DOI: 10.1111/j.1752-5209.2008.0019.x

It's not technical progress: empirical TFP determination and structural change

Markus Eberhardt[1]

St. John's College, Oxford & Centre for the Study of African Economies

Abstract: The process of economic growth is inseparable from the process of structural change. This paper examines the problem of measuring technological progress and structural change in the growth accounting literature. It uses a simple computation exercise to illustrate how dual economy considerations increase the difficulties associated with measuring technological change.

> '[There exists a] vast literature on dual economies. However, work by growth economists, whether at the research frontier or in writing textbooks, has only rarely acknowledged this tradition. Instead, the dominant focus on a one-sector model, and on steady states, has precluded some of the most interesting questions from even being asked.'
> Temple (2005, p. 468)

> 'The TFP growth rate obtained [from a growth accounting framework] need not represent only technological change and may not represent technological change at all.'
> Baier, Dwyer, and Tamura (2006, p. 27)

The questions why some countries produce so much more output per worker than others, and why certain economies have grown phenomenally while others have stagnated or even regressed over the same period of time go to the very heart of development economics. Since the emergence of a unique dataset for cross-country empirical analysis in the early 1990s – the Penn World Table – literally thousands of articles have used these data to investigate the 'deep determinants' of economic growth and development, the factors other than the 'proximate' labour, capital and non-material inputs, such as institutions, geography or trade openness that are deemed to drive the development process.

The empirical strategies in this literature are typically based on a version of the Solow-Swan growth model and usually involve the determination of Total Factor Productivity (TFP), the residual which accounts for the labour productivity variation that remains unexplained by proximate factor inputs. These strategies can be broadly divided into three categories: regressions, accounting frameworks, and frontier estimations. We have discussed the issues surrounding the first of these in detail elsewhere[2] and leave the latter, including data envelopment analysis, for future research.

In this paper we develop three major points: firstly, that the focus on aggregate economy data in the empirical literature has diverted attention from dual economy models pre-dating the standard Solow-Swan model – to the detriment of empirical growth analysis; secondly, that the growth accounting approach to TFP determination encounters serious conceptual and practical difficulties, and that consequently narrow definitions of TFP as a sound proxy for 'technical change' are misguided; and thirdly, that the difficulties of providing a meaningful interpretation of empirical TFP estimates are exacerbated once the notion of a dual economy is accounted for.

1 A brief history of dual economy models and aggregate growth empirics

In the early literature on developing countries a distinction was made between the processes of economic development and of economic growth. Economic development was seen to be a process of structural transformation whereby an economy which was 'previously saving and investing 4 or 5 percent of its national income or less, converts itself into an economy where voluntary savings is running at about 12 to

1. Correspondence: St. John's College, Oxford OX1 3JP; markus.eberhardt@economics.ox.ac.uk
Francis Teal, Charles Brendon, Dirk Bezemer and two anonymous referees have provided helpful comments and suggestions on an earlier draft. All remaining errors are my own. Financial support from the Economic and Social Research Council is gratefully acknowledged.
2. Eberhardt and Teal (2008).

Published by Blackwell Publishing Ltd

15 percent of national income' (Lewis, 1954, p. 155). An acceleration in the investment rate was only one part of this process of structural transformation; of equal importance was the process by which an economy moved from a dependence on subsistence agriculture to one where an industrial modern sector absorbed an increasing proportion of the labour force.[3]

In contrast to these models of 'development for backward economies' (Jorgensen, 1961, p. 309), where duality between the modern and traditional sectors was a key feature of the model, was the analysis of economic growth in developed economies. Here the processes of factor accumulation and technical progress occur in an economy which is already 'developed', in the sense that it has a modern industrial sector and agriculture has ceased to be a major part of the economy (Solow, 1956, 1957; Swan, 1956).

Much of the early growth modeling work proceeded without close connection to observed data. The models were in Solow's classic exposition of growth theory inspired by stylised 'Kaldor' facts. The dual economy models of structural transformation used case studies and facts at least as stylised as those in the Solow-Swan growth context. Empirical studies that followed in the period thereafter employed a vast array of explanatory variables of growth, while methodological, statistical, and conceptual problems made it difficult to draw reliable conclusions from the existing literature. The key papers which brought modeling and data together were the contributions of Barro (1991) and Mankiw, Romer, and Weil (1992). These initiated a major revival in the Solow-Swan model and effectively merged the concerns of economic development with those of growth. Crucially, the Penn World Table data supplies macro-data which ensure that the *aggregate* Solow-Swan model can be readily estimated. Empirical implementations of dual economy models were hindered by the lack of sectoral data from developing countries. In effect, valuable conceptual contributions from the dual economy models were lost to the mainstream empirical growth literature, which proceeded to neglect sectoral composition and structural change entirely in favour of an aggregate production framework.

