The implications of growth regressions for equality of opportunity

By Donal O'Neill

Economics Department, National University of Ireland Maynooth, Co. Kildare, Ireland; e-mail: donal.oneill@nuim.ie

In this paper we consider the usefulness of alternative measures of convergence in an equality of opportunity framework. In particular we use established results from the public finance and mobility literature to show that a form of β-convergence is both a necessary and sufficient condition for a reduction in inequality of opportunity for a wide range of popular inequality measures. We illustrate our approach using regional data from the United States, Japan, and Europe.

1. Introduction

The traditional approach to measuring cross-country income convergence involved regressing the growth rate of income on initial income to examine whether or not poor countries grew faster than richer countries (e.g. Barro and Sala-i-Martin, 1992; Sala-i-Martin 1996a). This approach has been the subject of much debate and has been criticized by many. At a fundamental level a number of authors, including Friedman (1992), Quah (1996), and Cannon and Duck (2000), argue that by itself, faster growth among poor countries, or β-convergence as it has become known, tells us little about the evolution of incomes. Friedman (1992), quoting Hotelling (1933), argues that 'the real test of a tendency towards convergence would be in showing a constant decline in the variance ... among individual enterprises.' In the growth literature this type of convergence has become known as σ-convergence.[1,2]

---

[1] For a recent discussion of the relationship between these two types of convergence, see Furceri (2005).

[2] Other critics of traditional growth regressions include Quah (1996), who argues that the speed of convergence estimated from growth regressions may simply reflect small-sample biases. However, he later acknowledges that the degree of precision reported in standard growth regressions casts doubt on this explanation. Lee *et al.* (1997) discuss the econometric problems that arise when using growth regressions to estimate the structural parameters of a growth model. Although important, this issue is distinct from the question addressed in our paper.

However, all of this debate has focused on the distribution of final outcomes. A number of economists have argued that greater emphasis should be placed on opportunities rather than observed outcomes (e.g. Fleurbay, 1995; Roemer, 1998, 2006). The equal-opportunity framework stresses the link between the opportunities available to an agent and the initial conditions which are inherited or beyond the control of these agents. For countries or regions one may be interested in knowing to what extent future opportunities are determined by initial income levels. Proponents of equality of opportunity accept inequality of outcomes that arise from choice or shocks that are independent of initial conditions.[3]

The question that we address in this paper is whether growth regressions can contribute in a meaningful way to studies that focus, not on the evolution of realized outcomes, but rather on the equality of opportunity across regions. We show that appropriate consideration of nonlinearities in the growth process is not only desirable when documenting the evolution of income over time but is also an essential component of a coherent equal-opportunity based framework. In particular, we use established results from the public finance and mobility literature to show that a form of β-convergence is both a necessary and sufficient condition for a reduction in inequality of opportunity for a wide range of popular inequality measures. We also show that this approach can be extended to facilitate welfare comparisons in ways that are not possible using some of the alternative convergence concepts that have been proposed. We examine the implications of our findings for a number of regional data sets previously examined in the growth literature.

2. Progressivity, growth, and equality of opportunity

Fields and Ok (2001) provide a comprehensive survey of the growing literature on the meaning and measurement of income mobility. A considerable portion of this literature has focused on the use of transition probabilities to measure mobility. A transition probability, $M_T(x|y)$, specifies the probability that a country with income y today will have at most income x at time T. In the growth literature associated transition matrices are almost always presented as descriptive tools for understanding the evolution of observed incomes over time (e.g. Quah, 1993). However, in his survey of welfare theoretic approaches to the measurement of mobility, Maasoumi notes that 'Mobility in any social hierarchy is an indication of opportunity' (Maasoumi, 1998, p.1). Benabou and Ok argue that many people care about mobility 'not because income movements are intrinsically valuable, but primarily because of the hope that it helps attenuate the effects of disparities in initial endowments on future income prospects' (Benabou and Ok, 2001, p.2). This view is common in welfarist studies of mobility, which tend to associate the origin-independence principle used when measuring economic mobility with the notion

---

[3]For further discussion see Fields (2004).

of equality of opportunity. Benabou and Ok thus define a 'process [as opportunity] equalising when it leads to ex-ante income prospects that are more evenly distributed than initial incomes or endowments' (Benabou and Ok, 2001, p.6). To study equality of opportunity they abstract from agents' aversion to risk and summarize future opportunities using conditional expectations.[4] In particular the opportunity set available at time T to an agent with current income y is summarized by $e_T(y)$, the expected income at time T given initial income y. Formally:

$$e_T(y) = \int_0^\infty x\,dM_T(x|y) \tag{1}$$

Given this framework they derive conditions on transition matrices by which a mobility process can be characterized as opportunity-equalizing, as well as establishing criteria to determine if one process is more equalizing than another.

