Topic 2: Empirical Evidence on Cross-Country Growth

Over the last few weeks, we examined the dynamics of output per worker in the Solow model using the decomposition

$$\frac{Y_t}{L_t} = A_t^{\frac{1}{1-\alpha}} x_t^{\frac{\alpha}{1-\alpha}} \tag{1}$$

where

$$x_t = \frac{K_t}{Y_t} \tag{2}$$

We will now use this framework to examine differences across countries in the level of output per worker as well as its growth.

Decomposing Levels of Output Per Worker

Equation (5) provides an accounting framework for explaining why some countries are rich and others are poor. Output per worker depends positively on the level of total factor productivity (TFP) and also on the capital-output ratio. Which of these factors matters most in determining which countries are rich and which are poor. One simple way of assessing this question is to draw a scatter plot of output per worker in a range of countries against the levels of TFP and the capital-output ratio.

Look at the first chart at the back of the handout. It uses the Penn World Tables. This dataset contains post-War data on PPP-equivalent GDP for a very wide range of countries, incorporating almost all of the global economy. It was first released in the 1980s as the result of detailed work by economists Robert Summers and Alan Heston and it has proved to be a major impetus for research on cross-country growth patterns. This chart compares the levels of output per worker and TFP in the year 2000 in 96 countries. Note that calculating TFP requires an assumption about α and these charts are based on the standard assumption that $\alpha = \frac{1}{3}$. The chart shows that there is a very strong correlation between output per worker and TFP. The second chart shows that there is also a positive correlation between output per worker and the capital-output ratio, but that this is a much weaker correlation.

One factor that could influence this calculation is that workers in rich countries tend to be more educated than workers in poor countries. One way to measure this is to construct measures of *human capital* based on estimates of the return to education. An important paper that did these calculations and used them to shed light on cross-country income

differences is the paper on the reading list by Hall and Jones (1999). The basis of the study is a "levels accounting" exercise that starts from the following production function

$$Y_i = K_i^{\alpha} \left(h_i A_i L_i\right)^{1-\alpha} \tag{3}$$

where h_i is the average level of human capital per worker. Note the slightly different treatment of TFP in this equation than the one we have been working with. The Hall-Jones production function can be re-expressed using our usual formulation as

$$\frac{Y_i}{L_i} = \left(\frac{K_i}{Y_i}\right)^{\frac{\alpha}{1-\alpha}} h_i A_i \tag{4}$$

Hall and Jones then constructed a measure h_i using evidence on levels of educational attainment and they also set $\alpha = 1/3$. This allowed them to use (4) to express all cross-country differences in output per worker in terms of three multiplicative terms: capital intensity, human capital per worker, and technology or total factor productivity. They found that output per worker in the richest five countries was 31.7 times that in the poorest five countries. This was explained as follows:

- Differences in capital intensity contributed a factor of 1.8.
- Differences in human capital contributed a factor of 2.2
- The remaining difference—a factor of 8.3—was due to differences in TFP.

The results from this paper show that differences in total factor productivity, rather than differences in factor accumulation, are the key explanation of cross-country variations in income levels.

Evidence on Growth: Do We See Convergence?

One of the key questions in macroeconomics is whether the international distribution of income displays *convergence*: Do poor countries tend to catch up with rich ones? There are some good reasons to expect that such a pattern may exist. For example, poor countries may be able to learn from rich countries, increasing their income levels by copying production techniques originally invented elsewhere. In contrast, there may be forces that work in the opposite direction. For instance, a common theme in development economics is that some countries can become stuck in "poverty traps" in which low income levels that only provide

for subsistence lead to low levels of savings and investment, thus limiting opportunities for growth.

