Oxford Review of Economic Policy, 2025, 41, 223–259
https://doi.org/10.1093/oxrep/graf042
Article

The rebuilding macroeconomic theory project part III: macroeconomic models of structural change

David Vines

Economics Department, Balliol College, and Oxford Martin School, University of Oxford, UK; Centre for Economic Policy Research, e-mail: david.vines@economics.ox.ac.uk

I am grateful to the authors of the papers in this issue the *Oxford Review* for the detailed conversations which I have had with them about their papers, and in particular to Joe Stiglitz and Tomohiro Hirano. I am also grateful to Olivier Blanchard, Wendy Carlin, Shamik Dhar, Gerry Holtham, John Llewellyn, and Paul Luk for helpful discussions which I have drawn upon when writing this paper. And I am grateful to the other editors of the *Oxford Review* for comments on earlier drafts of this paper. This is particularly true in the case of Chris Adam who helped me to restructure the paper in an important way.

Abstract

In this article I first set out a taxonomy of reasons for structural change within a macroeconomic model. There may be tipping points, either within one sector or due to strategic complementarities between sectors. There may be a process of gradual accumulation, within the model, of something that is particular to one sector which causes that sector to grow faster than other sectors. Or there may be structural changes in the economy that are external to the model whiMcKibbinch mean that the nature of the model needs to change. I use this taxonomy to help me to analyse the models of structural change that are presented in the papers in this issue of the *Oxford Review of Economic Policy*. First, I examine models of structural change that can occur in the process of economic growth. I then use the taxonomy to examine models of the kind of structural change that can also occur in the shorter term, as a result of the operation of stabilization policy. I also show how models can help us to understand structural change in the world economy. In addition, I examine how chaotic behaviour can cause structural change. And finally, I show how the models of structural change set out in the papers in this issue of the *Oxford Review of Economic Policy* can be of use to macroeconomic policy-makers.

Keywords: structural change, tipping point, Ramsey model, leading sector two-sector model, stabilization policy.

JEL codes: E10, E22, E23, E27, E61, F41, F42

I. Introduction

In macroeconomics, structural change refers to a fundamental shift in the long-term composition, organization, and functioning of an economy as it grows over time. More precisely, structural change is the process by which the relative importance of different sectors of the economy evolves, leading to changes in the patterns of production, employment, trade, and the distribution of income. Policy choices can lead to structural change. And structural change can drive changes in policy.

Structural change may include such elements as:

- technical changes: adoption of new technologies that transform production methods;
- demand shifts: with non-homothetic preferences, Engels-law effects induce a structural transformation away from agriculture;
- sectoral shifts: movement of labour and resources from low-productivity sectors (e.g. agriculture) to higher-productivity sectors (e.g. manufacturing, services, and technology);
- financial developments: changes in the size and composition of national savings;

– demographic and social changes: urbanization, education, and population dynamics;
– institutional changes: reforms in labour markets, or in governance structures;
– global integration: shifts driven by trade liberalization, globalization, and capital flows.

This issue of the *Oxford Review of Economic Policy* (*OxREP*) asks what causes structural change and why models of structural change matter for macroeconomic policy-making. We have brought together 16 contributions from leading scholars to discuss these questions. These papers all show why the canonical model of the macroeconomic growth process—the Ramsey model—is an inadequate way to understand the kinds of structural change which I have described above. The papers also help identify what models are needed in addition to that canonical model. It will become clear how the availability of these models might help improve both macroeconomic theorizing and macroeconomic policy-making.

This issue of *OxREP* is the third in a series in which distinguished authors have examined the future of macroeconomics as an intellectual discipline. The first of these was published in 2018 and examined the role of DSGE models (i.e. dynamic stochastic general equilibrium models) in providing advice for macroeconomic policy-makers. The authors were asked: 'Is the glass half empty or half full?' Although some authors challenged the simplicity of the assumptions that underly DSGE models, most valued the insight which they provide.[1] The second issue, published in 2020, examined whether models with multiple equilibria provide valuable advice for policy-makers. The answer was a modified 'yes'; such models change our understanding in ways which are often unexpected.[2] This third issue on modelling structural change broadens the agenda still further. It takes forward the multiple-equilibrium idea, and adds to this the ideas about composition, organization, and functioning that I have sketched above.

In the first section of this article, I provide a brief introduction. In section II, I outline a taxonomy of reasons why structural change can occur. Then in section III I go on to use this taxonomy to help me analyse the models of structural change to the growth process contained in papers in this *OxREP* issue. In section IV, I discuss the papers which show how structural change can also arise as a result of interactions between households, firms, and the policy-makers who are conducting stabilization policy in the shorter term, i.e. policy-makers at the central bank. In section V, I examine in some detail two papers that make clear how structural change can happen in the world economy system as a whole, rather than just in individual economies. In section VI, I discuss a paper which shows how structural change can emerge in models in which there is chaotic behaviour. Finally, in section VII, I describe two papers which show the way in which the models of structural change set out in this *OxREP* issue can be of use to macroeconomic policy-makers. In section VIII I offer some concluding remarks.

Some models of structural change, including a number of the models presented in papers in this *OxREP* issue, are somewhat technical. Where appropriate I have provided a non-technical explanation of the key arguments in these papers. This means that some sections of my article are rather dense; a reader of these sections will need to have the relevant paper open beside them as they read my article. As a result my article has become not just an essay on ideas but also a reader's guide for a set of interesting but challenging papers.

Some of the other models of structural change presented in this issue are much more straightforward, although they contain equally important ideas. I discuss the papers containing these models in much less detail, despite their significance.

Overhanging all that follows is the canonical model of economic growth: the Ramsey growth model. This model is normally set out as a simple, abstract, all-encompassing, one-sector model, of a representative agent deciding about savings and investment, using intertemporal optimization, and with a dynastic point of view that looks far into the future. That is how Olivier Blanchard and Stan Fischer described what they were doing in the Preface to their famous textbook (Blanchard and Fischer, 1989). Ever since then, textbook authors have aimed for something similar. What these accounts present is a model in which there is no structural change. The purpose of the model is to analyse the behaviour of an optimizing representative household which engages in a dynamic process of capital accumulation and growth. The model shows how the savings decisions of this

[1] See Vol. 34, Nos 1 and 2, available at https://academic.oup.com/oxrep/issue/34/1-2.
[2] See Vol. 36, No. 3, available at https://academic.oup.com/oxrep/issue/36/3.

household influence the level of the capital stock that emerges in the future—and thus determines how much output and consumption will be possible. More than this: the model shows that there is an optimal amount for this household to save.

The papers in this *OxREP* issue show—in many different ways—why that canonical model is an inadequate model of economic growth, precisely because of the possibility of structural change. It will become apparent that structural change is always possible, just around the corner.[3] This possibility should temper the confidence that one places in the policy recommendations which emerge from the canonical model. It is clear that a better understanding of the macroeconomics of structural change might have a beneficial influence on the way in which macroeconomics is practised as a discipline. And it is also clear—as I say in the conclusion—that such an understanding might change the way in which macroeconomics is taught to graduate students.

II. Macroeconomic models of structural change: a taxonomy

My taxonomy of structural change has three parts

First, some macroeconomic models contain tipping points which give rise to more rapid changes. In such models there is a region in which behaviour is locally unstable. Once the tipping point is reached, the economy tends towards a new equilibrium, quite different from the one which had previously been relevant. Such tipping points can come about within a single sector of an economy, for example due to increasing returns to scale. Or they can arise in multi-sectoral models due to strategic interactions between actors in different parts of an economy, for example between different members of a supply chain. The task of the model is then to analyse the reasons for such tipping points and the changes which they induce.

Second, some multi-sector macroeconomic models portray a process of gradual accumulation of something that is particular to one sector—for example capital, knowledge, or skills. This accumulation causes the output of that sector to gradually increase in size relative to output in the rest of the economy. The task of models of this kind is to display the reasons—to do with tastes, technology, or endowments—as to why such gradual changes might come about endogenously, within the model.

Third, there may be structural changes in the economy which are due to changes, external to the model, which mean that the nature of the model needs to change. Examples are market liberalization or industrial reorganization. Or there may be policy changes which cause structural change, such as the removal of protectionism in international trade. The task of the model is then to display the structural change which is caused by the change in circumstances being analysed.

In what follows I use this taxonomy to help me analyse the models of structural change presented in the papers in this issue. These are both models of structural change in the process of longer-term economic growth and models of structural change in relation to shorter-term macroeconomic policy-making. I consider these two sets of models in turn.

III. Longer-term economic growth

(i) Structural change due to tipping points in one-sector growth models

The first paper in Part I of this *OxREP* issue is by Tomohiro Hirano and Joe Stiglitz. In this paper the two authors discuss one way in which structural change can occur in the canonical Ramsey one-sector model of economic growth.

It is already well understood how structural change, to do with multiple equilibria, can occur in a Ramsey model for reasons related to the supply side of that model. There are at least two different reasons why this might happen. First, if increasing returns apply over a range, then the model can have two different equilibrium growth paths along which there are different levels of output. Second, it may be that the rate of technical progress is not exogenous but instead depends on how fast the economy is growing, perhaps because when growth is rapid, the expected rewards

[3] In fact, our argument goes beyond this. As we will see, an overlapping generations (OLG) set-up is a better representation of savings/accumulation decisions than that found in the Ramsey model. It will become apparent that a move to an OLG set-up, in which there is no longer the kind of transversality condition that is required to solve the Ramsey model, makes a big difference to all of the analysis which follows.

to invention are that much larger. In that case there may be one equilibrium growth path along which the rate of technical progress is high and another along which it will be low. This situation is analysed in Vines and Wills (2020), building on the work of Wendy Carlin and David Soskice (Carlin and Soskice, 2018) and on the earlier work of Stiglitz (Stiglitz, 1994). In their paper Hirano and Stiglitz do something different.

Growth paths with different levels of saving and investment

Hirano and Stiglitz show, instead, how multiple equilibria can also occur in a Ramsey model for reasons to do with the demand side of the model (Hirano and Stiglitz, 2025*a*, this issue—henceforth HS1). This amendment to the model which can bring this about involves replacing the assumption of an infinitely-lived representative-agent set-up with the assumption of an overlapping-generations (OLG) set-up in which each generation lives for just two periods. In such a set-up, one replaces the Euler equation for the infinitely-lived representative agent with an equation for each generation, showing how its members save when they are young and working to enable them to consume when they are old and retired. With that one amendment, the model can have two different stable equilibria. Savers and investors can converge onto an equilibrium with a high level of saving and investment which can be a self-fulfilling outcome, fully consistent with rational expectations. Or they can converge onto an equilibrium with a (much) lower level of savings and investment.

HS1 give a thought-provoking insight into the ideas about animal spirits that Keynes set forth in his *General Theory* (Keynes, 1936).In effect they show that animal spirits might be an endogenous self-fulfilling phenomenon.

The argument as to why that can happen is laid out in the following paragraphs.[4]

We know that in the normal Ramsey model the long-run equilibrium is unique. In that model the real interest rate, r, is tied down by the exogenously given rate of time preference. That long-run real interest rate then determines the long-run ratio of the capital stock to the effective labour force and so determines the level of output on the path along which the economy grows.

To begin their story, HS1 rely on the fact that, when the Euler equation in a Ramsey model is replaced by an OLG set-up in which agents live for two periods, the long-run real interest rate, r, is no longer tied down in this way. Instead, that long-run interest rate depends upon the amount which the young are willing to save. And so we can have a long-run equilibrium in which the real interest rate, r, will be equal to the marginal product of capital, which will be above the rate of time preference, δ. This gap between the r and δ is determined, on the one hand, by these savings preferences (i.e. the utility function of the representative household) and, on the other hand, by the reward for investing (i.e. by the nature of the production function). It turns out that there can be two stable long-run solutions for this interest rate, one in which consumers are happy with substantial savings, a low rate of interest, and a high stock of capital, and another, in which they are also happy, in which there are limited savings, a high rate of interest, and a low stock of capital.

Figure 1 in HS1 shows what can happen. Consumers work and save when they are young, and do not work but still consume when they are old. HS1 consider wide classes of utility and production functions for the representative consumer and the representative firm which include the class of constant-elasticity-of-substitution functions (or CES functions). They use certain simple parametrizations of these functions to discuss both what they call momentary equilibrium and long-run equilibrium in their model.

Momentary equilibrium is defined as the equilibrium which pertains in a short-run period in which the capital stock is given at some arbitrary initial level. We know that, given this initial capital stock, and given full employment of resources, the output of the economy will be pinned down because, at that instant, we also know the size of the labour force and the level of technology. Furthermore, the distribution of that output between labour and capital will be determined by whatever real wage is necessary to ensure that labour is fully employed. Consider now the savings behaviour of the young people in this economy. The SS curve in Figure 1 shows their supply

[4] HS1 present their model in a number of different ways. Here I provide a simple explanation of one version of their story.

of savings as a function of the rate of interest expected on their investments. HS1 consider, for instance, the case where consumers have a strong desire for intertemporal consumption smoothing, i.e. a low intertemporal elasticity of substitution, something which is consistent with the empirical evidence. (An extreme case of this is that of 'target' savings, where there is a certain amount that young individuals want to accumulate for the tine when they will be old.) This assumption means that, when the interest rate is high, they need to save less because the interest receipts which they will receive when they are old will be that much greater for any given amount of savings. And vice versa. So the SS curve is actually downward-sloping.