While Benabou and Ok (2001) characterize future opportunities using the underlying transition process, M, it is relatively straightforward to recast their results in terms of the underlying growth process. We adopt a general formulation in which regions draw a specific growth path from a set of possibilities summarized by $G_i(.)$. This set of possible growth paths is in turn determined by the regions's initial income draw.[5] The potential incomes available to region i at time T can then be summarized by the random variable Y_T^i, where,

$$Y_T^i = Y_0^i + G_i(Y_0^i) \tag{2}$$

Y_0^i is the random variable summarizing initial incomes and $G_i(Y_0^i)$ reflects the growth opportunities available to region i.

Some regions make choices that, *ex post*, lead to high growth rates; other regions make choices that are not so good. However, what matters in an equal-opportunity framework is the set of choices available to a region at time 0. Assuming that the choice of growth path is only constrained by initial income, ex-ante income prospects can be summarized using conditional expectations. Let $g(y_0) = E(G_i(Y_0^i| = y_0)) = E(Y_{Ti} - Y_{0i}|Y_{0i} = y_0)$ denote the expected change in income conditional on initial income level y_0. The opportunities available at time T, to a region with initial income y_0, can then be summarized by:

$$e_T(y_0) = y_0 + g(y_0) \tag{3}$$

This formulation allows different regions to choose different realizations of the growth schedule, while recognizing that the opportunities available as of time T are characterized by the expected schedule $g(.)$. Following Benabou and Ok (2001)

----

[4]This is discussed in more detail in their paper. For a general discussion of the issues that arise when evaluating opportunity sets see Sen (1985).

[5]Extending the model to allow for circumstances where the choice of the growth path is constrained by factors that are not captured by initial income requires a multi-dimensional approach which is not considered in the current paper.

we consider only monotone growth processes, such that for any y_1, y_2 with $y_2 > y_1$ then $g(y_2) > g(y_1)$. This implies $e_T(y_2) > e_T(y_1)$.

Writing the model in this way allows us to draw close parallels between the income convergence literature and the public economics literature on tax/benefit progressivity (see Lambert, 1993). Following the tax literature we define a growth process as progressive if $d(g(y_0)/y_0)/dy_0 < 0$, regressive if $d(g(y_0)/y_0)/dy_0 > 0$ and proportional if $d(g(y_0)/y_0)/dy_0 = 0$. Intuitively a growth process is progressive if low income countries experience faster expected growth rates than higher income countries.

We can now ask the following question: suppose we have a group of regions, indexed by i, with initial realized incomes y_0^i. Under what conditions on g will there be less inequality in opportunities (the e_T^i series) than in initial incomes (the y_0^i series)? Atkinson (1970) showed that if the Lorenz curve $(L(p))$ for one distribution lies everywhere below the Lorenz curve for another then inequality in the first distribution will be higher.[6] This is true for a range of popular inequality measures, including the variance, the coefficient of variation and the Gini coefficient. The following theorem establishes the conditions needed on $g(y_0)$ in order for the distribution of opportunities to Lorenz dominate the distribution of initial incomes.

Theorem 1   Let the opportunities available at time T, to a region with initial income y_0 be given as $e_T(y_0) = y_0 + g(y_0)$. The distribution of the e_T^i series Lorenz dominates the distribution of the y_0^i series for any realization of the initial distribution if and only if $d(g(y_0)/y_0)/dy_0 < 0$ for all y_0.

Proof   The proof of this theorem is a simple application of the Jakobsson-Fellman theorem (see Lambert, 1993).
See Appendix.

Theorem 1 states that progressivity is both a necessary and sufficient condition for Lorenz dominance of the opportunity set. The fact that progressivity is necessary warrants further discussion. If our only concern was whether the distribution of opportunities derived from a particular realization of initial incomes Lorenz dominated the initial income distribution, then progressivity, although sufficient, would not be necessary. Theorem 1 however, shows that if we want Lorenz dominance with respect to all possible initial income distributions that could have been realized then progressivity of the growth process is not only sufficient but is also necessary. This allows us to move beyond evaluating the growth process only in terms of realized distributions and to consider it in the context of initial distributions that could have been realized had circumstances been different at the start of the process.