Given that there are theoretical arguments on both sides, the strength of the forces for convergence is ultimately an empirical question. In the past, the debate about convergence has been complicated by data issues. For instance, while evidence was sometimes reported that there seemed to be some convergence in income levels among those countries that were well-off in the post-War period, others pointed out that reliable historical national income statistics tended only to be available in richer countries. This had the effect of biasing results in favour of finding that poor countries tend to catch up with rich ones.[1] Comparing real income levels across countries has also generally been tricky, because incomes are denominated in different currencies and it can be hard to compare absolute price levels across countries (i.e. calculation of Purchasing Power Parity equivalent incomes is hard).

Against this background, the publication in the 1980s of the Penn World Tables dataset, incorporating almost all of the global economy, was very valuable. With over 50 years of PWT data now available, the evidence on convergence is fairly clear: There is little evidence that poor countries tend to catch up, and some facts pointing the other way. The third figure at the back of this handout shows the growth in real PPP-adjusted GDP over the period 1960-2000 graphed against the value of real GDP in 1960. If convergence was operating, we would see a negative relationship with higher growth for countries that started out poorer. In fact, there is almost no correlation between growth and initial income levels. Other research, which attempts to dig back before 1950, points towards divergence in real incomes over longer time horizons.[2]

One caveat to these findings is reported in an important paper by Xavier Sala-i-Martin, "The World Distribution of Income: Falling Poverty and ... Convergence, Period" in the May 2006 *QJE*. He reports that when examined at the level of *individuals*, there has been a tendency for poor people's incomes to catchup. This is partly due to the fact that China and India have performed relatively well, and while these account for only two datapoints in our figure, they contain huge amounts of people. Still, it remains the case that, in general, a randomly-selected poor country should not be expected to systematically catch up with

[1] I am thinking here of the debate between William Baumol and Brad DeLong in the following articles in the *AER*: "Productivity Growth, Convergence and Welfare: What the Long Run Data Show" by Baumol (December 1986), with a reply by DeLong and rejoinder by Baumol and Edward Wolff in the December 1988 edition.

[2] See Lant Pritchett, "Divergence, Big Time," *Journal of Economic Perspectives*, Summer 1997.

income levels of richer countries.

Solow and Conditional Convergence

So why don't we see convergence? And should we really have expected to? One of the most important contributions to the debate on empirical growth was the 1992 *QJE* paper by Mankiw, Romer and Weil (henceforth, MRW). Where much of the previous research in this area had been descriptive or purely statistical, MRW argued for using the Solow model as an analytical framework for interpreting long-run growth patterns across countries. The model can be used to define the concept of *conditional convergence*, which we will now explain.

Consider a world in which there are separate countries, indexed by i, each having a Cobb-Douglas production function, so they can be written as:

$$\frac{Y_{it}}{L_{it}} = A_{it} x_{it}^{\frac{\alpha}{1-\alpha}} \tag{5}$$

Letting lower-case letters represent logged variables, this can be re-written as

$$y_{it} - l_{it} = a_{it} + \frac{\alpha}{1-\alpha} x_{it}. \tag{6}$$

So the dynamics of output per worker are given by

$$\Delta y_{it} - \Delta l_{it} = \Delta a_{it} + \frac{\alpha}{1-\alpha} \Delta x_{it}. \tag{7}$$

The steady-state path for output per worker is the level of output per worker consistent with the capital-output ratio being equal to its steady-state level, which we will denote as time-varying because determinants such as the investment rate can change over time:

$$y_{it}^* - l_{it} = a_{it} + \frac{\alpha}{1-\alpha} x_{it}^*. \tag{8}$$

So, the gap between output and its steady-state level can be written as

$$y_{it} - y_{it}^* = \frac{\alpha}{1-\alpha} \left(x_{it} - x_{it}^* \right) \tag{9}$$

Remember now that the capital-output ratio tends to close a fraction λ of the gap relative to the its steady-state level. In this case, this can be written as

$$\Delta x_{it} = \lambda \left(x_{i,t-1}^* - x_{i,t-1} \right) \tag{10}$$

Using equation (10), output per worker dynamics can then be expressed as

$$\Delta y_{it} - \Delta l_{it} = \Delta a_{it} + \lambda \left(y^*_{i,t-1} - y_{i,t-1} \right). \tag{11}$$

The convergence hypothesis can be formulated as a relationship between growth in output per worker and its past levels: Convergence obtains when the past level has a negative impact on subsequent growth. Equation (13) shows that the Solow model does indeed predict a negative coefficient on past levels of output. Crucially, however, it only predicts that such a relationship obtains once one has *conditioned on* the growth rate of technology and the steady-state level of output per worker. This idea of conditional convergence has received a lot of attention in the literature on empirical growth.