An aggregate function framework, however, only offers an appropriate construct in cross-country analysis if the economies investigated do not display large differences in sectoral structure (Temple, 2005), since a single production function framework assumes common production technology across all firms facing the same factor prices. Take two distinct sectors within this economy, assuming marginal labour product equalisation and capital homogeneity across sectors, and Cobb-Douglas-type production technology within each sector; if technology parameters differ between sectors, aggregated production technology *cannot* be of the Cobb-Douglas form. Finding differential technology parameters in *sectoral* production function estimation thus is potentially a serious challenge to treating production in form of an aggregated function.

An alternative motivation for a focus on sector-level rather than aggregate growth across countries is as follows: it is common practice to exclude oil-producing countries from any aggregate growth analysis, since 'the bulk of recorded GDP for these countries represents the extraction of existing resources, not value added' (Mankiw *et al.*, 1992, p. 413). The underlying argument is that sectoral 'distortions', such as resource wealth, justify the exclusion of the country observations. By extension of the same argument, we could suggest that given the large share of agriculture in GDP for countries such as Malawi (25–50%), India (25–46%) or Malaysia (8–30%) over the period 1970–2000, these countries should be excluded from any *aggregate* growth analysis since a significant share of their *aggregate* GDP derives from a single resource, namely land.[4] Sector-level analysis, in contrast, does not face these difficulties, since sectors such as manufacturing or agriculture are defined closely enough to represent a reasonably homogeneous conceptual construct.

Having already indicated the importance of agriculture for GDP for a number of countries, we complete this section by providing some more data to highlight the importance and dynamics of agriculture in a wider set of countries. As can be seen in Table 1 the shift away from agriculture has been most dramatic in the East Asia group, whereas the Sub-Saharan Africa has seen virtually no change over the same period.

Table 1 Share of Agriculture in GDP (in %)

	Decadal Medians			
	1960s	1970s	1980s	1990s
Canada & US		4.3	3.0	2.1
Europe		7.5	6.3	4.6
Latin America & Caribbean	20.4	19.2	13.9	12.2
Middle East & North Africa	13.6	8.8	6.3	11.4
Australia & New Zealand		9.4	6.7	3.7
East Asia & Pacific	35.8	28.2	21.4	18.3
Sub-Saharan Africa	38.5	34.1	32.9	32.5
South Asia	41.4	36.1	31.1	26.4

Source: World Bank (2008) World Development Indicators.

3. Kindleberger (1967); Kuznets (1961); Ranis and Fei (1961).
4. The quoted shares are from the World Bank World Development Indicators database (World Bank, 2008). For comparison, maximum share of oil revenue in GDP, computed as the difference between 'industry share in GDP' and 'manufacturing share in GDP' from the same database yields the following ranges for some of the countries omitted in Mankiw *et al.* (1992): Iran (12–51%), Kuwait (15–81%), Gabon (28–60%), Saudi Arabia (29–67%).

2 TFP determination via regression and growth accounting methods

In the following we investigate the growth accounting approach to empirical TFP determination in greater detail and discuss the difficulties arising from the approach. In order to provide clearer insights we however begin by discussing the parametric approach via econometric estimation of an aggregate production function. The definition of TFP holds in both cases: TFP represents the deviations of the actual, observed output growth from the growth rate implied by the growth of factor inputs. In the regression approach

$$Y_t = A_t K_t^{\beta} L_t^{1-\beta} \tag{1}$$

$$A_t = A_0 e^{\lambda t} \tag{2}$$

$$\Leftrightarrow \frac{\dot{Y}}{Y} = \lambda t + \beta \frac{\dot{K}}{K} + (1-\beta)\frac{\dot{L}}{L} \tag{3}$$

where the quotients represent percentage growth rates of output, capital and labour respectively, A_0 is TFP level and λ is the TFP growth rate.[5] In this approach TFP growth in the form of λ only captures disembodied, exogenous Hicks-neutral technological change which increases the efficiency of both capital and labour – the proverbial 'manna from heaven' (Chen, 1997). In this framework any technical progress embodied in the input factors (education, R&D, learning-by-doing) must be assumed to have been properly specified and accounted for when constructing the factor inputs. Finding a low value for λ does not confirm that technology has played no role in the growth process – it only suggests that disembodied TFP growth has not been substantial, but allows no judgement on the overall impact of technical progress.