....

[6]The Lorenz curve $(L(p))$ is a basic building block in the measurement of inequality and plots the proportion of income held by the poorest $p\%$ of the agents against p.

2.1 Progressivity, growth, and welfare

While the focus of this paper is on inequality of opportunity our framework also allows us to characterize the welfare properties of growth processes. To see this let $U(e(y))$ denote the utility accruing to a region with opportunities summarized by $e(y)$. We assume $U'(e) > 0$. Following an established tradition in public economics define social welfare as the average utility across initial income levels. That is

$$W = \int U(e(y)) \, f(y) dy \qquad (4)$$

where $f(y)$ is the distribution of initial incomes.[7] Although restrictions such as risk neutrality, the failure to model inequality within a region, the focus on final period opportunities and the specific choice of welfare function limits the nature of the welfare comparisons that can be provided by this approach, it is still instructive to show precisely how the nature of the growth process can impact on welfare comparisons within an equality of opportunity framework. The following theorem compares the change in welfare arising from the observed process with the change in welfare that would arise from an equal-yield proportional growth process i.e. a proportional growth process that generates the same overall level of opportunity as the observed process: [8]

Theorem 2 A monotone growth process increases welfare more than an equal yield proportional growth process applied to the same pre-growth income distribution for all strictly concave U and for all possible initial income distributions if and only if the growth process is progressive.

Proof See Appendix.

Theorem 2 states that progression in the growth process, over the entire range of income, is a necessary and sufficient condition for the resulting distribution of opportunities to welfare dominate both the initial distribution of opportunities and the subsequent distribution of opportunities that could have been derived from an equal-yield proportional growth process.

It is important to recognize the importance of progressive growth in these welfare comparisons. It is possible for the process to be monotonic, for mean income to rise and for inequality (as defined by the variance or Gini coefficient of opportunities) to fall and yet for Generalised Lorenz curves to cross so that unambiguous welfare rankings are not possible. In order to make unanimous welfare comparisons across distributions of opportunities it matters how the reduction in inequality is generated, and the necessary restrictions are captured by our measure of progressivity.

---

[7]See Lambert (1993, Ch. 4) for alternative rationalizations of this social welfare function.

[8]See also Corollary 3 of Benabou and Ok (2001).

3. Progressive growth and β-convergence

The above analysis shows how progressivity in the growth process can be used to facilitate inequality and welfare comparisons in an equality of opportunity framework. We now establish the relationship between progressivity in the growth process and measures of β-convergence derived from a traditional growth regression. To determine the progressivity of the growth process we need to establish whether $d\big(g(y_0)/y_0\big)/dy_0 \leq 0$ for all y_0. Since $d\big(g(y_0)/y_0\big)/dy_0 \leq 0$ for all y_0 if and only if $d\big(g(y_0)/y_0\big)/d\ln(y_0) \leq 0$ for all y_0, we can use the latter characterization of progressivity as a basis for our tests. The obvious approach to follow in this instance is simply to use a flexible form estimator which regresses observed growth rates on the log of initial income:

$$\frac{g_i(y_0)}{y_0} = m(\ln y_0^i) + \varepsilon_T^i \tag{5}$$

where ε_T^i is a realization of a mean zero error term. Progressivity of the growth process requires $dm(\ln y_0)/d\ln y_0 \leq 0$ everywhere. Therefore our proposed test for progressivity is nothing more than a negativity condition on the slope of a non-parametric cross-sectional growth-initial level regression. If we use the fact that $(g_i(y_0)/y_0) \approx \ln y_T^i - \ln y_0^i$ we can rewrite equation (5) as:

$$\ln y_T^i - \ln y_0^i = m(\ln y_0^i) + \varepsilon_T^i \tag{6}$$

Equation (6) is simply a flexible form growth regression. Thus the progressivity requirements needed for inequality and welfare comparisons in an equality of opportunity framework can be stated in terms of the β-convergence estimates obtained from a flexible specification of a growth regression. This highlights a potentially important role for non-linear growth regressions that extends beyond their ability to distinguish between competing theories of growth or their usefulness in summarising the evolution of realized outcomes.