MRW's paper assumed that technology levels across countries could be characterised as

$$a_{it} = gt + \epsilon_i \tag{12}$$

where ϵ_i is a random variable. This assumption allows for random differences in the *level* of technology across countries, but assumes that technology grows at the same rate everywhere. MRW (page 410) defended this assumption on the basis that "g primarily reflects the advancement of knowledge and this is not country-specific." In this case, equation (13) can be simplified to be

$$\dot{y}_{it} - \dot{l}_{it} = g + \lambda \left(y^*_{i,t-1} - y_{i,t-1} \right). \tag{13}$$

and one can expect to see a negative effect of lagged values of output once one has controlled for the factors that determine the steady-state capital output ratio, i.e. the investment share and the rate of growth of the labour force.

MRW estimated regressions of this sort and found that, once one conditioned on the determinants of the capital-output ratio (i.e. cross-country variations in investment rates and population growth rates) then one did indeed find evidence for convergence.[3] In this sense, the results were vindication for the Solow model. However, they also found that the estimated convergence speeds were lower than predicted by the model. Recall from the last handout that the model predicts a convergence speed of

$$\lambda = (1-\alpha)(\frac{g}{1-\alpha} + n + \delta) \tag{14}$$

[3]Specifically, MRW implemented this as a regression to explain growth rates over the period 1960-1985 for a large sample of countries from the PWT dataset. They assumed $\frac{g}{1-\alpha} + \delta = 0.05$ so that the determinants of the steady-state capital-output ratio were $\log(s_i) - \log(n_i + 0.05)$.

Plugging in reasonable values for these parameters, one would expect to observe conditional convergence of about six to seven percent per year. In fact, MRW found estimates of one to two percent. MRW's explanation for this result was that the parameter α might be higher than the standard value of about one-third, which implies a lower convergence speed. They developed a model in which "capital" corresponds not just to physical capital, but also human capital (skills accumulated through education), implying a model that operates like a standard Solow model except with a high value of α.

Constant TFP Growth Across Countries?

One drawback of MRW's approach to estimating conditional convergence is that it relies on a very restrictive assumption about the A_{it} technology or TFP term. How accurate is the assumption that this grows at the same rate everywhere? Rather than making an assumption about technology, an alternative approach is to make reasonable assumptions about the coefficient α, calculate A_{it} for each country based on this assumption, and then examine its behavior. How does A_{it} change over time and are there important variations in this growth across countries?

Look at Figure 3 at the back of McQuinn and Whelan (*Oxford Review of Economic Policy*, 2007)–a link to this paper is on the website. It shows percentage changes in TFP calculated in this manner, over the period 1960-2000 for a sample of 96 PWT countries, and graphs them against the percentage changes in output per worker. Not only are there very substantial variations across countries in TFP growth, but these are highly correlated with the growth in output per worker. Indeed, variations in TFP growth are the dominant source of variation across countries in their growth rates of output per worker, with a positive correlation between these two series of 0.79. In contrast, Figure 4 from this paper shows that there is a slight *negative* correlation between the growth rates of the capital-output ratio and output-per-worker. These calculations use the standard value of $\alpha = \frac{1}{3}$, but similar results can be obtained for a wide range of reasonable values of this parameter. The paper on the reading list by Easterly and Levine also emphasises the importance of TFP rather than capital accumulation, and is well worth reading.