The II curve in Figure 1 shows the demand for investment as a downward-sloping function of the rate of interest, which is a consequence of the relationship between the marginal product of capital and the rate of interest, the slope of which depends on the elasticity of substitution between labour and capital in the production function. HS1 assume that the marginal product of capital and the interest rate fall significantly when extra capital is accumulated; that is what will happen when the elasticity of substitution in production between labour and capital, σ, is small. HS1 provide evidence that this assumption may be realistic.

As drawn, Figure 1 shows three equilibria. What is new in the HS1 paper is that it turns out that sufficient non-linearity can be introduced into the model to get something like Figure 1, even if the utility function and the production function are constrained to take a straightforward CES form. Even in such a simple set-up there can be two stable solutions for the interest rate r, and thus for the level of investment, if $\sigma < (1—\vartheta)$, where ϑ is the intertemporal elasticity of substitution.[5] In a word, it appears that both young and old can be happy, or at least 'in equilibrium', either with a lot of savings, a low rate of interest, and a high stock of capital, or with not very large savings and a high interest rate and not very much capital.

Now consider long-run equilibrium. Since in this economy we know the rate of growth of the labour force, the rate of growth of technology, and the rate of depreciation, all of which are exogenous, we know how large investment will need to be in order to simply maintain the level of capital per effective worker at exactly its current level. Long-run equilibrium occurs when this level of investment is exactly the amount which young people are willing to save. HS1 show that it may well be possible to draw something like Figure 1 for the long run as well as for the short run. That is to say, it may well be possible, if capital per effective worker begins at a low level, for the savings–investment equilibrium to be like the one in the upper left of the diagram which causes investment to be simply such as to continue sustaining the capital stock per effective worker at precisely that low level. But it may also be possible, if capital per effective worker begins at a high level, for the savings–investment equilibrium to be like the one in the lower right of the diagram which causes investment to be high enough to sustain the capital stock per effective worker at this high level. In a word, it may also be true in the long run—as well as the short run—that both young and old are happy either with a lot of savings, a low rate of interest, and a high stock of capital, or with not very large savings and a high interest rate and not very much capital.[6]

Finally, it seems useful to consider the likely transition from the short run to the long run. It appears that, starting from any initial level of capital per effective worker, if savers and investors choose the high investment level of the two equilibria, then the economy will embark on a growth process which converges to the long-run equilibrium in which there is a high level of capital per effective worker, and vice versa. Given such a story, we can think of the short-run choice by investors and savers as an equilibrium-selection device, one which determines whether the economy converges onto a growth path with a high level of capital per effective worker, or whether, instead, it does the opposite. In particular, starting from any initial value of capital per effective worker, the economy can move to one of two transition lines. The upper-level transition path, starting with the higher level of savings and investment, depicts a path of transition to the long-run equilibrium, with a high level of savings, investment, and capital per effective worker. By contrast the lower-level transition path, starting with the lower level of savings and investment, depicts a path

[5] These equilibria are stable in the sense that when r is above its equilibrium value, S > I, so that r will be under pressure to fall, and vice versa. This condition is not sufficient for there to be multiple equilibria—some other conditions are also necessary, as HS1 show. But it seems plausible that in certain circumstances these conditions might hold.

[6] Indeed, HS1 show that the conditions under which there is multiplicity of long-run equilibria are even weaker than the conditions for multiplicity of short-run equilibria.

of transition to the long-run equilibrium, with a low level of savings, investment, and capital per effective worker. Selection of the momentary equilibrium selection will, in certain circumstances, determine to which of the long-run equilibria the system converges.

The support that these findings give to Keynes's ideas about animal spirits goes like this. If in any one short-run period, savers and investors converge on an equilibrium with a high level of investment, then that produces a self-fulfilling outcome, and vice versa. That is to say, the short-run choice by savers and investors might be able to determine to which of its possible long-run growth paths the economy will actually converge.

This kind of argument seems more powerful than the conventional assumption that changes in animal spirits are due to an exogenous shift in the beliefs of consumers or investors. It means, as HS1 maintain, that changes in animal spirits do not necessarily have to be based on irrational exuberance or pessimism. There can be more than one value of animal spirits which is a rational-expectations equilibrium. As a result, endogenous changes in animal spirits may be able—on their own—to determine whether the economy ends up on a more, or less, productive growth path, quite independently of any changes in tastes, technology, or endowments in the economy.

Of course, this argument about equilibrium selection will be weakened if animal spirits are volatile, and HS1 argue that in a more general model, where the economy is bombarded by non-stationary shocks, there are good reasons why they might be. In that case an equilibrium selection choice in the short run will not tie down the process of convergence to a particular long-run equilibrium, since, in future periods, the equilibrium-selection process may give rise to a different outcome.[7]

However, once any particular equilibrium has been selected, it may be difficult for actors in the economy to re-converge onto a different equilibrium. Blake, Dennis, Kirsanova, and Yates (BDKY) argue along those lines in their paper in this issue, as I discuss below. They show that it may be difficult to reverse any initial choice about which equilibrium is selected.[8] Nevertheless, HS1 provide a general parameterization for sticky expectations, showing that if they are not sticky enough, the short-run choice of momentary equilibrium does not necessarily tie down the long-run trajectory, and economic trajectories can be characterized by complex 'wobbles' without any fixed periodicity.[9]

The implications of this equilibrium selection argument for policy-makers are profound. It may be that, despite what has just been said, macroeconomic policy is able to nudge the choices made by savers and investors in the short run in such a way as to bring about a higher level of savings and investment. If they can do this, the long-run growth trajectory of the economy may turn out to be very different.

(ii) Structural change due to different sectoral growth rates in multi-sector growth models

I now discuss a number of papers, still in Part I of this *OxREP* issue, which show how structural change might occur in a model with two or more sectors. To understand what is happening, one needs an analysis of why it is that one of the sectors is likely to grow faster than others. The first of these papers by Douglas Gollin and Richard Rogerson sets out a number of reasons as to why this might happen, i.e. why an economy might come to have a 'leading sector'. This paper is followed by Dan Susskind's article showing how an endogenous process of technical progress can reinforce the way in which the leading sector of an economy grows faster than the other sectors by helping to reduce costs in the leading sector and thus further stimulate its relative advance.

[7] HS1 show that such volatile animal spirits can be *fully* consistent with rational expectations—unlike in the standard model, rational expectations *does not* tie down the economic path. There is an infinity of rational expectations trajectories.

[8] In fact, as we will see, BDKY are discussing a very different kind of multiple equilibrium, namely one to do with the conduct of macroeconomic policy, in which the kind of equilibrium chosen will depend on a choice made by a policy-making institution—namely the Central Bank. Nevertheless what they say in relation to their own model, about how difficult it may be to change the kind of equilibrium which is selected, is relevant here too.

[9] Of course, HS1 show that *not everything* is possible, i.e. we can still define bounds on the fluctuations that are consistent with rational expectations. Nevertheless, going beyond this, HS1 suggest that weakening the criteria for intertemporal consistency, say from fully rational expectations to some form of reasonably rational expectations, opens the door for a far larger range of economic fluctuations.

Susskind's paper is followed by a second paper by Tomohiro and Joe Stiglitz (HS2) which analyses the effects of speculation in land on the growth rate in a two-sector model containing a capital-intensive industrial sector and a land-intensive agricultural sector. The authors focus on the way in which this speculation is likely to reduce the overall growth rate of the economy by attracting savings away from the more progressive industrial sector.

Capital accumulation and technical progress

In their paper, Douglas Gollin and Richard Rogerson (2025, this issue) describe in broad terms the insights which have flowed from going beyond a one-sector model of economic growth to models with more than one sector.

We have understood ever since Adam Smith and David Ricardo that structural change can be studied in a two-sector model of industry and agriculture. The reason that this can happen is that capital accumulation and/or technical progress can be concentrated in the manufacturing sector in such a two-sector model.[10] The article and book by Arthur Lewis on economic growth (Lewis, 1954, 1955) became highly influential—winning him the Nobel Prize—precisely because they began to show how this two-sector process might be formalized.[11] Lewis set out institutional and demographic reasons why savings, capital accumulation, and technical progress are concentrated in the industrial sector of the economy, leading to an increase in the demand for labour in that sector. He thought that economic development involves a progressive movement of labour from a largely agricultural economy to a largely industrial one, which is why he considered the one-sector Solow growth model to be inadequate. If there was a surplus supply of labour in agriculture, he maintained, then withdrawal of labour from agriculture might not lead to an increase in the wage, or to a reduction in the availability of agricultural goods. But after the arrival of the economy at what became known as the Lewis turning point, wages would begin to rise in the manner suggested by the Solow growth model. And the withdrawal of labour from agriculture would give rise to an increase in the price of agricultural goods, further constraining the rate at which the economy might grow, just as Solow had suggested (Vines and Zeitlin *et al*., 2008).

Trevor Swan presented a fully worked out version of this Lewis model, almost immediately after it appeared,[12] in the paper in which he set out the same growth model as provided by Robert Solow (Swan, 1956). The last part of Swan's paper describes a set-up in which there is a constraint on growth due to a fixed supply of land. This constraint ensures that the capital accumulation process cannot achieve constant returns to scale, unlike what Solow assumed when presenting his model.[13] Swan examines such an economy in which capital accumulation and technical progress are concentrated in one of the sectors—industry—but in which the output in the other sector—agriculture—is constrained by the fixed supply of land.[14] The economy is also one which has passed the Lewis turning point so that there is no surplus labour available in agriculture. Swan showed that growth in such an economy would lead to an increase in the rents accruing to the owners of land and so to a constraint on the extent to which wages could rise as growth took place. In a subsequent analysis, Swan (1957) showed how this difficulty could be offset by technical change of a kind specifically designed to unbind the land constraint, using the ideas about directed technical change that are discussed by Daniel Susskind in his paper in this issue, which I consider below.

Gollin and Rogerson (henceforth GR) describe the wealth of empirical findings which have flowed from these Lewis-model insights, particularly in the last 20 years. As they explain, the central finding of the research work on both the one-sector growth model and models of the

[10] See Ricardo (1817) who built on the central idea in Book 2 of Adam Smith's *Wealth of Nations* (Smith, 1776). John Hicks provided a formal discussion of how he did this but not until in the mid-1960s (see Hicks, 1965, ch. 4).

[11] The Ricardian growth model was not formalized until some years after Lewis's article and book appeared (Pasinetti, 1960). I myself played a small part in that process of formalization; see Molana and Vines (1989).

[12] See Butlin and Gregory (1989) and Pitchford (2002).

[13] Solow said 'One can imagine the theory [presented here, in which there are constant returns to scale] as applying as long as arable land can be hacked out of the wilderness at essentially constant cost.' (Solow, 1956, pp. 66 and 67)

[14] In fact, Swan never set out a version of his model in which two sectors are explicitly identified, the formal model in Swan (1956) contains only one sector. But we know from other work of Swan's at the time that he was used to writing about a small open economy—like that of Australia—in which all tradeable goods could be sold in world markets at prices exogenous to the economy and could therefore be aggregated onto a single sector. See Swan, P. (2022). That appears to be the procedure which Swan adopted in that paper.

two-sector kind like the Lewis model, is that growth in output per person has been mainly caused by technical progress rather than by capital accumulation. So, of course, the obvious question becomes what is it that leads technical progress to be concentrated in one sector rather than another, and whether that concentration will continue to happen. As GR state, to provide a detailed answer to this question requires one to look at questions that go beyond the analysis of a one-sector model to models of two or more sectors. Only then can one begin to understand where the impetus for economic growth actually comes from and the way in which the resulting growth process is intrinsically connected with structural change.

One insight provided by a two-sector approach, as GR remind us, is that such an approach can help to explain the findings of Simon Kuznets on the relationship between economic growth and economic inequality. The growth process—according to Kuznets—is often associated with first an increase in inequality and then a reduction. The Lewis model provides a simple way of understanding how this might happen. If the real income per person is initially much lower in the agricultural sector than in the emerging industrial sector, then initially growth will cause inequality to increase as more of the population moves to the industrial sector and becomes better off. But there will come a point when so few people remain in the poorer agricultural sector that further movement of labour into the industrial sector begins to cause inequality to fall again.

GR show that models of more than one sector reveal interesting and important observations about the way in which labour relations differ between sectors. It is also possible to infer ideas about the way in which industrial organization will differ. Furthermore, as I will discuss below, different sectors are involved in trade in different ways. All of these factors will influence the way in which technical progress happens, capital is accumulated, and structural change takes place.

Directed technical change

The next paper in this part of the issue, by Daniel Susskind, continues the discussion begun by Gollin and Rogerson by asking: why is it that particular kinds of innovation take place? (Susskind, 2025, this issue). Endogenous growth theory has familiarized us with the idea that technical progress, in general, takes place because there are rewards to be made from innovating. But which *kinds* of innovation produce rewards? Susskind considers how technical progress is 'directed' to take certain forms by the incentives that are thrown up by each particular kind of innovation. He discusses in detail how these incentives should be modelled. And he considers how such models might be used to guide policy on innovation.

