LIFE DURING STRUCTURAL TRANSFORMATION

Jonathan R. W. Temple
Huikang Ying

Discussion Paper 14 / 650

28 November 2014
Revised 8 December 2014

Department of Economics
University of Bristol
8 Woodland Road
Bristol BS8 1TN
United Kingdom

Life During Structural Transformation *

Jonathan R. W. Temple†
University of Bristol

Huikang Ying‡
University of Bristol

December 8, 2014

Abstract

We examine whether structural transformation leads to a Kuznets curve. We present a dynamic general equilibrium model with heterogeneous workers, occupational self-selection and selective migration, and calibrate the model to survey data for Malawi. We show that structural transformation raises living standards unevenly. As development proceeds, the movement of workers from agriculture is associated with rising wage inequality, rather than a Kuznets curve. The increase in sectoral wage inequality is pronounced for agriculture. At the same time, structural transformation is associated with major reductions in rural poverty, and eventually in urban poverty.

JEL Categories: D31, O15, O41

Keywords: Structural transformation, inequality, Kuznets curve, Roy model

*We are grateful to Leandro De Magalhães for guidance on the survey data for Malawi. This research was funded by the DFID-ESRC Growth Research Programme, ESRC award no. ES/J009601/1, 'Structural Change and Productivity Growth in Africa', PI Margaret McMillan.

†Corresponding author. Department of Economics, University of Bristol, 8 Woodland Road, Bristol, BS8 1TN, UK. Email: jon.temple@bristol.ac.uk.

‡Department of Economics, University of Bristol, 8 Woodland Road, Bristol, BS8 1TN, UK. Email: huikang.ying@bristol.ac.uk.

1 Introduction

If the fame of a paper can be measured partly in the misconceptions that surround it, this is more than usually true of Kuznets (1955). It is often assumed that his paper established a non-monotonic relationship between inequality and the level of development, the inverted-U of the Kuznets hypothesis. But he had relatively little data to work with, and characterized his paper as 'perhaps 5 per cent empirical information and 95 per cent speculation, some of it possibly tainted by wishful thinking' (Kuznets, 1955, p. 26).

He argued that, as economies industrialize, the path of inequality will be driven by changes in sectoral structure. If two sectors differ in their distributions of income, movements of population from one sector to the other can generate an inverted-U for aggregate inequality. In particular, the Kuznets hypothesis has been associated with a transition from rural agriculture towards urban manufacturing and services. In many countries, this process is far from complete, and its implications for inequality remain uncertain.

A variety of mechanisms that could generate an inverted-U relationship have been analyzed. In this paper, we concentrate on the one that Kuznets originally proposed, which Anand and Kanbur (1993) call the *Kuznets process*. In the standard version of the process, structural transformation changes the aggregate wage distribution by changing the relative importance of the sectoral wage distributions. The urban sector is usually assumed to have a higher mean income than the rural sector, and a greater degree of inequality. As the employment share of the urban sector increases, overall wage inequality changes over time. This process can easily give rise to an inverted-U relationship between wage inequality and the urban sector's share of total employment, of the form that Kuznets hypothesized; see, in particular, Robinson (1976) and Anand and Kanbur (1993). But it is often assumed that neither sectoral wage distribution is affected by migration, which is hard to justify. Further, the standard approach imposes a fixed ratio of the average urban wage to the average rural wage, which remains constant throughout a lengthy growth process. The reasons for this differential in average wages are left unexplained.

We develop a general equilibrium model that overcomes these objections. Building on the work of Lagakos and Waugh (2013), the model incorporates worker heterogeneity and occupational self-selection, but remains simple enough to calibrate to survey data, so that dynamic paths can be simulated. The model generates a rich account of the dynamics of living standards, wage inequality and poverty rates. Selective migration plays a central role: workers sort across sectors according to their comparative advantage, as in the Roy (1951) model. We show that this change in assumptions radically alters the nature of the Kuznets process. As the employment share of agriculture falls, changes in the composition of the agricultural and non-agricultural labor forces can generate a path of steadily increasing wage inequality, rather than an inverted-U.

We explore these effects in quantitative terms using a simple dynamic model. We calibrate the model to survey data for Malawi in 2010, and simulate the time paths that would arise under a hypothetical structural transformation. Our main result is that wage inequality rises over time. This finding does not rely on use of a particular inequality measure: under our maintained assumptions, Lorenz curves show an unambiguous increase in inequality, at the sectoral and aggregate levels. Our result will therefore hold for any summary measure of inequality that is Lorenz-consistent. Moreover, the result arises while assuming that workers can move freely between sectors, in line with the findings of Young (2013). Unlike many previous studies of the Kuznets process, our findings do not rely on unexplained barriers to mobility, or a fixed wage differential.

We also examine a range of other outcomes. We carry out simple welfare comparisons using stochastic dominance criteria, and show that structural transformation will be welfare-improving under our assumptions. But due to composition effects, the model gives rise to equilibrium differences in average living standards across sectors, even though workers are mobile. We also examine the time profiles of poverty measures. Although structural transformation reduces rural poverty quickly, similar reductions in urban poverty take longer to emerge. Structural transformation is therefore associated with the urbanization of poverty (Ravallion (2002)). Put differently, if we allow for selective migration, urban poverty does not fall to a low level until relatively late in the process.

It is important to emphasize that our paper is about the Kuznets process, rather than the Kuznets curve. There are many influences on inequality that we do not analyze, and the paper does not aim to give a complete account of the different forces that could give rise to a Kuznets curve, on which there is a large literature.[1] Instead, we focus on a single mechanism, and show that it does not readily generate the inverted-U that Kuznets anticipated.

This finding can be linked to the contested history of cross-country empirical studies. It has been known since the mid-1970s that a Kuznets curve sometimes appears in cross-country data for a single point in time (for example, Ahluwalia (1974, 1976a,b)). But there is no consensus on whether a Kuznets curve arises in longitudinal data: see Barro (2000) and Deininger and Squire (1998) for different views. Kanbur (2012) observes that a number of countries have undergone secular increases in inequality. The cases he cites include Bangladesh, China, Ghana, India, and South Africa, and also some of the Latin American countries, at least before 2000. These findings suggest that, as Kuznets acknowledged, his inverted-U hypothesis might have involved some wishful thinking.

The paper is structured as follows. Section 2 relates our contribution to the literature. Section 3 sets out the model, and section 4 describes the parameter choices

[1]For example, there is a literature on wealth inequality in one-sector models where households differ in their assets: see Álvarez Peláez and Díaz (2005) and Caselli and Ventura (2000).

that will be adopted in the simulations. Section 5 describes some of the features of the simulated structural transformation, before section 6 analyzes its consequences for wages, welfare, wage inequality and poverty. Section 7 examines the robustness of our findings, while section 8 concludes.

2 Discussion

As the introduction noted, the traditional models of the Kuznets process have major drawbacks. The difference in average incomes across sectors is often imposed exogenously, rather than derived from microeconomic foundations. But for an inequality analysis, it seems risky to assume that an income differential is independent of growth and structural transformation. More fundamentally, the differential is often assumed to reflect unspecified barriers to the intersectoral mobility of labor, an assumption which is not supported by recent work. In a detailed analysis, Young (2013) finds that survey data for 65 countries point to considerable mobility. Rural-urban income gaps appear to be driven mainly by the higher skill intensity of urban production, a composition effect. Further, the high incidence of migration from urban to rural areas is suggestive of sorting on the basis of unobservable skill, rather than barriers to mobility that maintain a higher wage in urban areas. In our analysis, we assume that workers are perfectly mobile between sectors, and rural-urban differences in average wages are driven by differences in skill composition, consistent with Young (2013).

The standard accounts of the Kuznets process are restrictive in other ways. It is often assumed that workers are homogeneous, and that the sector-specific distributions of wages do not change as individuals move from one sector to the other. This rules out selective migration. It also raises the question of precisely how the within-sector distributions can be invariant to population movements. As characterized by Anand and Kanbur (1985, p. 43), the migration process is one in which 'a representative subsample of the rural distribution shifts to the urban sector and relocates itself representatively across the urban distribution'. This is not an easy outcome to achieve, unless it is assumed that workers in a given location simply receive, at each instant, a random draw from the wage distribution for that location. But then it is not clear why wages differ across homogeneous workers, or why workers remain in a sector with a less favorable distribution.

Our approach avoids these drawbacks, building on the work of Lagakos and Waugh (2013). In their two-sector model, individuals have a level of productivity that would apply if they work in agriculture, and a separate level that would apply in non-agriculture. Given these distinct skill levels, each individual chooses, at every date, whichever sector will maximize their income at that date. Depending on their choice of sector, one of their sector-specific productivity levels will be relevant to production or 'realized', while the other is unobserved or 'latent'. The analysis is based on occupational self-selection according to comparative advantage, along the lines sketched by Roy (1951). If the sectors are considered to be spatially distinct, we can think of the Roy equilibrium as a spatial equilibrium in which no worker prefers another location to their current location.[2]

We show that, in this framework, time paths for wage inequality can look very different to the standard Kuznets process. This is because composition effects now play a central role: structural transformation will change the skill composition of the labor force in each sector. When a single worker moves, one productivity level changes from realized to latent, and the other from latent to realized. The sectoral wages per efficiency unit of labor also evolve over time. Hence, a structural transformation will generate changes in wage inequality and other outcomes, without needing to assume either an exogenous income differential across sectors, or time variation in the individual-specific distributions of productivity.

To the extent that this mechanism is relevant, it becomes easier to explain why the evidence on the Kuznets curve has remained ambiguous and contested. Among recent analyses, Frazer (2006) highlights the differences in the paths of inequality within countries. Angeles (2010) concentrates on the relationship between the Gini coefficient and either the employment share of non-agriculture, or the share of population living in urban areas. He writes that 'when we examine countries individually we do not find a systematic pattern linking employment outside agriculture or urban population with inequality' (p. 469). His Table 3 suggests that the turning points in estimated quadratic relationships often arise for high shares of non-agricultural employment: this is more suggestive of monotonic increases or declines in inequality, than of non-monotonic relationships.

Despite the ambiguous evidence on the inverted-U hypothesis, it remains much discussed. It is sometimes seen as reassuring in two different ways. First, it suggests that rising inequality is a transitory phase of development, to be reversed later by market forces or public policy. Second, to the extent that inequality arises from a rural-urban income differential, policy interventions could be effective; an intervention that reduces the rural-urban differential will typically reduce inequality. But neither of these consolations is available in the Lagakos-Waugh framework. An economy may find itself on a development path in which inequality will continue to rise as development proceeds, asymptotically approaching a constant; and this happens even in an undistorted economy, in which workers can move freely between sectors. In that case, as Young (2013) emphasizes, policy responses would have to look beyond eliminating distortions.