Convergence Estimated via the Capital-Output Ratio

One might expect that the reliance on an inaccurate assumption about TFP may make MRW's method for estimating conditional convergence speeds a bit suspect. In my opinion

it does. Think about it: MRW assume that output is converging everywhere to a path described by a common growth rate for TFP. If TFP is actually growing at different rates around the world, or even if it grows at approximately the same rate everywhere but exhibits random country-specific shocks, then the actual steady-state that countries are converging towards will be quite different from what MRW assume. MRW's finding of slow convergence may just reflect that fact these economies are not really converging towards their assumed steady-state paths at all.

It turns out that estimates of convergence speeds that do rely on the incorrect assumption of constant common TFP growth point to speeds very similar to those predicted by the original Solow model. In research I have carried out with Kieran McQuinn, we use the fact that one can estimate the coefficient λ from the dynamics of the capital-output ratio to provide new empirical estimates.[4] In other words, our approach is to use equation (10) to directly estimate λ. We find estimates in the range of six to seven percent per year, about as predicted by Solow. One important advantage of this approach is that it does not rely on *any* assumption about the process for TFP across countries.

Of course, one can question the importance of the λ parameter. Conditional convergence relates to the dynamics of the capital-output ratio and the McQuinn-Whelan Figure 4 shows that these are not a very important factor in determining long-run economic growth. So this suggests that, despite having received a lot of attention from growth researchers, conditional convergence is not a very important topic for the big picture. Instead, the bigger question that should be examined is what determines TFP. We turn to this question next.

[4] Kieran McQuinn and Karl Whelan. "Conditional Convergence and the Dynamics of the Capital-Output Ratio", *Journal of Economic Growth*, June 2007. Available at www.karlwhelan.com.

