Growth and Inequality:

Dependence on the Time Path of Productivity Increases (and other Structural Changes)†

Manoj Atolia*
Florida State University

Santanu Chatterjee**
University of Georgia

Stephen J. Turnovsky***
University of Washington

Revised Version
September 2011

Abstract

This paper examines the significance of the time path of a given productivity increase on growth and inequality. Whereas the time path impacts only the transitional paths of aggregate quantities, it has both transitional and permanent consequences for wealth and income distribution. Hence, the growth-inequality tradeoff generated by a given discrete increase in productivity contrasts sharply with that obtained when the same productivity increase occurs gradually. The latter can generate a Kuznets-type relationship between inequality and per-capita income. Our results suggest that economies with similar aggregate structural characteristics may have different outcomes for income and wealth inequality, depending on the nature of the productivity growth path.

†This paper has benefited from presentations at the 2009 Midwest Macroeconomic Meetings, held in Bloomington IN, May 2009, the 8th Journées Louis-André Gérard-Varet Conference in Public Economics, Marseille, France, June 2009, the 2010 AEA/ASSA Annual Meetings in Atlanta, GA, the 2010 Conference of the Society for Computational Economics in London, UK, and at the University of Hamburg, Germany. We are grateful to Shankha Chakraborty for comments on an earlier draft.

*Department of Economics, Florida State University, Tallahassee, FL 32306. Phone: +1-850-644-7088. Email: matolia@fsu.edu
**Department of Economics, Terry College of Business, University of Georgia, Athens, GA 30602. Phone: +1-706-542-3696. Email: schatt@terry.uga.edu
***Department of Economics, University of Washington, Seattle, WA 98195. Phone: +1-206-685-8028. Email: sturn@u.washington.edu.

1. Introduction

The relationship between income inequality and economic growth has been extensively discussed since Kuznets' (1955) pioneering work first appeared over half a century ago. Since then, the question of whether these two key economic variables are positively or negatively related has been extensively debated, although no definitive conclusion has been reached. Early growth regressions by Alesina and Rodrik (1994), Persson and Tabellini (1994), Perotti (1996), and others, yield a negative growth-inequality relationship. But more recent studies obtain a positive, or at least more ambiguous, relationship; see for example, Li and Zou (1998), Forbes (2000), and Barro (2000).[1] From a theoretical perspective, this empirical controversy should not be surprising. Because an economy's growth rate and its income distribution are both equilibrium outcomes, the growth-inequality relationship – whether positive or negative – will reflect the underlying set of forces to which both are simultaneously reacting. To understand these linkages it is necessary to examine this relationship using a consistently specified general equilibrium framework.

In this paper we employ such a framework to consider the impact of one of the major determinants of the growth-inequality relationship, namely an increase in productivity.[2] The key result we shall establish is that the effects of a productivity increase of a given magnitude on wealth and income inequality depend crucially upon the time path along which the productivity increase accrues. This in turn has important consequences for the growth-inequality tradeoff, and further, may help explain why economies with similar aggregate structural characteristics may nevertheless have very different income and wealth distributions. While we focus on a productivity increase as being particularly salient, it will become evident that the argument in fact applies to any structural change that occurs over time. Hence the issue we are addressing is quite general, and therefore highly significant for understanding the dynamics of the growth-inequality tradeoff.

In a completely general setup, where the equilibrium growth rate and income distribution are mutually dependent, their joint determination and the analysis of their relationship becomes intractable; see Sorger (2000). This paper, on the other hand, is related to a growing body of

[1] Various explanations for these results and their differences are summarized by García-Peñalosa and Turnovsky (2006).
[2] This is one of the key factors influencing the growth-inequality relationships; see Solimano (1998) and Piketty (2006).

research that exploits the fact that if the underlying utility function is homogeneous in its relevant arguments, the aggregate economy can be summarized by a representative agent, as a result of which aggregate behavior becomes independent of the economy's distributional characteristics. Rather, the distributions (e.g. of wealth and income) reflect the evolution of the aggregate economy; see e.g. Caselli and Ventura (2000), García-Peñalosa and Turnovsky (2006, 2007), Turnovsky and García-Peñalosa (2008), Kraay and Raddatz (2007), Carroll and Young (2009) and Barnett et al. (2009). While awareness of this aggregation property dates back to Gorman (1953), by rendering the analysis tractable, it assumes particular importance in studying the growth-inequality relationship. Moreover, the class of utility functions to which this aggregation applies includes the constant elasticity utility function that dominates contemporary growth theory.

Inequality is necessarily associated with heterogeneous agents. Recently, differential initial endowments of capital across economic agents have received a lot of attention as an underlying source of heterogeneity.[3] A crucial mechanism generating the endogenous distribution of income is the relationship between agents' relative capital stock and their relative allocation of time to leisure. In the long run, this relationship is positive, as wealthier agents who have a lower marginal utility of wealth increase their consumption of all goods, including leisure.[4] In the short run, however, this relationship is conditioned by the time path a given productivity change is expected to follow, and the differences in the consumption-smoothing motives it generates for rich and poor agents.

A key feature of this labor allocation-relative wealth mechanism is that it introduces *hysteresis* in the dynamic adjustment characterizing the relative holdings of capital. This occurs because the impact of any structural change on the long-run evolution of wealth inequality, and subsequently on income inequality, depends critically upon the initial response of leisure (labor supply) to the underlying shock. This initial response, in turn, depends upon the entire (known) time

[3] By identifying agents' heterogeneity with their initial physical asset endowments, we are embedding distributional issues within a more traditional growth-theoretic framework. Indeed, the role of the return to capital, which is essential in that literature, has tended to receive less attention in much of the recent discussions of income inequality, which have emphasized other aspects such as human capital and growth; see e.g. Galor and Zeira (1993), Bénabou (1996b), and Viaene and Zilcha (2003), among others. The argument that the return to capital is essential to understanding distributional differences has, however, been addressed by Atkinson (2003), and is supported by recent empirical evidence for the OECD [see Checchi and García-Peñalosa, 2010].

[4] This long-run negative relationship between wealth and labor supply is supported by empirical evidence, obtained from a variety of sources. This is discussed by García-Peñalosa and Turnovsky (2006).

path that the structural change is expected to follow. That is, not only the dynamic evolution (which is expected) but also the long-run distributions of wealth and income inequality become *path dependent*. Thus, a central insight of this paper is that the effects of a productivity increase of a *given* magnitude on the *long-run distributions* of both wealth and income are crucially dependent upon the time path that the productivity increase is assumed to follow. This is in sharp contrast to the dynamics of the aggregate economy. In this case, the time path of a productivity change affects only the transitional path of the aggregate economy and has no impact on its steady state.

To illustrate the role of path dependence we compare the consequences of two alternative specifications of the productivity increase. The first is the conventional one, where the full productivity change occurs instantaneously as an unanticipated permanent *discrete* increase in the level of productivity. The second is where the same overall increase in the level of productivity occurs, but is acquired *gradually* over a known time path, and is therefore anticipated after the first instant.[5] These two specifications of the productivity increase have exactly opposite consequences for the initial responses of leisure. In the first case leisure initially declines, while in the latter case it initially increases, leading to profoundly different distributional consequences.

The path dependence of the distributions of wealth and income inequality helps provide some key insights into the ambiguous empirical relationship between growth and inequality. Given the analytical complexity of the theoretical framework, all our results are derived using numerical simulations. The main findings are summarized below:

1. Whereas a discrete productivity increase always leads to a monotonic decline in wealth inequality, its gradual introduction leads to a non-monotonic adjustment, with an initial increase followed by a gradual decline after some period of time. In the long run, a gradual increase is likely to lead to *more*, rather than less, wealth inequality unless the flexibility of production is extremely high. Furthermore, whereas a discrete productivity increase leads to an initial increase in income inequality, followed by a monotonic decline to below its initial

[5] More specifically, consider the case of a closed economy using antiquated production techniques that decides to open up to the adoption of more modern technology. In this context the question is: should this economy adopt policies that attain the new level of technology instantaneously or should the adoption process be more gradual?

level, its gradual introduction generates essentially the opposite time profile.

2. The model permits a diversity of distributional equilibria for structurally similar countries: countries with similar structural conditions may end up with very different levels of inequality, depending on the time path of productivity changes. This result is consistent with the experiences of countries in East Asia and Latin America, who have similar levels of per-capita income but very different levels of income inequality.[6]
3. A gradual productivity increase generates a Kuznets-type inverted-U relationship between inequality and per-capita income. In contrast, a discrete change in productivity can generate only an inverse monotonic relationship. The fact that the much-debated Kuznets relationship can be generated in the context of a simple one-sector Ramsey model is a salient feature of this paper, providing a simple theoretical justification for a controversial empirical relationship; see Ray (1998).
4. A country's distance from the technological frontier has important implications for both inequality and its persistence.
5. Even though the long-run outcomes for the aggregate economy are independent of the time-path of the underlying productivity change, the transitional responses can be quite sensitive to it. For example, while a discrete productivity shock leads to an instantaneous increase in output, a continuous change leads to exactly the opposite response, where both employment and output decline. This may help explain why some countries go through difficult transitions following structural changes. Mexico's short-run experience with trade liberalization in the mid 1990s following the adoption of the NAFTA is a good example.

Since the model is solved numerically, we conduct extensive robustness checks by varying the two key parameters of the model, namely the elasticity of substitution between capital and labor in production, and the speed of the underlying productivity change. Our main results are robust to

[6] The diversity of distributional equilibria prevents Deininger and Squire (1996, 1998) from finding a statistically significant relationship between the *level* of income and inequality in over 75 per cent of their cross-country sample even after controlling for initial differences in inequality. There may be other reasons that economies structurally similar in the aggregate may have different distributional characteristics, the most obvious being differential fiscal policies; see e.g. García-Peñalosa and Turnovsky (2007) where this issue is discussed in the context of an endogenous growth model.

variations in these parameters. Motivated by the differences in the productivity gaps of different countries from the technology frontier, we also present robustness results for the size of the productivity shock.

The standard procedure of assuming that productivity increases occur fully on impact, rather than gradually, while convenient analytically, is arguably less realistic than positing some form of continuing adjustment. Several examples can be given to support this view. First, to the extent that productivity increases reflect government investment in infrastructure, the notion that the increases occur gradually over time, rather than instantaneously, seems more plausible. Budget restrictions, installation costs, and bureaucratic impediments inevitably force governments to spread their investments over time as multi-year projects. The US interstate highway system, initiated in the 1950's, is a good example and so are the recent public infrastructure policies of China and India. Foreign aid to developing countries, especially when "tied" to investment projects as schools, roads, hospitals, etc. is likely to be granted over time, and is yet another good example of a gradual productivity change. Second, productivity increases generally reflect the assimilation of new productive techniques that may require learning for complete adaptation, and this too takes time. An example of this is the general purpose technologies (GPT) such as steam, railroads, lasers and, more recently information technology; see Aghion and Howitt (2009, ch 9). Finally, from the perspective of a developing economy, one can interpret the productivity increase as representing a closing of the 'productivity gap' which again is likely to take years to eliminate.