We also study the evolution of rural and urban poverty as structural transformation proceeds. In an influential paper, Ravallion and Datt (1996) studied the relationship between poverty rates and the sectoral composition of growth, using data on sectoral poverty rates and average incomes for India. A potential objection is that, in a multi-

[2] An implication of the Roy model is that workers will generally strictly prefer their current location to other locations. Contrary to many spatial models with homogeneous workers, there is no presumption that workers are indifferent between locations in equilibrium.

sector model, average incomes at the sector level are general equilibrium outcomes. For example, if labor mobility equalizes wages and thereby marginal products, this links average products across sectors. The average product might rise in one sector because TFP increased in the other sector; movements in average products can be misleading about the underlying sources of growth.[3] By simulating an (admittedly simple) structural model, we can investigate the effects at work, and study counterfactuals, in ways that are ruled out by reduced-form empirical approaches.

Finally, we note that a wider literature has considered the nature of recent structural changes in selected African countries, partly associated with the fast growth of the early 2000s. Work in this vein includes Jedwab et al. (2014), McMillan et al. (2014), and McMillan and Harttgen (2014). Our work is less empirically driven: we do not investigate historical structural transformations or the experience of particular countries. Instead, we introduce a framework in which structural transformation influences wage inequality, sectoral income differentials, and poverty rates. In principle, it could be combined with some of the ideas in the wider literature, such as the analysis of home production in Gollin et al. (2004) or the 'consumption cities' of Jedwab et al. (2014); more ambitiously, the analysis could seek to follow Helpman et al. (2010) in combining firm heterogeneity, unemployment and match-specific heterogeneity. We leave these possibilities to future work.

3 The model

The underlying structure of the model is simple. We consider a small open economy with two sectors, in which the production equilibrium at a given point in time is essentially that of Lagakos and Waugh (2013). We then study how the equilibrium changes over time in response to technical progress, factor endowments and population growth. In effect, we allow these forces to generate a sequence of Lagakos-Waugh equilibria, with factor prices changing over time. Since the model does not have a closed-form solution, our main findings will be derived from simulations, where certain key parameters are inferred from detailed survey data for Malawi.

3.1 Model assumptions

In more detail, we consider a two-sector or dual economy model, the structure traditionally used to analyze the Kuznets process.[4] We assume there is a rural sector which produces an agricultural good, and co-exists with an emerging urban 'non-agricultural' sector that produces a composite good. The composite good can be interpreted as a

[3]Ravallion and Datt acknowledge this interdependence in their paper, but their preferred interpretations tend to equate 'rural growth' and 'urban growth' with changes in mean incomes, rather than with underlying forces such as TFP improvement or capital accumulation.

[4]For reviews of the literature on dual economy models, see Kanbur and McIntosh (1988), Temple (2005) and Gollin (2014).

bundle of manufacturing goods and services; to keep the analysis simple, we abstract from the distinction between manufacturing and services. We treat the urban good as the numéraire and the unit of account. For brevity, we treat the terms 'rural' and 'agricultural' as interchangeable, and likewise 'urban' and 'non-agricultural'.

Also in keeping with tradition, we use simplifying assumptions to focus on the mechanism that Kuznets introduced. We assume that both goods are traded on world markets, and there are no international flows of labor or capital. The economy we consider is too small to influence world prices through its consumption or production decisions. A major advantage of these assumptions is that we do not have to specify household preferences. Instead, the price of the agricultural good relative to the price of the non-agricultural good will be determined exogenously, by world prices. Given that we rule out international borrowing and lending, trade will be balanced at each date. At a given date, the economy will import one good and export the other, but for the issues we consider here, the trade pattern is a secondary concern. In the terminology of trade theory, we are considering a production equilibrium, and this is sufficient for our purposes.[5]

There are three classes of agents in the economy: landowners, capitalists and workers. Within these classes, we abstract from most forms of heterogeneity. Landowners consume the entirety of the rents from land, and do not supply any labor. Capitalists own physical capital, and again do not supply any labor. They save a constant fraction, identical across capitalists, of their capital income, and consume the rest. The saving by capitalists will be the only source of investment in the economy. Finally, workers own neither land nor capital, but supply labor and consume their wage income. The landowners are infinitely-lived and live in identical households that grow at an exogenous rate n; this is also true of capitalists.[6]

The production technologies in the two sectors are Cobb-Douglas and have constant returns to scale, but with output elasticities that differ across sectors. In per worker terms, they are given by

$$
\begin{aligned}
y_a &= k_l^{\alpha} \cdot k_a^{\beta} \cdot (A_a \cdot e_a)^{1-\alpha-\beta} \\
y_m &= k_m^{\mu} \cdot (A_m \cdot e_m)^{1-\mu} \qquad (1)
\end{aligned}
$$

where inputs are normalized by the total (i.e., economy-wide) number of workers.[7] k_l denotes land usage per worker. k_a and k_m are the stocks of physical capital in the two sectors, divided by the total labor force. e_a and e_m represent the per capita effective labor employed in the two sectors. A_a and A_m are technology indices, assumed to grow at the exogenous rates g_a and g_m respectively.

In both sectors, profit-maximizing firms choose factor inputs under conditions of

[5] See Graham and Temple (2006) for a similar approach.

[6] We can think of newborns as inheriting the type of their parents; see Song, Storesletten and Zilibotti (2011) for a related approach in the context of China.

[7] The time subscript t is omitted here and henceforth when no confusion is caused.

perfect competition, and make zero profits in equilibrium. Capital and (effective units of) labor receive their marginal products. We use y_a and y_m to denote the levels of output in the two sectors divided by the total labor force. The relative price of the agricultural good is denoted by p, and we treat it as fixed over time. Aggregate output per worker y is then given by $p \cdot y_a + y_m$.

These assumptions are simple enough that we can introduce variation in the sector-specific productivity levels of workers. We describe the productivity levels of worker i in the two sectors by a vector $\{z_a^i, z_m^i\}$. We abstract from human capital decisions or firm-specific productivity levels, and hence this vector will be treated as time-invariant. When workers are employed in the agricultural sector, we can think of their productivity z_a as 'realized' and z_m as 'hidden', and *vice versa* for an individual employed in non-agriculture. This will generate non-trivial changes in the (realized) productivity distributions, within sectors and for the whole economy, as the economy develops.

Our labor market assumptions follow Lagakos and Waugh (2013). We assume that workers can move freely between sectors, and choose whichever sector maximizes current labor income at a given date. The labor market in each sector clears, and so all workers are fully employed at each date. Importantly, although workers are heterogeneous in productivity, they are perfect substitutes at fixed ratios. Let Ω_a and Ω_m denote the sets of workers who are employed in agriculture and non-agriculture respectively, at a given date. The supplies of efficiency units of labor, relative to total employment, can then be defined as:

$$e_a \equiv \frac{1}{L} \int_{i \in \Omega_a} z_a^i di$$

$$e_m \equiv \frac{1}{L} \int_{i \in \Omega_m} z_m^i di \tag{2}$$

To study the production equilibrium at a given date, we will need to understand how the sets Ω_a and Ω_m are determined. Given that, within a sector, workers are perfect substitutes at fixed ratios, we can think of an individual worker's wage in a given sector as the product of a wage per efficiency unit of labor for that sector, and their sector-specific productivity level. The agricultural and non-agricultural wages per efficiency unit of labor play a central role; we refer to them as 'effective wages' for brevity.

Profit-maximizing employers in a given sector will employ workers until the marginal product of an efficiency unit of labor is equal to the effective wage in that sector. Hence, we can write the two effective wages as

$$\zeta_a = (1 - \alpha - \beta) \frac{p y_a}{e_a}$$

$$\zeta_m = (1 - \mu) \frac{y_m}{e_m}$$

and define their ratio,

$$q \equiv \frac{\zeta_a}{\zeta_m} \tag{3}$$

Hence, q is the ratio of effective wages. Importantly, there is no expectation that $\zeta_a = \zeta_m$ (or $q = 1$) even under perfect mobility. Instead, the equilibrium q will reflect the decisions of workers on where to locate, and q will be determined jointly with the allocation of workers in which each worker is maximizing their income.

Hence, as in Lagakos and Waugh (2013), we consider occupational self-selection by income-maximizing workers. The necessary and sufficient condition for worker i to select the agricultural sector is

$$\zeta_m \cdot z_m^i \leq \zeta_a \cdot z_a^i$$

or more compactly,

$$\frac{z_m^i}{z_a^i} \leq q$$

Put differently, to select agriculture, the worker must have a comparative advantage in agriculture. Conversely, the cut-off rule for a worker to enter non-agriculture is $\frac{z_m^i}{z_a^i} > q$.

The equilibrium labor income of worker i is given by

$$w^i = \max\left\{\zeta_a \cdot z_a^i, \zeta_m \cdot z_m^i\right\} \tag{4}$$

This mechanism determines the equilibrium labor composition of each sector, namely the sets Ω_a and Ω_m, at a given date. We can define the associated employment shares as

$$l_a = \text{Prob}\left\{i \in \Omega_a\right\} = \text{Prob}\left\{\frac{z_m^i}{z_a^i} \leq q\right\}$$

$$l_m = \text{Prob}\left\{i \in \Omega_m\right\} = \text{Prob}\left\{\frac{z_m^i}{z_a^i} > q\right\} = 1 - l_a \tag{5}$$

and the levels of effective labor supplies are then defined as in (2).

Finally, we assume that capital is perfectly mobile between sectors, and hence the returns to capital are equal across the two sectors at each date:

$$r + \delta = \beta \cdot \frac{y_a}{k_a} = \mu \cdot \frac{y_m}{k_m} \tag{6}$$

where r denotes the interest rate and δ is the depreciation rate of physical capital. This completes the description of the equilibrium at a given date.

3.2 Dynamic paths

Now we depart from Lagakos and Waugh, by embedding the self-selection mechanism in a simple growth model. Given that we are studying a small open economy facing a fixed relative price, structural transformation will be driven by a combination of technical progress in each sector and capital accumulation, with no role for non-homothetic preferences. Since the production technologies differ across sectors, technical progress, capital accumulation and population growth will alter the marginal product of an efficiency unit of labor in each sector. Some workers will therefore see their comparative advantage change from one sector to the other. Since mobility is costless and workers maximize their current income at each date, workers will switch sectors until equilibrium is restored at a new allocation and a new set of factor prices. Each sector's share of total employment will change over time, and the individual-level productivity distributions within each sector will also change, with implications for wage inequality and poverty within sectors, and for the whole economy.

We are primarily interested in the consequences of structural transformation rather than its causes, so we keep the growth model deliberately simple. We assume that workers consume the entirety of their labor income. Owners of capital save a fixed fraction of their capital income, and we denote this fraction by s_k.[8] Since there are no international capital flows, and no government, domestic saving implies domestic investment. The transition equation for the capital stock is given by

$$\begin{aligned} k_{t+1} &= \frac{1}{1+n}\left[s_k \cdot (r_t+\delta)\cdot k_t + (1-\delta)\cdot k_t\right] \\ &= \frac{1}{1+n}\left[s_k \cdot \eta_{kt}\cdot y_t + (1-\delta)\cdot k_t\right] \end{aligned} \tag{7}$$

where η_{kt} is the share of capital income in total income at date t. This can be written as $\pi_{at}\cdot\beta + (1-\pi_{at})\cdot\mu$, where π_{at} denotes the share of agricultural value added in total value added at date t.