In the growth accounting approach TFP is derived as a residual from observed output growth:

$$Y = F(K, L, t) \Rightarrow$$

$$\frac{\partial Y/\partial t}{Y} = \frac{\partial F/\partial K}{Y} K \frac{\partial K/\partial t}{K} + \frac{\partial F/\partial L}{Y} L \frac{\partial L/\partial t}{L} + \frac{\partial F/\partial t}{Y} \tag{4}$$

where the final term represents technological shift of the production function (TFP). Under standard neoclassical assumptions the factor shares of capital and labour in output (η_k; η_L) can be employed to yield a 'deterministic' equation for TFP growth:

$$\frac{\dot{Y}}{Y} = \eta_K \frac{\dot{K}}{K} + \eta_L \frac{\dot{L}}{L} + \Delta TFP \tag{5}$$

$$\Leftrightarrow \Delta TFP = \frac{\dot{Y}}{Y} - \eta_k \frac{\dot{K}}{K} - \eta_L \frac{\dot{L}}{L} \tag{6}$$

The growth rate of TFP is at least in theoretical terms identical to the above λ – disembodied, 'Hicks-neutral' exogenous technical progress. In practice however, one needs to keep in mind that TFP is a residual, such that it represents a 'catch-all' of output growth that cannot be explained by factor accumulation. Crucially, if TFP growth is recovered via growth accounting, its coefficient 'need not represent only technological change and may not represent technological change at all' (Baier *et al.*, 2006, p. 27).[6] Firstly, any measurement error in output, labour or capital enters the residual term and thus TFP growth. Secondly, violations of the standard assumptions of constant returns to scale and private and social marginal product equality can add to further accounting error. Thirdly, the question over which types of capital to include in the aggregate stock variable (how to account for inventories, durable goods, land), and the appropriate treatment of capital depreciation or of capacity utilisation rates pose further hurdles in the accurate construction of factor variables. Changes over time such as in the number of hours worked, property rights or the economic regime all can result in apparent TFP changes when the residual method is used. Fourthly, violations of the imposed factor parameter homogeneity across countries lead to additional terms in the residual – if the capital coefficient β in country X is considerably lower than the imposed 'common' parameter, the residual is artificially inflated. Finally, growth accounting provides no mean to test the statistical significance of TFP growth values obtained.

Thus since measurement error, violations of assumptions and incorrect variable construction can cause considerable bias, the conclusion reached by Baier *et al.* (2006) is hardly surprising. As many of the above points equally apply to the regression framework one would suggest that both regression and accounting approaches face equally large problems in determining TFP growth. Nevertheless, growth accounting has been the instrument of choice, due to a well-known empirical paradox: aggregate production function estimations commonly fail to obtain a capital coefficient β reasonably close to .3, the figure obtained from macro data on income share of labour and capital. Any mismatch between inflated capital parameter estimates and macro data does not seem to strike many authors as important.[7]

5. For simplicity we have assumed that TFP evolves in a linear fashion. We discuss production function regression and TFP estimates in greater detail in Eberhardt and Teal (2008).
6. The following discussion draws mainly on Baier *et al.* (2006), Barro (1999), Easterly and Levine (2001) and Lipsey and Carlaw (2004).
7. It is not uncommon to find papers with implied negative capital coefficients in the underlying production function – with the policy implication that destruction of the capital stock would further growth – although these results are typically glossed over.

A review of the empirical findings for TFP highlights the vast difference of results obtained. Early growth accounting exercises by Abramowitz (1956) and Solow (1957) found the overwhelming share of output growth was attributable to TFP growth. More recent work by Easterly and Levine (2001) did result in reduced residuals, nevertheless far from zero. These authors strengthen their argument regarding the importance of TFP growth by pointing to empirical studies investigating the timing of growth and savings/capital formation, which suggest that the causality runs from the former to the latter but not vice-versa.