The rest of the paper proceeds as follows. Section 2 sets out the analytical model and the evolution of the aggregate dynamics, while Section 3 discusses the distributional dynamics of wealth and income. Since the focus is on the transitional dynamics, which are too complex to solve analytically, we employ numerical simulations. These are reported in Section 4. Section 5 discusses implications of our analysis for the empirical literature on growth and inequality, focusing specifically on the Kuznets curve, and Section 6 presents concluding remarks.

2. Analytical framework

We consider a decentralized economy having a single representative firm and heterogeneous households. The source of heterogeneity among consumers is the initial distribution of capital endowments. For simplicity, we assume a completely laissez-faire economy which operates in the absence of a government or social planner.

2.1 Technology and factor payments

Aggregate output is produced by a single representative firm according to a standard neoclassical production function[7]

$$Y(t) = A(t)F\left(K(t), L(t)\right) \qquad F_L > 0, F_K > 0, F_{LL} < 0, F_{KK} < 0, F_{LK} > 0 \qquad (1)$$

where, K, L, and Y denote the per-capita stock of capital, labor supply, and output. In addition, $A(t)$ represents the level of productivity, which is exogenous to the firm's decisions.

The key feature of our analysis is that the level of productivity is assumed to increase gradually from its initial level, A_0, to a higher long-run level, $\tilde{A}$, both of which are publicly known. This is specified by the (known) deterministic growth path

$$A(t) = \tilde{A} + \left(A_0 - \tilde{A}\right)e^{-\theta t}, \;\; \theta \geq 0 \qquad (2)$$

The parameter θ thus defines the time path followed by the increase in productivity. The conventional approach to specifying productivity increases is to assume that they occur instantaneously. This is obtained (or at least approximated) as a special case by letting $\theta \to \infty$ in (2), so that the new productivity level is achieved virtually instantaneously. However, the more general specification introduced in (2) is clearly important. This is because, as we will demonstrate subsequently, there is a sharp contrast between how θ affects the behavior of *aggregates* and *distributions*. As one would expect, it affects the transitional path of the aggregate economy, but not the aggregate steady state. In contrast, it has profound impacts on *both* the time paths and the steady-state levels of both wealth and income inequality.

The assumption that the increase in technology occurs at a constant proportionate rate, and is

[7] Both factors of production have positive, but diminishing, marginal physical products and the production function exhibits constant returns to scale, with $F_{KL} > 0$ being a consequence of the latter assumption.

completed only asymptotically, is made purely for analytical convenience. It is straightforward to generalize (2) to the case where the new level of productivity is reached in finite time, *T*.[8] The analysis could also be modified to allow for the technology increase to follow a more general functional form, and the same general qualitative conclusions would emerge.

The wage rate, *w*, and the return to capital, *r*, are determined by the marginal physical products of labor and capital:

$$w(t) \equiv w(K,L) = AF_L(K,L) \tag{3a}$$

$$r(t) \equiv r(K,L) = AF_K(K,L) \tag{3b}$$

where we have dropped the time notation from the variables. Note that both the wage rate and the return on capital reflect the current level of productivity, $A(t)$.

2.2 Households

At time 0, the economy is populated by N_0 households, represented as a continuum between 0 and N_0, each indexed by *i*. Population grows uniformly across households at an exponential rate, *n*, so that at time *t*, household *i* has grown to e^{nt} and the total population of the economy is $N(t) \equiv N_0 e^{nt}$. Households are identical except for their given initial endowments of capital, $K_{i,0}$, so that the average initial stock of capital in the economy is $K_0 = (1/N_0)\int_0^{N_0} K_{i,0}di$. At time *t*, with the growing population and accumulation of capital, the average per-capita amount of capital is

$$K(t) = \frac{1}{N_0 e^{nt}}\int_0^{N_0} K_i(t)e^{nt}di = \frac{1}{N_0}\int_0^{N_0} K_i(t)di .$$

where $K_i(t)$ is the per capita capital owned by household *i*. From a distributional perspective, we are interested in household *i*'s relative share of the total capital stock in the economy, $k_i(t)$, namely

$$k_i(t)\frac{K_i(t)e^{nt}}{(1/N_0)\int_0^{N_0} K_i(t)e^{nt}di} = \frac{K_i(t)}{(1/N_0)\int_0^{N_0} K_i(t)di} = \frac{K_i(t)}{K(t)}$$

[8] This could be done by specifying the productivity growth function, $A(t) = \tilde{A} + (A_0 - \tilde{A})[e^{-\theta t} - e^{-\theta T}](1 - e^{-\theta T})^{-1} \quad t \le T$, $A(t) = \tilde{A} \;\; t \ge T$, or alternatively $A(t) = \tilde{A} + ((T-t)/T)(A_0 - \tilde{A}),\ t \le T$.

At all points of time, the mean of the distribution is normalized to unity, while the the initial (given) standard deviation of relative capital (the coefficient of variation of the level of capital) is $\sigma_{k,0}$.[9]

We now consider household *i*, which, like all others, is endowed with a unit of time that it can allocate to either leisure, l_i, or labor, $L_i \equiv 1 - l_i$. The household chooses its rates of consumption, C_i, and leisure to maximize lifetime utility represented by the iso-elastic function:

$$\max \int_0^\infty \frac{1}{\gamma}\left(C_i(t)l(t)_i^\eta\right)^\gamma e^{-\beta t} dt, \quad \text{with } -\infty < \gamma < 1, \eta > 0, \gamma\eta < 1 \tag{4}$$

where $1/(1-\gamma)$ equals the intertemporal elasticity of substitution.[10] This maximization is subject to the household's initial endowment of capital, $K_{i,0}$, together with its capital accumulation constraint

$$\dot{K}_i(t) = (r(t) - n)K_i(t) + w(t)(1 - l_i(t)) - C_i(t) \tag{5}$$

2.3. Macroeconomic equilibrium

Summing over all households, equilibrium in the capital and labor markets is described by

$$K(t) = \frac{1}{N_0}\int_0^{N_0} K_i(t)di, \quad L(t) = 1 - l(t) = \frac{1}{N_0}\int_0^{N_0}(1 - l_i(t))di \tag{6}$$

Note that in equations (3a) and (3b), we have expressed the wage and the return to capital, *w*, *r*, as functions of average capital, *K*, and employment, L. We can equivalently write them as functions of aggregate leisure time, $(1-l)$, namely, $w = w(K,l)$ and $r = r(K,l)$.

The key elements facilitating the aggregation across households are the homogeneity of the utility function and perfect factor markets. The first-order conditions (see Appendix) can be used to derive the following relationships:[11]

[9] We should emphasize that this formulation does not impose any particular distributional form, other than assuming the existence of a mean and an arbitrary measure of initial dispersion, $\sigma_{k,0}$. As will become clear later, the distributional dynamics of wealth and income we derive will reflect that of the arbitrary initial endowments, $\sigma_{k,0}$.

[10] The preponderance of empirical evidence suggests that this is relatively small, certainly well below unity, so that we shall restrict $\gamma < 0$. See, the discussion of the empirical evidence summarized and reconciled by Guvenen (2006).

[11] See Section A.1 of the Appendix for the derivation of (7). It requires taking the time derivative of (A.1a) and combining with (A.1b); see also Turnovsky and García-Peñalosa (2008).

$$\frac{\dot{K}_i(t)}{K_i(t)} = r(t) - n + \frac{w(t)}{K_i(t)}\left[1 - l_i(t)\left(\frac{1+\eta}{\eta}\right)\right] \quad (7a)$$

$$\frac{\dot{C}_i(t)}{C_i(t)} = \frac{\dot{C}(t)}{C(t)}; \quad \frac{\dot{l}_i(t)}{l_i(t)} = \frac{\dot{l}(t)}{l(t)} \quad \text{for all } i \quad (7b)$$

Eq. (7a) describes the evolution of agent *i*'s capital stock and (7b) states that all agents choose the same growth rate for consumption and leisure, implying further that average consumption, *C*, and leisure, *l*, will also grow at the same common growth rates.

Summing over all households, the macro-dynamic equilibrium is described by the following equations, expressed in terms of average capital, leisure (labor supply) and productivity, where the time notation has been dropped for the core dynamic variables, *K*, *l*, and *A*:

$$\dot{K} = AF(K,1-l) - \frac{A(t)F_L(K,1-l)l}{\eta} - nK \quad (8a)$$

$$\dot{l} = \frac{1}{G(K,l)}\left\{AF_K(K,1-l) - \beta - n - (1-\gamma)\frac{F_{Kl}(K,1-l)}{F_l}\left[AF(K,1-l) - \frac{AF_l(K,1-l)l}{\eta} - nK\right]\right\} \quad (8b)$$

$$\dot{A} = \theta\left(\tilde{A} - A\right) \quad (8c)$$

where, $G(K,l) \equiv \frac{1-\gamma(1+\eta)}{l} - (1-\gamma)\frac{F_{LL}}{F_L} > 0,$ and $L = 1 - l.$

With this aggregation, the complete dynamics of the economy can be represented by the core dynamic system consisting of l_i, l, k_i, K, A, the evolution of which is described by (7a), (7b), and (8a)-(8c). There are two key points to note about the macroeconomic equilibrium of the economy. First, the aggregate dynamics are entirely *independent* of any distributional characteristics. This is a consequence of the homogeneity of the underlying utility function and, as previously acknowledged, has been known since Gorman (1953). Second, the dynamics of both capital and leisure depend on the current level of productivity, $A(t)$. At any time, this in turn depends upon the anticipated long-run change, $(\tilde{A} - A_0)$, together with its growth rate along the transitional path, θ.

The steady-state equilibrium is attained when $\dot{K} = \dot{l} = \dot{A} = 0$ in (8a)-(8c), and determines the steady-state values of $\tilde{K}$ and $\tilde{l}$, in addition to the given and known steady-state level of productivity

$\tilde{A}$. Since the steady state is independent of θ, the long-run effects of an increase in productivity are *independent* of the time path by which it is achieved. Thus, a productivity increase of a given magnitude, whether it occurs instantaneously as a discrete jump, or is attained only gradually, will lead to identical steady-state changes for the aggregate economy. However, the transitional responses may be significantly different (as illustrated later in the numerical experiments). In contrast, as we will see below, the rate of productivity increase, θ, has fundamental consequences for wealth and income inequality, both in transition and across steady states.[12]

3. Distributional dynamics

To derive the distribution dynamics, we must solve for the remaining variables, namely the relative stock of capital and leisure (k_i and l_i), in the core-dynamic system. A key difference between the distributional dynamics from the aggregate dynamics described previously is that it generates *hysteresis*, i.e., the dependence of long-run outcomes on initial conditions. This is a direct consequence of (7b): the proportionality of individual and aggregate consumption and leisure along the transition path.[13] Note that, having already solved for the path of *l*, the path of individual leisure, l_i, is determined except for the initial jump at $t = 0$. This initial jump, therefore, determines the labor supply for a household in the final steady state, which also determines its steady-state relative income and capital stock. As we will show below, the initial response of the labor-leisure choice, in turn, depends critically on the time path of the underlying shock (discrete versus continuous).