When combined with occupational self-selection, the dynamic equilibrium of this economy is a time series sequence

$$\left\{y_a, y_m, y, k_a, k_m, k, q, l_a, l_m, e_a, e_m, r, w^i\right\}_{t=1,2,3,\ldots}$$

that satisfies the law of motion (7) and the intratemporal (static) equations (1) to (6) at each date, given a certain specification of initial endowments, including a strictly positive initial capital stock k_0.

Under the above assumptions, the model is dynamically recursive: given factor endowments and efficiency levels, we can solve for the Lagakos-Waugh equilibrium at date t, update the factor endowments and efficiency levels, solve for the equilibrium at date $t+1$, and so on. Hence, our simulations are essentially studying a sequence of Lagakos-Waugh economies, and the numerical analysis is straightforward.[9] Ying (2014) studies a model with Ramsey consumers, and simplifies the problem by ruling

[8] Our approach to modelling aggregate investment, based on factor incomes, is therefore in the tradition of Kaldor and Pasinetti. See Bertola et al. (2006) for a discussion of the relevant literature.

[9] Analysing a forward-looking saving/investment decision is computationally demanding in this context, because of the need to solve numerically for the Lagakos-Waugh equilibrium at each date for each candidate set of paths for the state variables. We have experimented with a shooting algorithm, but have not found an approach that keeps the computational time manageable.

out dependence between individuals' agricultural and non-agricultural productivity levels, while considering a two-parameter family for the sectoral wage distributions. These assumptions lead to closed-form solutions and some analytical results, and also allow simulations to be implemented easily, but the simplicity comes at a price: wage inequality is necessarily constant over time.

4 Calibration

In the following sections, we analyze the paths of welfare, wage inequality and poverty rates during structural transformation. The specification of a joint distribution for individual-specific sectoral productivity levels will play a key role. We calibrate its parameters based on survey data for Malawi, using an approach similar to the one Lagakos and Waugh adopt for US data. We fix some of the remaining model parameters based on precedents in the literature. We then calibrate the model's initial conditions to match selected characteristics of Malawi in 2010, and simulate the transitional dynamics for the following 100 years, for given paths of sectoral technology levels. We should clarify our objectives at the outset: we are not seeking to match historical data for Malawi, or to make serious predictions about its future growth path. Instead, we aim to isolate the effects of structural transformation on wage inequality and other outcomes. Among the many forces that we omit are schooling decisions, intergenerational flows through bequests, estate taxation, progressive income taxation, and major shocks such as wars.[10]

4.1 Productivity distribution

In generating time paths for sectoral and economy-wide wage inequality and poverty rates, the assumptions on the joint distribution for individual productivity levels $\{z_a^i, z_m^i\}$ will be critical. Following the quantitative analyses in Lagakos and Waugh (2013), we assume that workers' productivity endowments are drawn from a continuous joint distribution with dependence between the two sectoral productivity levels; for simplicity, this same joint distribution also applies to newborns.

The cumulative density function of our chosen distribution is given by

$$H(z_a, z_m) = C\left[H_a(z_a), H_m(z_m)\right]$$

where

$$H_a(z_a) = e^{-z_a^{-\theta_a}}$$

$$H_m(z_m) = e^{-z_m^{-\theta_m}}$$

[10]On inequality in human capital, see Morrisson and Murtin (2013); on the potential relevance of the other forces mentioned, see Piketty (2006).

are the two marginal distributions, given by Fréchet distributions with sector-specific shape parameters θ_a and θ_m respectively. In the case of agriculture, a lower θ_a corresponds to higher average productivity and higher dispersion across individuals within agriculture, and similarly a lower θ_m corresponds to higher average productivity and higher dispersion for non-agriculture.

The expression $C(u, v)$ denotes a Frank copula, which models the potential dependence between the sectoral productivity levels as follows:

$$C(u, v) = -\frac{1}{\rho} \log \left[1 + \frac{(e^{-\rho u} - 1) \cdot (e^{-\rho v} - 1)}{e^{-\rho} - 1} \right]$$

The copula forms the joint distribution with the required marginals, while allowing for a dependence parameter, $\rho \in (-\infty, \infty) \setminus \{0\}$. A positive ρ implies that workers who are relatively productive in agriculture will also tend (on average) to be relatively productive in non-agriculture. The case where $\rho = 0$ corresponds to independence in the sectoral productivity endowments: the Frank copula then reduces to $C(u, v) = u \cdot v$. This is similar to the case analyzed extensively, with Ramsey consumers, in Ying (2014). That case is analytically more tractable, but as noted previously, it implies that wage inequality is necessarily constant over time.

The above specification requires us to choose values for the parameters θ_a, θ_b and ρ. We follow an approach similar to Lagakos and Waugh (2013), in which the three parameters are chosen to match three moments of the sectoral wage distributions: the variances of log wages in agriculture and non-agriculture, and the ratio of average wages across the two sectors. To do this, we draw on micro-level data from the Third Integrated Household Survey of Malawi (IHS3) conducted by the World Bank and the National Statistical Office of Malawi. De Magalhães and Santaeulalia-Lopis (2014) use these data to construct a variety of statistics, and we use their figures as the basis for the relevant moments. The survey provides sectoral labor income for individuals over the period March 2010 to March 2011.

One limitation of these data is that we do not have a panel, and hence we cannot follow Lagakos and Waugh in separating the permanent component of wages from the transitory component. To address this, we conduct two calibration experiments based on different assumptions. In the first experiment, we assume that the wages in the survey data do not have a transitory component, so that the observed distribution of wages in a given sector maps straightforwardly into the distribution of sector-specific productivity levels. When the three moments are simultaneously matched, we obtain the parameters $\theta_a = 2.15$, $\theta_m = 1.18$ and $\rho = 5.24$. For the second experiment, we assume that the relative importance of the transitory component is the same as that found in Lagakos and Waugh.[11] In this experiment, we obtain the parameters $\theta_a = 3.40$, $\theta_m = 1.53$ and $\rho = 6.93$.

[11] For the details of their calculation, see the online appendix of Lagakos and Waugh (2013).

The two experiments have similar implications for the productivity endowments. As in Lagakos and Waugh, the variance in individual productivity is greater in non-agriculture than in agriculture. The finding that $\rho > 0$ implies that workers with relatively high productivity in agriculture tend (on average) to have relatively high productivity in non-agriculture, which is reasonable if some skills are transferable across sectors. In the remainder of the paper, we focus on the simulation results from our second experiment. The first experiment, with more dispersion in productivity levels across individuals, yields similar predictions on growth and structural transformation, and also qualitatively similar results for outcomes such as wage inequality and poverty rates.

4.2 Parameters and initial values

The other parameter assumptions are shown in Table 1. The output elasticities of land and physical capital in agriculture follow the estimates in Mundlak and Hellinghausen (1982). The output-capital elasticity for non-agriculture is calibrated so that the economy-wide capital share matches the average value for sub-Saharan African countries in the data of Caselli and Feyrer (2007). For the rate of technical progress in agriculture, we adopt the estimate of Martin and Mitra (2001) for low-income countries, while the rate for non-agriculture is chosen to be slightly higher. The depreciation rate for physical capital is set to 6% a year, which follows Caselli (2005), among others.

[Table 1 about here.]

The initial state of the economy will be calibrated to match data for Malawi in 2010. For this purpose, we draw on two data sets: the Africa Sector Database (ASD, de Vries, Timmer and de Vries, 2013) and the Penn World Table 7.1 (PWT, Heston, Summers and Aten, 2012). We use the ASD for sectoral data on productivity and employment shares, but we follow some other authors in assuming that a certain fraction of agricultural output is not measured in the national accounts. This is a natural assumption in the context of subsistence agriculture, and has been discussed by Parente et al. (2000) in particular.[12]

To infer the 'true' share of agricultural value added in total value added, we use the data on employment shares to solve for the migration cut-off, q, and thus e_a and e_m given the endowment distributions. We can then infer the share π_a by noting an implication of the model:

$$\frac{1-\alpha-\beta}{1-\mu} \cdot \frac{e_m}{e_a \cdot q} = \frac{1-\pi_a}{\pi_a} \tag{8}$$

Some of the remaining parameters are calculated from data for Malawi over 1965-2010. The population growth rate that we project forwards is assumed constant, and

[12] A further complication, which we do not allow for, is that some off-farm employment is classified as employment in agriculture, as Schmitt (1989) discusses. Herrendorf and Schoellman (2013) show that a significant measurement problem appears to be present even in the data for the USA.

based on the mean for 1965 to 2010. We use the PWT to obtain the ratio of investment to GDP, and infer the saving rate of capitalists that determines the evolution of the capital stock. In estimating the initial stock of physical capital, for 2010, we face the problem that a perpetual inventory calculation, taken at face value, yields an implausibly high capital-output ratio in 2010. This is perhaps not surprising given the cautions about such calculations emphasized by Pritchett (2000). We therefore construct an alternative capital stock series which assumes that a fraction of each year's investment is wasted, implying a more plausible capital-output ratio in 2010. The precise calculation draws on the Public Investment Management Index of Dabla-Norris, Brumby, Kyobe, Mills and Papageorgiou (2012).

5 Model dynamics

5.1 Structural transformation

Before we turn to describing the consequences of structural transformation in detail, we first provide an overview of the simulation results. Some of the relevant outcomes are shown in figure 1. Panel (a) shows the time path of the ratio of effective wages, or $q \equiv \zeta_a/\zeta_m$. As we saw previously, it is this quantity which plays a central role in the allocation of labor.

Given our assumptions on production technologies and capital accumulation, q declines over time. An increasing proportion of workers will see their comparative advantage change from agriculture to non-agriculture. As a result, some of those initially employed in agriculture will switch to non-agriculture, generating the decline in the agricultural employment share shown in panel (b). Under our assumptions, the transformation is slow: it takes a century to reduce the employment share in agriculture from 65 percent to 10 percent. This is not necessarily unrealistic. According to the ASD data set, between 1965 and 2010, agricultural employment in Malawi decreased from 84 percent of total employment to 65 percent, or less than twenty percentage points.

Some of the explanations for the slow pace of the structural transformation are direct consequences of our parameter assumptions: the rate of investment is low, and we have assumed a relatively small gap between sectoral productivity growth rates. A less familiar argument is that the extent of the dispersion in productivity levels across individuals also plays a role in the rate of transformation. This result was first derived in Ying (2014) in the context of a model without dependence ($\rho = 0$ in our notation).

[Figure 1 about here.]