4. Empirical application

In this section we illustrate our approach using regional data sets taken from Barro and Sala-i-Martin (1995). These data correspond to those used by Sala-i-Martin (1996b) to study regional cohesion. Sala-i-Martin estimated linear growth regressions for the regions of the United States, Japan, and Europe. In order to apply Theorems 1 and 2 however we must allow for possible nonlinearities. To do this we extend Sala-i-Martin's empirical analysis by estimating flexible non-parametric growth equations for each of these data sets. In particular we estimate the following flexible form growth equation:

$$\ln\left(\frac{y_T^i}{y_0^i}\right)/N = m[\ln(y_0^i)] + \epsilon_T^i \tag{7}$$

In each case we use the Nadaraya-Watson kernel estimator to obtain a flexible estimate of $m[\ln(y_0)]$. For the estimates presented in this paper we use a Gaussian
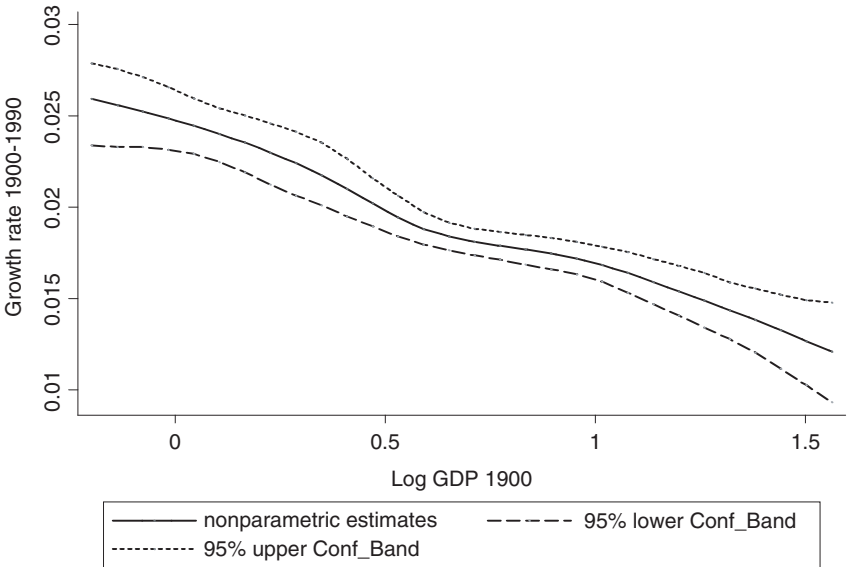
Fig. 1. Nonparametric estimate of the growth process across the US states

kernel to weight the data and choose bandwidths using the least squares cross-validation approach, though similar results were obtained using different kernels and a range of bandwidths. The dates for which the analysis is conducted depend on data availability and differ across data sets. The data for the US refer to real annual personal income *per capita* for each of the 48 contiguous states from 1900–90. The Japanese data measure real *per capita* income between 1955 and 1990 for the 47 prefectures. Finally the European data measure GDP *per capita* in each of 90 regions of Europe from 1950–90, covering Germany (11 regions), United Kingdom (11 regions), Italy (20 regions), France (21 regions), The Netherlands (four regions), Belgium (three regions), Denmark (three regions), and Spain (17 regions).[9]

The nonparametric estimates, $m[\ln(y_0)]$, for the US states, the Japanese prefectures, and the European regions are given in Fig. 1–3 respectively. The confidence bands reported in these figures were constructed by applying the Bonferroni correction to bootstrapped pointwise intervals (see Hardle, 1990). Our principal concern is the extent to which the growth process exhibits progression or regression over the income range; equivalently the extent to which the slope of $m[\ln(y_0)]$ is