3.1 Distribution of capital (wealth)

Wealth inequality is characterized in terms of household *i*'s capital stock relative to the average, namely by the evolution of $k_i(t) \equiv K_i(t)/K(t)$. Combining (7a) and (8a) leads to the following dynamic equation for the *i*-th agent's relative capital stock:

---

[12] The steady-state conditions, together with the homogeneity of the production function, imply $\tilde{l} > \eta/(1+\eta)$; see Turnovsky and García-Peñalosa (2008). This inequality yields a lower bound on the steady-state allocation to leisure that is consistent with a feasible equilibrium. As we shall see below, this condition is critical in characterizing the distributional dynamics.

[13] Hysteresis arises because of the relationships $\dot{l}_i/l_i = \dot{l}/l$ and $\dot{C}_i/C_i = \dot{C}/C$, as a result of which l_i and C_i are proportional to *l* and *C*, respectively. With continuous-time, this introduces a zero root into the individual-level dynamics.

$$\dot{k}_i = \frac{AF_L(K,l)}{K}\left[1-\nu_i l\left(1+\frac{1}{\eta}\right)-\left\{1-l\left(1+\frac{1}{\eta}\right)\right\}k_i\right] \tag{9}$$

where K, l, A evolve in accordance with (8a)-(8c) and the initial relative capital $k_{i,0}$ is given from the initial endowment. Since $\dot{l}_i/l_i = \dot{l}/l$ in (9), we may write $l_i = \nu_i l$, where $\left(1/N_0\right)\int_0^{N_0}\nu_i di = 1$ and ν_i (relative leisure) is constant for each *i*, and to be determined. Setting $\dot{k}_i = 0$, and using the fact that $\tilde{l} > \eta/(1+\eta)$ [see footnote 12], leads to the following positive long-run relationship between relative leisure and relative capital:

$$\tilde{l}_i - \tilde{l} = \left(\tilde{l} - \frac{\eta}{1+\eta}\right)\left(\tilde{k}_i - 1\right) \text{ for each } i \tag{9'}$$

While our simulations employ shooting algorithms to solve (9) for the time path of the relative stock of capital, in conjunction with the aggregate dynamics specified in (8a)-(8c), the intuition underlying the dynamic structure can be better understood by characterizing a linear approximation. To do this, we linearize (9) around the steady state. In the Appendix we show that the resulting bounded solution for the relative stock of capital is:

$$k_i(t) - 1 = \delta(t)(\tilde{k}_i - 1) \tag{10}$$

where,

$$\delta(t) \equiv \left[1 + \frac{\tilde{A}F_L}{\tilde{K}}\int_t^\infty \left(1 - \frac{l(\tau)}{\tilde{l}}\right)e^{-\beta(\tau - t)}d\tau\right].$$

Setting $t = 0$ in (10), we can solve for agent *i*'s steady-state relative capital stock:

$$k_{i,0} - 1 = \delta(0)(\tilde{k}_i - 1) = \left(1 + \frac{\tilde{A}F_L}{\tilde{K}}\int_0^\infty \left(1 - \frac{l(\tau)}{\tilde{l}}\right)e^{-\beta\tau}d\tau\right)(\tilde{k}_i - 1) \tag{11}$$

where $k_{i,0}$ is given from the initial distribution of relative capital endowments.

Equations (10) and (11) characterize the evolution of relative capital. First, given the time path of the aggregate economy, in particular $l(\tau)$, and the distribution of initial capital endowments, (11) determines the steady-state distribution of capital, $(\tilde{k}_i - 1)$. Once this is known, (10) then describes the time path of relative capital, which can be expressed in the convenient form[14]

$$k_i(t)-\tilde{k}_i=\left[\frac{\delta(t)-1}{\delta(0)-1}\right]\left(k_{i,0}-\tilde{k}_i\right) \tag{12}$$

Because of the linearity of (10)-(12), we can immediately transform these expressions into corresponding relationships for the standard deviation of the distribution of relative capital across agents, which serves as a convenient measure of wealth inequality:

$$\sigma_k(t)-\tilde{\sigma}_k=\left(\frac{\delta(t)-1}{\delta(0)-1}\right)(\sigma_{k,0}-\tilde{\sigma}_k) \tag{13}$$

where, $\sigma_k(t)=\delta(t)\tilde{\sigma}_k$ and $\sigma_{k,0}=\delta(0)\tilde{\sigma}_k$.

The crucial difference between this analysis and previous work lies in the evolution of the productivity shock $A(t)$, which is reflected in the time path of $\delta(t)$. In the case where the complete productivity increase occurs instantaneously, $l(\tau)-\tilde{l}=\left(l(0)-\tilde{l}\right)e^{\mu\tau}$, and $\delta(t)$, $\delta(0)$ simplify to

$$\delta(t)=1+\left(\frac{1}{\beta-\mu}\right)\frac{\tilde{A}F_L(\tilde{K},\tilde{L})}{\tilde{K}}\left(1-\frac{l(t)}{\tilde{l}}\right);\ \delta(0)=1+\left(\frac{1}{\beta-\mu}\right)\frac{\tilde{A}F_L(\tilde{K},\tilde{L})}{\tilde{K}}\left(1-\frac{l(0)}{\tilde{l}}\right)$$

where μ is the negative (stable) eigenvalue corresponding to the linearized aggregate dynamic system specified in (8a)-(8c). Note that only the current allocation of time to leisure relative to its steady-state allocation is relevant in determining current wealth inequality relative to its long run level. When the productivity increase occurs gradually over time, the entire time profile of $A(t)$, as reflected in $l(t)$, needs to be taken into account; see equation (A.5) in the Appendix.

In general, the term $\delta(t)$ in (10) highlights the role played by the time path of leisure in determining the long-run change in wealth inequality. For example, if during the transition $l(\tau)<\tilde{l}$, so that leisure approaches its long-run steady state from below, then $\delta(t)>1$ and wealth inequality will decline over time; see (13). As our simulations show, this is the case for a discrete productivity increase, where leisure increases (following an initial drop) and wealth inequality declines monotonically over time. On the other hand, a gradual productivity increase leads to an initial increase in leisure, taking it initially above its new (lower) steady-state level. But since the

[14] Note also that the constant $\nu_i=l_i/l$ can be determined from (9'), and is given by $\nu_i=1+\left(1-(1/\tilde{l})(\eta/(1+\eta))\right)(\tilde{k}_i-1)$.

transitional path is U-shaped, eventually approaching $\tilde{l}$ from below, whether inequality rises or falls over time depends upon the extent to which $l(\tau) > \tilde{l}$ during the early phase of the adjustment.

The other point to observe is that the closer $l(\tau)$ is to its steady state, $\tilde{l}$, the smaller is the subsequent adjustment in $l(t)$, and hence the smaller is the overall change in the distribution of wealth. This is because if the economy and therefore all individuals fully adjust their respective leisure times instantaneously, they will all accumulate wealth at the same rate, causing the wealth distribution to remain unchanged.

3.2 Distribution of income

Defining household *i*'s per capita income as $Y_i(t) = r(t)K_i(t) + w(t)\left(1 - l_i(t)\right)$, and average economy-wide per capita income as $Y(t) = r(t)K(t) + w(t)\left(1 - l(t)\right)$, we define relative income by $y_i(t) = Y_i(t)/Y(t)$. This leads to the following equation of motion for relative income:[15]

$$y_i(t) - 1 = \varphi(t)\left[k_i(t) - 1\right] \tag{14}$$

where, $\varphi(t) = 1 - \left(1 - s(t)\right)\left[1 + \frac{l(t)}{1 - l(t)}\left(1 - \frac{1}{\tilde{l}}\frac{\eta}{1+\eta}\right)\frac{1}{\delta(t)}\right]$

$s(t)$ represents the share of capital in total output. Again, because of the linearity of (14) in $(k_i(t) - 1)$, we can express the relationship between relative income and relative capital in terms of corresponding standard deviations of their respective distributions, namely

$$\sigma_y(t) = \varphi(t)\sigma_k(t) \tag{14'}$$

4. Numerical analysis

The model set out in Sections 2 and 3 will be solved and analyzed numerically, using the following functional forms and parameterization:[16]

[15] See Turnovsky and García-Peñalosa (2008) for details regarding the derivation of the equations of motion for relative income and capital.
[16] These parameters are generally standard in the literature and noncontroversial. For an extensive discussion of the calibration of the Ramsey model, see Cooley (1995).

Utility function:	$\gamma = -1.5,\ \beta=0.04,\ \eta = 1.75$
Production function:	$Y = A(\alpha K^{-\rho} + (1-\alpha)L^{-\rho})^{-1/\rho}$; $A_0 = 1,\ \alpha = 0.4$
Elasticity of substitution:	$s = 0.75, 1, 1.25;\ s \equiv 1/(1+\rho)$
Productivity level:	$A_o = 1,\ \tilde{A} = 1.5$
Productivity growth:	$\theta = 0.1$
Population growth:	$n = 0.015$

Preferences remain specified by a constant elasticity utility function with an intertemporal elasticity of substitution of 0.4, while the elasticity of leisure in utility is 1.75. The production function is of the Constant Elasticity of Substitution (CES) form, where we allow the elasticity of substitution to vary between $s = 0.75$ (low substitution), $s = 1$ (Cobb-Douglas), and $s = 1.25$ (high substitution).

We adopt the following strategy. We consider the aggregate and distributional consequences of a 50% increase in productivity (A increases from its benchmark value of $A_0 = 1$ to $\tilde{A} = 1.5$), which we allow to take effect in two alternate ways: (i) an immediate one-time unanticipated jump in productivity from 1 to 1.5. This represents a *discrete* increase in *A*, and corresponds qualitatively to much of the previous literature, and (ii) the same increase in A ($A_0 = 1$ to $\tilde{A} = 1.5$) taking place *gradually* over time, where $A(t)$ adjusts at the (known) rate $\theta = 10$ % per period (year). In the latter case, the higher productivity level is achieved asymptotically. As a result, the instant it starts to increase, the subsequent levels of productivity are fully anticipated along the transition path. For each of these scenarios, we numerically characterize the economy's aggregate and distributional dynamics for the three specified values of the elasticity of substitution in production to demonstrate the robustness of our results for this important parameter of the model.

4.1. Solution Algorithm

Intertemporal models grounded in optimizing behavior typically give rise to saddle-point solutions, the exact numerical computation of which is often difficult. Many papers, therefore, obtain linear approximations to the "true" dynamics. One alternative to deriving exact solutions for non-linear dynamic systems is to use some type of "shooting" algorithm (forward or reverse) to locate the path that lies on the stable manifold.[17] The choice between forward and reverse shooting

[17] See Atolia and Buffie (2009b) for other alternatives.

depends on many factors, including the nature of the dynamic system and the type of shock under consideration (Atolia and Buffie, 2009b). Forward shooting computes the equilibrium path by searching over the initial values of the jump variables, whereas in reverse shooting the search is conducted over the terminal values of the state variables. For a unit-root system, where final steady-state values are not known, the forward shooting algorithm is the appropriate solution technique.