Much of Susskind's attention is focused on the idea that technical progress might be biased towards economizing on the use of particular factors of production. That is, technical progress might be labour-saving or capital-saving. Or, in more detail, it might be focused on saving particular kinds of labour—whether unskilled, routine, or skilled. He describes a number of ways in which technical progress might be directed in this way. All of this can be modelled, in a valuable way, in a one-sector model of the economy. In which case, structural change might be thought of as being the innovations that happen in the process of producing 'output-in-general'.

However, in the later part of his article Susskind goes on to discuss how and why technical change might be directed towards particular sectors of the economy. Thus, Susskind maintains, the idea of directed technical change can help one to understand why it was that the incentives facing producers caused 'technological progress.. [to be] biased towards the labour-saving, energy-intensive innovations that started the Industrial Revolution—the steam engine, spinning jenny, roller spinner, power loom, and so on' (Susskind, 2025).

It was profitable to carry out this innovation in the industrial sector rather than the agricultural sector—in the way which Ricardo and Lewis had described but not explained—because the industrial sector was where it was most profitable to innovate. Susskind shows how this can be modelled.

Susskind carries his historical discussion over into an examination of the contemporary world, considering how and why it is now profitable to innovate in particular sectors, be they energy, automotive, or IT. And in the last part of the paper he discusses the way in which policy—taxes, subsidies, regulation—might be used to influence the way in which innovation is directed. There is a discussion of how those whose livelihoods are displaced by innovation may resist this process and of what policy might do in the face of this resistance.

Endogenous growth, financial development and structural change

The next article in Part I of this issue is a second paper by Tomohiro Hirano and Joe Stiglitz (2025*b*). Here, they show how Lewis-model ideas can be developed by adding three additional factors to the set-up: the effects of endogenous technical change in the industrial sector; the effects of an overlapping-generations population structure on the savings behaviour of those who work in the industrial sector; and the effects of financial developments which enable those who save in the industrial sector to also invest in land as well as investing in industrial capital (Hirano and Stiglitz, 2025*b*, henceforth HS2). In this article, the authors demonstrate how the structure of the economy will be profoundly altered by these three factors.

The basic ideas in HS2 can be laid out as follows.[15] The model has, at its core, a specific-factors set-up in which labour is employed in both the industrial sector and the agricultural sector, in combination with capital in the industrial sector and with land in the agricultural sector. Consumers who are employed in the industrial sector consist of overlapping generations of workers who, while working, save for their retirement. HS2 suppose that growth in the industrial sector is subject to what they call a standard form of Marshallian-external-increasing-returns-to-scale in capital. This makes the model an endogenous growth model, since the rate of technical progress will depend on the rate of growth of the economy. HS2 initially assume a particular parameterization of this form of technical progress such that, although each individual firm experiences diminishing returns in the application of labour to the capital which it employs, in the aggregate the economy experiences constant returns to scale in capital.[16] The model is thus able to act as an AK model (the kind of model in which the level of output is related to the level of capital in a linear way) without experiencing any diminishing returns in the growth process. Nevertheless, the two-sector nature of the model influences the way in which this growth takes place, since the productivity of labour in the agricultural sector, and the relative price of the agricultural good, which in turn depends on the nature of the consumer demand for that good and thus on the consumers' utility function, will all influence the size of the A parameter in the AK model.

This result is a clever story about how the kind of two-sector economy imagined by Lewis might perpetually grow, once the Lewis turning-point has been passed.[17] The growth process becomes one in which two conflicting forces are held in balance: the shortage of land as economic growth takes place is perpetually compensated for by the technical progress in the industrial sector. The proportion of the labour force employed in the industrial sector now no longer grows but remains constant, even though capital accumulation is taking place. For each individual firm, the marginal product of labour is being pulled up by means of the economy-wide capital accumulation and technical progress. But at the same time, the real wage is perpetually rising at a rate which ensures that the demand for labour by the representative firm continues to be equal to the unchanging amount of labour which is employed in the industrial sector. The speed of the pulling-up process is determined endogenously by the model and—because of the AK structure of the model—remains constant over time.

[15] Hirano and Stiglitz present a number of different versions of this model, as is the case in their HS1 paper. In particular they also set out a version of their model in which there is only one sector but in which there is a constraint on growth due to a fixed supply of land. (As noted above, that is what Swan did when presenting his version of the Lewis model.) The two-sector version of their model that I am explaining here fits in well with the two-sector story that I am telling in this part of my essay. The one-sector version of their model enables them to study—in the same way that Swan did—how economic growth increases the price of land and the rents which accrue to those who own it. I will not discuss that model here, but I recommend it because it says something about the place of cities in the overall process of economic development.

[16] They show later in their paper that their main results do not depend on this.

[17] The workings of the economy (before any land reform) are as follows. Given the production function in the two sectors and the utility function of the workers in the industrial sector, we can work out the relative price of food and thus the real incomes of those who work in industry and so how much they are able to save. HS2 make use of a simple utility function, namely $u_t = \log(c_{1t}) + \beta\log(c_{2t+1})$, where t is time, c_{1t} is the consumption of industrial goods, and c_{2t} is the consumption of agricultural goods, or food. This utility function ensures that the savings rate is independent of the rate of return which these workers receive from their holdings of industrial capital, a simplification which makes the model easier to solve. This savings rate then determines the growth rate of the economy. But given that land remains scarce and that production in agriculture does not experience technical progress, the price of food will need to perpetually rise as this growth takes place. A simple trick is deployed, by supposing that the agricultural production function has a Cobb–Douglas form. This together with what is assumed about the utility function of the workers, ensures that the proportion of national income given over to land rents, and the proportion of consumer expenditure devoted to the purchase of food, both remain constant, even though agricultural rents and the price of food are perpetually rising as growth takes place.

The model's AK structure also ensures that, following on any shock to the system, the economy immediately moves onto its new equilibrium growth path. What happens, say, after an earthquake destroys half of the industrial capital is that the economy immediately adapts, in a general-equilibrium manner, to the shock. The outcome will be one in which, immediately, the wage and the price of food are both lower. And it will be one in which, immediately, the process of economic growth continues at the same rate as before (and continues forever), but from a lower base.

The second set of insights from the AK model's structure concerns the effects of land reform on the overall growth process. These effects are very remarkable.

As a result of any such land reform, industrial workers are now able to invest their savings not just in the entitlements to income which come from their holding of shares in physical capital. They can also hold titles to the ownership of land as part of their asset portfolio. By doing this they receive rents from their ownership of land. They also benefit from the fact that the price of the land which they own will be perpetually rising, because of the shortage of land.

There is again an equilibrium growth path for this post-reform economy along which the growth rate remains constant. The AK nature of the model set-up goes on being helpful in ensuring that this can be the case. But now, along this growth path, workers will be investing only some of their savings in industrial capital, because they are also investing in holdings of land. The value of those land holdings keeps pace with their holdings of industrial capital precisely because of the appreciation in the value of land which is continually taking place. The authors call this process of holding claims to land, whose price is perpetually rising, a process of 'speculation in land'.

The effect of this speculation is to *lower* the growth rate of the economy.

The fact that such a financial liberalization can lower the growth rate might at first sight seem paradoxical. As the authors say, it goes against the convention that financial reform can be useful in helping to promote the process of economic growth. But the finding can be straightforwardly explained. Savings are being diverted from investment in industrial capital because they are being invested in holdings of land. In effect, part of the savings of industrial workers are being diverted away from investment in industrial assets—investment which would enable the capital stock to grow and so would enable an increase in the rate of technical progress—and are being directed towards the investment in assets which give the asset holder the ownership of some land. Part of the value of these land holdings arises not because this land is doing anything useful, and thereby generating rents, but simply because its price is rising. Once it is explained this way, the fact that land reform causes a reduction in the growth rate of the economy becomes intuitively obvious, even although at first sight it might seem to be counter-intuitive.

The AK structure of this version of the model, in which there is speculation in land, is still one that ensures that the solution for the growth rate of the economy is very simple. Once again—following any shock to the system—the economy immediately moves to its new equilibrium growth path. What happens after, say, an earthquake destroys half of the industrial capital is again that the economy immediately adapts to what has happened, in a general-equilibrium manner. The outcome will again be one in which, immediately, the wage and the price of food are both lower. And it will again be one in which the economy immediately moves to its new general-equilibrium structure, which continues the same forever.

HS2 go on to examine extensions to this AK model in which the fixity of the general-equilibrium structure associated with their AK model disappears. This happens because—in the generalizations of the model which are examined—the Marshallian external economies in the industrial sector are no longer able to fully offset the effects of diminishing returns caused by the shortage of land. In this case, the growth rate of the economy gradually converges to a rate of growth which is determined by whatever exogeneous process of technological change is happening, and the level of the capital stock gradually converges to the one at which the growth of capital stock happens at just this rate. This will be amplified by whatever benefits arise from (the limited) endogenous growth (in their model, again, arising from the limited Marshallian increasing returns). In such a case, after an earthquake destroys some of the industrial capital, the level of capital per effective worker will—eventually—return to its previous position. As part of this process, the relative price of land will at first collapse, because there is no longer, in the short run, such a large demand for land (as the real incomes of workers will have fallen since they are working with much less capital). But this relative price of land will gradually return to its long-run level. As HS2 show, the solution of the model will, in these circumstances, have saddle-path properties. The initial jump down in the relative price of land will involve a jump to a level from which it can gradually recover, along its saddle-path trajectory, to its final resting position. That is something which happens gradually as the stock of capital is slowly rebuilt.

HS2 also consider an extension of the model in which savers can leverage their holdings of land, making money out of the fact that the price of land is perpetually increasing. An earthquake in such an extended model is capable of leading to a major initial collapse in the price of land. That is because the earthquake will damage the collateral which the borrowers have used to support their loans. But in this model—providing there is not a crisis which causes the whole economic system to disintegrate—the price of land will, in due course, move along a saddle path back towards equilibrium as the capital stock is gradually adjusted back to its long-run equilibrium level and as the collateral of the borrowers is gradually rebuilt.

This model brings the Lewis model into the twenty-first century in showing how we can incorporate endogenous technical change in the aggregate industrial sector, and in showing how speculation in land can alter the growth properties of the economy in a striking manner. The authors examine a number of significant policy questions thrown up by the model, including how land reform should be managed, and whether the leverage described above should be constrained by regulatory policy.

Structural change due to involvement in international trade

Part I of this issue continues with a paper by Peter Warr and Arief Anshory Yusuf describing the way in which economic growth can lead to structural change in any single economy, as economic growth leads that economy to export more of the goods in which it has a comparative advantage, with the result that the growth in other sectors falls behind growth in the exporting sectors (Warr and Yusuf, 2025*a*, this issue, henceforth WY). WY use a small, comparative-static, empirically-based, computable-general-equilibrium (CGE) model of the Thai economy to provide an empirical, model-based study of the way in which there have been long-term changes in the composition of economic output in Thailand during the period from 1990 to 2014. The exercise tells us much about the kinds of structural change that we might observe in other emerging market economies that are relatively open to international trade in the way that the Thai economy has been.

We have had models for a long time showing the way in which involvement in international trade leads to structural change. Examples are the Ricardian model, the Heckscher–Ohlin (HO) model, and the oligopolistic models of new trade theory. And there are also the data-intensive CGE models which—unlike the HO model—contain many sectors.

WY use an applied version of the classical Australian model, which is, in effect, a three-good-three-factor variant of the HO model.[18] In the original version of the Australian model it was assumed, as in the HO model, that only one good was exportable—commodities—and only one good would be importable—manufactures. It was also assumed that a third good could be added to the HO set-up which was neither exported nor imported but was non-tradeable. This was also, unlike the HO model, a specific factors model; that is, although labour was assumed to be used in the production of all three goods, one of the tradeable goods—manufactures—was assumed to be produced only with labour and capital and the other—commodities—was assumed to be produced by land and capital, while the non-traded good was assumed to be produced by labour alone.[19]

WY deploy the same small-open-economy, three-factor, three-good structure, as the Australian model, but allow labour and capital to be used in the production of all goods[20] and also allow—in a CGE manner—for the use of intermediate inputs into production. Manufactures are the most capital-intensive good, agricultural goods are the most land-intensive, and the services sector is the most labour-intensive. And for WY tradability is a matter of degree. Within the Australian

[18] See Jones (1965), Anderson (1987), Devarajan *et al.* (1993), and Menzies and Vines (2008). WY use data from the full GTAP database (Global Trade Analysis Project) but aggregate these data into three sectors.

[19] See Brigden *et al.* (1929) and Reddaway (1937).

[20] But, as in the Australian model, land is used only in the production of agricultural goods.

model, domestically produced importable goods are perfect substitutes for imported goods—implying that the Armington elasticity of substitution between these goods is effectively infinite in absolute value—and the demand for exports is infinitely elastic. This implies that, in the Australian model, the domestic prices of both importable goods and exportable goods are fully determined by international prices. By contrast, within the WY model, there are finite Armington elasticities of substitution between imported goods and domestically produced goods, and export demand elasticities are finite; the values of these elasticities are empirically determined. These assumptions imply that there can be considerable domestic price flexibility within the WY model. The set-up of the model also—unlike in the Australian model—allows all three goods to be exported or imported. But the calibration of the model is such that manufactures form a larger proportion of exports than they do of imports. And relatively little of the demand for services is met from abroad and relatively few services are exported.