---

[9]Following Sala-i-Martin (1996b) the European GDP figures are expressed as deviations from country specific means. Thus the estimated growth process we present for Europe should be interpreted as a common, within-country growth process. More details on these data are available in Barro and Sala-i-Martin (1995).
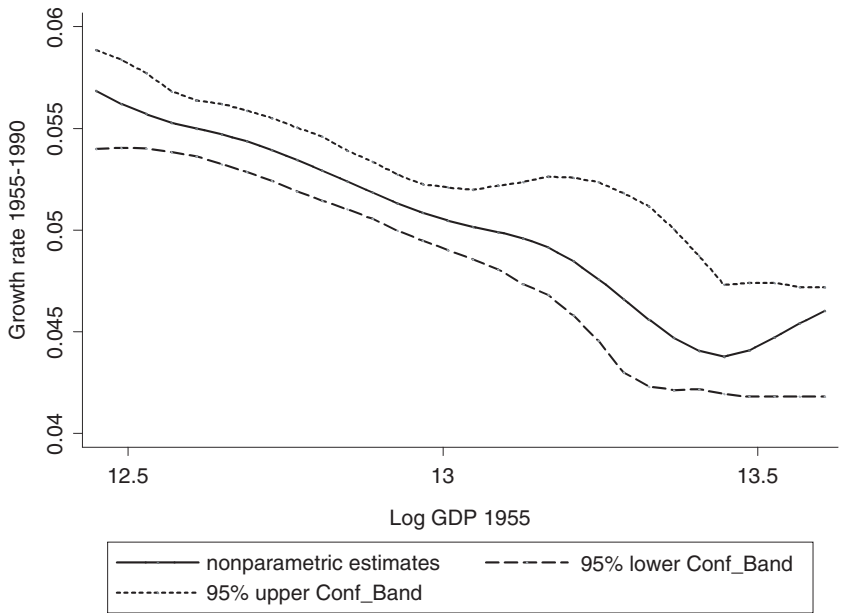
Fig. 2. Nonparametric estimate of the growth process across the japanese prefectures
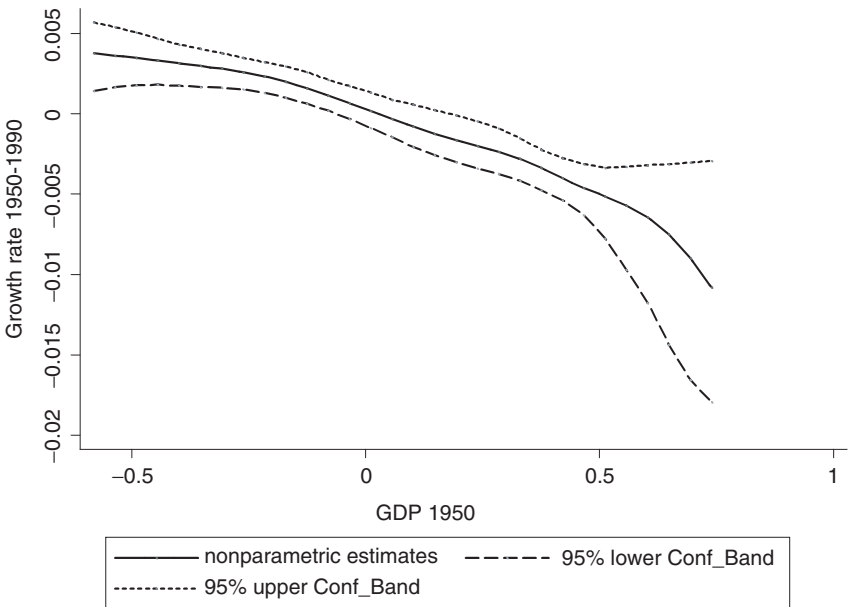

Fig. 3. Nonparametric estimate of the growth process across the european regions

negative or positive over the entire range of y_0. Figures 1–3 show that all the regional growth processes exhibit progressivity over almost all of their respective income ranges. Indeed the only evidence of regressive growth for these data occurs among high income Japanese prefectures. However, even then the confidence bands are such that we cannot rule out progressive growth over this income range. Combining these estimates with the theoretical results in Theorem 1 we can thus conclude that, for a wide range of popular inequality measures, there is less inequality in the opportunities afforded by the observed regional growth processes than was evident in the initial distribution of incomes. However, the implications of Theorem 1 go further and state that inequality of opportunity would have fallen no matter what the realized distribution of initial incomes had been. Since average income has also risen over this time period we can also conclude that, within the welfare framework specified in earlier, the distribution of opportunities available today within each of these regions welfare dominates that available at the start of the period.