In a series of papers, Atolia and Buffie (2009a, 2009b) have developed a set of shooting algorithms that identify the global saddle path much more efficiently by using a new distance mapping. Atolia and Buffie (forthcoming) develop a set of innovative forward-shooting algorithms for unit-root systems that combine the new distance mapping mentioned above with the insight that there exists a one-to-one relationship between guesses for the initial values of jump variables in the dynamic system and the consistent guess for the new steady state. As our complete dynamic system consisting of $l, l_i, K, k_i,$ and A has two jump variables, we use the circle-search algorithm of Atolia and Buffie (forthcoming) that underlies their *UnitRoot-Circle* program to obtain an exact solution to the dynamics of our unit-root system.[18] As this algorithm allows solving for unit-root problems with two jump variables, we have the benefit of solving the complete dynamic system consisting of both the aggregate and the individual-level dynamics in a single step.[19]

4.2 An increase in productivity: Discrete versus continuous adjustment

4.2.1. Aggregate dynamics

Fig. 1 depicts the transition paths for the aggregate variables, K and l, corresponding to the two specifications of the productivity increase. To demonstrate robustness, paths are depicted for

[18] Although we employ non-linear solution techniques for our numerical analysis, the first-order linearization procedure, albeit an approximation, is useful in guiding our intuition, particularly in the role of the initial response in leisure. Therefore, the linearized solutions for both the aggregate economy and the distributions are set out in the Appendix. This exposition and equation (A.7) in particular, highlights the role played by the initial response of leisure in the aggregate dynamics, as it internalizes the information regarding the time profile of the productivity increase. Elsewhere, we have investigated the accuracy of conventional linearization procedures in characterizing the dynamics of a standard aggregate Ramsey model; see Atolia, Chatterjee, and Turnovsky (2010). For a model having the structure here, linearization can accommodate quite large structural changes without committing unacceptably large errors, at least for moderate values of the elasticity of substitution (less than unity). For large values (around 1.25) the errors become more significant.

[19] Alternatively, because of the block-recursive structure one can solve the problem sequentially, first solving for the aggregate dynamics using a reverse-shooting procedure in Atolia and Buffie (2009b) and then using this solution to solve for the "individual-level" dynamics using a unit-root forward-shooting algorithm.

the three values of the elasticity of substitution in production, s. Although the long-run responses of leisure (labor supply), capital, and output are identical for both specifications of the productivity increase, their short-run responses and transitional paths are dramatically different. Irrespective of the elasticity of substitution, we see that a discrete productivity increase causes leisure to *decline* instantaneously (and labor supply to increase), after which it immediately reverses and increases monotonically to its new steady-state level, which lies above or below its original pre-shock level, depending upon whether $s \gtrless 1$. By contrast, a continuous (gradual) productivity increase causes an immediate *increase* in leisure, which is then quickly reversed, causing leisure to overshoot its long-run equilibrium during the subsequent decline to steady state. Similar differences are displayed in the initial phases of the transitional path for capital. While a discrete productivity increase leads to a gradual monotonic accumulation of capital to the new steady-state, a continuous increase actually leads to a short-run *decumulation* (for about 10 years), before capital accumulation begins. This gives the time-path of the capital stock a U-shaped trajectory, the depth of which increases as s declines. These differences in the adjustments of leisure (labor supply) and capital translate directly into differences in the dynamic adjustment of output (not shown). A discrete productivity increase causes output to increase instantaneously followed by a further rise during transition, while a continuous shock causes output to fall on impact. For the reasons discussed in Section 3.1 [see eq. (11)], and as we shall illustrate soon, these transitional differences, and in particular the adjustment of leisure, have a critical impact on the economy's distributional dynamics.

Why does the dynamic response of the aggregate economy differ so dramatically for the two types of productivity change? The explanation lies in the information being revealed to the agent on impact of the shock, relative to the time path of the higher productive capacity associated with the long-run realization of the shock. For a one-time discrete increase in A, the enhanced long-run productivity is fully realized instantaneously by the agent, and immediately raises the marginal product of both labor and capital. Consequently, labor supply immediately increases on impact of the shock, and the enhanced productivity of capital generates immediate incentives for capital accumulation, and the stock of capital begins to rise. Output increases instantaneously, and can accommodate the increase in consumption associated with the higher level of permanent income resulting from the productivity increase.

In contrast, if the productivity increase occurs only gradually, the enhanced productive capacity necessary to support the increase in consumption will take effect only over time. In the short run, the long-run change in the level of productivity is fully anticipated by the agent, thereby increasing permanent income, and raising aggregate current consumption. But the increase in productive capacity is immediately reflected only as an increase in its growth rate, $\dot{A}(0) = \theta(\tilde{A} - A_0)$. Thus, since the instantaneous *level* of productivity remains unchanged, current output cannot rise, and the increase in consumption resulting from the anticipation of higher future income is achieved through reduced investment and a decline of the capital stock. In fact, the increase in short-run consumption and lower productivity (relative to the long-run) causes the agent to increase leisure, which causes output to also decline on impact of the shock.

Fig. 1 also provides some insights as to why certain countries have difficult transitions when they adopt policies such as liberalization that induce gradual structural changes. One particular example to which the results from figure 1 can be related to is Mexico's trade liberalization through the adoption of NAFTA in 1994, which led to a collapse in output and an increase in unemployment (proxied by leisure in our setup).

4.2.2 Distributional dynamics: Diversity of transitory and long-run outcomes

Fig. 2 illustrates the dynamic responses of the distributions of capital (wealth) and income to the two specifications of the productivity increase, corresponding to the three values of the elasticity of substitution, $s = 0.75, 1, 1.25$. Specifically, we plot the evolution of the standard deviation of wealth and income relative to their respective initial (pre-shock) standard deviations. The most striking feature of these distributional time paths is that not only do the two specifications of productivity increases have contrasting effects on the short-run distributions of capital and income, but contrary to the aggregate economy in fig. 1, the long-run effects are also dramatically different. In other words, while the aggregate economy reaches identical steady-states, irrespective of whether the productivity change occurs discretely or gradually, the distributions do not. This reflects the fact that the long-run distributions of wealth and income are *path dependent*, depending critically on the underlying process through which the steady-state equilibrium is attained (i.e., whether the productivity increase occurs discretely or gradually).

The contrasts between the dynamic adjustments in wealth and income distributions in response to the two types of productivity increase are sharpest for low values of s. Focusing first on $s = 0.75$ and 1 (in fig. 2), we see that a discrete productivity increase generates a gradual monotonic decline in wealth inequality over time. Income inequality increases instantaneously in the short run, before declining monotonically to an equilibrium value that is below its pre-shock level. By contrast, a gradual increase in productivity of the same magnitude increases wealth inequality in transition. However, the time path of wealth inequality is non-monotonic and follows an inverted U-shaped trajectory. On the other hand, income inequality falls instantaneously on impact of the shock, before it rises in transition to an equilibrium that exceeds the pre-shock benchmark. Like wealth inequality, the transitional adjustment of income inequality also follows a non-monotonic inverted U-shaped trajectory, peaking at around 15 years. For high values of the elasticity of substitution, such as $s = 1.25$ (the third panel of fig. 2), the differences in the responses are less pronounced. The instantaneous and transitional responses of wealth and income inequality remain as above. But now in the long-run, both wealth and income inequality decline for the two types of productivity increases, though the continuous productivity shock still leads to higher levels of long-run wealth and income inequality relative to the discrete shock.

An important dimension of variation of the growth and development experience across countries is the rate of productivity growth θ (e.g. arising from the speed of reforms). Fig. 3 presents sensitivity analysis for the dynamics of wealth and income inequality for alternative values of θ, given $s = 0.75$.[20] In particular, the following patterns can be detected:

(i) The slower the rate at which a given increase in productivity is achieved (longer is the "catch-up" time), the greater are the long-run increases in wealth and income inequality.

(ii) Our previous result that wealth and income inequality worsens in the short and the medium run is robust to the rate of productivity growth, but slower productivity growth

[20] As in fig. 2, we plot the evolution of the standard deviation of wealth and income relative to their respective initial (pre-shock) standard deviations.

increases the *persistence* of both wealth and income inequality.

The experiments illustrated in fig. 3 indicate that countries that experience a faster "catch-up" process will also experience smaller increases in inequality, which in turn are also less persistent over time, compared to countries that experience slower catch-up speeds. These results are consistent with the fact that Asian economies that have been developing at faster rates than Latin American economies also have substantially less wealth and income inequality.[21,] We should note that we have performed these experiments for much larger values of θ and do indeed find that as θ becomes large, the time paths of wealth and income distributions converge to those associated with the discrete productivity increase.

More generally, we can conclude that a diverse range of distributional outcomes are possible along transitional paths as well as in the long run, depending on the time path followed by the underlying productivity shock. In particular, two countries starting with the same initial distribution of wealth, and having the same per-capita income today may have very different degrees of wealth and income inequality, if they have experienced the same overall productivity increase but at different rates. Thus, the diversity of growth experiences of different countries may be reflected in the cross-sectional diversity of wealth and income inequality.

As stressed above, the critical element in determining the evolution of the distributions is the time path followed by average leisure, $l(t)$ and its implications for the rates of factor return. To assist in understanding the intuition underlying these diverse time paths, it is convenient to recall the steady-state condition (9'), which when combined with (10) at $t=0$, yields

$$l_i - \tilde{l} = \frac{1}{\delta(0)}\left(\tilde{l} - \frac{\eta}{1+\eta}\right)\left(k_{i,0} - 1\right) \qquad (9')$$

A *discrete* increase in productivity raises the return to both labor and capital. On impact, average leisure, $l(0)$, falls as agents substitute toward labor supply. This decrease in average leisure increases $\delta(0)$ in (11), implying an overall monotonic reduction in wealth inequality over time. This happens because the amount of leisure time chosen by agents with above (below) average wealth

[21] Gini coefficients for income are typically 10 or more points higher for Latin American countries than they are for Asian countries.

declines (increases). That is, wealthier people initially increase their work time, while poorer people work less, and income inequality increases. Over time, as average leisure increases, the relative income of agents having above-average wealth declines, and income inequality declines accordingly.

In contrast, a *continuous* increase in productivity leads to an initial increase in $l(0)$, taking it above $\tilde{l}$. This increase in the initial average leisure decreases $\delta(0)$, implying an increase in wealth inequality in the long-run; see (13). This initial increase in leisure is, however, immediately reversed and falls below $\tilde{l}$ during the subsequent transition. The net effect of this on the evolution of wealth inequality depends on whether the positive amounts of $l(\tau)-\tilde{l}$ during the early phase dominate the negative amounts in the latter phase. This depends upon the elasticity of substitution. If $\sigma \leq 1$, we see that the early excesses dominate $\left(\delta(0) < 1\right)$, implying a non-monotonic inverted-U time path for the distribution of wealth, as it approaches its higher steady-state level.

As a consequence of the increase in long-run wealth inequality, the amount of leisure time chosen by people with above (below) average wealth increases (decreases), causing income inequality to decline. Over time, as average leisure decreases, the relative income of agents having above average wealth increases, and income inequality increases accordingly, although it also reflects the inverted-U time path of wealth inequality. In the case of a high elasticity of substitution, $s = 1.25$, most of the time $l(\tau) < \tilde{l}$, so that the downward pressure on wealth inequality dominates and except for a brief period at the start of the transition, wealth inequality declines over time. As a consequence of this, income inequality also declines, albeit slightly over time.