As is standard, structural transformation involves a reallocation of capital as well as labor. This is shown in panel (c), which shows the share of capital that is employed in agriculture. As workers and capital leave agriculture, its share of aggregate output declines. In panel (d) we show two measures of the agricultural output share. The solid line is the 'true' measure generated from the simulation. The dashed line corresponds to the output share that would be observed in the data, given that some agricultural output is not measured in the national accounts; in principle, it is this measure which could be compared against output share data.

The results presented thus far are conventional. In panels (e) and (f) we show results that are more distinctive. The panels show the average individual productivity for each sector and how it changes over time. The average individual productivity increases from 1.23 to 1.59 in agriculture over the transformation, while it decreases from 5.31 to 2.70 in non-agriculture. This pattern arises naturally from changes in comparative advantage. In the long run, only those workers who are highly productive in agriculture relative to non-agriculture will remain in agriculture; meanwhile, the non-agricultural sector will have seen workers enter whose comparative productivity levels initially led them to work in agriculture. This effect leads to a decline in the average level of individual productivity in non-agriculture — but, as we now discuss, not a major decline in the average output per worker of the sector.

5.2 Relative labor productivity

In this subsection, we briefly discuss the ratio of average products of labor, often denoted by *RLP* in the literature. This is conventionally defined (for example, Bourguignon and Morrisson (1998)) as the ratio of output per worker in non-agriculture to output per worker in agriculture:

$$\text{RLP} = \frac{y_m/l_m}{p \cdot y_a/l_a}$$

where l_a and l_m correspond to the employment shares of agriculture and non-agriculture respectively. Temple and Woessmann (2006) calculated *RLP* for a large number of countries. Consistent with earlier work, they found that *RLP* is greater than one in most countries, and has declined over recent decades. Hence, the average product of labor is higher in non-agriculture, but grows faster in agriculture. We should not expect average products to be equalized across sectors. In a market economy with homogeneous labor, the usual efficiency condition is that marginal products are equalized, but then average products will typically differ.[13] The significant complication in the heterogeneous labor model of Lagakos and Waugh (2013) is that *RLP* will be influenced by the distributions of individual productivity levels within each sector.

[Figure 2 about here.]

Figure 2 shows the time path of *RLP* in our simulation. The solid line, representing the 'true' measure when agricultural output is correctly measured, indicates that *RLP*

[13] Put differently, equilibrium differences in average products across sectors are a condition of efficiency, rather than a sign of its absence.

is around two. The dashed line constructs the *RLP* measure that a researcher would see in the data, if some agricultural output is not measured in the national accounts. This corresponds to a higher (measured) level of *RLP* and amplifies its absolute downwards movement.[14] Nevertheless, in common with some other work in this field, the model clearly understates the downwards trend in *RLP*, compared to the large declines sometimes observed in the data. One possible explanation for those declines is that agricultural output becomes better measured as countries develop and subsistence agriculture becomes less relevant. Alternatively, the predictions of the current model could be improved by introducing home production, perhaps along the lines of Gollin et al. (2004). We treat these variations as beyond the scope of the current paper.

6 Wages and welfare

Now that we have described the overall patterns of a hypothetical structural transformation, we can turn to the implications for wages, welfare, inequality and poverty. We look not only at the wage distributions directly, but also use stochastic dominance criteria to make distributional comparisons. Throughout, we consider the welfare only of workers, not landowners or capitalists, and all our inequality results concern inequality in wage income. As we emphasized earlier, we are not seeking to compare possible drivers of inequality.

6.1 Productivity and wage distributions

When examining changes in wage distributions, it is useful to bear in mind that there are two forces at work. A worker's maximized income has two components: their 'realized' productivity level, which varies with their sector, and the effective wage which applies in that sector. Therefore, structural transformation reshapes the wage distributions in several different ways. First and most obviously, the relative importance of the sector-level distributions will change, as in the standard analyses of the Kuznets process. Technical progress and capital accumulation induce changes in effective wages, which have a direct effect on labor incomes, for a given composition of employment within each sector. But there is also an indirect general equilibrium effect. For equilibrium to apply at each date, the effective wages and the allocation of individuals across sectors will be jointly determined. Technical progress and capital accumulation will tend to induce selective migration, altering the composition of employment in each sector.

The individual-level productivity distributions in the two sectors are derived from the distribution of productivity endowments, $H(z_a, z_m)$, conditional on the self-selection

[14] The change in the level brings it closer to the data: Temple and Woessmann (2006) report that the average *RLP* in sub-Saharan African countries declined from 11.8 to 6.1 over 1960-1996. These high numbers suggest that output mismeasurement could indeed be relevant.

cutoff, q. To be more specific, the distribution of individual-level productivity in agriculture is evaluated as the product of the two marginal individual-level productivity densities for z_a and z_m, integrated over the range $z_m \in (0, z_a \cdot q]$, implied by the selection condition $\left\{ \frac{z_m^i}{z_a^i} \leq q \right\}$; see Ying (2014).

[Table 2 about here.]

In Table 2, we present some data on the productivity distributions for the years $t+1$, $t+50$ and $t+100$ in our simulation, where t corresponds to 2010. As discussed earlier, average individual productivity increases in agriculture but decreases in non-agriculture, due to selective migration. For our calibrated joint distribution for the productivity levels, structural transformation makes little difference to the median productivity in agriculture, but halves the median productivity in non-agriculture.

The productivity distributions become more unequal in both sectors. This can be seen from the log variance, and also from the ratio of the 90th to the 10th percentile. This effect can arise easily in the model, given the dependence between agricultural and non-agricultural productivity levels ($\rho > 0$). Some relatively unskilled workers move to non-agriculture, increasing the dispersion of ability within the sector; at the same time, some workers with low productivity remain in agriculture, which leads inequality in that sector to increase.

If we combine these individual productivity levels with the sectoral effective wages, we can obtain the wage distributions for the two sectors and the economy as a whole. Figure 3 plots the distributions of wages (and log wages) for the sectors and the whole economy, for the years $t+1$, $t+50$ and $t+100$, where t is 2010. Some important features of these distributions are also highlighted in Table 3.

[Figure 3 about here.]

[Table 3 about here.]

The figures and table show a clear pattern: the locations of the distributions shift to the right over time, and their variances increase. The average wage is substantially higher in non-agriculture than in agriculture for all three dates considered. Hence, even when labor mobility is costless, composition effects are powerful enough to ensure that the average urban wage is higher than the average rural wage, and increases faster. Median wages for the two sectors are closer to one another, but again move apart over time. The log-variance and the ratio of the 90th percentile to the 10th percentile, shown in table 3, suggest that wage inequality rises over time, but these measures are not Lorenz-consistent. We defer a more rigorous analysis of inequality until later in the paper.

6.2 Welfare

In this section, we consider the implications of the calibrated model for welfare. For this analysis, it helps to recall that a given worker's wage is the product of their 'realized' productivity level, which varies depending on their chosen sector, and the effective wage which applies in that sector. Our assumptions on how different workers combine in production are important here. Since we are assuming that, within a sector, workers are perfect substitutes at fixed ratios, we rule out variation in wages within sectors driven by changing relative supplies.

One consequence of this approach is that structural transformation is necessarily welfare-improving when ζ_a and/or ζ_m are both non-decreasing over time, as they are in our simulation. If they grow at the same rate, all workers are paid higher wages and there is no reallocation; all workers are then better off. If one of (ζ_a, ζ_m) grows at a faster rate than the other, selective migration will take place, but those who have moved must be better off than previously (otherwise they would not have moved) and those who remain in the sector that is declining in relative size are no worse off, given that wages within sectors are not affected by relative supplies. Hence, if neither ζ_a nor ζ_m is decreasing over time, structural transformation will lead to a Pareto-improvement in this model.

We should therefore expect the wage distribution at any given date to stochastically dominate the distributions at earlier dates. We show this in figure 4, which plots the cumulative density functions of log wages against the population percentiles for three dates ($t+1$, $t+50$, $t+100$). For both sectors, as well as the whole economy, the cumulative density functions shift strictly to the right over time, so that the distributions in the later years first-order dominate those of earlier years. In more concrete terms, the person at a given percentile of the wage distribution in a given year is always better off than the (perhaps different) person at that percentile in earlier years.

[Figure 4 about here.]

If we take wages to be our welfare measure, we can also use stochastic dominance criteria to compare welfare between the two sectors at a given point in time. For the years $t+1$, $t+50$ and $t+100$, figures 5 and 6 compare welfare between the sectors in terms of first-order stochastic dominance and second-order stochastic dominance respectively. These figures show that we cannot rank the sectors in terms of welfare, even in the weaker second-order sense of the Generalized Lorenz curve of Shorrocks (1983). Put differently, although there has been some recent interest in comparisons of average life satisfaction across regions or sectors, these comparisons can be misleading once heterogeneity and occupational self-selection are acknowledged. Differences in average life satisfaction will often emerge in equilibrium but, as with differences in the average product of labor, these should not necessarily be a policy concern.[15]

[15]See Breinlich et al. (2014) for related discussion and references.

[Figure 5 about here.]

[Figure 6 about here.]

So far, our results indicate that structural transformation raises welfare, but the benefits are distributed unevenly. A useful way to illustrate these uneven benefits is to plot the average growth rate of wages against the percentiles of the wage distribution: this is the growth incidence curve used in Ravallion and Chen (2003). Figure 7 presents growth incidence curves for the first fifty years of the simulation, the second fifty years, and the whole period (2010-2110). The lines all lie above zero, consistent with first-order stochastic dominance. For the whole economy, the curves are upward sloping, suggesting an increase in overall inequality. For the agriculture sector, the curve is modestly decreasing across the first three quartiles, but then rises sharply for the richest quartile. The curve for the first half-century is flatter than for the second half-century, which suggests that the largest effects on inequality arise later in the structural transformation. In non-agriculture, the growth incidence curve has a slight U-shape for the first half-century, when the richest and the poorest benefit almost equally from economic growth. For the century as a whole, the curve for non-agriculture is modestly upward sloping, suggesting a modest increase in inequality for this sector.

[Figure 7 about here.]

6.3 Wage inequality

We now study wage inequality in more depth. We will look at the path taken by wage inequality at the sector level, and for the economy as a whole. The path of economy-wide wage inequality is partly shaped by changes in employment shares, as in classic analyses of the Kuznets process due to Robinson (1976) and Anand and Kanbur (1993). But our model adds further considerations. The evolutions of the sectoral wage distributions will be influenced by composition effects (selective migration) while economy-wide inequality will also be influenced directly by differential growth rates for the sectoral effective wages. We will see that these additional considerations can overturn the prospect of an inverted-U in the path of wage inequality.

We initially examine inequality in terms of Lorenz dominance, rather than less general summary measures.[16] Figure 8 shows the Lorenz curves for the two sectors and the whole economy for $t+1$, $t+50$ and $t+100$. We can see that the Lorenz curve barely moves for the non-agricultural sector, but shifts out substantially for agriculture. The outwards movement for agriculture is greater for the second half-century, but the effect on the whole-economy Lorenz curve is larger for the first half-century, consistent with the decline in the agricultural employment share over time. Overall, it should be clear that wage inequality increases for agriculture and the whole economy.

