share (with industry & year dummies)									
Entry	$\Delta\pi_{it}$	.165	1.07	.217	1.11	.175	.828	.237	.751
	(s.e.)	(.099)	(.506)	(.085)	(.410)	(.132)	(.711)	(.111)	(.544)
	N	182	182	572	572	182	182	572	567
Exit	$\Delta\pi_{it}$	.061	.094	.025	-.077	-.125	-.545	.017	-.245
	(s.e.)	(.099)	(.507)	(.081)	(.415)	(.125)	(.706)	(.105)	(.569)
	N	182	181	572	571	182	181	572	569

Notes: N = number of industry x year observations. An occasional observation is lost when taking the ln of a zero entry or exit rate. Partial correlation = correlation of residuals from regression on industry and year dummies. Regressions = regression of entry or exit rates on industry & year dummies and the change in the share of non-agricultural employment ($\Delta\pi_{it}$).

relatively consistent across the years, and is similar to the detail used by McLaughlin & Bils.[46] For a given industry i, examined in period t, workers are classified as stayers if they were in the same industry i in period t-1, entrants if they were in a different industry j in period t-1, and leavers if they worked in industry i in period t-1 but work in industry j in period t. To be in the sample a worker needs to both report industry of employment and allow the calculation of ln wage per hour in consecutive years. This eliminates unknown industry and workers who were completely out of employment in one year or the other. Every worker who is an entrant in

[46] 9 sectors: 1 agriculture, forestry & fishing; 2 mining; 3 manufacturing; 4 construction; 5 transport, communications & utilities; 6 wholesale & retail trade; 7 finance, insurance and real estate; 8 other services (except gov't); and 9 government & armed forces. 24 sectors: 1 agriculture, forestry & fishing; 2 mining; 3 metal industries; 4 machinery (inc. electrical); 5 motor vehicles & other transportation equipment; 6 food & kindred products (inc. tobacco); 7 textile mill products, apparel & other fabricated textile products, plus shoes; 8 paper & allied products; 9 chemical & allied products, petroleum & coal products, and rubber & misc. plastic products; 10 printing & publishing; 11 other manufacturing; 12 construction; 13 transportation; 14 communication; 15 public utilities; 16 wholesale trade; 17 retail trade; 18 finance, insurance and real estate; 19 business services; 20 personal services; 21 health; 22 education; 23 other services (except gov't); 24 government & armed forces.

industry i in period t is a leaver from some industry j in period t-1. Although I use all 9 or 24 sectors to categorize workers, I focus on entry/exit rates in the 7 or 22 private non-agricultural sectors.[47] Overall I have about 61500 individual x year observations (a little over half in representative sample households) in these industries, with about 15% of these being entrants or leavers according to the broad sectoral definitions and 23% according to the narrow sectoral definitions.

I begin by reporting, in the top panel of Table C-1, the correlation between the sample fractions, at the industry x year level, of entrants (in entrants and stayers) and leavers (in leavers and stayers). As shown, there is a very strong positive correlation between the fraction of the sample that enters an industry between period t-1 and t and the fraction that leaves the same industry between the same two periods. This holds true whether the measures are in levels or in lns, using both the low income and representative sample or just the representative sample alone. The partial correlation of the entry and exit rates, after removing industry and year fixed effects, is weaker but still generally highly significant. In contrast, in the bottom panel of the table I report the regression of the industry entry and exit rates on the change in the sector's share of non-agricultural employment, with industry and year fixed effects. As shown, the regression coefficients are almost universally insignificant, the only exception being entry rates for the 22 industry measure, and this result largely disappears when the sample is restricted to PSID representative households alone.