5. The growth-inequality relationship

As noted in the introduction, the relationship between growth and inequality has been extensively discussed. This paper belongs to a growing strand of research that contends that these processes are endogenous outcomes in the course of economic development. The critical issue then concerns the underlying mechanisms or shocks that affect the joint evolution of growth and inequality. A comprehensive survey by Solimano (1998) identifies several factors, such as the national savings rate, investment in physical or human capital, productivity growth, education, capital markets, and public policy that can significantly influence the growth-inequality relationship.

Viewed in this context, can we address the consequences of productivity growth (technological change) for the growth-inequality relationship, both during the transition and in the steady-state?

In considering this issue, we focus on two aspects: (i) the relationship between the evolution of per-capita income and income inequality in the face of a productivity increase. In particular, we ask the question what are the consequences (if any) of the nature of productivity change for the well-known empirical Kuznets' curve? (ii) Whether the distance from the technology frontier has implications for inequality and its persistence. In other words, do countries that require a lot of "catching-up" with regard to technology also generate more inequality in the process?

5.1. Per-capita income and inequality: The Kuznets curve revisited

The celebrated Kuznets curve is an inverted U-shaped relationship between income inequality and a country's level of development, say the level of per-capita income. The idea underlying this relationship is that in the initial stages of development the accumulation of physical capital is important and, therefore, capital-rich agents gain disproportionately relative to the capital-poor. Hence, at low levels of per-capita income, income inequality rises. After a certain income level is attained, physical capital becomes less important for development (possibly due to the emergence of human capital and knowledge), and income inequality declines with further increases in per-capita income.[22] Clearly, such a relationship cannot be generated by an unanticipated *discrete* increase in productivity. This is because after the initial jumps in output and income inequality, the subsequent increases in income are associated with a monotonic decline in income inequality.

On the other hand, our more flexible specification where the level of productivity increases gradually is more promising. As we have already noted in figs. 2 and 3, after an initial decline on impact, income inequality does indeed follow an inverted U-shaped trajectory over time. Therefore, fig. 4 plots takes the same transition paths as in figs. 2 and 3, but plots wealth and income inequality against the (normalized) per-capita income, for the three values of the elasticity of substitution in production.[23] As is evident, elements of a Kuznets curve emerge for both wealth and income

[22] Galor and Moav (2004) develop an elegant theory of growth in which human capital accumulation replaces physical capital accumulation as the prime engine of growth in the process of development. They argue that their theory offers a unified explanation for the effect of income inequality on the process of economic growth.

[23] The time range considered in fig. 4 starts from the period after the instantaneous adjustments are complete.

inequality during the transition, after the initial adjustments have been completed. Thus, our flexible, more general specification of productivity change not only allows a diversity of distributional outcomes, depending on the growth experiences, but also generates a Kuznets type relationship during transition that is robust to the ease of substitution between factors of production.

It is important to stress that the inverted U-shaped relationship between inequality and per-capita income generated by our model relies on a mechanism that is very different from the traditional explanations of the Kuznets curve. While the sectoral composition of capital, labor, and knowledge is the traditionally understood mechanism behind the Kuznets curve, we focus on the role of differences in the consumption-smoothing responses of the rich and the poor to a gradual, but anticipated, change in productivity, and its effects on investment and the choice of labor and leisure. Recently, Piketty (2006) has discussed the role of "waves" of technological change (such as general purpose technologies) that can generate waves of an inverted U-shaped relationship between inequality and income over time. Our findings fit nicely into that story.

To consider the robustness of the Kuznets curve to other parameters, fig. 5 reports further sensitivity analysis with respect to the speed with which the productivity change is introduced. Specifically, we consider the following experiment: for a given increase in the *level* of productivity (50% increase in A from 1 to 1.5), we consider three alternative *rates* of change, i.e., $\theta = 0.05$, 0.1, and 0.2. These three cases characterize scenarios where a given "catch-up" in productivity occurs slowly ($\theta = 0.05$), at the benchmark rate ($\theta = 0.1$), or quickly ($\theta = 0.2$). Wealth and income inequality are plotted relative to normalized per-capita income. The corresponding time paths are drawn in fig. 5 for the case of an elasticity of substitution in production equal to 0.75. As we can see, the essence of the Kuznets curve that is generated by a continuous productivity change is robust with respect to the speed of productivity growth during transition.

5.2. Productivity gaps and inequality

An important area of differences in the development experiences of different countries relates to the size of productivity increases, due, for example, to the range or magnitude of their economic reforms or the possibility of "catch up" due to the distance from the technology frontier. While there are differences in magnitude of the responses, this section shows that the short-run and medium-run rises in wealth and income inequality (and the concomitant emergence of the Kuznets' curve) is a robust feature of the development experience, irrespective of the size of the necessary "catch up".

This can be seen from fig. 6 which depicts the sensitivity of the distributional dynamics for different levels of productivity "catch-up," for any given rate of productivity growth. Specifically, for a given productivity growth rate ($\theta = 0.1$), we consider three magnitudes of "catch-up," i.e., $(\tilde{A} - A_0)/A_0 = 0.25$, 0.5, and 1. In other words, these characterize cases where the difference between the initial and steady-state levels of productivity are small (25%), at the benchmark (50%), or large (100%). The following patterns can immediately be identified:

(i) The larger the initial gap in productivity, the greater is the long-run rise in wealth and income inequality.

(ii) The higher the initial productivity gap, the larger is the short-run increase (following the instantaneous responses) in wealth and income inequality, relative to their initial (pre-shock) levels.

These results indicate that the further away a country is from its long-run technology or productivity frontier, the larger will be the inequality generated in converging to the frontier. Countries that are closer to their long-run technology frontier will have less persistent inequality than countries that are further away.

6. Conclusions

The relationship between growth and inequality is one of the most fundamental (and elusive) ones in development economics. We employ a general-equilibrium heterogeneous-agent growth model with certain well-known aggregation properties that generates hysteresis in the dynamics of wealth and income inequality, but not in the aggregate dynamics. This is manifested in dynamic adjustments of distributional variables being dependent on the initial response of leisure (labor supply), following a structural change. Since with forward-looking agents, the initial response is dependent upon future anticipations of these structural changes, this implies further that their long-run effects on the distributions of wealth and income are path-dependent.

This paper, therefore, examines the consequences of the time path of a productivity change on the distributions of wealth and income. As a benchmark, we have considered an increase in productivity of a given magnitude, and compared the distributional implications when (i) it is introduced gradually, with (ii) the more conventional situation where it all occurs instantaneously. The main conclusion is that the time path along which a productivity increase of a given magnitude is introduced has dramatic consequences for both wealth and inequality. In general we find that the gradual introduction of a given productivity increase has adverse distributional consequences, and certainly much more adverse than when they are introduced instantaneously.

This has important consequences for a range of issues. First, it suggests that if the government introduces some productivity-enhancing policy, such as investment in infrastructure, with the objective of stimulating economic growth, it should do so rapidly. While delay and gradual implementation will have no adverse permanent effects on the aggregate performance of the economy, they will generate expectational effects on the labor-leisure choice that will lead to a worsening of wealth and income inequality in the long-run.

The paper also has some interesting implications for the empirical relationship between inequality and growth. We show that a gradual productivity change can indeed generate a Kuznets-type inverted U-shaped relationship between inequality and per-capita income, with a diverse set of possible long-run outcomes for inequality across structurally similar countries. This is also shown to be a very robust finding. As Ray (1998, chapter 7) points out, one problem with the generally inconclusive empirical literature on the Kuznets curve is the lack of an underlying theory that can generate a testable specification of this relationship. By articulating an explicit mechanism through which per-capita income and inequality can be linked – namely the differences in the dynamic responses of the labor-leisure choice between rich and poor in conjunction with the time profile of structural changes – and in the context of the simple one-sector neoclassical growth model, our results provide a step in that direction, one that may provide a useful basis for future empirical work.

Further, we show that a country's distance from the technology frontier has important implications for inequality and its persistence. This relates to the issue of "economic backwardness" pioneered by Gerschenkron (1962), who in his seminal work hypothesized various consequences of increasing backwardness. Our analysis contributes to this important discussion by suggesting that, at a given rate of convergence, the more backward an economy initially is, the more inequality, both short-run and long-run, will be generated in the process of catching up to the technology frontier.

In light of the significance of the path-dependence characteristic and the fact that we have relied on numerical simulation methods, it is important to return to the question of robustness. This issue has several dimensions, including the assumed functional forms, the source of the underlying heterogeneity, the nature of the shocks, and the chosen parameter values. The key factor generating the path dependence of distributions is the proportionality of individual and aggregate consumption and leisure during the transition [eq. (7b)]. This in turn is a consequence of the homogeneity of the utility function, and as Turnovsky and García-Peñalosa (2008) show, this generalizes beyond the constant elasticity utility function employed here. Any homogeneous utility function will generate the path-dependence property. Within that widely-employed class of model, our sensitivity analysis has shown that our results are robust with respect to a wide range of parameter values. Although we choose to focus on the time path followed by a productivity increase, the issue is in fact a generic one, applying to any form of structural change. Moreover, though we have chosen to focus on one source of heterogeneity, namely the initial endowments of capital, work by Caselli and Ventura (2000) and more recently García-Peñalosa and Turnovsky (2011) indicates that a similar structure can emerge for other sources of heterogeneity as well, such as differential skill levels in labor. The point is that as long as the time path for wealth inequality depends upon the initial response of leisure, its subsequent time path, and in turn that of income inequality, will depend on the path followed by the structural change. On the other hand, if the utility function does not have the homogeneity property we are assuming, so that aggregate behavior and distributions are simultaneously determined rather than having the recursive structure that we are exploiting here, we cannot in general characterize the transitional dynamics. The extent to which, if at all, path-dependence may exist in this case remains an open question.