During the 1990s a debate raged over the importance of factor accumulation vis-à-vis 'technical progress' (equated with TFP) in the development of the East Asian Tigers. In a series of papers Alwyn Young (1992, 1994, 1995) concluded that based on his growth accounting exercises factor accumulation played a vital and 'technical progress' a minor role in these nations' spectacular economic growth. Implications for East Asia's future growth were thought to be pessimistic: economic growth without technical progress was judged unsustainable in the long-run. Using alternative data, in particular for human capital, Klenow and Rodriguez-Clare (1997a,b) in contrast found that TFP growth can explain up to 90% of the variation in growth rates of output per workers.

These contradictory findings were reconciled by Chen (1997), who pointed out that more recent work has gone to great lengths to adjust for *quality* changes in the process of factor input aggregation. The correct conclusion to be drawn from Young's work is thus not that economic development in the East Asian Tigers was based on factor accumulation alone and therefore would implode similarly to that of the Soviet Union, but that *disembodied* technical progress has not been important in their growth, while *embodied* technical change has gone unmeasured.

In the same vein, Lipsey and Carlaw (2001, p. 14) stress that '[w]e cannot in practice distinguish empirically between genuine disembodied technical change and technological change that is embodied in a new machine'. They conclude that 'whatever TFP does measure . . . it emphatically does not measure all of technological change' (Lipsey and Carlaw, 2001, p. 43).

3 TFP and growth accounting with structural change

The multiple concerns regarding measurement of TFP in an aggregated fashion aside, it is insightful to analyse growth accounting in the presence of structural transformation. The following simple computational exercise can show that sectoral structure and marginal product differentials between sectors have a significant bearing on 'expenditure-weighted' calculations of change, such as those implemented in TFP growth accounting. Recalling the standard neoclassical assumptions for the growth accounting framework (5) it must hold that

$$\eta_K = \frac{rK}{Y}, \eta_L = \frac{wL}{Y}, \eta_K + \eta_L = 1 \qquad (7)$$

where equation (7) states that factor shares in production accruing to capital and labour are such that equation (5) represents a 'Divisia index' of inputs, made up of *the percentage change in input weighted by the relative share in input cost*.

Now assume a dual economy model with stylised agriculture (a) and manufacturing (m) sectors in two time periods

$$Y_a = A_a K_a^{\alpha} L_a^{1-\alpha} \qquad (8)$$

$$Y_m = A_m K_m^{\beta} L_m^{1-\beta} \qquad (9)$$

Assume that TFP-levels differ across sectors ($A_m > A_a$) – for simplicity there is no change in factor inputs in the aggregate economy (only shift between sectors) or TFP growth in either sector between the two periods. We can show that during a period of structural transformation where labour moves from the agricultural to the manufacturing sector, in which equality of marginal products between sectors is violated, the relative expenditure shares (on labour and capital) do not provide appropriate weights for aggregating percentage changes in labour forces. The marginal product inequality between sectors can be thought of as a wage difference between the agricultural and the manufacturing sector. Aggregation to a macro production function in this 'transitional equilibrium' is not possible without bias. The reason for this is that the expenditure weights α and β are not free from influence of marginal products.

Our computational exercise can illustrate this point: in Table 2 we consider an 'agrarian' economy (total population 100), where in the base period 90% of workers are employed in agriculture. We consider three different scenarios: (A) large marginal product differences between agriculture and manufacturing; (B) intermediate-sized differences; (C) no differences. We then consider the migration of 10 input units (workers) from the agricultural to the manufacturing sector. As can be seen in column [6], aggregated *quantity*-based calculations show no overall change in inputs. Columns [7] to [9] provide *expenditure*-weighted input changes (such as employed in a growth accounting exercise) under the three scenarios.