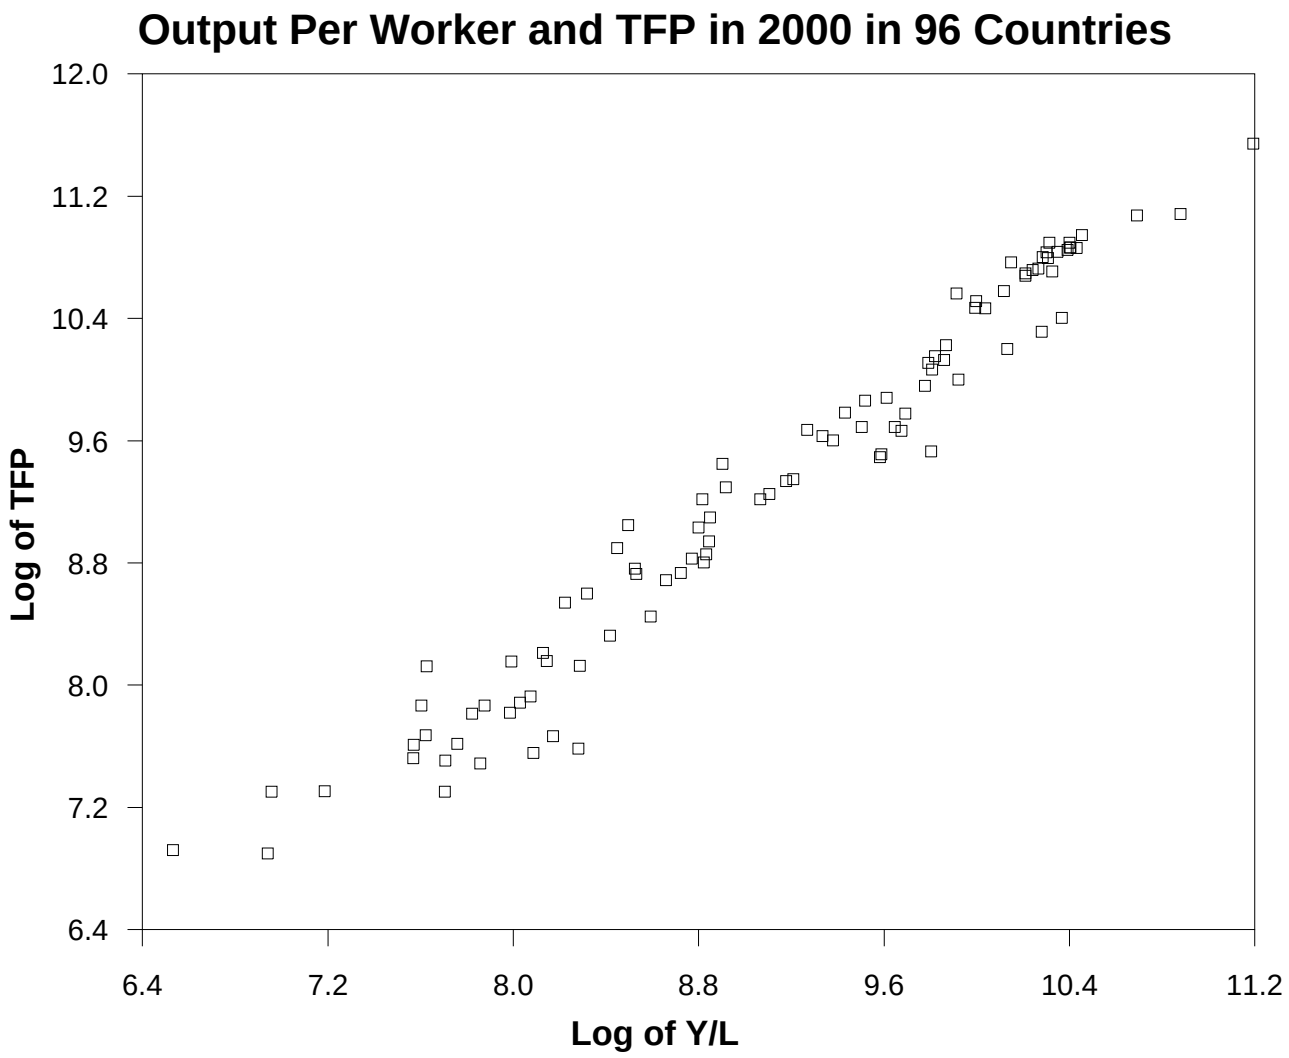
Output Per Worker and TFP in 2000 in 96 Countries
Log of TFP
12.0
11.2
10.4
9.6
8.8
8.0
7.2
6.4
6.4
7.2
8.0
8.8
9.6
10.4
11.2
Log of Y/L