Appendix: Linearized Solution for Dynamic System for a Gradual Increase in Productivity

A.1 Dynamics of aggregate economy

The first-order conditions corresponding to the household maximization problem in (4) and (5) give rise to the following key relationships:

$$\eta \frac{C_i(t)}{l_i(t)} = w(t) \tag{A.1a}$$

$$(\gamma-1)\frac{\dot{C}_i(t)}{C_i(t)} + \eta\gamma\frac{\dot{l}_i(t)}{l_i(t)} = \frac{\dot{\lambda}_i(t)}{\lambda_i(t)} = \beta + n - r(t) \tag{A.1b}$$

where λ_i is agent i's shadow value of capital. Equation (A.1a) equates the marginal rate of substitution between consumption and leisure to the price of leisure, while (A.1b) is the Euler equation modified to take into account the fact that leisure changes over time. Using (A.1a), we may write the individual's accumulation equation, (5), in the form

$$\frac{\dot{K}_i(t)}{K_i(t)} = r(t) - n + \frac{w(t)}{K_i(t)}\left[1 - l_i(t)\left(\frac{1+\eta}{\eta}\right)\right] \tag{A.1c}$$

Aggregating (A.1c) over the individuals and using (3a) and (3b), we can derive the macroeconomic equilibrium and the dynamics of the aggregate economy, reported in (8a)-(8c):

$$\dot{K} = A(t)F(K,L) - \frac{A(t)F_L(K,L)l}{\eta} - nK \tag{A.1d}$$

$$\dot{l} = \frac{1}{G(K,l)}\left\{A(t)F_K(K,L) - \beta - n - (1-\gamma)\frac{F_{KL}K,L)}{F_L}\left[A(t)F(K,L) - \frac{A(t)F_L(K,L)l}{\eta} - nK\right]\right\} \tag{A.1e}$$

$$\dot{A}(t) = \theta\left(\tilde{A} - A(t)\right) \tag{A.1f}$$

where: $G(K,l) \equiv \frac{1-\gamma(1+\eta)}{l} - (1-\gamma)\frac{F_{LL}}{F_L} > 0, \qquad l + L = 1$

Linearizing (A1.d) and (A1.e) around the steady state equilibrium and $\tilde{A}$, yields

$$\begin{pmatrix} \dot{K} \\ \dot{l} \\ \dot{A} \end{pmatrix} = \begin{pmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ 0 & 0 & -\theta \end{pmatrix} \begin{pmatrix} K-\tilde{K} \\ l-\tilde{l} \\ A-\tilde{A} \end{pmatrix} \tag{A.2}$$

where,

$$a_{11} = \tilde{A}\left[F_K - \frac{F_{KL} l}{\eta}\right] - n\,; \;\; a_{12} = -\tilde{A}F_L\left[1+\frac{1}{\eta}\right] + \tilde{A}F_{LL}\frac{\tilde{l}}{\eta}\,; \;\; a_{13} = \frac{n\tilde{K}}{\tilde{A}}$$

$$a_{21} = \frac{1}{G}\left[\tilde{A}F_{KK} - (1-\gamma)\frac{F_{KL}}{F_L}a_{11}\right]; \;\; a_{22} = \frac{1}{G}\left[-\tilde{A}F_{KL} - (1-\gamma)\frac{F_{KL}}{F_L}a_{12}\right]; \;\; a_{23} = \frac{1}{G}\left[-F_K - (1-\gamma)\frac{F_{KL}}{F_L}a_{13}\right]$$

This is a third order system with two stable eigenvalues; (i) μ, where $\mu < 0$ is the negative root to $\mu^2 - (a_{11} + a_{22}) - (a_{11}a_{22} - a_{12}a_{21}) = 0$ and (ii) $-\theta$. Using standard solution procedures, the general form of the stable solution is

$$K(t) = \tilde{K} + Ce^{\mu t} + \varphi_1(A_0 - \tilde{A})e^{-\theta t} \tag{A.3a}$$

$$l(t) = \tilde{l} + \left(\frac{\mu - a_{11}}{a_{12}}\right)Ce^{\mu t} + \varphi_2\left(A_0 - \tilde{A}\right)e^{-\theta t} \tag{A.3b}$$

where C is arbitrary and

$$\varphi_1 \equiv [a_{12}a_{23} - a_{13}(a_{22} + \theta)]\Delta^{-1},\; \varphi_2 \equiv [a_{21}a_{13} - a_{23}(a_{11} + \theta)]\Delta^{-1} \tag{A.4}$$

and $\Delta \equiv (a_{11} + \theta)(a_{22} + \theta) - a_{12}a_{21}$. As long as the two stable eigenvalues are distinct – a very weak restriction – $\Delta \neq 0$.[24] Imposing the initial condition, $K(0) = K_0$ by setting $t = 0$ in (A.3a) yields

$$C = (K_0 - \tilde{K}) - \varphi_1(A_0 - \tilde{A})$$

so that starting from $K(t) = K_0, A(t) = A_0$ the time paths for $K(t)$ and $l(t)$ are

$$K(t) = \tilde{K} + [(K_0 - \tilde{K}) - \varphi_1(A_0 - \tilde{A})]e^{\mu t} + \varphi_1(A_0 - \tilde{A})e^{-\theta t} \tag{A.5a}$$

$$l(t) = \tilde{l} + \left(\frac{\mu - a_{11}}{a_{12}}\right)\left[(K(t) - \tilde{K}) - \varphi_1\left(A_0 - \tilde{A}\right)e^{-\theta t}\right] + \varphi_2\left(A_0 - \tilde{A}\right)e^{-\theta t} \tag{A.5b}$$

[24] The case where $\theta = -\mu$ can also be easily solved.

The solution (A.5a)-(A.5b) represents a general form, which covers all ranges of θ, including the conventional case where the full change in productivity occurs as a discrete change at time zero. This is obtained by letting $\theta \to \infty$ in which case $\varphi_1, \varphi_2 \to 0$, and the solution reduces to:

$$K(t) = \tilde{K} + (K_0 - \tilde{K})e^{\mu t} \tag{A.5a'}$$

$$l(t) = \tilde{l} + \left(\frac{\mu - a_{11}}{a_{12}}\right)\left[K(t) - \tilde{K}\right] \tag{A.5b'}$$

In the long-run, as $t \to \infty$, $K(t) \to \tilde{K},\, l(t) \to \tilde{l}$ independently of the path as defined by θ.

To see the role of the initial adjustment in leisure, consider (A.5b) at time $t = 0$, namely

$$l(0) = \tilde{l} + \left(\frac{\mu - a_{11}}{a_{12}}\right)\left[(K_0 - \tilde{K}) - \varphi_1\left(A_0 - \tilde{A}\right)\right] + \varphi_2\left(A_0 - \tilde{A}\right) \tag{A.6}$$

Thus the initial response of leisure to a productivity increase at that time is:

$$\frac{dl(0)}{dA} = \frac{d\tilde{l}}{dA} - \left(\frac{\mu - a_{11}}{a_{12}}\right)\frac{d\tilde{K}}{dA} + \left(\frac{\mu - a_{11}}{a_{12}}\right)\varphi_1 - \varphi_2 \tag{A.7}$$

where the role of the time path is contained in φ_1 and φ_2.

(i) If the productivity increase occurs as a discrete jump at time zero, $\varphi_1, \varphi_2 \to 0$, and the instantaneous response of leisure is the standard expression

$$\frac{dl(0)}{dA} = \frac{d\tilde{l}}{dA} - \left(\frac{\mu - a_{11}}{a_{12}}\right)\frac{d\tilde{K}}{dA} \tag{A.8}$$

which implies an initial decline in $l(0)$.

(ii) If the increase in productivity A occurs gradually, at the rate θ, as in (A.1f), we find that $\left((\mu - a_{11})/a_{12}\right)\varphi_1 - \varphi_2 < 0$, implying an initial increase in $l(0)$, as suggested by our simulations.

A.2 Dynamics of the relative capital stock

To obtain the dynamics of individual capital we linearize equation (9) around the steady-state $\tilde{K}, \tilde{l}, \tilde{k}_i, \tilde{l}_i$. This is given by

$$\dot{k}_i(t) = \frac{\tilde{A}F_L}{\tilde{K}}\left[\left(1+\frac{1}{\eta}\right)(\tilde{k}_i - v_i)(l(t)-\tilde{l}\,) + \left[\left(\tilde{l}\left(1+\frac{1}{\eta}\right)-1\right)\left(k_i(t)-\tilde{k}_i\right)\right]\right] \tag{A.9}$$

Combining the steady-state conditions corresponding to (8a) and (8b) we get

$$\frac{F_L(\tilde{K},\tilde{L})}{\tilde{K}}\left[\tilde{l}\left(1+\frac{1}{\eta}\right)-1\right] = \beta$$

and rewriting equation (9') as

$$v_i = \frac{(1-\tilde{k}_i)}{\tilde{l}\,(1+1/\eta)} + \tilde{k}_i$$

enables us to express (A.9) in the more compact form

$$\dot{k}_i(t) = \beta\left(k_i(t)-\tilde{k}_i\right) + \frac{\tilde{A}F_L}{\tilde{K}}\frac{\left(\tilde{k}_i-1\right)}{\tilde{l}}\left(l(t)-\tilde{l}\,\right) \tag{A.9'}$$

The stable solution to this equation is

$$k_i(t)-1 = \left(\tilde{k}_i-1\right)\left[1+\frac{\tilde{A}F_L}{\tilde{K}}\int_t^\infty\left(1-\frac{l(\tau)}{\tilde{l}}\right)e^{-\beta(\tau-t)}d\tau\right] \tag{A.10}$$

Setting $t=0$ in (A.10) and noting that $k_{i.0}$ is given, we obtain

$$k_{i,0}-1 = \left(\tilde{k}_i-1\right)\left[1+\frac{\tilde{A}F_L}{\tilde{K}}\int_0^\infty\left(1-\frac{l(\tau)}{\tilde{l}}\right)e^{-\beta\tau}d\tau\right] \tag{A.11}$$

Thus, having determined $\tilde{K},\tilde{L}$, and the time path for $l(t)$ from (A.5), equation (A.11) determines $\tilde{k}_i$, and knowing $\tilde{k}_i$, (A.10) in turn determines the entire time path for $k_i(t)$.

In the case of the discrete productivity shock, when the aggregate economy follows (A.5'), we find $l(\tau)-\tilde{l} = (l(0)-\tilde{l})e^{\mu\tau}$ and (A.10), (A.11) reduce to

$$k_i(t)-1 = (\tilde{k}_i-1)\left[1+\left(\frac{1}{\beta-\mu}\right)\frac{\tilde{A}F_L(\tilde{K},\tilde{L})}{\tilde{K}}\left(1-\frac{l(t)}{\tilde{l}}\right)\right] \tag{A.10'}$$

$$k_{i,0}-1 = (\tilde{k}_i-1)\left[1+\left(\frac{1}{\beta-\mu}\right)\frac{\tilde{A}F_L(\tilde{K},\tilde{L})}{\tilde{K}}\left(1-\frac{l(0)}{\tilde{l}}\right)\right] \tag{A.11'}$$

References

Aghion, P., Howitt, P., 2009. The Economics of Growth. MIT Press, Cambridge, MA.

Alesina, A., Rodrik, D., 1994. Distributive politics and economic growth. Quarterly Journal of Economics 109, 465-490.

Atkinson, A.B., 2003. Income inequality in OECD countries: Data and explanations. CESifo Economic Studies 49, 479-513.

Atolia, M., Buffie, E.F., 2009a. Smart forward shooting. Computational Economics 33, 1-30.

Atolia, M., Buffie, E.F., 2009b. Reverse shooting made easy: Automating the search for the global nonlinear saddle path. Computational Economics 34, 273-308.

Atolia, M., Buffie, E.F., 2011. Solving the unit root problem in models with an exogenous world market interest rate. Macroeconomic Dynamics (forthcoming).

Atolia, M., Chatterjee, S., Turnovsky, S.J., 2010. How misleading is linearization? Evaluating the dynamics of the neoclassical growth model. Journal of Economic Dynamics and Control 34, 1550-1571.

Barnett, R., Bhattacharya, J., Bunzel, H., 2009. Choosing to keep up with the Joneses and income inequality. Economic Theory (forthcoming).

Barro, R.J., 2000. Inequality and growth in a panel of countries. Journal of Economic Growth 5, 5-32.

Bénabou, R., 1996a. Heterogeneity, stratification, and growth: Macroeconomic implications of the community structure and school finance. American Economic Review 86, 584-609.

Bénabou, R., 1996b. Equity and efficiency in human capital investment: The local connection. Review of Economic Studies 63, 237-264.

Carroll, D., Young, E., 2009. The stationary wealth distribution under progressive taxation. Review of Economic Dynamics 12, 469-478.