The original Australian model depicted an economy that had already passed the Lewis turning point. As a result, all three factors were assumed to be fully employed. There was also no international mobility of capital. These two assumptions are also deployed by WY.

WY simulate their model to study the way in which four different factors have affected the structure of the Thai economy. These are: differential growth rates of aggregate supplies of physical capital, labour, and land (they call this the Rybczynski effect); differential growth rates of total factor productivity between sectors; changes in relative international prices; and changes in sectoral rates of trade protection. They also examine the effect of differences in the expenditure elasticities of demand for the output of the three sectors, which they call the 'Engel's law' effect.

The authors conclude that, in Thailand's case, the second factor—differential rates of productivity growth between sectors—has been by far the most important determinant of structural change. This finding can be easily understood given the nature of their model. First, they show that the productivity growth which occurred in Thailand was concentrated in the manufacturing sector. Their simulation of the effects of this productivity growth in manufacturing shows that the growth which it caused was not impeded by a shortage of agricultural goods, and a consequential rise in the price of these goods, since these goods were available as imports or as a diversion from exports. Second, more subtly, their model also suggests that, because of the increase in real wages which this simulation caused—unimpeded by the very small rises in the price of agricultural goods—the price of non-tradeable goods rose relative to the price of tradeable goods (both agricultural goods and manufactures) which had the effect of shifting demand away from non-traded goods towards traded goods, creating an additional reason for why the manufacturing sector was able to grow relatively rapidly as a result of the technical change which had occurred.

WY also find that the very rapid accumulation of capital in the Thai economy in the period being studied does not explain as much of the increase in the relative size of the manufacturing sector as one might have initially expected. First, although the growth rate of the capital stock relative to the labour force is large,[21] their model allows for the fact that capital is used in the production of services and agricultural goods as well as in the production of manufactures. The solution to their model assumes that capital is mobile between sectors in such a way as to equalize the rate of return on capital wherever it is applied in the economy. This means that some of the capital that was accumulated over the period studied was—in effect—spread out across the agricultural sector and the services sector, rather than merely causing an increase in the rate of growth of the manufacturing sector.

The structure of the WY model also explains why the Engel's law effects were much smaller than one might have expected. Put simply, that law suggests that, since the income elasticity of demand for food is much less than one, rising real income, as the economy grows, is likely to cause a gradual change in the proportions of demand devoted to different goods, away from primary commodities and towards manufactures and services. In the kind of small open economy depicted in the Australian model the prices of the two tradeable goods are exogenous, relative to one another. But the price of the non-tradeable, services, is endogenous, relative to these two tradeable prices. The shifts in demand resulting from Engel's law effects raise the domestic price of services relative to the two tradeable prices. Resources move from some combination of the two tradeables

21 WY naturally assume that the stock of land did not grow at all during this period.

industries to services. The effects on production of the two tradeables depends on factor intensities in the three industries. Contraction of agriculture would release very little capital. If the services industry uses capital, the capital required for its expansion will presumably come primarily from manufacturing. Manufacturing output will contract, relative to what would otherwise have been the case. If production of services uses only labour, as assumed in the original Australian model, the labour required for expansion of services must come primarily from agriculture. Agricultural output will contract. But in the WY model, the endogeneity of the two tradeables' domestic prices relative to one another—a consequence of the Armington assumptions and the export demand assumptions—moderates these impacts, lessening the magnitude of these Engel's law effects.

In a follow-up paper, WY explore the role played by domestic price flexibility in determining the structural-change consequences of the kind of exogenous shocks that are analysed in this paper (Warr and Yusuf, 2025*b*). As already noted, the domestic prices of both importable goods and exportable goods *do* respond to domestic demand shocks, as well as the prices of non-traded goods, because of the assumptions noted above. By contrast, in the Australian model the domestic relative prices of both importable goods and exportable goods are fully determined by international prices, although the price of non-traded goods is endogenous. Furthermore, in the standard Heckscher–Ohlin–Samuelson (HOS) small country trade-theoretic model, as used by Rybczynski, there is not even this degree of flexibility since there are no non-traded goods. WY show that this difference in assumptions about domestic price determination explains the difference between the empirically inaccurate structural-change predictions that can be derived using the Australian model, or the HOS model, and the more accurate results that are obtained in this paper.

The model which WY use also reveals an important fact about the effects on the Thai economy of the large reduction in protectionism that occurred during the period being studied. One might have expected that the reduction in the protection of import-competing production which happened would have caused a significant move in the Thai economy away from the production of importable goods and towards exports.[22] However, simulations of the WY model suggest otherwise. The main reason for this appears to be that the increase in exports to which this process led caused a significant fall in the price of Thai exports on world markets. That fall in price clearly dampened the degree to which resources are reallocated in response to protection.

As the authors maintain, their investigation of the Thai economy may well be relevant to how one might come to understand structural change in other countries that are experiencing a similar growth path. Of course, the details will be different. But the authors have demonstrated a method by which such studies might be carried out, using a simple comparative-static CGE model of the kind which now exists for many countries.

(iii) Structural change caused by the transition to net zero carbon emissions

The final two papers in Part I of this issue deal with a very particular form of structural change, namely the move of the global economy towards net-zero emissions of greenhouse gases. This move will inevitably involve structural change of a multisectoral kind, as the renewable energy sector grows at the expense of sectors producing or using fossil fuels. A paper by Rick van der Ploeg and Tony Venables (2025*a*, this issue) shows how this process is likely to involve tipping points caused by strategic complementarities between different sectors of the economy in a number of different ways. And a paper by Ross Garnaut (2025, this issue) shows how such policies will necessarily lead to a changing structure of international trade as some countries—for example, Australia—come to specialize in the production and export of commodities using non-fossil-fuel energy in which they have a comparative advantage. Two papers, each by several authors, describe the kinds of models which should be constructed to help manage the policy decisions which will be needed as the global economy moves towards a net-zero-emissions position. I postpone my discussion of these two papers until section VII of this article.

It is first desirable to begin with macro big-picture issues, and examine the place of integrated assessment models, or IAMs, in the transition to net zero. The IPCC (Intergovernmental Panel on

[22] Some of the reduction in protection which occurred was, in fact, a reduction in the protection of exportable goods (manufactures) some of which had in fact been imported. This was also important in reducing the reallocation of resources which happened.

Climate Change) makes use of IAMs to investigate how macroeconomic outcomes are likely to be damaged by global warming and how policy-making should respond to these looming problems (Intergovernmental Panel on Climate Change, 2022). The best known of these models is the DICE (dynamic, integrated climate economy) model developed by William Nordhaus at Yale University (Nordhaus, 1992, 2018). This model has proved an influential research vehicle for the past quarter of a century, and Nordhaus was awarded a Nobel Prize in Economics in 2018 for his work in this field. These models—and their continuing development—cast such a large analytical shadow over this whole area of policy-making that it seems important to begin our discussion of this topic by considering how they work and the insights they provide.[23]

IAMs contain—at their core—a rather simple Ramsey growth model. To this is added a model of the evolution of the stock of atmospheric carbon and the effects of this stock in causing global economic damage. One can think of such models as a two-sector generalization of the Ramsey growth model, in which one of the sectors produces normal goods and services and the other sector uses these goods and services to carry out abatement activity, including carbon capture and storage, and the production of energy in a carbon-friendly manner using solar and/or wind power.

Consumers save, which causes the stock of capital to grow. This, in turn, enables the output of normal goods and services to grow. But the continuing increase in output leads to more carbon emissions, which causes the stock of atmospheric carbon to continue growing, leading to an increase in the global temperature and so to an increase in climate-related damage. In such a set-up there are two slowly moving state variables: not just the capital stock, as in the Ramsey model, but also the stock of atmospheric carbon.

In order to stop damage growing without bound, resources must be devoted to abatement activity, such as carbon capture and storage (CCS). Resources must also be devoted to investment in capital to produce energy in a climate-friendly manner, by means of solar panels and wind turbines, and the construction of nuclear power stations, so as to reduce the emissions intensity of output. To analyse the role of this additional form of capital accumulation, IAMs need to include a third, slowly moving state variable in their model set-up. Here I simplify my discussion by concentrating exclusively on abatement. But the paper by van der Ploeg and Venables, discussed below, considers the accumulation of such climate-friendly capital.

When a position of net zero emissions (NZE) has been reached, abatement activity will need to exactly compensate for the fact that the economy is still growing, something which would otherwise lead to a growing stock of atmospheric carbon. Let us suppose that—as in the normal use of a Ramsey growth model to study macroeconomic policy-making—the task of policy-makers is to maximize an appropriately discounted flow of consumption over time. But now this is to be achieved while the economy is in transition towards NZE, along which emissions are being gradually curtailed. IAMs enable one to study how policy should be configured to do this.

In short, two things are necessary if the transition to NZE is to be well managed.[24]

First, when choosing their overall level of savings, households need to recognize the fact that each addition to the stock of capital will generate more carbon emissions and thus cause additional damage to the economy. Compared with the outcome in a standard Ramsey growth model, households will therefore need to constrain the amount of capital which they accumulate. IAMs enable one to investigate the extent to which a carbon tax, or environmental regulations, could be used to push the economy towards this desirable outcome. Susskind (2024) provides a clear analysis of the trade-off which this constraint imposes as the economy transitions to NZE. This is a managed-growth solution rather than the kind of no-growth solution which has been popularized by Kate Raworth in her discussion of 'doughnut economics' (Raworth, 2017).

Second, along any such desirable path to NZE, the amount of resources which policy-makers devote to abatement[25] must be optimal at each and every point in time. In particular, resources devoted to abatement have a social cost—they are resources which are not being used for consumption—but they also have a social benefit—a reduction in economic damage caused by atmospheric carbon. A desirable outcome is one in which, at every point in time, the marginal

[23] I am grateful to Shamik Dhar for discussions of the material that follows.

[24] I owe my understanding of both of these ideas to my conversations with Shamik Dhar.

[25] And in a more general treatment, to the investment in capital used to produce energy in a climate-friendly manner.

cost of using resources for abatement exactly equals the marginal benefit from reducing economic damage.

The virtue of IAMs is that they enable one to examine how the behaviour of consumers and policy-makers needs to be coordinated in order bring an outcome with the two desirable features outlined in these last two paragraphs.

Those building IAMs have been able to use them to examine what would need to happen to policy if circumstances change so that, for example, the process of abatement became more efficient or if—to go in the other direction—the damage caused by an increasing stock of atmospheric carbon became more significant than had originally been expected. They have also examined what would be a desirable response, by households and policy-makers, if the stock of atmospheric carbon suddenly increases, as a result, say, of the sudden melting of polar ice caps.

It is true that IAMs can help us study what happens when things get bad, but what they do not do, by and large, is study tipping points. Nor do they focus on the kinds of unstable processes which can *cause* tipping points.

Tipping points in the transition to net zero

Rick van der Ploeg and Tony Venables show how important such unstable processes are likely to be in the process of transition to net zero (van der Ploeg and Venables, 2025*a*, this issue, henceforth VV). In particular, they first show that the presence of strategic complementarities stemming from technological spillovers, from peer effects in demand, or from the kind of propagation and amplification mechanisms depicted in IAMs, are all likely to hasten the process of transition to fossil-free energy sources in the global economy. Learning by doing is also really important.

VV describe how the existence of such tipping points can be used to increase the effectiveness of climate policies. They suggest that it may be possible to achieve a transition to NZE with much smaller policy interventions than has been thought to be necessary. Nevertheless, if there are such tipping points, costly climate policies may well be needed to initially shift the economy from its current high-emissions path to a low-emissions path—what they call ‘radical and ambitious plans’. It may be that it is only possible for these costly policies to be withdrawn once such tipping points have been passed. More generally, VV show that radical policies can set in motion social, technological, and political tipping points. The rationale for such policies will then be strengthened by key households, corporations, and institutions which are at the centre of propagation networks. Thus, they argue, climate policies should identify those key agents and leverage them. Their proposals offer a complementary perspective to that offered by scholars concentrating on early warning signals. This leads VV to identify what they call ‘sensitive intervention points’ in order to implement more effective and transformative climate policies.

The version of their model which VV publish in this *OxREP* issue is a comparative-static model. But in a companion paper they present a dynamic model in which heterogeneous consumers make forward-looking choices between brown and green durable goods and then establish conditions under which peer or network effects create lock-in to an initial brown steady state, preventing any transition to an alternative green steady state (van der Ploeg and Venables, 2025*b*) The central question then becomes how to make such a transition path possible. A path might become sustainable because of peer effects among the consumers, each of whom recognizes that it is actually happening. Nevertheless, they also show how, in such circumstances, there may be a need for a supportive policy, such as a green subsidy, which exceeds a critical threshold level, in order to make a move to such a transition path possible.