Table C-2 follows the McLaughlin & Bils methodology, examining the average relative ln wages of different groups. Without adjustment for worker characteristics, the wages of entrants or leavers are found to be between 11 and 19 percent lower than those of stayers (first four columns). With adjustment for worker characteristics (last four columns), these mean differences are greatly reduced and, in many cases, rendered statistically insignificant. Moreover, in all cases the vast majority of the estimates that underlie the calculation of these averages are insignificant. Thus, for example, while the relative wages of entrants to stayers in expanding sectors are on average 3.1 percent (7 sectors) or 2.9 percent (22 sectors) lower among the representative PSID sample, only about 1/10th of the industry x year differences that underlie the calculation of these means are, by themselves, statistically significant at the 5% level. Unlike McLaughlin & Bils, Table C-2 reports relative wages in both expanding and contracting sectors for all measures. As shown, while the relative wages of entrants are lower than stayers in

[47] I relate these rates to the BLS Current Employment Statistics based historical SIC measures of employment, which exclude agriculture, while the focus on private sector activity is consistent with the measures examined earlier in the paper.

Table C-2: Mean Wage Differences Between Industry Entrants or Leavers vs. Stayers in the PSID								
	Average ln Wages Differences				Adjusted for Worker Characteristics			
	7 Private Sectors		22 Private Sectors		7 Private Sectors		22 Private Sectors	
	$\pi_{it} \uparrow$	$\pi_{it} \downarrow$	$\pi_{it} \uparrow$	$\pi_{it} \downarrow$	$\pi_{it} \uparrow$	$\pi_{it} \downarrow$	$\pi_{it} \uparrow$	$\pi_{it} \downarrow$
Entrants vs. Stayers (representative & low income PSID sample)								
Mean Dif.	-.141	-.189	-.128	-.161	-.031	-.065	-.033	-.050
(s.e.)	(.014)	(.014)	(.010)	(.010)	(.009)	(.010)	(.007)	(.007)
Significant/N	27/90	44/92	65/241	84/331	12/90	25/92	25/241	41/331
Leavers vs. Stayers (representative & low income PSID sample)								
Mean Dif.	-.144	-.164	-.128	-.134	-.012	-.032	-.008	-.019
(s.e.)	(.014)	(.015)	(.011)	(.010)	(.010)	(.011)	(.008)	(.007)
Significant/N	25/92	32/89	53/243	68/328	7/92	6/89	21/243	12/328
Entrants vs. Stayers (representative PSID sample)								
Mean Dif.	-.140	-.154	-.118	-.117	-.031	-.059	-.029	-.022
(s.e.)	(.019)	(.019)	(.014)	(.014)	(.012)	(.014)	(.009)	(.010)
Significant/N	18/90	34/91	34/241	59/326	10/90	20/91	22/241	35/326
Leavers vs. Stayers (representative PSID sample)								
Mean Dif.	-.120	-.162	-.108	-.112	-.007	-.045	-.004	-.008
(s.e.)	(.019)	(.019)	(.015)	(.014)	(.014)	(.013)	(.010)	(.011)
Significant/N	11/92	28/89	29/242	50/327	7/92	9/89	21/242	28/327

Notes: Observations are industry x year measures of wage differences. Adjusted for Worker Characteristics = the coefficients on entrant (or leaver) yearly dummies in industry level regressions with controls for sex, age, age2, race (African-American), education (8 categories) and year (dummies), with random effects for PSID individuals. $\pi_{it} \uparrow$ ($\pi_{it} \downarrow$): observations in industries whose share of total employment increased (decreased) in that year. Mean Dif: mean year x industry difference for observations with $\pi_{it} \uparrow$ or $\pi_{it} \downarrow$; s.e. = standard error of the mean difference; N = number of industry x year observations; Significant = number of such observations which are, individually, significantly different from 0 at the 5% level.

expanding industries, the difference is, generally, even larger in *contracting* industries. Similarly, while the relative wages of leavers are lower than stayers in contracting industries, the difference is generally almost as large in *expanding* industries. These results completely undermine the interpretation of these wage differences as reflecting the relative efficacy of entrants in expanding industries and leavers in contracting industries.

In sum, changes of industrial sector in the PSID appear to reflect a form of "churning", whereby both entry and exit simultaneously occur within industries. It is not hard to motivate such movement, either with models of creative destruction within sectors or with a more general idiosyncratic destruction of existing jobs and appearance of new opportunities. Workers with systematically lower human capital appear to play a disproportionate role in this churning, as adjustment for observable characteristics eliminates most of the relative wage differences. While these facts are interesting in and of themselves, they provide little insight into the impact of the expansion or contraction of industry employment shares on average worker efficacy.