Caselli, F., Ventura, J., 2000. A representative consumer theory of distribution. American Economic Review 90, 909-926.

Checchi, D., García-Peñalosa, C., 2010. Labour market institutions and the personal distribution of income in the OECD. Economica 77, 413-450.

Cooley, T.F., 1995. Frontiers of Business Cycle Research. Princeton University Press, Princeton, NJ.

Deininger, K., Squire, L., 1998. New way of looking at old issues: inequality and growth. Journal of Development Economics 57, 259-287.

Deininger, K., Squire, L., 1996. A new data set measuring income inequality. World Bank Economic Review 10, 565-591.

Forbes, K., 2000. A reassessment of the relationship between inequality and growth. American Economic Review 90, 869-887.

Galor, O., Moav, O., 2004. From physical to human capital accumulation: Inequality and the process of development. Review of Economic Studies 71, 1001-1026.

Galor, O., Zeira, J., 1993. Income distribution and macroeconomics. Review of Economic Studies 60, 35-52.

García-Peñalosa, C., Turnovsky, S.J., 2006. Growth and income inequality: A canonical model. Economic Theory 28, 25-49.

García-Peñalosa, C., Turnovsky, S.J., 2007. Growth, income inequality and fiscal policy: What are the relevant tradeoffs? Journal of Money, Credit and Banking 39, 369-394.

García-Peñalosa, C., Turnovsky, S.J., 2011. Income inequality, mobility, and the accumulation of capital: The role of heterogeneous labor productivity, unpublished ms.

Gerschenkron, A., 1962. Economic Backwardness in Historical Perspective. Harvard University Press, Cambridge MA.

Gorman, W., 1953,. Community preference fields. Econometrica 21, 63-80.

Guvenen, F., 2006. Reconciling conflicting evidence on the elasticity of intertemporal substitution: A macroeconomic perspective. Journal of Monetary Economics 53, 1451-1472.

Kraay, A., Raddatz, C., 2007. Poverty traps, aid, and growth. Journal of Development Economics 82, 315-347.

Kuznets, S., 1955. Economic growth and income inequality. American Economic Review 45, 1-28.

Li, H.Y., Zou, H.F., 1998. Income inequality is not harmful to growth: Theory and evidence. Review of Development Economics 2, 318-334.

Perotti, R., 1996. Growth, income distribution, and democracy: What the data say. Journal of Economic Growth 1, 149-187.

Persson, T., Tabellini, G., 1994. Is inequality harmful for growth? American Economic Review 84, 600-621.

Piketty, T., 2006, The Kuznets curve: Yesterday and tomorrow. In Banerjee, A.V., R Bénabou, and D. Mookherjee (eds.), Understanding Poverty, Cambridge University Press, Cambridge UK.

Ray, D., 1998. Development Economics, Princeton University Press, Princeton, NJ.

Solimano, A., 1998. The end of the hard choices? Revisiting the relationship between income distribution and growth," in A. Solimano (ed.), Social Inequality: Values, Growth, and the State. University of Michigan Press, Ann Arbor MI.

Sorger, G., 2000. Income and wealth distribution in a simple growth model. Economic Theory 16, 23-42.

Turnovsky, S.J., García-Peñalosa, C., 2008. Distributional dynamics in a neoclassical growth model: The role of elastic labor supply. Journal of Economic Dynamics and Control 32, 1399-1431.

Viaene, J.-M., Zilcha, I., 2003. Human capital formation, Income inequality, and growth. In T.S. Eicher and S.J. Turnovsky (eds.), Inequality and Growth: Theory and Policy Implications. MIT Press, Cambridge MA.