Modelling the role of international trade in the transition to net zero

The paper by Ross Garnaut contains a review of the many global issues facing the international community in the endeavour of creating a net zero environment. Garnaut reminds us of the ingredients that have been necessary in the past for successful economies to create wealth and well-being, namely the coming together in one place of supplies of energy, materials, new ideas, capital, and labour (and, ideally, cheap labour). Historically, those places which were fortunate to possess all of these assets became wealthy. Those that did not possess them needed to resort to acquiring them from elsewhere by trade, and a large amount of it.

Looking ahead, this process will need to happen again. Garnaut shows that the building of a world economy with net zero emissions will involve major changes in the location of global production and trade. There will be transformative shifts in comparative advantage in production of energy, chemical inputs to industry, and industrial output. Different commodities will enter international trade from different sources and with different destinations. This will transform the location of manufacturing of iron for steel-making, aluminium, and other metals for a variety of uses, green fuels for transport, and urea and other manufactured chemicals. Comparative advantage in these commodities will be especially strong where relative abundance of renewable energy resources sits alongside relatively rich endowments of complementary resources—mineral oxides for iron, aluminium, and other metals production, together with sustainably harvested biomass for transport fuels and other chemical manufactures.

Garnaut's paper is not about modelling. But it contains a very clear indication of the kind of sectoral decomposition which will be needed if the process of international transformation which he describes is to be properly modelled. Such modelling will be essential for the purposes of global economic policy-making. The paper in this *OxREP* issue by Liu *et al*., discussed below, describes how such a model has actually been built, namely the G-Cubed model. Warwick McKibbin and I have discussed how this model can be used to study the direction and magnitude of the trade flows which will be necessary as the economy makes a transition to NZE (McKibbin and Vines, 2023).

IV. Shorter-term macroeconomic policymaking

The four papers in Part II of this *OxREP* issue discuss the way structural change can emerge in an economy in the short term, as a result in particular of the interaction between the macroeconomic policy-makers who set the interest rate and those in the private sector who determine the inflation rate. As we shall see, multiple equilibria and tipping points can emerge in a number of different ways.

Central to the discussion which follows is what I will call the *extended* version of the canonical theory of economic growth. The canonical model to which I referred in the Introduction is the Ramsey model of an optimizing representative household that engages in a dynamic process of capital accumulation and economic growth. In this 'extended' version of that model, two frictions constrain the growth process, in such a way as to make macroeconomic policy-making necessary. First, investment is undertaken by firms (not by the representative household as in the Ramsey model) which face costs in adjusting the capital stock; the level of investment by firms is determined by the size of Tobin's *Q*. Second, prices are slow to adjust, as in the new Keynesian model, so that the price signals coming from Tobin's *Q* are not—in general—sufficient to encourage firms to invest exactly what is required to ensure that resources are always fully employed. There may be deficient demand, or the opposite, putting price stability at risk. In response to this, the central bank follows a Taylor rule—raising the nominal interest rate when inflation is above target and vice versa. This extended canonical model is, in effect, the well-known DSGE model due to Christiano *et al*. (2005) and Smets and Wouters (2007).[26]

In long-run equilibrium, the level of output in this extended canonical model is determined by the size of the effective labour force.[27] Capital accumulation is driven by the representative consumer's decision to save. Financial intermediation ensures that investors invest what is saved by households, as in the basic Ramsey model.

However, shocks to the labour force and the level of technology can disturb such an equilibrium growth path, as can shocks to the desire to save. The canonical model can nevertheless be solved for the 'neutral' real interest rate after any such shock. It is the real interest rate which ensures that both capital and labour remain fully employed at all periods of time, so that there is no inflationary pressure, despite the adjustment that will becomes necessary after any shock.

[26] A straightforward description of the workings of that model can be found in Vines and Wills (2018, 2020) and Vines and Luk (2025).

[27] This is determined by the level of the actual labour force and the level of disembodied labour-saving technology, both of which grow at an exogenous rate.

Nevertheless, because the actual interest rate deviates from the neutral rate in the short run, output can be different from the level at which capital is fully employed, and labour is willingly supplied. As a result, there can be either inflation, or disinflation, which, in turn, establishes a role for monetary policy, which is assumed to follow a Taylor rule. This rule sets the *nominal* rate of interest so that the *real* interest rate is above its neutral rate when inflation is above its target, and vice versa, which is what causes inflation to converge to its target. The model employs a standard Calvo (1983) price-setting mechanism in which firms are able to change their prices in response to both demand pressure and expected future inflation, but can only do so sequentially, and gradually. The movement of the model away from long-run equilibrium will depend on both of the frictions added to the model. In the case, for example, of a cost-push shock to inflation, both the size of adjustment costs to capital and the degree of nominal rigidities will help determine how quickly the interest rate change brought about by the Taylor rule will get the economy back to its equilibrium growth path.

(i) Strategic complementarity between monetary policy-makers in the present and the future

Andrew Blake, Richard Dennis, Tanya Kirsanova, and Antony Yates show how multiple equilibria can arise in the extended canonical model described above, if there is only one amendment to the model, to do with the behaviour of the monetary policy-maker (Blake *et al.*, 2025, this issue, henceforth BDKY). Their analysis provides a further set of striking implications in relation to Keynes's ideas about animal spirits. It seems that beliefs about the nature of policy responses to inflation shocks may become self-fulfilling, in a way which I now describe.

In the BDKY paper, the authors' one amendment to the extended canonical model is as follows: instead of supposing that the monetary authorities follow a simple Taylor rule, they suppose—in a way which is very familiar[28]—that the central bank minimizes an intertemporal loss function. The components of this loss function are the discounted weighted sum of $(\pi—\pi^T)^2$ and y^2 into the future, where $(\pi—\pi^T)$ is the difference of inflation from its target and y is the output gap, which, of course, depends on the way in which the accumulation of capital changes the productive potential of the economy.

BDKY show that multiple equilibria can emerge if the policy-maker acts with discretion, re-optimizing policy in each period of time. They demonstrate that, as a result of this behaviour, there may be two very different ways in which the monetary authorities deal with an inflation shock. The existence of multiple discretionary equilibria of this kind opens the door to expectations-based coordination failures between current and future policy-makers. If policy is re-optimized in every period, then outcomes depend not only on fundamentals and preferences, but also on the beliefs of current policy-makers about the behaviour of future policy-makers.

The fact that there may be two different ways of dealing with an inflation shock can be explained as follows. In this set-up, in which monetary policy is made with discretion, a persistent cost-push shock will, naturally, give rise to the conventional disinflation process as one possible story: the central bank raises the interest rate, depressing output and exerting downward pressure on wages and prices. Inflation will fall gradually, at a speed which depends on the nominal rigidities. These rigidities exist because, in this extended canonical model, it is assumed that only a proportion of those who fix prices are able to alter them in each period of time. But as BDKY show, there may also be another, unconventional, kind of equilibrium in which, following the inflation shock, the central bank actually *lowers* the interest rate, stimulating investment. The higher level of investment leads to a higher capital stock and so to a higher level of potential output in the future and thus to a disinflationary environment. That is something which is so effective in bringing down inflation in the future that the policy-maker never has to raise the interest rate in order to bring inflation down. The result is an exaggerated stimulus to inflation in the initial

[28] Galí (2015) discusses the way in which a Taylor rule may be thought of as a short-hand representation of the behaviour of a policy-maker carrying out flexible inflation targeting in which output loss and inflation loss are traded off against each other. BDKY and others suggest that this popular way of thinking about things is itself a simplification of how a policy-maker might behave whose objectives are the maximization of the present discounted value of consumption.

period, but nevertheless an outcome in which inflation rapidly falls back towards target. At the same time the capital stock will gradually return to its initial level.

This second unconventional outcome is made possible by strategic complementarities between current and future policy-makers, which are transmitted to the economy through their effects on the expectations of the private sector and—through that channel—on the amount which the private sector decides to invest in capital. Of course, this second outcome clearly delivers a lower level of welfare than the first one. But nonetheless it is an equilibrium.

BDKY show that the economy may become stuck in a suboptimal outcome of this kind, in which responses to inflation shocks happen in this unsatisfactory manner. Although transitions between equilibria are theoretically possible, coordination on a different equilibrium may be practically difficult. In the BDKY model, policy-makers and the private sector may find it hard to learn how to coordinate on an alternative equilibrium once they have settled into one. As I have already noted in my discussion of the first of the papers by Hirano and Stiglitz, this equilibrium-selection issue is a challenging one. Unless supported by a process of learning in the private sector, or commitment on the part of the policy-maker, or some other kind of coordination mechanism, the economy may not be able to get itself away from being one in which responses to inflation are of this unconventional kind.

It seems that the private sector may, in some circumstances, be correct in believing that the economy's policy-makers may not deal with inflation shocks in the normal way by raising the interest rate. This is a worrying outcome.

(ii) Inferential expectations and sudden belief changes

I have just shown that the model presented by BDKY shows how multiple equilibria can emerge because of the way in which the *policy-maker* behaves. Timo Henckel, Gordon Menzies, and Daniel Zizzo put forward a different story about the interactions between the policy-maker and the private sector to do with the behaviour of the *private sector*. This idea involves a theory of abrupt shifts in expectations of inflation (Henckel *et al*., 2025, this issue, henceforth HMZ). As I note below, this model could also be used to explore the origins of multiple equilibria, although the authors do not use their model in this way.

HMZ formalize an idea whose origins can be found in a very early paper by Bob Rowthorn (Rowthorn, 1977). Rowthorn was writing a decade after Milton Friedman presented the natural-rate hypothesis and the vertical Phillips curve. In response, Rowthorn suggested that as inflation gradually began to build in the 1960s, people in the private sector failed at first to notice this fact in the way that Friedman had claimed they would. As a result, Rowthorn said that they would set their wages and prices in a way consistent with an upward-sloping Phillips curve. This behaviour would continue, he said, until people suddenly noticed the rise in inflation and took action, causing the Phillips curve to become vertical—even if not immediately, then very soon.[29] HMZ provide a microeconomic foundation for this kind of behaviour, arguing that at least some actors behave in this way because they conduct a hypothesis test on the expected rate of inflation. The authors maintain that such people hold firmly to their initial beliefs until the weight of contradictory evidence reaches a critical threshold, prompting a sudden and seemingly disproportionate revision of their expectations.

HMZ use these ideas to build and simulate a simple 3-equation macro model to provide qualitative insights into the recent surge in inflation and the policy dilemma which it presented. Their model—except for one detail—is an entirely standard kind of simple New Keynesian model which is linear in form. We can think of their model as a simplified version of the extended canonical model described at the beginning of section IV: simplified because it omits any analysis of longer-term processes of capital accumulation and growth and focuses entirely on the need—in the presence of nominal rigidities—to control the rate of inflation. The model's components are: an IS curve, which makes the output gap depend on the real interest rate; a Taylor rule which, in effect, makes the real interest rate depend on the inflation rate; and a Phillips curve which makes the inflation rate depend on both the output gap and the expected rate of inflation. The expected rate of inflation in the model is partly backward-looking and partly forward-looking.

[29] See also Rowthorn (2024) and Michl and Rowthorn (2024).

That much is familiar and standard. If those who have forward-looking expectations of inflation form them in a manner which is rational, and model-consistent, then it is well known that such a linear model will produce a unique response to any shock, essentially because the model is linear as well as simple (see Galí, 2015). For example, a positive inflation shock will cause the central bank to raise the interest rate, leading to a gradual return of inflation to its target level (Galí, 2015). And the more people there are who are forward-looking in this way, the more quickly will any inflation shock be dealt with.

The non-standard element of the model which HMZ use is this. Those who are forward-looking are supposed to have beliefs which are anchored on the target rate of inflation—which is low. This means that they will not initially move their inflation expectations upward after any positive shock to inflation. HMZ suppose that these people will only change their expectations of inflation after a hypothesis test convinces them that the actual inflation rate has veered too far off track for too long. (The hypothesis test formalizes the idea of 'too far' and 'too long'.) In these circumstances, HMZ suppose that the inflation expectations of these people suddenly switch, when they become convinced, to being rational, and model-consistent.

Given this set-up, HMZ use their model to explain why the recent inflation shock—largely caused by the war in Ukraine—was initially slow to take hold, but then led to a very rapid increase in both expected inflation and actual inflation. HMZ are also able to suggest something similar about the way in which the inflation rate was slow to fall back in many countries after the Ukraine-war shock. They argue that this subsequent outcome might have been observed because those who are forward-looking came to believe that the inflation rate would remain at its high post-shock level, until they had sufficient evidence that this was not the case.