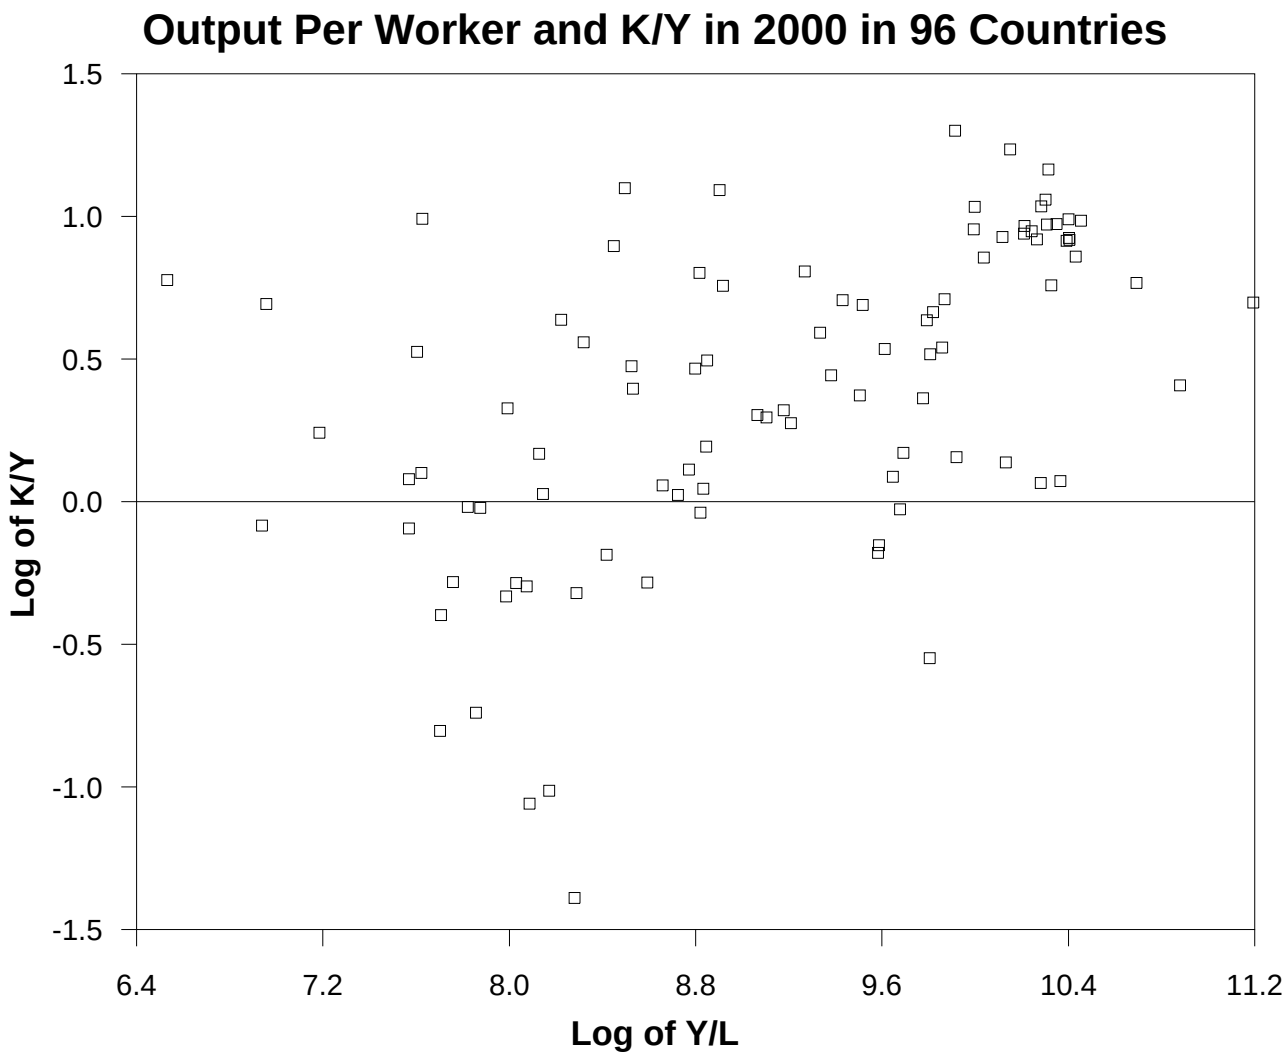
Output Per Worker and K/Y in 2000 in 96 Countries
1.5
1.0
0.5
0.0
-0.5
-1.0
-1.5
Log of K/Y
6.4
7.2
8.0
8.8
9.6
10.4
11.2
Log of Y/L