Figure 1 : Aggregate Dynamics

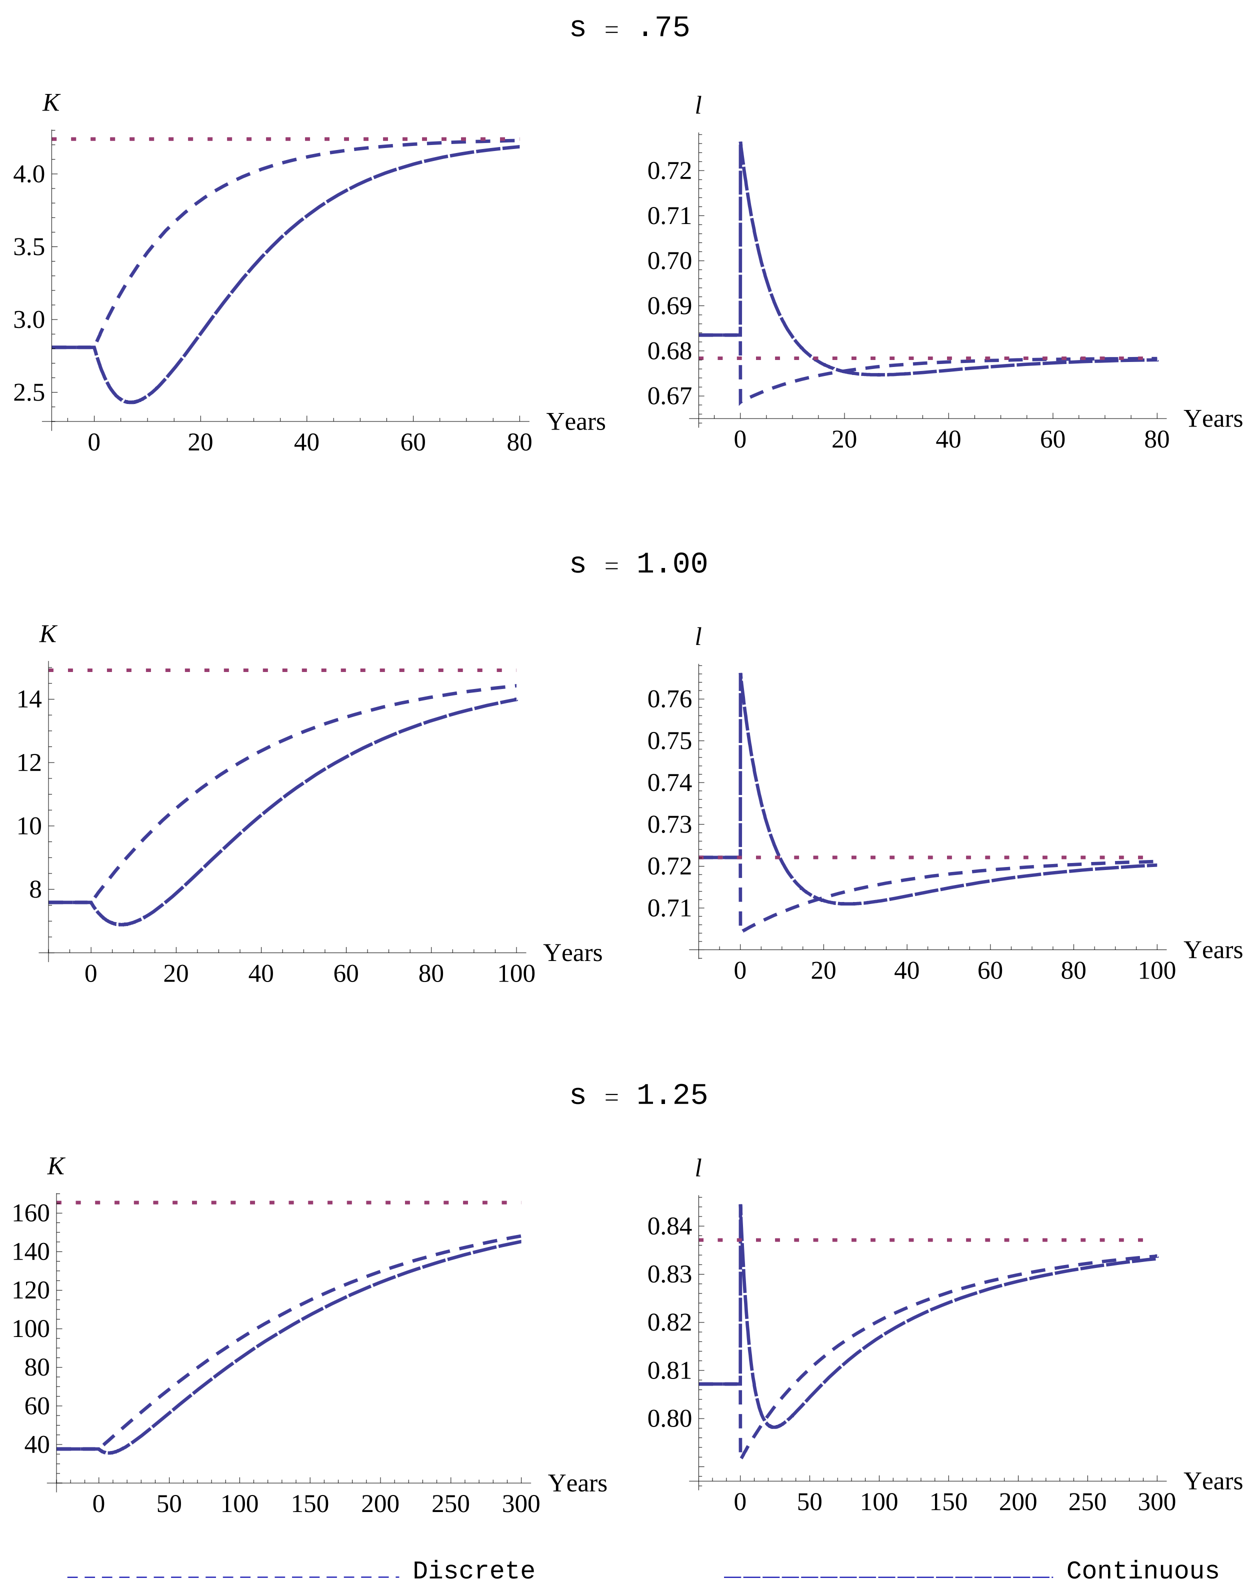

Figure 2 : Distributional Dynamics

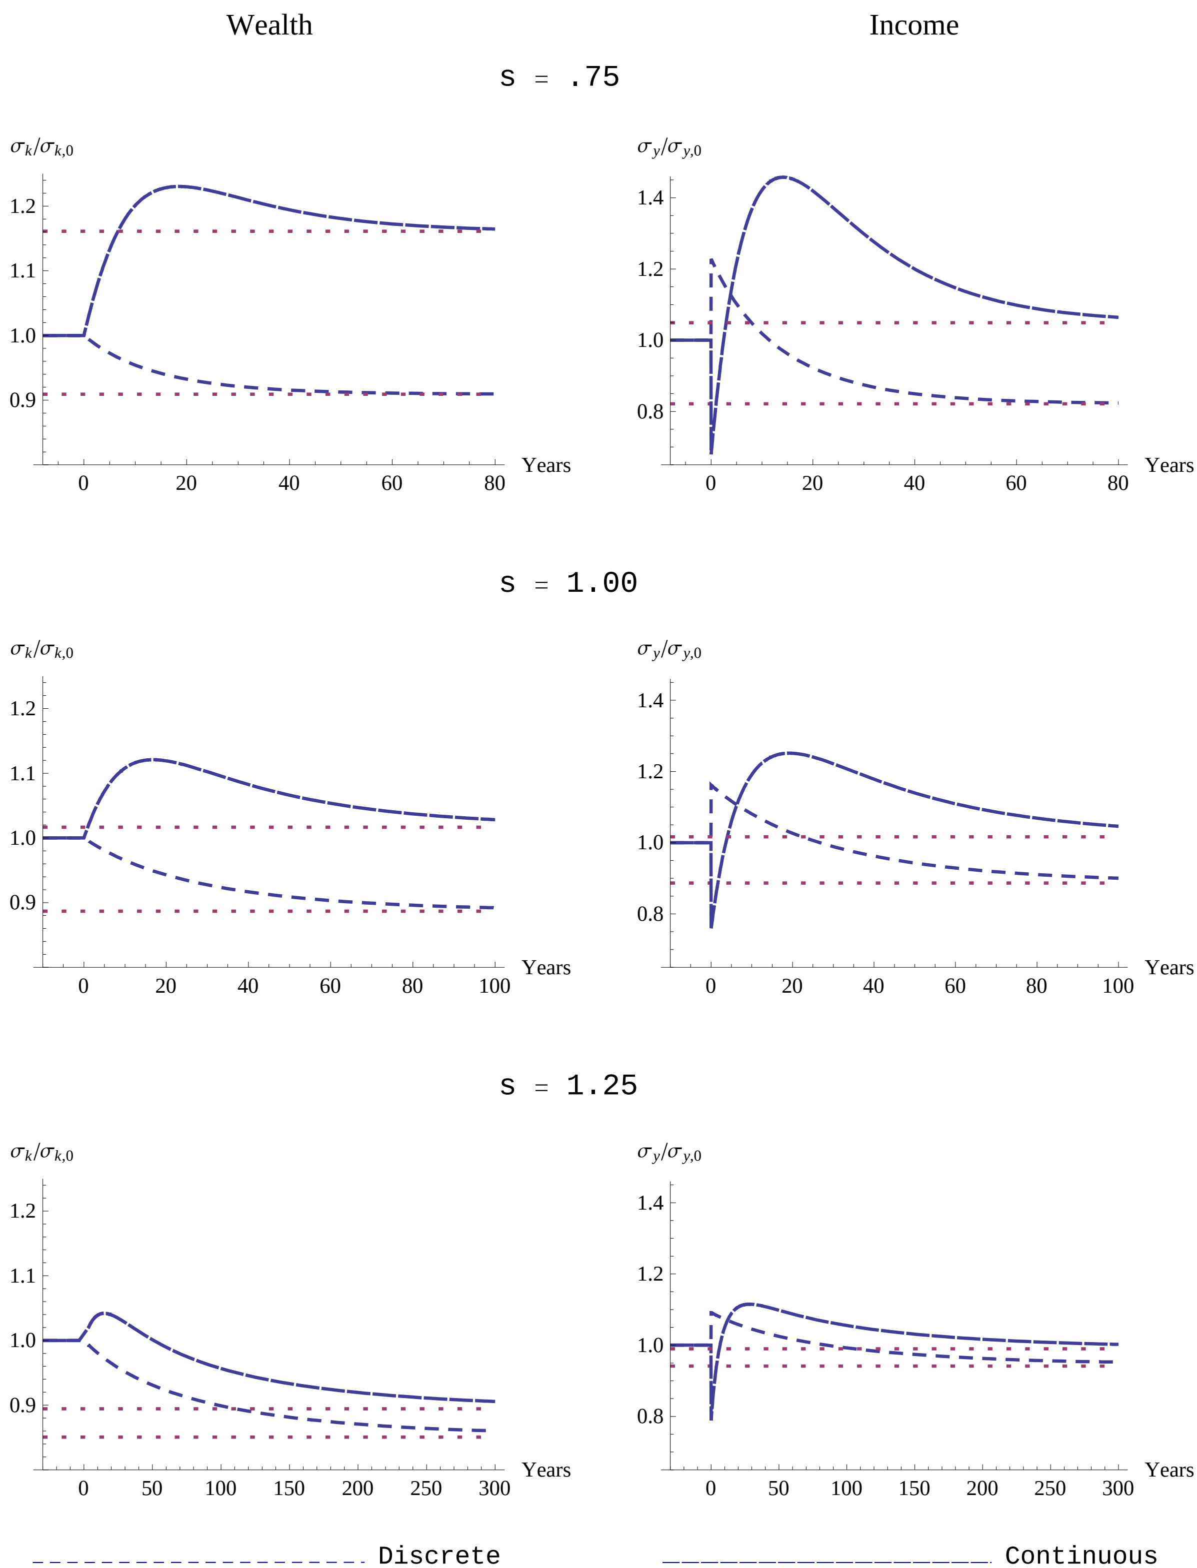

Figure 3 : Robustness of Distributional Dynamics

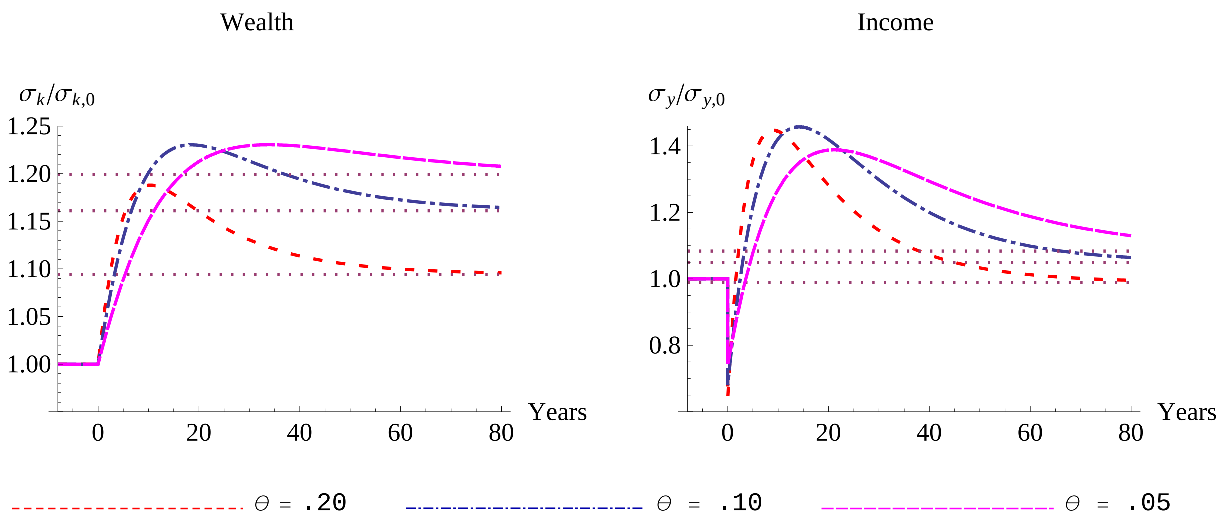

Figure 4 : Kuznets' Curve

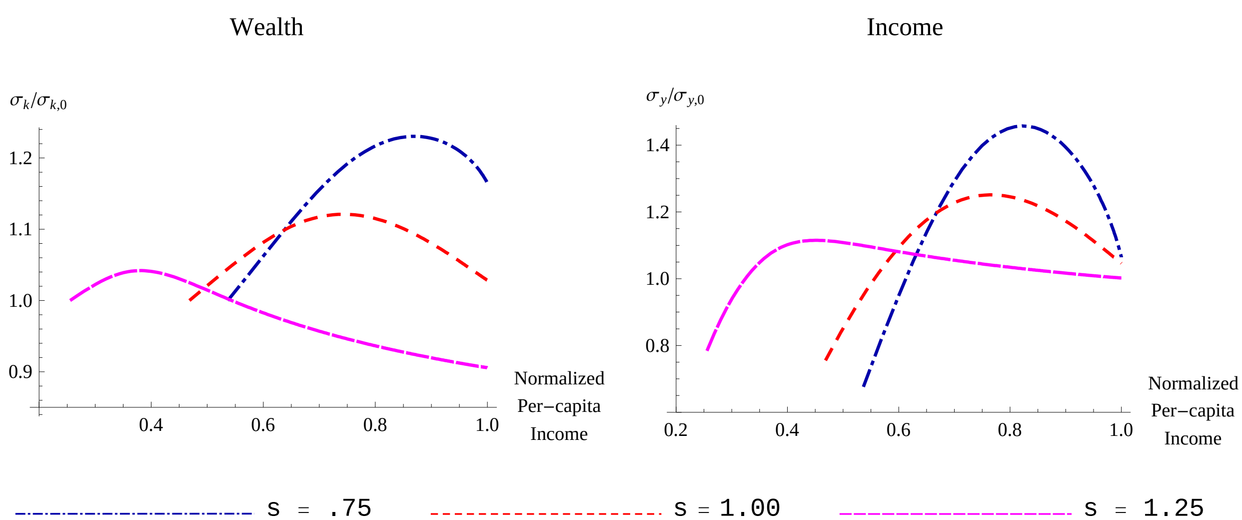

Figure 5 : Robustness of Kuznets' Curve

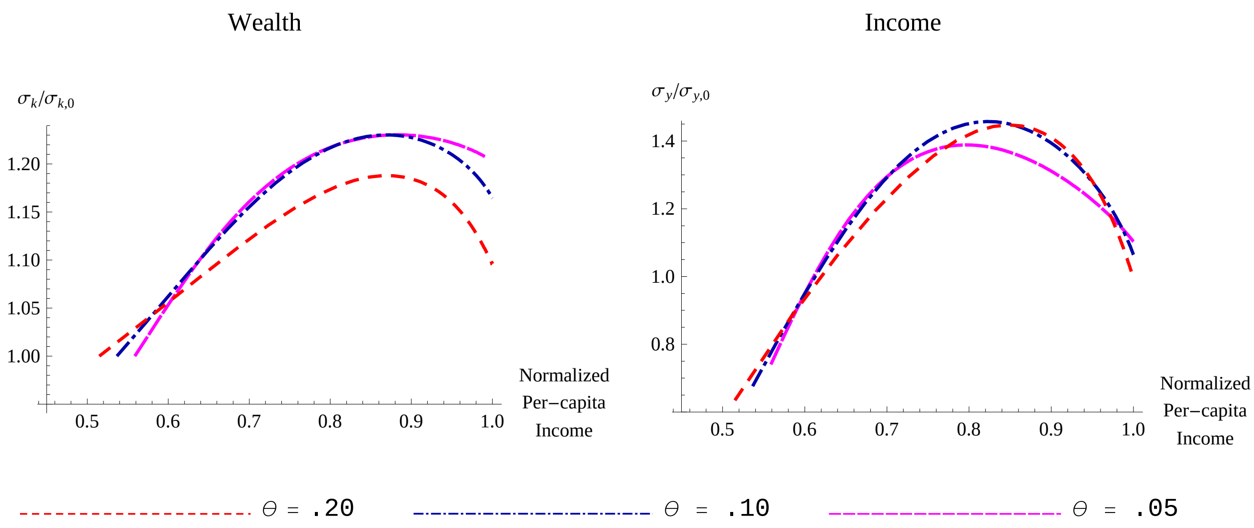

Figure 6 : Productivity Gaps, Catch up, and Distributional Dynamics