Finally it is interesting to compare our non-parametric results on regions to earlier work using cross-country data sets (e.g. Liu and Stengos, 1999; Kalaitzidakis *et al*. 2000; Okada, 2006). The cross-country analysis tends to find strong evidence of nonlinear growth processes which is consistent with the multiple-regime equilibrium view advocated by Durlauf and Johnson (1995). This is in contrast to our results for regional convergence. To understand the differences between the regional and cross-country analysis we note that the multiple steady state models (e.g. Azariadis and Drazen, 1990) often invoke thresholds in physical or human capital which limit the diffusion of technology across countries. This produces multiple steady states and gives rise to nonlinear growth patterns. The linear model on the other hand is consistent with a single-steady state equilibrium without threshold externalities. Our findings suggest that barriers to technology diffusion may be less evident in regional data sets than in the cross-country data, so that even the relatively poor regions are able to benefit from technology spillovers in a way that relatively poor countries are not.

5. Conclusion

Cross-section regressions in which growth rates are regressed on initial values are widespread in growth economics. These regressions have been criticized on the grounds that they do not fully incorporate the range of restrictions imposed by alternative growth models and for failing to provide much insight into the evolution of income distributions over time. Accepting these criticisms, our paper draws on similarities between the convergence literature and the tax progressivity literature in public economics to show that appropriately specified growth regressions may, nevertheless, facilitate inequality and welfare comparisons in an equal-opportunity framework that are not possible with alternative measures of convergence. We illustrate this using regional data for the US, Japan, and Europe.

The growth processes estimated are consistent with a single equilibrium linear model yielding unambiguous reductions in inequality of opportunity and welfare improvements in each of the regions considered.

Acknowledgements

References

Atkinson, A. (1970) On the measurement of inequality, *Journal of Economic Theory*, **2**, 244–63.

Azariadis, C. and Drazen, A. (1990) Threshold externalities in economic development, *Quarterly Journal of Economics*, **CV**, 501–26.

Barro, R. and Sala–i-Martin, X. (1992) Convergence, *Journal of Political Economy*, **100**, 223–51.

Barro, R. and Sala-i-Martin, X. (1995) *Economic Growth*, McGraw-Hill, New York.

Benabou, R. and Ok, E. (2001) Mobility as progressivity: ranking income processes according to equality of opportunity, Working paper No. 8431, NBER, Cambridge, MA.

Cannon, E. and Duck, N. (2000) Galton's Fallacy and economic convergence, *Oxford Economic Papers*, **52**, 415–19.

Durlauf, S. and Johnson, P. (1995) Multiple regimes and cross-country growth behaviour, *Journal of Applied Econometrics*, **10**, 365–84.

Fields, G. (2004) Does income mobility equalise longer-term incomes?, New measures of an old concept, ILR Working paper, Cornell University, Ithaca, NY.

Fields, G. and Ok, E. (2001) The measurement of income mobility an introduction to the literature, in J. Silber (ed.) *Handbook on Income Inequality Measurement*, Kluwer Academic Press, Boston, MA.

Fleurbay, M. (1995) Equality of opportunity or equal social outcome?, *Economics and Philosophy*, **11**, 25–55.

Friedman, M. (1992) Do old fallacies ever die?, *Journal of Economic Literature*, **XXX**, 2129–32.

Furceri, D. (2005) β and σ-Convergence: a mathematical relation of causality, *Economic Letters*, **89**, 212–15.

Hardle, W. (1990) *Applied Nonparametric Regression*, Cambridge University Press, New York.

Hotelling, H. (1933) Review of the 'Triumph of Mediocrity in Business' by Horace Secrist, *Journal of the American Statistical Association*, **28**, 463–5.

Kalaitzidakis, P., Mamuneas, T., and Stengos, T. (2000) A non-linear sensitivity analysis of cross-country regressions, *Canadian Journal of Economics*, **33**, 604–17.

Lambert, P. (1993) *The Distribution and Redistribution of Income*, Manchester University Press, Manchester.

Lee, K., Pesaran, M., and Smith, R. (1997) Growth and convergence in a multi-country empirical stochastic solow model, *Journal of Applied Econometrics*, **12**, 357–92.

Liu, Z. and Stengos, T. (1999) Non-linearities in cross-country growth regressions: a semi-parametric approach, *Journal of Applied Econometrics*, 14, 527–38.

Maasoumi, E. (1998) On mobility, in D. Giles and A. Ullah (eds) *Handbook of Applied Economic Statistics*, Marcel Dekker, New York.

Okada, T. (2006) What does the solow model tell us about economic growth?, *Contributions to Macroeconomics*, 6.

Quah, D. (1993) Galton's Fallacy and tests of the convergence hypothesis, *Scandanavian Journal of Economics*, 95, 427–43.