HMZ go on, in the last part of their paper, to talk about circumstances which might give rise to a multiple-equilibrium situation for a very different sort of reason than that put forward by BKD. Suppose that, after an inflation shock has happened, forward-looking people remain fixed on the idea that inflation will remain high. And suppose that the central bank does indeed follow a Taylor rule, but that it is free to choose the size of the coefficient on inflation in the Taylor rule, i.e. whether it is large or small or in between. The central bank could decide to act 'firmly' and bring inflation down quickly—with the benefits which that will bring—by operating according to a Taylor rule with a large coefficient on inflation, causing a large output loss, and inflation reduction, in the short-to-medium term, so that the forward-looking people change their minds quite quickly—using their hypothesis test—and so come to the view that inflation will fall back to its target level. It may seem worthwhile for the policy-maker to act in this way because, after a lag, inflation comes down fast. Alternatively, the central bank could decide to act in a 'timid' manner, causing a small output loss, both initially and subsequently. As a result, inflation will come down only slowly, and those who hold forward-looking inflation expectations will take a long time to form the view—using their hypothesis test—that inflation has come down enough for them to, once again, expect inflation to be back on target. Nevertheless it may seem worthwhile to the central bank for it to act in this way because the output losses which it causes will always be small. However, the model which HMZ deploy suggests that it may not be desirable for the central bank to follow a policy which is somewhere in between: action that is firm enough to lead to quite large output losses initially, but not firm enough to cause the forward-looking part of the private sector to change its mind quickly, meaning that the process of inflation reduction is slow, but nevertheless still costly. This suggests that if the central bank is free to choose which kind of path to follow, the political economy of the inflation-reduction process may well have two kinds of equilibria—the fast one and the slow one—but not an equilibrium which lies somewhere in between. As I have already said, HMZ do not explore this possibility, but it seems like an interesting idea to investigate. It is the non-linearity introduced by their hypothesis test which seems to make this idea worth looking at.

(iii) Credibility

The next paper in Part II continues the discussion of how multiple equilibria can arise in the short term; the authors of this paper argue that different self-sustaining equilibria can arise, depending on whether the central bank is credible—or not. The paper was written in response to

Ben Bernanke's review of monetary policy making at the Bank of England (Bernanke, 2024), and the authors—Douglas Laxton, Haykaz Igityan, and Shalva Mkhatrishvili—focus on the way in which the private sector's expectations of inflation depend on central bank credibility. They build a model of how this credibility can be won, or lost, leading to a multiple-equilibrium outcome (Laxton *et al*., 2025, this issue, henceforth LIM). They then use this model to compare the credibility of the European Central Bank (the ECB) and the US Federal Reserve (the Fed), and argue that the lower credibility of the ECB led to poorer economic outcomes in the euro area. They go on to claim that, if the Bank of England were to follow their proposals for strengthening credibility, the economic performance of the UK might well be much improved. That is what makes their paper an important—and constructive—response to the Bernanke Review.

The starting point of LIM's analysis is a rather simple, and conventional, open-economy model, of a linear kind, which LIM estimate empirically in a two-country version for the US and the EU. They present this model in section III of their paper.[30] Like the closed-economy model used by HMZ which I discussed above, the model deployed by LIM contains, for each country, the three standard equations for a short-run macro-model: an IS curve that determines the level of output in a way depending on the real interest rate; an expectations-augmented Phillips curve that determines the rate of inflation in a partly forward-looking manner; and a form of Taylor rule which depicts how the nominal interest rate responds to the inflation rate. The open-economy features which they add to this model are of a Dornbusch-model kind. International capital movements ensure that—given the expected future exchange rate—the spot exchange rate finds a level which is consistent with uncovered interest parity (UIP), at least in a modified form. The exchange rate which emerges in this way has an influence on both the IS curve and the Phillips curve. It is clear from my earlier discussion that such a linear model will produce a unique response to any shock. For example, a positive inflation shock will cause the central bank to raise the interest rate, and a gradual return of inflation to its target level, partly because there will be an appreciation of the exchange rate. A negative demand shock will lead to a lower interest rate, a depreciated real exchange rate, and a gradual return of output to its non-inflationary level.

The strategy of LIM's research has been to augment this simple linear model with a non-linear sub-model of how credibility evolves, following shocks to the economy like those which I have just mentioned. This work is then designed to show how the resulting level of credibility affects both inflation expectations in the Phillips curve and the expected exchange rate in the exchange-rate equation. LIM then demonstrate that a loss of credibility can lead to outcomes in these equations which result in a worsening of macroeconomic performance. And they argue that this worsening of performance can be a self-fulfilling multiple-equilibrium outcome in which credibility, once lost, is hard to regain.

The determination of credibility in their sub-model is, in fact, not discussed until towards the end of their paper—in section IV(iv). But it is useful to explain how this sub-model works at the beginning of any attempt to understand what the authors are doing.

LIM calculate credibility, and its effects, in the following way.

LIM take inflation expectations to be a mixture of a backward-looking component and a forward-looking component; the latter is—unlike in the paper by HMZ—assumed to be always rational and model consistent. But the weight on the forward-looking component is not constant; it is denoted by a variable called *Cred,* which is used as a measure of credibility, and which lies between zero and one. The idea is that, the more credible the policy-maker, the larger will be the weight, *Cred*, which is placed on the forward-looking model-consistent determinant of inflation expectations, and so the larger will be the weight which the private sector places on what the overall macro-model predicts that inflation outcomes will be, when forming their expectations of the inflation rate. And vice versa. In order to calculate *Cred*, the model first computes, in equations 18 and 19, using actual past data, the proportion of time in which actual inflation has turned out to be either well above, or well below, the level which the model has forecast.[31] These proportions are then subtracted from unity in equation 20 to obtain the *Cred* variable. We can thus interpret *Cred* as saying that the credibility of the policy-maker will be close to unity if the previous pre-

[30] Rather confusingly they call their model 'canonical', using this word in a very different way from that in which I have been using it.
[31] The details of these equations need not concern us here.

dictions of inflation, using their model, have been more or less on track for much of the time, and vice versa. This *Cred* variable is then used—in equation 17—to determine the extent to which inflation expectations are forward-looking, and model consistent, in the Phillips curve, or simply backward-looking, in which case they follow a gradual partial-adjustment process.

The chain of reasoning about multiple equilibria then works as follows. The extent to which the model would have been able to forecast the actual inflation outcome in the recent past determines the *Cred* variable. Good forecasting performance will cause the *Cred* variable to be close to unity causing inflation expectations to be primarily forward-looking, and that—so the argument goes—will improve the forecasting performance of the model going forward. And vice versa: poor forecasting performance in the past will cause the *Cred* variable to be well below unity causing inflation expectations to be largely backward-looking, and that—it is said—will degrade the forecasting performance of the model going into the future. That is to say, a model which has forecasted inflation well in the past will compute a high value for *Cred*, so that, going forward, expectations will be close to what that model predicted, making it likely that the Phillips curve will predict inflation well into the future, which will keep the *Cred* variable at a high level. And vice versa.

Such a set-up suggests that large, uncontrolled movements in the rate of inflation are in danger of damaging credibility. One possibility of this kind, which LIM discuss in detail in sections IV(iv) and IV(v) of their paper, can be caused by a zero lower bound to interest rates. That lower bound of course prevents the Taylor rule from operating—hence introducing a non-linearity into the model. This means that, when the zero lower bound is reached, deflation can set in, causing inflation to remain below target, perhaps for some period of time. Such deflation will interact with the credibility-propagation mechanism which I have just described, leading to a loss of credibility. And that credibility will be hard to revive.

The authors use Figure 8 to suggest that something like their credibility-propagation mechanism was recently at work in the US and the EU. In the figure they depict, for the period from January 2018 to the early 2020s for each of these regions, the actual data for inflation expectations, and beneath this their calculated value for credibility. Commenting on their figure, the authors state,

> The main takeaway is that inflation was better anchored in the US heading into the pandemic where Fed credibility was near perfect. Meanwhile, the euro area was still mired in a low inflation regime with inflation expectations persistently below the ECB target and therefore entered the pandemic with worse credibility.

LIM then depict, in section V(v) of their paper, how this credibility-propagation mechanism works within their full macroeconomic model. They do this by simulating the effects of the onset of the Covid pandemic in both the US and the euro area, starting in the first quarter of 2021. The results are depicted in Figure 10, which is the culmination of the analytical work presented in their paper.

The starting position for their simulation is the one which LIM had previously depicted in Figure 8—namely that the US entered the pandemic period with a much higher level of credibility. The simulation shows the onset of the pandemic causing both output and inflation to plummet in both regions. In the US the simulation supposes that the Fed quickly cuts rates to reach the lower bound, while in Europe the interest rate was already at that lower bound. What happens next in the simulation depends on the difference in credibility between the two regions. As a result of the downturn the Fed still spends a long time with interest rates at the lower bound, but not as long as the ECB. There are two reasons for this. In the case of the US, the stimulus coming from the initial decline in the policy rate—and the extended period at the lower bound, combined with the higher level of credibility—is compelling enough for output to recover and for financial markets to expect inflation to return to levels that are more consistent with the Fed's 2% target. That is to say, in the US, inflation is expected to quickly return to target, precisely because of the higher level of credibility. The euro area enters the pandemic with interest rates at the zero lower bound already, since the ECB had previously allowed inflation to drift downwards, and, as a result, with a lower level of credibility. So there is, when the pandemic hits, no interest rate cut which the EU authorities can carry out. Furthermore, as the authors say: 'in the case of the euro area, no one believes the ECB is potent enough to respond to the shock and so inflation and inflation expectations drift even lower from a lack of credibility.'

In addition, there is an extra twist to this story, due to the effect of credibility on the expected future exchange rate. I mentioned the role of the expected future exchange rate on the current exchange rate when setting out LIM's original simple model. But I have not yet discussed the effects that this can have. LIM continue on from the quotation above with the following statement:

> This divergence in inflation outlook widens the expected price level between the US and the euro area and the EURUSD exchange rate appreciates, reinforcing the deflationary forces and amplifies the negative demand shock and needlessly prolongs the recovery in the euro area.

This idea might not immediately be familiar—it was not familiar to me. I say this because interest rates in the UK and the US are both projected to be at the zero lower bound for some period of time in the simulation and so—recalling that the spot exchange rate is determined in order to make UIP hold—one might not have thought that the simulation would suggest any change in the nominal exchange rate. But the fact that inflation in the EU area is projected to remain below target for many time periods means—as LIM say in the quotation above—that the price *level* in the EU area is projected to fall relative to that in the US, in a way which is not expected to be reversed. This will cast a shadow over the spot exchange rate causing it to *appreciate*. As LIM say, this will further dampen the recovery that the simulation predicts for the EU area.

Sections V and VI conclude the paper by turning to the implications of their analysis for policy-making at the Bank of England, and at other central banks. Here I focus on just two of the many valuable points which they make.

LIM argue that risk management should play a part in choosing a strategy for macroeconomic policy-making. One reason for their concern about risk stems from another possibility which I have not yet mentioned: the idea that the Phillips curve might be non-linear. In section IV(i) of their paper LIM discuss the possibility that, once output passes a certain level, inflation might begin to rise very rapidly. Such a situation might mean that, in the presence of a normal kind of Taylor rule, inflation would rise well above target in response to any unexpectedly large levels of output in the economy. Their analysis suggests that this would lead to a loss of credibility, making the economy harder to control in future periods. As a consequence, they suggest in section IV(iii) that any such simple Taylor rule should be replaced with a monetary policy that is based on the minimization of a loss function which heavily penalizes high rates of inflation. This would ensure, they say, that the response to any such outcome would be firm and rapid. It would also ensure that it became widely known that the response would, in fact, be firm and rapid, enabling credibility to be preserved.

This particular example leads on to the second, much more general, proposal made by LIM, which has two parts. First, LIM agree with Ben Bernanke that central banks need to adopt a scenario-based approach to monetary policy. But second, they believe there is a need for policy-makers to make clear the way in which their policy would respond to whatever scenario is being portrayed. The authors argue that monetary authorities should publicly display what their interest rate response would be to each and every scenario. They suggest that this is important precisely because of the analysis of credibility which they have undertaken.

Thus, taking the example of an inflation-surge scenario, LIM argue that the credibility of the central bank's inflation target requires the bank to make clear the way in which interest-rate increases would be deployed to deal with the challenge which such a surge would present. They do not sidestep either the problems which any monetary policy committee would encounter in reaching agreement as to what such a hypothetical policy-stance might look like, or the difficulty there might be in communicating this information to the general public. Nevertheless, they make clear what the risk would be if this potential policy response is not made explicit. It is clear to them that the presentation of any scenario which does not make the policy response explicit would run the risk of undermining credibility. Their modelling of the credibility-propagation mechanism supports that claim.

This key proposal—built as it is upon a formal model of credibility—is what makes the LIM paper a powerful response to the Bernanke Review of policy-making at the Bank of England.

(iv) Structural change caused by contemporary changes in policy frameworks

The final paper in Part II of the *OxREP* issue is by Andrew Hodge, Zoltan Jakab, Jesper Lindé, and Vina Nguyen. In this paper the authors assess the extent to which reforms to the operational aspects of the monetary frameworks deployed by the Fed and the ECB contributed to the surge in core inflation that followed the Covid pandemic in 2021 (Hodge *et al*., 2025, this issue, henceforth HJLN). As I describe below, they then go on to examine the welfare cost of what was done.

These structural reforms to the conduct of monetary policy (MP) involved a form of institutional change that was independent of the Covid pandemic, and thus relate to something that is directly relevant to the theme of this *OxREP* issue.[32] They included a move from flexible inflation targeting to flexible *average* inflation targeting in the US (FAIT) and from an asymmetric target of 'below, but close to 2%' to *symmetric* flexible inflation targeting (FIT) in the euro area (my italics). This is in contrast to the fiscal policy (FP) events which occurred at the same time, designed to provide support to aggregate demand in the face of Covid-19, in particular the huge fiscal stimulus which occurred in the US. The challenge which the authors faced was that of finding a plausible decomposition of cause and effect, since the fiscal stimulus was implemented at the same time as the change to the monetary policy regimes.