[16] See Bourguignon (1990) for an earlier paper which uses Lorenz comparisons in the context of a dual economy.

[Figure 8 about here.]

For analysing the time paths for wage inequality, a summary measure such as the Gini coefficient is helpful. Figure 9 plots the Gini coefficients over time, and shows a clear increase in the Gini coefficient for agriculture; this drives an increase in the economy-wide Gini coefficient, which approaches the (barely changing) Gini coefficient for the non-agricultural sector. Again, this evidence for convergence of the economy-wide and the non-agricultural levels of inequality is natural given the declining employment share of agriculture. Much the same pattern emerges from Theil's L index, also known as the mean logarithmic deviation (MLD), which is plotted in figure 10.

[Figure 9 about here.]

An advantage of MLD is that we can decompose this measure into the respective contributions of within-sector and between-sector components, as in figure 10(b). Within-sector inequality increases over time, largely due to the increased inequality within agriculture. It is noticeable that between-sector inequality makes only a modest, and declining, contribution to aggregate wage inequality. We can see its modest changes as the outcome of two offsetting forces: changes in sectoral effective wages drive actual wages further apart, but also induce reallocation, giving rise to composition effects. The increasing average individual-level productivity of those workers who remain in agriculture offsets the changes in effective wages, limiting the extent of between-sector inequality.

[Figure 10 about here.]

To investigate this further, figure 11 plots wages in non-agriculture relative to agriculture, including the ratio of average wages, and the ratios at various percentiles. The ratio of average wages shows a small decline, implying that between-sector inequality declines slightly. The ratio of median wages eventually shows a slight increase. The wage ratio for the 90th percentile declines substantially, so that the 'rich' workers in agriculture see their living standards converge to the 'rich' in non-agriculture. The wage ratio for the 10th percentiles barely changes over time, and so the urban poor remain somewhat worse off than the rural poor. By construction, these are workers whose productivity is low, but who would be even worse off in agriculture.

[Figure 11 about here.]

We end this section by comparing wage inequality across sectors at a given date. The sectoral Lorenz curves are shown in figure 12, for the years $t+1$, $t+50$ and $t+100$. As shown in all three panels, wage inequality in non-agriculture is unambiguously greater than in agriculture. But it is also clear that the gap narrows over time, consistent with our earlier finding that agricultural wage inequality rises in the wake of selective out-migration.

[Figure 12 about here.]

6.4 Poverty

Under our assumptions on the labor market and the growth process, we would expect structural transformation to be associated with declines in rural and urban poverty. Anand and Kanbur (1985) showed that poverty indicators of the Foster et al. (1984) type will change monotonically during the standard Kuznets process. Unexpectedly, however, poverty indicators of the Sen (1976) type can display an inverted-U shape, because structural transformation may be associated with increased inequality among the urban poor. Given our earlier discussion, this result may be an artifact of the random-draw aspect of the sectoral wage distributions.

The FGT class of measures introduced by Foster, Greer and Thorbecke (1984) take the form:

$$P_{\epsilon} = \frac{1}{N} \int_{i \in \Omega_{poor}} \left[\left(w_{poor} - w_i \right) / w_{poor} \right]^{\epsilon} di$$

where w_{poor} is the selected poverty line, N is the measured population and Ω_{poor} denotes the set of people in poverty. Three cases are of particular interest. When $\epsilon = 0$, the FGT measure is the proportion of the population below the poverty line. When $\epsilon = 1$, the measure corresponds to the average gap between the poor's income and the poverty line. When $\epsilon = 2$, the measure gives more weight to the gaps between the poverty line and wages for the poorest, and hence takes account of the severity of poverty.

[Figure 13 about here.]

First, we select a poverty line. Based on the IHS3 survey, the National Statistical Office of Malawi defined a consumption poverty line as 37,002 Malawi Kwacha, which corresponds to roughly 2/3 of the average consumption level, 54,568 Malawi Kwacha. We calibrate the poverty line in our simulation to match this ratio for the first year of the simulation. The calibrated poverty line implies that the head-count ratio is slightly higher in non-agriculture than in agriculture for the first year of the simulation. This conflicts with the evidence that poverty rates are higher in rural areas than urban areas for most countries; see, for example, Lipton and Ravallion (1995) and Ravallion (2002), and on Malawi, Mussa and Pauw (2011). Matching this feature of the data better would be an obvious area for future work.

Figure 13 shows the three measures of poverty over time. It is clear that poverty declines substantially over the course of a structural transformation. This finding accords with the stochastic dominance analysis we presented in figure 4: whatever the poverty line we select, later years will have less poverty than earlier years. Perhaps more interesting is that, in the structural transformation we consider, improvements in rural poverty emerge ahead of those in urban poverty. The analysis therefore associates a structural transformation with the urbanization of poverty: an increasing share of the poor are located in urban areas. Also note that, for an interval of the simulation, economy-wide poverty rates (head-count ratios) decline relatively slowly: to reduce poverty to a strong extent, the structural transformation has to be substantial, and urban poverty rates remain high until late in the process.

7 Robustness

In this section, we briefly discuss robustness to alternative assumptions, and in particular alternative choices of the distribution of individual-level productivity levels. The assumed joint distribution plays a key role, given the importance of composition effects to the analysis. We first present some findings for values of θ_a that are lower than our baseline value, $\theta_a = 3.40$. As θ_a declines, this increases the mean level of productivity in agriculture, and also its dispersion. This also has an indirect effect via our calibration procedure: as θ_a is reduced, the 'true' share of agricultural value added in total value added, π_a, will increase.

[Table 4 about here.]

Table 4 shows the values of θ_a that we adopt, and the implied shares of agriculture in total value added. Figure 14 plots the *MLD* inequality measure for the two sectors and the whole economy, and also decomposes inequality into within-sector and between-sector components. It is clear that lower values of θ_a correspond to a higher level of economy-wide wage inequality. Moreover, the dynamics change, and the lower values of θ_a give rise to a Kuznets curve, although the extent of the ultimate decline in economy-wide inequality is quite modest. The Kuznets curve arises because, once θ_a becomes sufficiently small, wage inequality in the non-agricultural sector declines over time. In principle, composition effects can give rise to a Kuznets curve, even labor can move freely between sectors.

[Figure 14 about here.]

8 Conclusions

The idea that structural transformation influences wage inequality has a long history. Analysis of the 'Kuznets process' usually suggests that, as countries develop, wage inequality will follow an inverted-U. The quantitative analysis in our paper leads us to question this conclusion. If we allow for worker heterogeneity and occupational self-selection, an inverted-U is far from inevitable. Instead, in our calibrated model, wage inequality usually rises over the course of a structural transformation, mainly driven by rising wage inequality in the agricultural sector. Under our maintained assumptions, this result holds for any Lorenz-consistent inequality measure. But we also show that a structural transformation represents a Pareto improvement, and is associated with major reductions in poverty. Rural poverty shows the steepest drop initially, and structural transformation is associated with an urbanization of poverty: given selective migration, urban poverty rates remain high until relatively late in the process.

These findings show what can be learnt from models of this type. The analysis generates a rich account of the dynamics of living standards, wage inequality and poverty rates. But there are a number of issues that remain to be addressed. Although we have looked at outcomes under alternative assumptions, more needs to be done in this direction. Extensions to the model could seek to match aspects of the data better, such as relative poverty rates in the urban and rural sectors, and the declines in the relative productivity of non-agriculture that are often observed in the data. It also seems important to investigate whether our overall findings are robust to alternative choices for the assumed joint distribution for individual sector-specific productivity levels. One interpretation of our paper is that future analyses of these questions cannot afford to ignore worker heterogeneity and occupational self-selection. Once these are taken into account, it becomes harder to argue that the Kuznets process leads to a Kuznets curve.

References

Ahluwalia, M. S. (1974). Income inequality: Some dimensions of the problem, *in* H. B. Chenery (ed.), *Redistribution With Growth*, Oxford University Press, London.

Ahluwalia, M. S. (1976a). Income distribution and development: Some stylized facts, *The American Economic Review* **66**(2): pp. 128–135.

Ahluwalia, M. S. (1976b). Inequality, poverty and development, *Journal of Development Economics* **3**(4): 307 – 342.

Álvarez Peláez, M. J. and Díaz, A. (2005). Minimum consumption and transitional dynamics in wealth distribution, *Journal of Monetary Economics* **52**(3): 633–667.

Anand, S. and Kanbur, R. (1985). Poverty under the Kuznets process, *Economic Journal* **95**(1): pp. 42–50.

Anand, S. and Kanbur, S. (1993). The Kuznets process and the inequality—development relationship, *Journal of Development Economics* **40**(1): pp. 25 – 52.

Angeles, L. (2010). An alternative test of Kuznets' hypothesis, *Journal of Economic Inequality* **8**(4): pp. 463–473.

Barro, R. J. (2000). Inequality and Growth in a Panel of Countries, *Journal of Economic Growth* **5**(1): 5–32.

Bertola, G., Foellmi, R. and Zweimüller, J. (2006). *Income Distribution in Macroeconomic Models*, Princeton University Press, Princeton.

Bourguignon, F. (1990). Growth and inequality in the dual model of development: The role of demand factors, *Review of Economic Studies* **57**(2): pp. 215–228.

Bourguignon, F. and Morrisson, C. (1998). Inequality and development: the role of dualism, *Journal of Development Economics* **57**(2): 233–257.

Breinlich, H., Ottaviano, G. I. P. and Temple, J. R. W. (2014). Regional growth and regional decline, *Handbook of Economic Growth*, 1 edn, Vol. 2, Elsevier, chapter 4, pp. 683–779.

Caselli, F. (2005). Accounting for cross-country income differences, *in* P. Aghion and S. N. Durlauf (eds), *Handbook in Economics 22*, Vol. 1, Part A of *Handbook of Economic Growth*, Elsevier: North-Holland, Amsterdam, pp. pp. 679 – 741.

Caselli, F. and Feyrer, J. (2007). The marginal product of capital, *Quarterly Journal of Economics* **122**(2): pp. 535–568.

Caselli, F. and Ventura, J. (2000). A representative consumer theory of distribution, *American Economic Review* **90**(4): 909–926.

Dabla-Norris, E., Brumby, J., Kyobe, A., Mills, Z. and Papageorgiou, C. (2012). Investing in public investment: an index of public investment efficiency, *Journal of Economic Growth* **17**(3): pp. 235–266.

De Magalhães, L. and Santaeulalia-Lopis, R. (2014). Consumption, income and wealth of the very poor: The cross-sectional facts for Malawi, Tanzania and Uganda, *Manuscript*, University of St. Louis.

de Vries, G., Timmer, M. and de Vries, K. (2013). Structural transformation in Africa: Static gains, dynamic losses, *GGDC research memorandum 136*.