It can be seen that the larger the marginal product differential between sectors, the larger the computed change in factor input. In a growth accounting exercise based on *aggregate* data we have not witnessed any change in factor inputs (total population is still 100), but our productivity increase has been sizeable. Refer back to equation (5), the fundamental growth accounting equation:

Table 2 Structural change and TFP in an agrarian economy

Expenditure-weighted computations of input-change during sectoral transition: Three scenarios for marginal product differences between sectors									
	Base period ($t = 0$)				Following period ($t = 1$)				
	[1] Labour	[2] Unit price of labour (MPL)	[3]	[4]	[5] Labour	[6] Change in quantity*	[7] Expenditure-weighted changes*,**	[8]	[9]
Scenario		A	B	C		All	A	B	C
Sector									
Agriculture	90	1	1	1	80	−11.8%	−6.4%	−8.7%	−10.0%
Manufacturing	10	5	2	1	20	66.7%	30.4%	17.2%	10.0%
Aggregate	100	1.40	1.10	1	100	0.0%	24.0%	8.4%	0.0%

Notes:
* Aggregate quantity change = (change in manu*avg. quantity share in manu) + (change in agri*avg. quantity share in agri); Aggregate expenditure weighted change = (change in manu*avg. expenditure share in manu) + (change in agri*avg. expenditure share in agri).
** Uses the two period average expenditure weights.

$\dot{L}/L$ did not change if we measure *aggregate* labour inputs, $\dot{K}/K$ did not change ... but due to the reallocation of workers into the more productive sector of the economy, aggregate output has changed: $\dot{Y}/Y > 0$. As is now very clear, the *aggregate* growth accounting approach would attribute all of this increase to ΔTFP.

Table 3 carries out the same calculations for a shift of 10 input units (under scenarios A, B or C), but assumes different structures of the economy in the base year. Column (1) shows five distinct cases, where case (I) repeats the result from the previous table. It can be seen that the more industrialised the economy in the base period (cases II–V), the lower the computed expenditure-weighted input change in the given period.

These numerical examples show that sectoral transformation which under pure quantity aggregation would show a zero change in factor inputs lead to considerable bias if expenditure-weighted input measures are used. Aggregate growth accounting, which applies expenditure-weighted input measures, cannot detect the source of observed productivity increase resulting from more (less) labour being employed in the relatively more (less) productive sector and attributes the change to TFP. Narrow interpretations of TFP as technical progress are therefore inappropriate.

Summary and concluding remarks

In this paper we charted the recent history of growth empirics and suggested that the ready availability of aggregate production data in form of the Penn World Table has been to the detriment of a dual economy approach to the analysis of growth and development. We highlighted the importance and dynamic change of the agricultural sector in many developing countries and provided a number of salient reasons why the aggregate economy approach may lead to biased empirical results. We then turned to a discussion of TFP growth accounting, one of the three standard approaches to 'empirical' TFP determination,

Table 3 Structural change and TFP at different stages of development

Expenditure-weighted computations of input-change during sectoral transition: Five examples of a dual economy setup						
	Base period	Following period				
Structure of the economy	[1] Labour split	[2] Labour split	[3] Quantity Changes	[4] Expenditure-weighted changes	[5]	[6]
Scenario			All	A	B	C
Agriculture : Manufacturing (I)	90 : 10	80 : 20	0.0%	24.0%	8.4%	0.0%
Agriculture : Manufacturing (II)	70 : 30	60 : 40	0.0%	16.5%	7.3%	0.0%
Agriculture : Manufacturing (III)	50 : 50	40 : 60	0.0%	12.4%	6.4%	0.0%
Agriculture : Manufacturing (IV)	30 : 70	20 : 80	0.0%	10.0%	5.7%	0.0%
Agriculture : Manufacturing (V)	10 : 90	0 : 100	0.0%	8.2%	5.0%	0.0%

which was shown to face considerable conceptual as well as practical difficulties. Our fundamental criticism is that TFP is too readily equated with technical progress, rather than being treated as a 'measure of our ignorance' (Abramowitz, 1956) containing not only the impact of variables omitted from the model but also errors arising from variable mismeasurement, model misspecification or violation of the constant returns to scale assumption. Our discussion of the empirical debate surrounding the importance of TFP in the rise of the East Asian Tigers highlights that these issues make it difficult to rely on TFP estimates for theoretical reasoning and policy formulation. Our final section then carried out a small computational exercise to highlight the distorting impact of structural change (the shift of labour and/or capital from one sector to another) on TFP derived from aggregate growth accounting. We showed that simple quantity changes between sectors resulted in apparent positive TFP change – this finding questions any interpretation of TFP as a solid proxy for technical change, rather than the measure of ignorance label that we prefer.