To meet this challenge, HJLN employ a two-country New Keynesian DSGE model of the US and the euro area to assess how much of the 2021 inflation surge in these two places can be explained by the changes in monetary policy frameworks and by the fiscal stimuli. The model which HJLN use is carefully set out and explained in Appendix A of their paper. Although the model might seem highly elaborate, it is moderately easy to understand since it is essentially a two-country version of the extended canonical model described at the beginning of section IV. That is to say, the model for each country is a Ramsey growth model to which have been added nominal rigidities and adjustment costs of capital, in a way which makes its framework overlap with that of New Keynesian models. The model is therefore more complex, and finely calibrated, than the two-country model of the US and the EU area used by LIM, that I discussed above. That model abstracted from the separate modelling of consumption and investment, from the micro-founded modelling of production and consumption decisions, and from the detailed modelling of trade between the two countries, all of which are carefully dealt with here. It is also a TANK model (i.e. a Two-Agent New Keynesian model) in which households are divided into two groups, namely those which are liquidity-constrained and those which are not. This last development is essential if the model is to be able to discuss the effects of fiscal stimulus without becoming bogged down by an assumption of Ricardian equivalence. The HJLN model also builds on such a framework in that it contains a financial accelerator. In fact the model used here is—in some ways—a two-country, single-good version of the many-region many-sector model operated by Liu, McKibbin, Shuetrim, and Wilcoxen which I describe in section V below.

The authors' task is to use this model to explain how much of the surge in inflation in the US and the euro area can be explained by the changes in monetary policy frameworks and fiscal stimuli which occurred. They demonstrate that the explanation depends on how monetary and fiscal policies interact both within and across countries, which is why the two-country two-agent structure of their model is important. They also demonstrate that the explanation depends on whether the slope of the Phillips curve is thought to have shifted in recent years. Model simulations using the initial calibration of the Phillips curve (i.e. using the pre-pandemic calibration of its slope) suggest that the monetary-policy reforms and the fiscal stimulus together explain rather little of the surge in inflation. But if the slope of the Phillips curve is assumed to have been steeper—which a calibration using pandemic-period data suggests is appropriate—then the combined effect of the two sets of policies (i.e. MP under a new regime and fiscal policy stimulus) explain about two-thirds and nearly half of the increase in core inflation in US and EU, respectively, during 2021.

[32] They involved a change in policy strategy rather than anything to do with a tipping point.

Furthermore, in Panel B of Figure 4 the authors suggest that fiscal stimulus had a greater impact on inflation than the change in the monetary-policy regime, in both the US and the euro area. In particular, the model simulations indicate that the introduction of FAIT in the US and FIT in the euro area can together explain around 0.7 percentage points of the difference between US core inflation and the pre-2021 IMF staff forecasts. For the euro area, the amount of unexpected core inflation in 2021 explained by the new monetary regimes (both in the US and in the euro area taken together) is smaller, at around 0.2 percentage points. By contrast, the US fiscal stimulus can explain almost 1 percentage point of the unexpected core inflation in 2021 domestically and 0.3 percentage points in the euro area.[33]

The measured welfare effects are large, particularly highlighting the adverse impacts on the US, regardless of whether a flat or steep Phillips curve is applied. These results are driven by the central bank loss function, which is assumed to depend on the output gap and the inflation rate equally, in line with the dual mandate given to the Fed which requires it to care about both inflation and employment. Nevertheless, the authors note that even if they were to set the weight on the output gap to zero, their calculations still report a significant economic loss.

Of course it would have been helpful to examine the outcomes over a longer period of time. However the baseline forecast—with which the effects of the change in monetary frameworks and the fiscal stimulus were compared—needed to be taken from the IMF staff projections that were made *before* the changes in monetary policy and fiscal policy had had their influence. And, in any case, from February 2022 any such exercise would have been tangled up with the effects of the war in Ukraine.

The HJLN paper has an important implication. Conditional on a prior belief that the slope of the Phillips Curve was invariant, then the reforms to the monetary framework were entirely coherent and were not going to dissipate into a higher rate of inflation. What policy-makers seem to have got wrong was believing that the flat Phillips Curve would persist. This paper highlights the importance of designing policies that are robust to potential changes in the parameters of the model that is being used. This is a strong conclusion, which mirrors that reached by LIM in their paper, which I have discussed above, and which is taken up in the paper by Zenghelis *et al.* which I discuss in section VII below.

V. Models of structural change in the global economy

Part III of this *OxREP* issue contains two papers which are, in effect, the culmination of the discussion of how modelling structural change requires us to go beyond simple one-sector models. These two papers take the analysis beyond the level of the individual economy to the global level, analysing models of the world economic system. The first of these papers is by Seth Benzell, Laurence Kotlikoff, Maria Kazakova, Guillermo Lagarda, Kristina Nesterova, Victor Yifan Ye, and Andrey Zubarev, who use the Global Gaidar Model. The second is by Weifeng Larry Liu, Warwick McKibbin, Geoffrey Shuetrim, and Peter Wilcoxen, who present simulations using the G-Cubed model. In global models such as these it is not just that some sectors grow faster than others, but also that some countries grow faster than others. These two papers seek to understand how that might happen. Many of the ideas about modelling structural change which I have discussed earlier play an essential role in the modelling of structural change at the global level that is presented in these two papers.

The comparison of the results from these two different models raises an important research question. The findings on structural change coming from simulations of the Global Gaidar model—that used by Benzell *et al.*—are much easier to understand. But the G-Cubed model simulated by Liu *et al.* is able to present results that are much more detailed. How much detail is necessary? What detail cannot be omitted?

[33] Notice from the figures that the combined impact of monetary and fiscal policy changes falls below the sum of the individual components. That is because the inflation triggered by the fiscal policy responses causes the monetary policy stance to be that much tighter than it would have been if there has simply been a move in each country to a new monetary-policy framework.

(i) Projecting the future of the world economy to 2100

The first of these two papers aims to describe how the relative position of countries in the world economy will have changed radically by the time this century ends (Benzell *et al*., 2025, this issue, henceforth BKKLNYZ). Which regions will come to dominate the world economy in 2100? Will high saving rates in fast-growing regions lead to a global capital glut? And, does ageing indicate a need for far higher tax rates in particular regions? The authors attempt to answer these three questions—and some others as well—by carrying out a number of simulations using the Global Gaidar Model.

Their findings are striking. Their baseline projection suggests that by 2100 China and India will become the world's two largest economies, with, respectively, 27.0% and 16.2% of world GDP. At this time, the shares in world output of the US and Western Europe (including the UK) are projected to be 12.3% and 11.9%. By contrast, in 2017—the year in which the model simulations begin—the US and China both accounted for roughly 16% of world GDP, lagging behind Western Europe (plus the UK) which accounted for almost 20%. At that time, India produced only 7.0% of world output. The baseline projections of the authors' model also suggest that by 2100 there will be a significant global savings glut, major reductions in the world interest rate, and large increases in ageing-related taxation.

BKKLNYZ are to be congratulated for producing a modelling strategy which enables them to analyse their questions in a transparent manner. They model the world as an overlapping set of regional models, 17 in all. For each of these regions, their model contains an amended version of the canonical Ramsey model, in which the assumption of a single infinitely-lived representative-agent set-up is replaced by the assumption of an overlapping-generations (OLG) set-up. As Figure 2 makes clear, that OLG set-up is one in which there are many generations of individuals living in each region of the world, each of which is carrying out intertemporal optimizing decisions about its savings, consumption, and employment behaviour, decisions that are carefully calibrated in the way that BKKLNYZ describe.[34] But since there are assumed to be no intentional bequests in the model, these decisions made by each generation only concern what will happen in their own lifetime, rather than having regard to events which will stretch far out into the future, as in the Ramsey model. There is perfect mobility of capital across all regions, so that there is only one global market for capital. As a result, labour is fully employed in all regions at all times and capital is accumulated everywhere until its expected return is exactly equal to the world rate of interest.

The models for each region have the following crucial features: the OLG structure for each region has a region-specific age/skill composition; there is only one kind of good produced anywhere in the world[35] so that trade balances reflect gaps between regions' and countries' optimal consumption and production decisions (but sum to zero); and government budgets are balanced by means of lump sum taxes.

The model is solved by BKKLNYZ for a baseline growth-path along which projections are made for five exogenous variables: population growth, population age structure, productivity catch-up growth relative to the US, fiscal adjustment, and automation. The task which the authors set themselves is to explain which of these five things is most important in explaining their striking findings about relative size listed above.

There are two essential parts to an explanation of their findings.

First, the authors find that region-specific productivity catch-up is the main determinant of future economic size—what they call global economic power—with potential demographic change representing the second biggest factor. Automation and fiscal policy all play second fiddle to these forces. 'Short of a near-term, sustained, and substantial drop in China's and India's catch-up rates.. this is Asia's economic century with China or India rapidly becoming the world's economic kingpin and the other taking second place.'

[34] By way of comparison, in the OLG version of the Ramsey model used by Hirano and Stiglitz in their one-sector growth model, which I disused above at the beginning of Section III, each agent simply lived for two periods. In the first period the agent worked and in the second period he or she was retired. The begin examined here model is much more complex than that, for reasons which include the fact that the model makes allowance for the existence of children.

[35] There is, however, a separate treatment of energy.

The reason underpinning this finding becomes obvious from Table 2.1 as soon as one understands the way in which the authors' OLG version of the Ramsey model actually works. In 2017, GDP per capita (a rough measure of productivity) for India was only one-tenth of that in the US, and for China it was only one-quarter. Thus, although one might initially think that demographic changes in the world are enormous and so are likely to be hugely important, they are not big enough. Although China's population is likely to fall by 400 million, the projected change in the size of the Chinese population relative to that in the US is much smaller than the projected change in China's productivity relative to that in the US, since China's population will still be roughly two-thirds of what it was initially. The same calculations are also true for India. The only region for which this is not true is Sub-Saharan Africa which begins from a very small base. These quantitative results clearly depend in part on the OLG/Ramsey structure of the model, since the model shows that an accumulation of capital in each region will accompany the effects of the exogenous changes in population and productivity.

Second, the authors state that:

> [the] general equilibrium transition path [which they present] entails major global capital deepening. This reflects the global ageing process combined with the growing relative economic importance of China, India, and other high-saving-rate regions. Capital deepening spells, in turn, a decline in the world interest rate and a rise in global wages. Indeed, in general equilibrium, the world interest rate falls from 5.98 per cent in 2017 to 1.18 per cent in 2100.

The reason underlying this finding is also very clear. On the one hand, savings will depend on age structure. On the other hand, investment in each of the regions will depend on returns to capital in that region so that capital will flow to regions where productivity is improving (creating balance of payments deficits and accumulation of asset imbalances on which interest must be paid). The extent to which the world interest rate falls will depend on the balance between, on the one hand, the tendency for savings to be high in some regions and, on other hand, the speed at which technology is transferred to the world's poorer countries. In other words, given any particular savings behaviour, a more rapid rate of technology transfer would tend to increase the demand for capital, stimulate the level of investment and raise the global interest rate.

It is worth noting the following sobering observation which the authors make.

> [O]ur model provides a warning to regions, including the US and China, that will come under major economic stress due to the interaction of deficit/pay-as-you-go-financed fiscal policies and population ageing. Unless these regions can limit fiscal outlays, they face dramatically higher future average and marginal tax rates. This could limit labour supply, induce tax avoidance and evasion, and promote emigration, particularly of high-skilled workers.

This conclusion obviously follows from a combination of factors: on the one hand, demographic assumptions along with, on the other hand, the assumption that governments cannot run deficits. A rising proportion of the aged population will clearly induce fiscal stress in these circumstances. These quantitative insights evidently depend on the Ramsey-model structure of the Global Gaidar Model. That is because the future ratio of pension payments to future output depends on the level of output which will emerge in the future. And that will, in turn, depend on the extent to which capital has been accumulated.

(ii) Modelling structural change at both the national level and the sectoral level

The second of the two papers in Part III, by Weifeng Larry Liu, Warwick McKibbin, Geoffrey Shuetrim, and Peter Wilcoxen, examines three scenarios very similar to those considered by BKKLNYZ. These three scenarios show the effects of demographic change, the technological catch-up of developing countries towards the global productivity frontier, and a process of global productivity growth in manufacturing driven by automation technology (Liu *et al*., 2025, this issue, henceforth LMSW).

LMSW present results for the effects of exogenous changes between now and 2050, rather than for the rest of the century as BKKLNYZ do.[36] They use the G-Cubed model, which, like the

[36] In fact they simulate the model up to 2150, but only present results running up to 2050 in this paper.

Global Gaidar Model is a multi-region model with 18 regions, but which, in addition, separately identifies six different sectors in each economy. As a result, the model captures not only the direct channel of different productivity growth across sectors, but also the indirect channels through sectoral and international linkages. The analysis is able to highlight how asymmetric productivity growth across sectors, as well as countries, can reshape global economic structure over time, with important implications for world economic growth, international trade, and the distribution of global economic power.