Deininger, K. and Squire, L. (1998). New ways of looking at old issues: inequality and growth, *Journal of Development Economics* **57**(2): 259–287.

Foster, J., Greer, J. and Thorbecke, E. (1984). A class of decomposable poverty measures, *Econometrica* **52**(3): pp. 761–766.

Frazer, G. (2006). Inequality and development across and within countries, *World Development* **34**(9): pp. 1459 – 1481.

Gollin, D. (2014). The Lewis Model: A 60-Year Retrospective, *Journal of Economic Perspectives* **28**(3): 71–88.

Gollin, D., Parente, S. L. and Rogerson, R. (2004). Farm work, home work, and international productivity differences, *Review of Economic Dynamics* **7**(4): pp. 827–850.

Graham, B. and Temple, J. (2006). Rich Nations, Poor Nations: How Much Can Multiple Equilibria Explain?, *Journal of Economic Growth* **11**(1): 5–41.

Helpman, E., Itskhoki, O. and Redding, S. (2010). Inequality and unemployment in a global economy, *Econometrica* **78**(4): 1239–1283.

Herrendorf, B. and Schoellman, T. (2013). Why is measured productivity so low in agriculture?, *Working paper*, Arizona State University.

Heston, A., Summers, R. and Aten, B. (2012). Penn World Table, *Version 7.1*, Center for International Comparisons of Production, Income and Prices at the University of Pennsylvania.

Jedwab, R., Gollin, D. and Vollrath, D. (2014). Urbanization with and without Industrialization, *Working Papers 2014-01*, The George Washington University, Institute for International Economic Policy.

Kanbur, R. (2012). Does Kuznets still matter?, *Working paper*, Cornell.

Kanbur, S. M. R. and McIntosh, J. P. (1988). Dual Economy Models: Retrospect and Prospect, *Bulletin of Economic Research* **40**(2): 83–113.

Kuznets, S. (1955). Economic growth and income inequality, *American Economic Review* **45**(1): pp. 1–28.

Lagakos, D. and Waugh, M. E. (2013). Selection, agriculture, and cross-country productivity differences, *American Economic Review* **103**(2): pp. 948–980.

Lipton, M. and Ravallion, M. (1995). Poverty and policy, *in* H. Chenery and T. Srinivasan (eds), *Handbook of Development Economics*, Vol. 3 of *Handbook of Development Economics*, Elsevier, chapter 41, pp. 2551–2657.

Martin, W. and Mitra, D. (2001). Productivity growth and convergence in agriculture versus manufacturing, *Economic Development and Cultural Change* **49**(2): pp. 403–422.

McMillan, M. and Harttgen, K. (2014). What is driving the African 'growth miracle'?, *Working Paper 20077*, NBER.

McMillan, M., Rodrik, D. and Verduzco-Gallo, I. (2014). Globalization, structural change, and productivity growth, with an update on africa, *World Development* pp. 11–32.

Morrisson, C. and Murtin, F. (2013). The Kuznets curve of human capital inequality: 1870–2010, *Journal of Economic Inequality* **11**(3): pp. 283–301.

Mundlak, Y. and Hellinghausen, R. (1982). The intercountry agricultural production function: another view, *American Journal of Agricultural Economics* **64**(4): pp. 664–672.

Mussa, R. and Pauw, K. (2011). Poverty in Malawi: current status and knowledge gaps, *Policy Note 9*, IFPRI.

Parente, S. L., Rogerson, R. and Wright, R. (2000). Homework in development economics: Household production and the wealth of nations, *Journal of Political Economy* **108**(4): pp. 680–687.

Piketty, T. (2006). The Kuznets curve: yesterday and tomorrow, *in* R. B. Abhijit V. Banerjee and D. Mookherjee (eds), *Understanding Poverty*, Oxford University Press, Oxford.

Pritchett, L. (2000). The Tyranny of Concepts: CUDIE (Cumulated, Depreciated, Investment Effort) Is Not Capital, *Journal of Economic Growth* **5**(4): 361–84.

Ravallion, M. (2002). On the urbanization of poverty, *Journal of Development Economics* **68**(2): 435–442.

Ravallion, M. and Chen, S. (2003). Measuring pro-poor growth, *Economics Letters* **78**(1): pp. 93–99.

Ravallion, M. and Datt, G. (1996). How Important to India's Poor Is the Sectoral Composition of Economic Growth?, *World Bank Economic Review* **10**(1): 1–25.

Robinson, S. (1976). A note on the U hypothesis relating income inequality and economic development, *American Economic Review* **66**(3): pp. 437–440.

Roy, A. D. (1951). Some thoughts on the distribution of earnings, *Oxford Economic Papers* **3**(2): pp. 135–146.

Schmitt, G. (1989). Simon Kuznets' 'Sectoral shares in labor force': a different explanation of his (I+S)/A ratio, *American Economic Review* **79**(5): 1262–1276.

Sen, A. (1976). Poverty: An ordinal approach to measurement, *Econometrica* **44**(2): pp. 219–231.

Shorrocks, A. F. (1983). Ranking income distributions, *Economica* **50**(197): pp. 3–17.

Song, Z., Storesletten, K. and Zilibotti, F. (2011). Growing like China, *American Economic Review* **101**(1): pp. 196–233.

Temple, J. R. W. (2005). Dual economy models: a primer for growth economists, *The Manchester School* **73**(4): pp. 435–478.

Temple, J. R. W. and Woessmann, L. (2006). Dualism and cross-country growth regressions, *Journal of Economic Growth* **11**: pp. 187–228.

Ying, H. (2014). Growth and structural change in a dynamic Lagakos-Waugh model, *Manuscript*, University of Bristol.

Young, A. (2013). Inequality, the urban-rural gap, and migration, *Quarterly Journal of Economics* **128**(4): pp. 1727–1785.

Table 1: Parametric assumptions

α	β	μ	δ
0.16	0.38	0.29	0.06
g_a	g_m	n	s_k
0.014	0.018	0.029	0.75

Table 2: Some features of the productivity distributions

	Agriculture			Non-agriculture		
Year	1	50	100	1	50	100
Mean	1.23	1.26	1.59	5.31	3.34	2.70
Median	1.02	0.98	1.01	2.76	1.68	1.34
Log-variance	0.16	0.20	0.35	0.59	0.62	0.67
90th-10th ratio	2.50	2.67	4.07	6.08	6.37	7.07

Source: Authors' simulation

Table 3: Some features of the wage distributions

	Agriculture			Non-agriculture			Overall		
Year	1	50	100	1	50	100	1	50	100
Mean	1.47	2.93	5.89	2.24	4.21	8.20	1.74	3.82	8.04
Median	1.22	2.27	3.72	1.16	2.12	4.07	1.21	2.20	4.02
Log-variance	0.16	0.20	0.35	0.59	0.62	0.67	0.31	0.49	0.64
90th-10th ratio	2.50	2.67	4.07	6.08	6.37	7.07	3.38	5.17	6.86

Source: Authors' simulation

Table 4: Assumptions for robustness experiments

	θ_a	Imp. π_a
Baseline	3.40	0.47
Robustness	2.80	0.49
	2.40	0.52
	2.00	0.58
	1.60	0.70