The purpose of this article was to highlight these issues and give an indication of their importance for the empirical analysis of growth and development. A next step would naturally be a more formal treatment of the matters discussed, providing simulation results for *aggregate* economy empirical approaches to data characterised by an underlying dual economy nature and vice-versa. These simulations could incorporate not only the investigation of the TFP growth accounting approach that was the focus in this article, but also the production function estimation and frontier estimation approaches.

References

Abramowitz, M. (1956). 'Resource and output trend in the United States since 1870'. *American Economic Review*, 46 (2), 5–23.

Baier, S. L., Dwyer, G. P., & Tamura, R. (2006). 'How important are capital and total factor productivity for economic growth?' *Economic Enquiry*, 44 (1), 23–49.

Barro, R. J. (1991). 'Economic growth in a cross-section of countries'. *Quarterly Journal of Economics*, 106 (2), 407–443.

Barro, R. J. (1999). 'Notes on growth accounting'. *Journal of Economic Growth*, 4 (2), 119–137.

Chen, E. K. (1997). 'The Total Factor Productivity debate: Determinants of economic growth in East Asia'. *Asian-Pacific Economic Literature*, 11 (1), 18–38.

Easterly, W., & Levine, R. (2001). 'It's not factor accumulation: Stylised facts and growth models'. *World Bank Economic Review*, 15 (2), 177–219.

Eberhardt, M., & Teal, F. (2008). *Modeling Technology and Technological Change in Manufacturing: How do Countries Differ?* (CSAE Working Paper, WPS/2008-12). Centre for the Study of African Economies.

Jorgensen, D. W. (1961). 'The development of a dual economy'. *Economic Journal*, 72 , 309–334.

Kindleberger, C. P. (1967). *Europe's post-war growth: the role of labor supply.* Cambridge, Mass.: Harvard University Press.

Klenow, P. J., & Rodriguez-Clare, A. (1997a). 'Economic growth: A review essay'. *Journal of Monetary Economics*, 40 (3), 597–617.

Klenow, P. J., & Rodriguez-Clare, A. (1997b). 'The neoclassical revival in growth economics: Has it gone too far?' *NBER Macroeconomics Annual*, 12 , 73–103.

Kuznets, S. (1961). 'Economic growth and the contribution of agriculture: notes on measurement'. *International Journal of Agrarian Affairs*, 3 (1), 56–75.

Lewis, W. A. (1954). 'Economic development with unlimited supplies of labour'. *The Manchester School*, 22 , 139–191.

Lipsey, R. G., & Carlaw, K. I. (2001). *What does Total Factor Productivity measure?* (Unpublished manuscript)

Lipsey, R. G., & Carlaw, K. I. (2004). 'Total factor productivity and the measurement of technological change'. *Canadian Journal of Economics*, 37 (4), 1118–1150.

Mankiw, G., Romer, D., & Weil, D. N. (1992). 'A contribution to the empirics of economic growth'. *Quarterly Journal of Economics*, 107 (2), 407–437.

Ranis, G., & Fei, J. (1961). 'A theory of economic development'. *American Economic Review*, 51 (4), 533–556.

Solow, R. M. (1956). 'A contribution to the theory of economic growth'. *Quarterly Journal of Economics*, 70 (1), 65–94.

Solow, R. M. (1957). 'Technical change and the aggregate production function'. *Review of Economics and Statistics*, 39 (3), 312–320.

Swan, T. W. (1956). 'Economic growth and capital accumulation'. *Economic Record*, 32 (2), 334–361.

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

World Bank. (2008). *World Development Indicators*. Online database.

Young, A. (1992). 'A tale of two cities: factor accumulation and technical change in Hong Kong and Singapore'. In O. Blanchard & S. Fischer (Eds.), *NBER Macroeconomic Annual 1992.* Cambridge, Mass.: MIT Press.

Young, A. (1994). 'Lessons from the East Asian NICs: a contrarian view'. *European Economic Review*, 38 (3–4), 964–973.

Young, A. (1995). 'The tyranny of numbers: confronting the statistical realities of the East Asian growth experience'. *Quarterly Journal of Economics*, 110 (3), 641–680.

Markus Eberhardt is a DPhil student in Economics at St. John's College, Oxford, where he also completed an MPhil in Economics. He also holds a Masters in Development Economics from the University of East Anglia. His first degree was in Modern Chinese Studies.