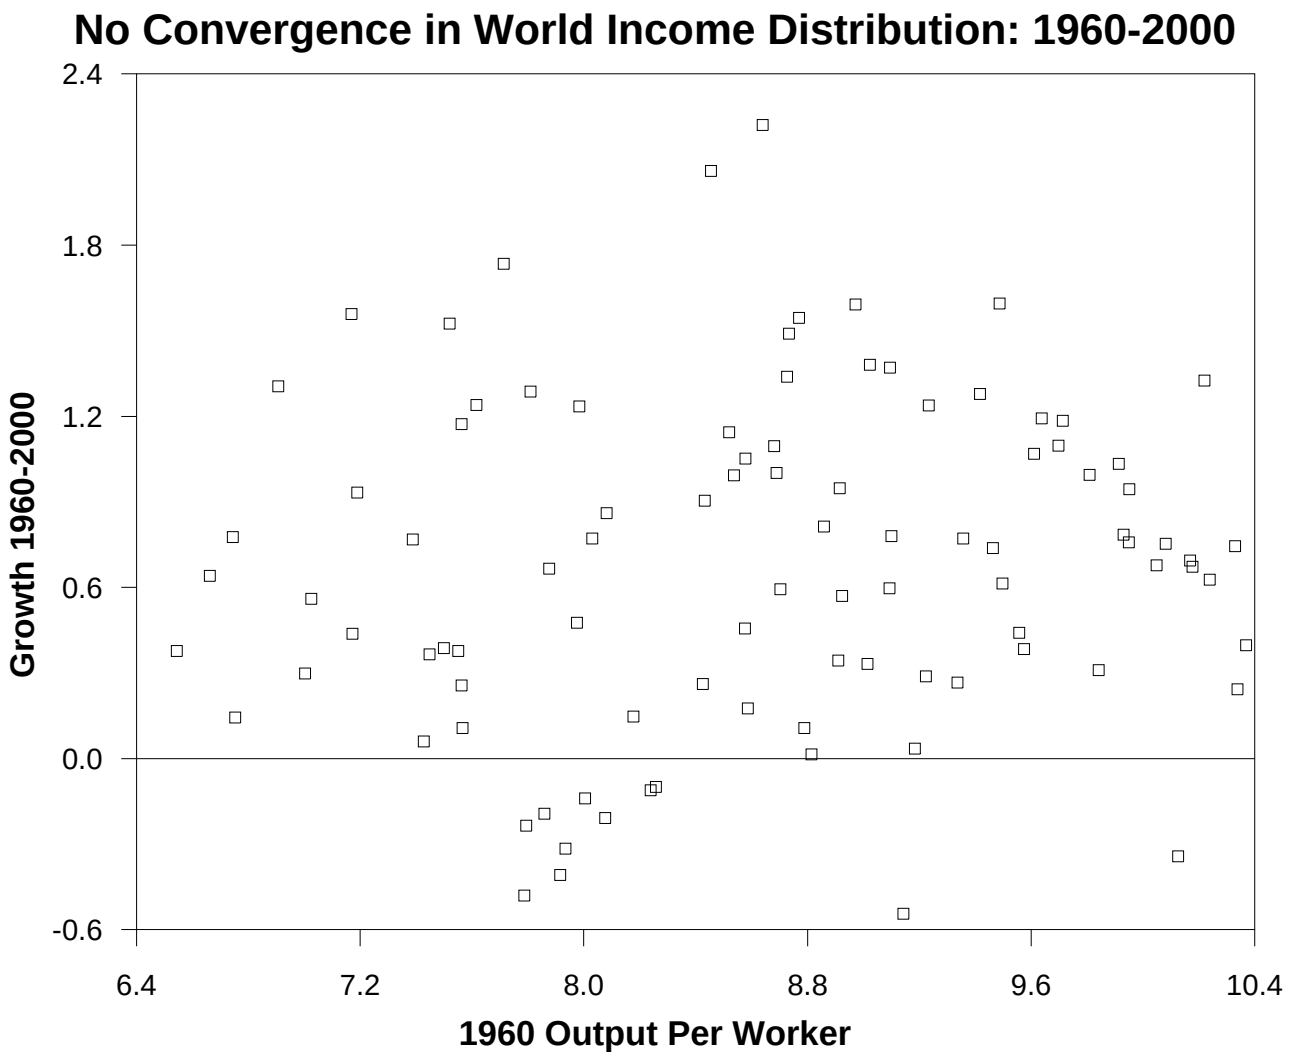
No Convergence in World Income Distribution: 1960-2000
Growth 1960-2000
2.4
1.8
1.2
0.6
0.0
-0.6
6.4
7.2
8.0
8.8
9.6
10.4
1960 Output Per Worker