Source: Authors' simulation

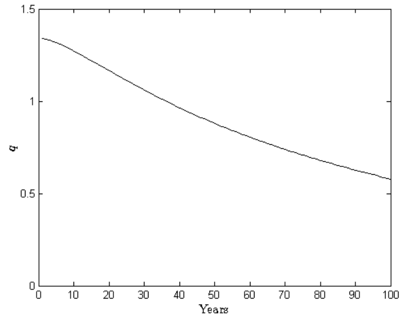

(a) Differential in effective wages

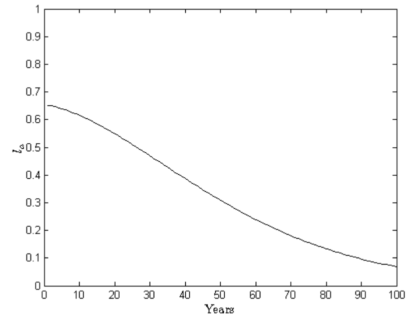

(b) Employment share in agriculture

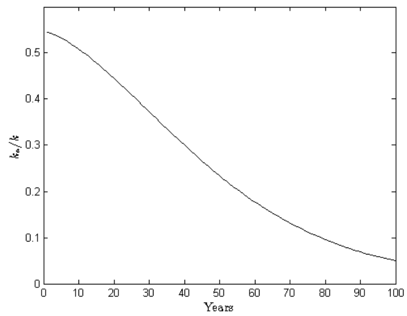

(c) Capital share of agriculture

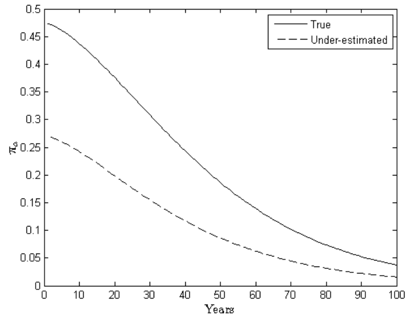

(d) Output share of agriculture

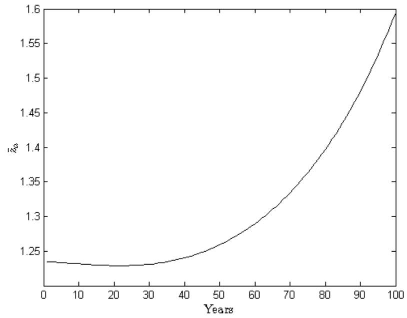

(e) Average individual productivity in agriculture

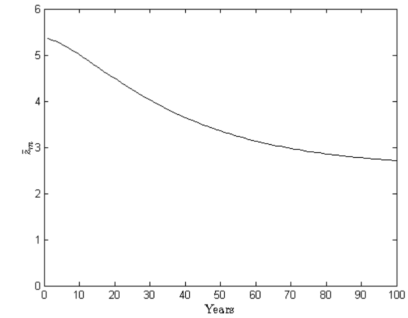

(f) Average individual productivity in non-agriculture

Figure 1: Structural transformation

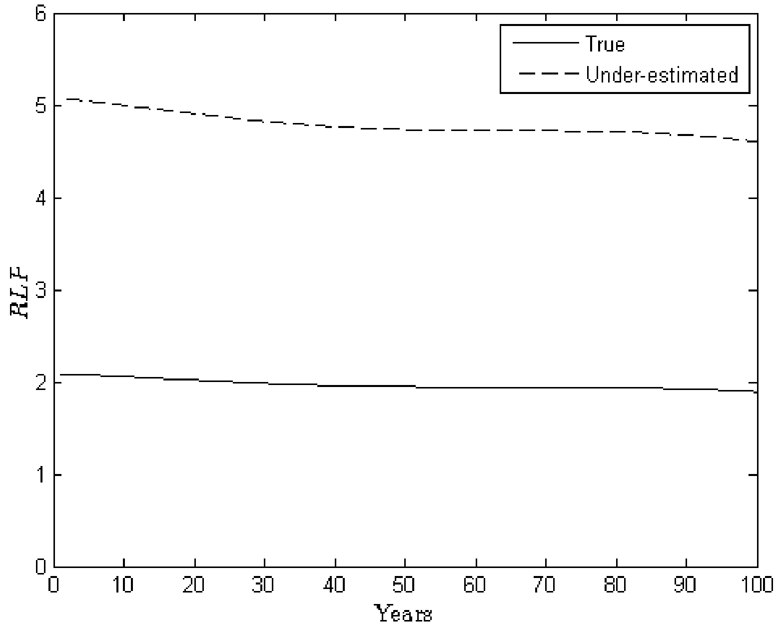

Figure 2: Relative labor productivity

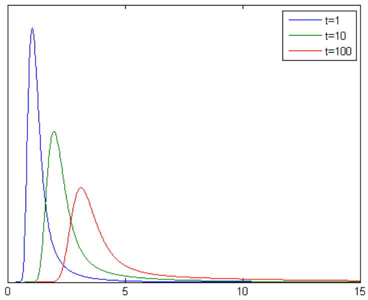

(a) Distributions of wages in agriculture

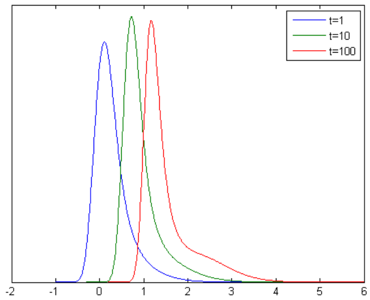

(b) Distributions of log wages in agriculture

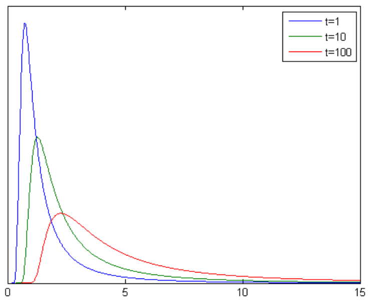

(c) Distributions of wages in non-agriculture

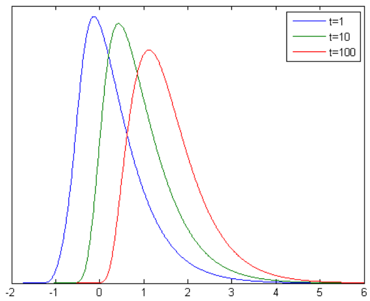

(d) Distributions of log wages in non-agriculture

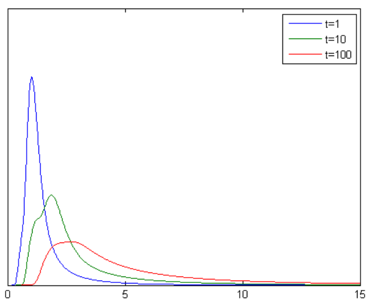

(e) Distributions of wages for the economy

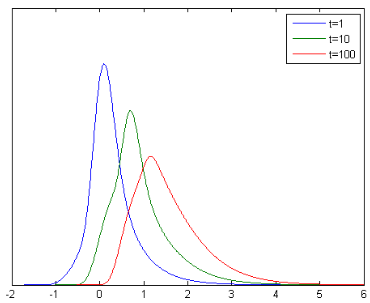

(f) Distributions of log wages for the economy

Figure 3: Wage distributions over time

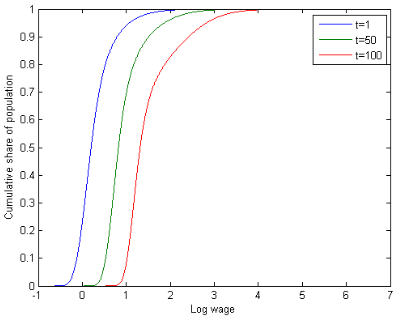

(a) CDFs for agriculture

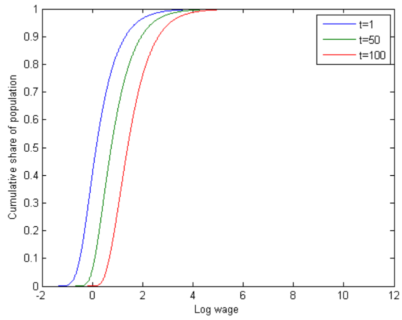

(b) CDFs for non-agriculture

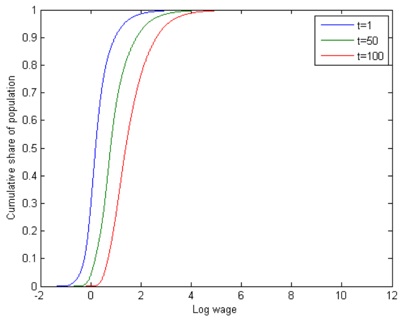

(c) CDFs for the economy

Figure 4: Cumulative density functions over time

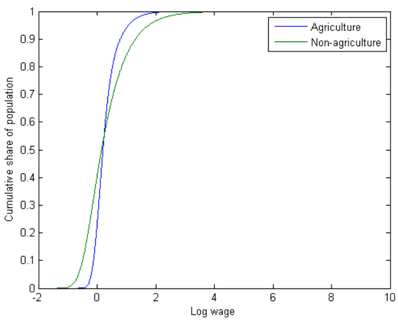

(a) CDFs at Year 1

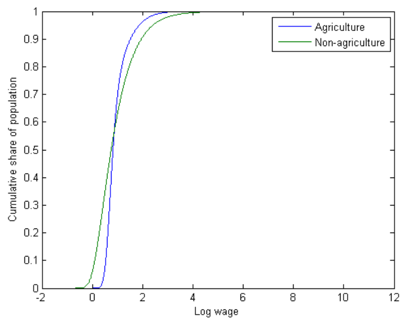

(b) CDFs at Year 50

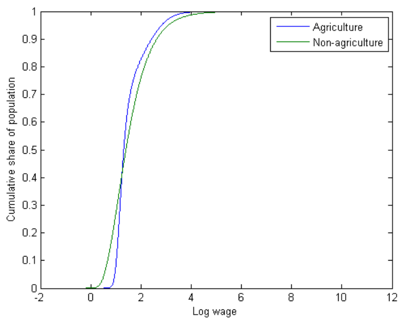

(c) CDFs at Year 100

Figure 5: Cumulative density functions across sectors

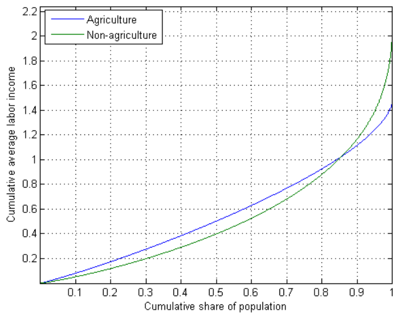

(a) Generalized Lorenz curves at Year 1

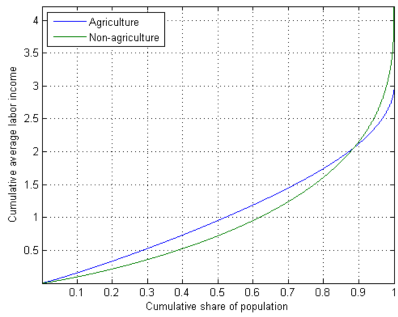

(b) Generalized Lorenz curves at Year 50