Quah, D. (1996) Empirics for economic growth and convergence, *European Economic Review*, 40, 1353–75.

Roemer, J. (1998) *Equality of Opportunity*, Harvard University Press, Cambridge, MA.

Roemer, J. (2006) Economic development as opportunity equalization, Cowles Foundation Discussion Paper No. 1583, Yale University, New Haven, CT.

Sala-i-Martin, X. (1996a) The classical approach to convergence analysis, *Economic Journal*, 106, 1019–36.

Sala-i-Martin, X. (1996b) Regional cohesion: evidence and theories of regional growth and convergence, *European Economic Review*, 40, 1325–52.

Sen, A. (1985) *Commodities and Capabilities*, North-Holland, Amsterdam.

Shorrocks, A. (1983) Ranking income distributions, *Economica*, 50, 1–17.

Appendix

Proof of Theorem 1

$\Longrightarrow$By definition $e_T(y_0) = y_0 + g(y_0)$.

Define the expected growth rate for person with initial income y_0 as $\gamma(y_0) \equiv (g(y_0)/y_0)$. If our observed growth process is progressive, that is $\gamma'(y) < 0$ for all y, then we can use the Jakobsson-Fellman theorem (Lambert, 1993, p.150) and our assumption of monotonicity to conclude that :

$L_{e_T}(p) \geq L_{y_0^i}(p)$ all $p \in [0,1]$ and all initial income distributions.

$\Longleftarrow$ Suppose $L_{e_T}(p) \geq L_{y_0^i}(p)$ all $p \in [0,1]$ for all possible realized distributions of initial income.

Suppose there was a range of initial incomes $[y_0^*, y_0^{**}]$ over which $\gamma'(y) > 0$.

By changing the inequalities in the foregoing $\Longrightarrow$ proof, we can show that $L_{e_T}(p) \leq L_{y_0^i}(p)$ for any realized distribution of initial incomes whose support is contained within $[y_0^*, y_0^{**}]$. This would provide a contradiction of our original assumption. $\square$

Proof of Theorem 2

$\Longrightarrow$ Under proportional growth the Lorenz curve of opportunities at time T, (e_T^i), must equal the Lorenz curve for initial incomes or opportunities ($e_0 \equiv y_0$).

That is $L_{e_{prop,T}}(p)=L_{e_0}(p)$ for all $p \in [0,1]$

By definition $e_T(y_0) = y_0 + g(y_0)$.

Taking the average across agents we get that $\bar{e}_T=\bar{e}_0(1+\bar{\gamma})$, where $\bar{\gamma}$ is the average expected growth rate across agents ($\bar{\gamma} = \sum (g(y_0)/y_0)/N$) and $\bar{e}_0$ is average initial income or opportunities.

Hence the Generalised Lorenz Curve for future opportunities derived from a growth process characterized by $\gamma(y)$ can be expressed as :

$GLC_{e_{\gamma T}}(p)=\bar{e}_0(1+ \bar{\gamma})L_{e_T}(p),$

where $L_{e_T}(p)$ is the Lorenz curve of future opportunities.

If our observed growth process is progressive, that is $\gamma'(y) < 0$ for all y, then we can use the Jakobsson-Fellman theorem (Lambert, 1993, p.150) and our assumption of monotonicity to conclude that :

$GLC_{e_{\gamma T}}(p)=\bar{e}_0(1+\bar{\gamma})L_{e_T}(p)\geq \bar{e}_0 \ (1+\bar{\gamma})L_{e_0}(p)= \bar{e}_0 \ (1+\bar{\gamma})L_{e_{prop,T}}(p)$ all $p \in [0,1]$.

The first inequality follows from our assumptions of monotonicity and progressivity and the last equality follows from step 1 of the proof.

By definition this implies that:

$GLC_{e_{\gamma T}}(p)\geq GLC_{e_{prop,T}}(p)$ all $p \in [0,1]$.

Referring to Shorrocks (1983) completes the proof in this direction.

$\Leftarrow$ Suppose $GLC_{e_{\gamma T}}(p)\geq GLC_{e_{prop,T}}(p)$ for all p and any pre-growth income distribution.

Then following the logic above we can establish that

$L_{e_T}(p)\geq L_{e_0}(p)$ for all p and all pre-growth income distributions.

From the Jakobsson-Fellman theorem we can then conclude that the mobility process is progressive for all y_0.   □