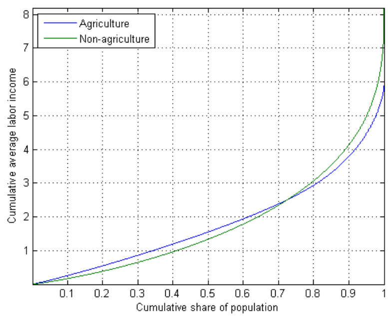

(c) Generalized Lorenz curves at Year 100

Figure 6: Generalized Lorenz curves across sectors

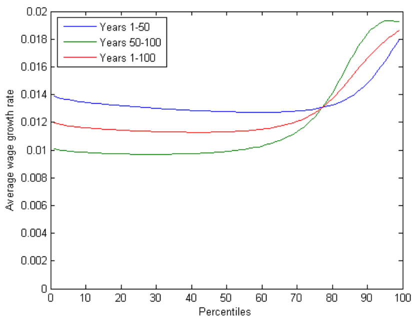

(a) GICs for agriculture

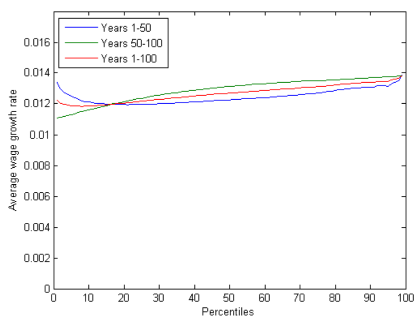

(b) GICs for non-agriculture

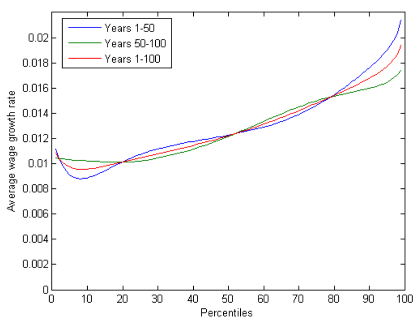

(c) GICs for the economy

Figure 7: Growth incidence curves over time

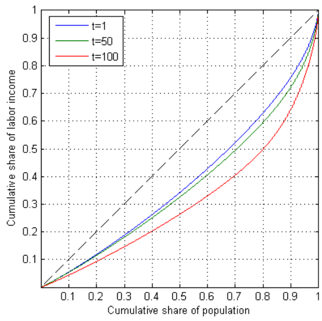

(a) Lorenz curve for agriculture

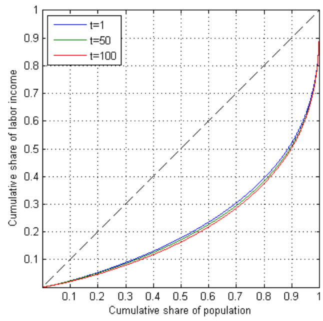

(b) Lorenz curve for non-agriculture

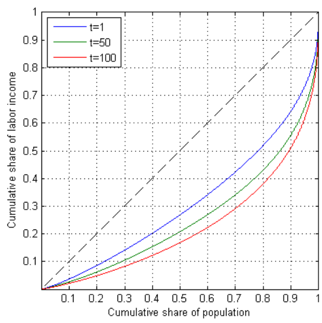

(c) Lorenz curve for the economy

Figure 8: Lorenz curves over time

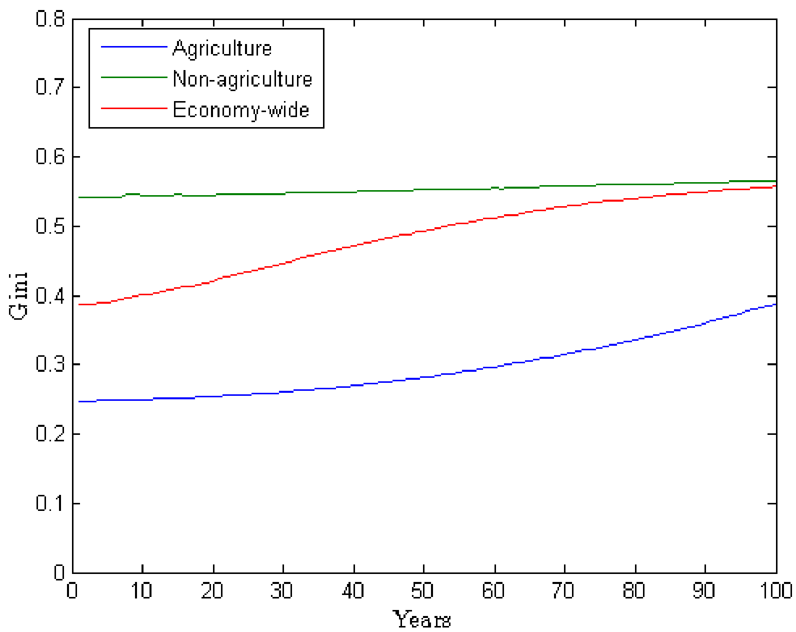

Figure 9: Gini coefficients

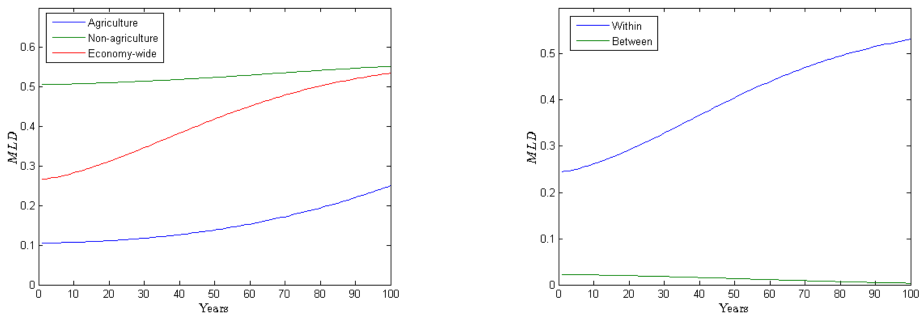

(a) MLD

(b) MLD decomposition

Figure 10: Mean log deviation

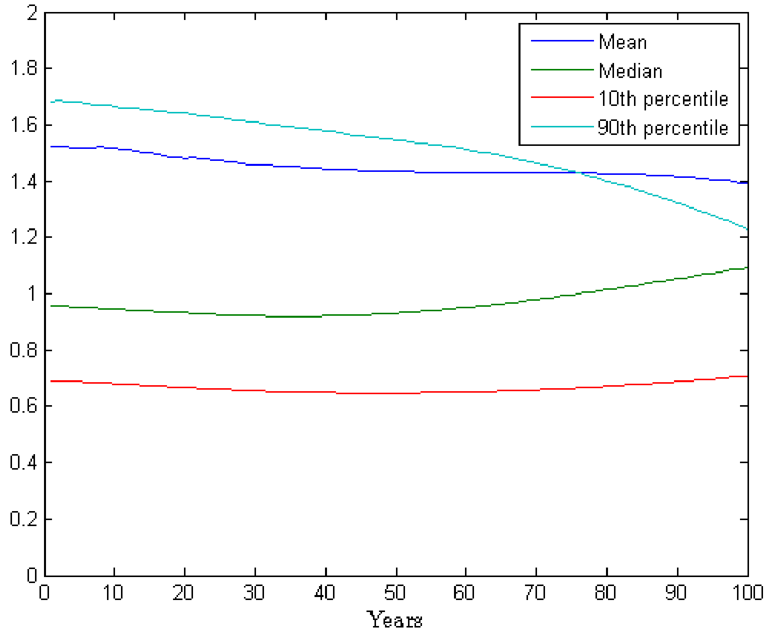

Figure 11: Wage ratios

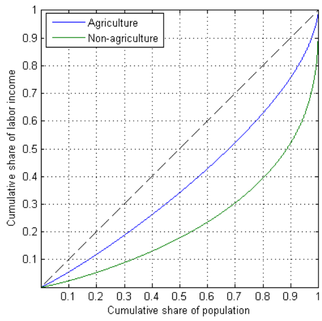

(a) Lorenz curves at Year 1

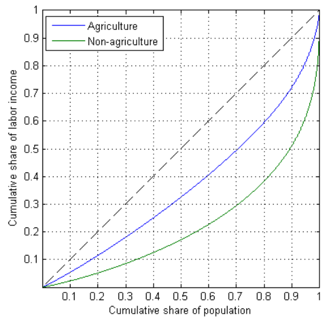

(b) Lorenz curves at Year 50

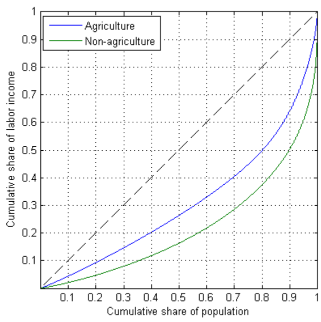

(c) Lorenz curves at Year 100

Figure 12: Lorenz curves across sectors

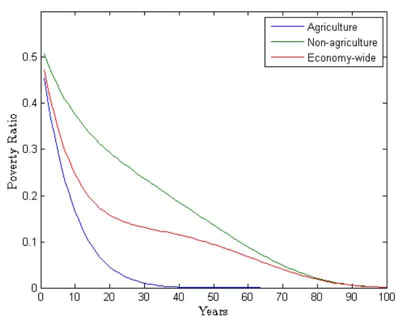

(a) Poverty rate

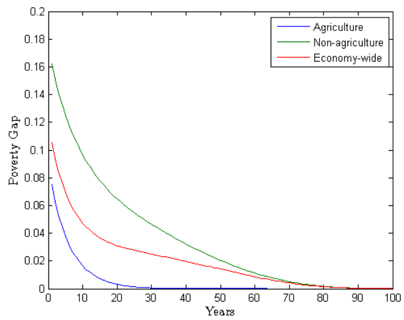

(b) Poverty gap

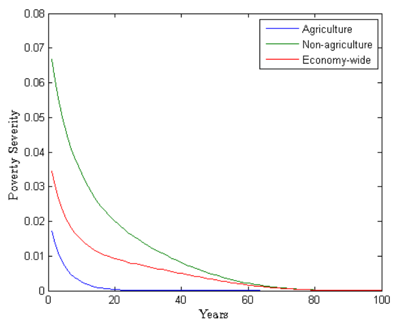

(c) Poverty severity

Figure 13: Poverty measures

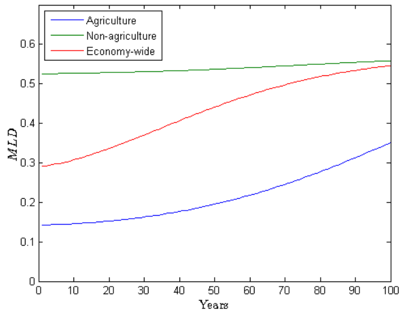

(a) $\theta_a = 2.8$

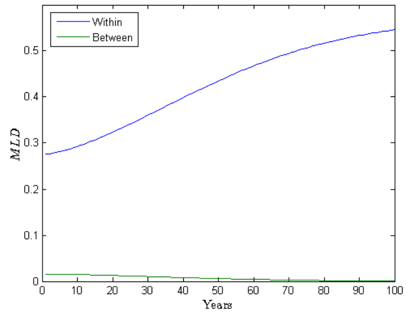

(b) $\theta_a = 2.8$

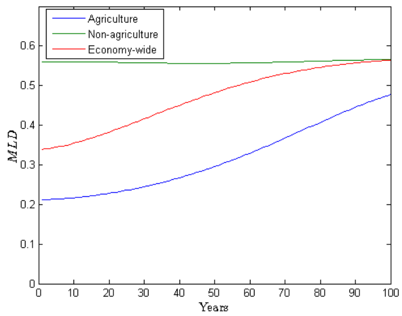

(c) $\theta_a = 2.4$

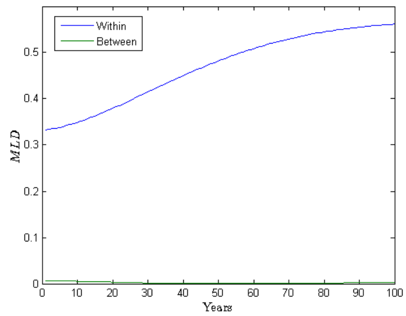

(d) $\theta_a = 2.4$

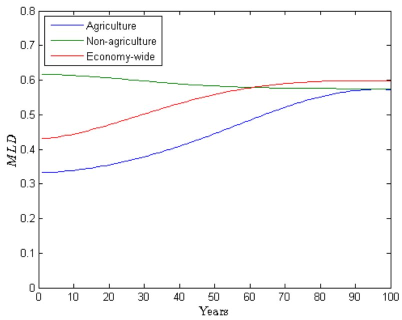

(e) $\theta_a = 2.0$

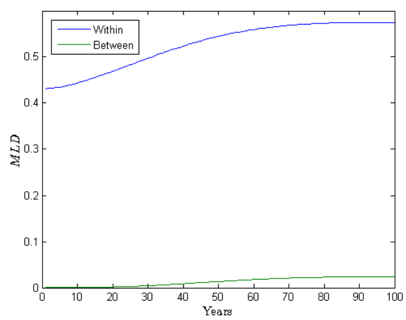

(f) $\theta_a = 2.0$

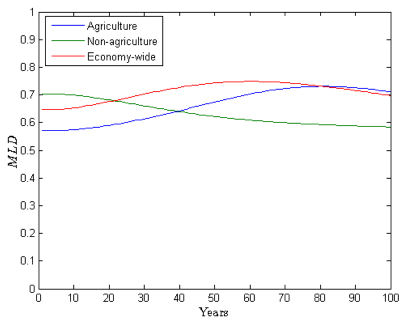

(g) $\theta_a = 1.6$

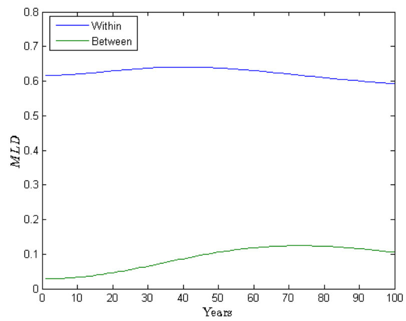

(h) $\theta_a = 1.6$

Figure 14: Robustness for inequality dynamics