At first sight this might seem like an extremely complex modelling exercise.[37] Nevertheless, what is going on is also reasonably straightforward.[38] At the core of the model for each region is, in effect, an overlapping generations (OLG) version of the extended canonical model described at the beginning of section IV above, i.e. an OLG version of the Ramsey growth model to which both nominal rigidities and adjustment costs of capital have been added. Here, I explain how the G-Cubed model deploys the OLG version of the extended canonical model in a multi-regional setting.

Capital accumulation in each region is driven by the decision to invest by the firms that exist in each sector in that region. But because this is an extended canonical model, what is invested is not determined by what is saved. Instead, there is an explicit investment function for each sector. That investment decision is forward-looking, depending on the expected future need for capital, and the extent of investment is governed by the costs of adjustment of capital. The long-run equilibrium growth path for this model is a Ramsey equilibrium growth path in which the capital accumulation to which these investment functions give rise exactly keeps pace with population growth and technical progress.[39]

In the G-Cubed model, OLG features are introduced by adopting the idea of perpetual youth in which, in each period, individuals have a given probability of death. The representative consumer in each region is thus based on a forward-looking Euler equation[40] but with a probability of death included, as in Yaari (1965), Blanchard (1985), and Weil (1989). This means that each individual will discount the future at a rate higher than his or her rate of time preference. A consequence is that the long-run equilibrium value for the interest rate in each region must be higher than the rate of time preference of the representative individual in that region. The existence of non-Ricardian equivalence caused by the positive probability of death means that the long-run real interest rate can deviate from the sum of the pure rate of time preference plus the long-run growth of productivity. There are also risk premia incorporated into the model, the effect of which is to further raise the interest rate.

In the short run, shocks to the level of technology, and to its expected rate of change, disturb this growth path, while disturbance is also caused by shocks to the desire to save, and to the financial intermediation process. In principle, there exists an endogenous 'neutral' real rate of interest in each region which would—despite such shocks—cause resources to remain fully employed in each period.

However, because in each region the actual interest rate deviates from the neutral rate in the short run, output in the region can be different from the level at which capital is fully and labour is willingly supplied. As a result, there can be excess demand for labour or unemployment, and so either inflation or disinflation. The model employs a Calvo-like inflation mechanism, but specifically for wage setting. As a result, economy-wide wages respond to both demand pressure and expected future inflation in the manner of Calvo (1983), but do so sequentially, and only gradu-

[37] The six sectors in the version of the model used here are: energy, mining, agriculture, manufacturing durables, manufacturing non-durables, and services. The first three sectors are important for their links to resource endowments that vary widely across the globe. Separating manufacturing into durables and non-durables is important because interest rate changes have differential effects. Also, there is a capital-goods-producing sector in each region but the output of this sector is not traded. The 18 countries/regions in the version of the model used there are identified in the paper. More aggregation of sectors and regions to make the model smaller is possible, but the authors believe that the grouping used here is the most useful if one seeks to understand the effects of the scenarios reported in their paper.

[38] The core reference for the G-Cubed model is McKibbin and Wilcoxen (2013). See also https://documentation.gcubed.com/. The main features of the model are discussed in detail in McKibbin and Vines (2000).

[39] In the short run there is a distinction between the optimizing firms and rule-of-thumb firms in each sector of each economy. In the long run both kinds of firms follow the same decision rules.

[40] In the short run there is also a distinction between the optimizing consumers and rule-of-thumb consumers. In the long run both kinds of consumers follow the same decision rules.

ally. Firms, in turn, set prices in a way which responds to both wage levels and demand pressures. The existence of such inflationary pressures establishes a need for monetary policy, which is assumed to set the interest according to what the authors call a Henderson–McKibbin–Taylor rule (see Henderson and McKibbin (1993) and Taylor (1993)). This rule means that the real interest rate in each region is above its natural rate when inflation in that region is above its target, and vice versa. Fiscal policy can stabilize demand, but over time government deficits lead to public debts which, to ensure fiscal solvency, lead to higher levels of taxes to pay the higher debt interest, but only gradually.

Being global, the G-Cubed model needs to capture the effects of international trade on output in each of the six sectors in each region: the model joins all regions in the world together through international trade in the goods and services produced by each of the sectors. Trade flows are determined by carefully modelled export and import functions, and these map consistently into the equations for imports and exports in other regions.

Countries are also joined together through asset-market arbitrage. The model supposes perfect international mobility of capital between regions, and the exchange rate is determined, *à la* Dornbusch, by the uncovered interest parity (UIP) condition, except for regions with pegged exchange rates and for those countries within the European Monetary Union. But there is explicit allowance for risk premia in these UIP equations. Because inflation pressures differ between regions there are large real interest rate differences between countries which, in turn, drive real exchange rates in important ways over time. Changes in real exchange rates in each region have significant influences on trade flows between the regions, and also on the levels of investment in each of the sectors in each of the regions.

The results of the authors' investigations are set out explicitly in their paper and are just as striking as those produced by BKKLNYZ.

In the productivity catch-up scenario, which relates most closely to the analysis by BKKLNYZ, the initial productivity gaps and the catch-up rates can differ across sectors, resulting in different productivity growth rates. As a result, the model captures not only the direct channel of different productivity growth across sectors, but also the indirect channels through sectoral and international linkages. The simulation of demographic change shows that demographic asymmetries between regions are large enough to mean that they are likely to reshape the distribution of global economic power, even between now and 2050, let alone between now and 2100. This demographic change is also likely to significantly change economic structure, even in the next 25 years, because of different factor intensities across sectors. Finally, the simulation of the G-Cubed model concerning manufacturing productivity growth due to automation shows that sector-specific differences in productivity are likely to lead directly to structural change, which is also propagated indirectly through input–output linkages. Also, as the manufacturing share varies across economies, even a uniform growth rate in manufacturing across regions is not likely to result in identical structural adjustments, but is instead likely to generate asymmetric changes in manufacturing sectors across regions.

The two papers which I have just described in effect enable a model-comparison exercise in relation to the modelling of global structural change. This exercise leads to an overall conclusion with four parts.

- (i) It seems important to make use of some kind of OLG version of the canonical Ramsey model when modelling structural change at the global level. That is because such structural change depends in a fundamental way on the effects of capital accumulation and technological innovation which a Ramsey-type model enables one to portray. Both the Global Gaidar Model and the G-Cubed model are models of this kind.
- (ii) Using a one-commodity model to carry out an analysis of structural change at the global level enables one to produce some important results. Use of the Global Gaidar Model has made this clear.
- (iii) Sectoral details do seem to matter. Thus confining one's modelling of global structural change to the use of a one-commodity model causes one to miss some things that obviously matter.

(iv) Because of this omission, it does seem that a proper analysis of structural change at the global level requires one to make use of an OLG version of the extended canonical Ramsey model in which both adjustment costs of capital and nominal rigidities are added to the basic Ramsey model. The G-Cubed model is a model of this kind. The reason for the need for a thorough analysis is as follows. One needs adjustment costs of capital in the model because, to investigate how sectoral change is propagated sector-by-sector, one needs equations describing the gradual response of investment in each sector to changes in productivity in that sector, and to changes in the demand for that sector's output. These are gradual responses: the existence of adjustment costs of capital makes the redeployment of capital across sectors take time. The pictures of simulation results displayed by LMSW makes this very clear. One also needs nominal rigidities in the model, because when wages are set in the Calvo-like manner that I have described, this will have important implications, of two kinds, for the outcomes in each sector, both of which are apparent in the pictures of simulation results displayed by LMSW. First, the resulting wage costs will help to determine the relative profitability of each sector and so influence investment in that sector. And second, the resulting wage costs which—in conjunction with the Taylor rule for monetary policy that determines the interest rate and so the (real) exchange rate—will influence the exports and imports of the goods in each of the sectors, in ways which differ between sectors because of the different elasticities that are found in the trade equations for each sector. This too will influence investment in each of the sectors.

This is a challenging conclusion. Structural change appears important in the global economy. Modelling the causes of this structural change is a demanding exercise. Some short-cuts, of the kind deployed in the Global Gaidar Model, are useful, but not in all circumstances.

VI. Models in which there is chaotic behaviour

(i) Structural change due to chaotic behaviour in the global growth process

The approach to modelling structural change that I have taken so far in this paper is to consider departures from the canonical model and its extension, for the three kinds of reasons that I have identified: namely, things that cause tipping points to emerge, factors which cause one sector to grow faster than other sectors, and external developments, or changes in policies, which mean that the model needs to change. However it is possible to take a more radical view and to argue that, taken together, these three reasons mean that structural change is going on all the time, all over the place. That is effectively the approach adopted by Doyne Farmer in his paper, which is the only article in Part IV of this *OxREP* issue (Farmer, 2025, this issue). I now discuss the Farmer paper in some detail precisely because of the radical position that he adopts.

Farmer criticizes the canonical model and its extension by attacking the two fundamental ideas which underpin them. The first is that economic agents carry out intertemporal optimization, using a true model of the economy, in the absence of any concerns about uncertainly, so that they have a 'rational expectation' of what will emerge. The second is that they do this in the absence of any frictions that might constrain optimizing behaviour, apart from the particular frictions which are explicitly identified in the model.

Many of the papers which I have considered in this paper proceed in the way which Farmer criticizes. Thus, for example, the analysis of short-run macroeconomic policy-making which I discussed in section IV was carried out using models which are set up in a way which Farmer explicitly criticizes, namely they assume intertemporal optimization in models in which there are no frictions at all other than the nominal rigidities and the adjustment costs of capital that are explicitly identified. Furthermore, the paper by van der Ploeg and Venables presents models of the move to net zero emissions which are grounded in the idea that economic agents carry out intertemporal optimization: the authors argue that a move to non-fossil-fuel technology can become self-fulfilling as and when enough producers in the economy understand that this process is under way and so come to expect the reduction of costs that will accompany the 'learning by doing' inherent in the process.[41] Such analysis can be very revealing.

Nevertheless, Farmer will have none of this. He insists that we begin with an agent-based-modelling approach in which each of the actors follows a simple rule which does not come from any process of optimization, explicit or implicit.

According to Farmer, there are two reasons for proceeding in this way.

First, there may be many local optima in the kinds of non-linear models which he considers—and each individual agent will need to form a view as to which of the local optima all the other agents will have settled on. In these circumstances, it will be very difficult to determine which of the many possible outcomes the model should depict. Farmer maintains it is trivial to find an optimum for a convex function. But for a non-linear problem of this kind, with more than a few agents, it becomes literally impossible to figure out what all these optimizing agents will end up doing. Indeed, if the kind of non-linearities which Farmer identifies in his Basel leverage cycle model are pervasive everywhere, then the kind of tipping points, and boom–bust behaviour, which he identifies in that model might also be found everywhere. In that case, the kind of rational-expectations solution obtainable from a conventional model might not turn out to be very useful. The solution will depict the behaviour of the model until a tipping point is reached, at which point the model will come to have a new attractor, for which a new rational-expectations solution will be relevant. The position of any such tipping point may turn out to be difficult to predict, or even to identify.

Second, if one sets out in the way that Farmer suggests, instead of taking a rational-expectations-perfect foresight approach, one can instead ask how it is that people come to learn about what they are doing, and how they learn about what might happen in the future, given all of their past experience.

Despite this criticism, it is possible that agent-based models and DSGE models converge towards a commonly shared modelling practice. In particular, as individuals learn about the past, in the way studied in agent-based models, two things might happen. First, the simple-rules behaviour of individuals may well, through a process of trial and error, increasingly become optimal behaviour. And second, the expectations of each individual about what happens in the future may well come closer to what would happen in a rational expectations solution. Both of these factors would lead to some convergence of modelling practice. Indeed, all sensible applications of the rational-expectations optimizing procedure are increasingly investigating the extent to which this is the case.

Nevertheless, it is clear from Farmer's paper that the kinds of models which he advocates are able to identify reasons for structural change that emerge from the simple-rules behaviour that underpins these models. This means that—until the convergence of modelling approaches has proceeded much further—there will still be room for the conventional approach and the agent-based-modelling approach, side by side.

VII. Choosing and using macroeconomic models of structural change

The final part of this the *OxREP* issue, Part V, is about choosing and using the right models to guide policy-making in relation to structural change.

(i) The need for scenario analysis

The first paper, by Dimitri Zenghelis, Hector Pollitt, Jean-Francois Mercure, and Frank Geels, analyses the implications of structural change for the kinds of models which should be used by economic policy-makers (Zenghelis *et al*., 2025, this issue, henceforth ZPMG). The conclusion which ZPMG reach is that use of such models should involve scenario analysis rather than following the conventional model-simulation approach.

[41] The version of their model which they publish in this *OxREP* issue is a static model in which the intertemporal optimization is implicit. But, as noted above, they have produced a paper which makes the dynamics explicit. See van der Ploeg and Venables (2025*b*).

Much of their discussion relates to models of the structural change caused by global warming. But the points which they make have wider implications than this.

In the first section of their paper, after the introduction, ZPMG provide an explicit account of what they mean by 'structural change'. It is clear they think that structural change involves a move across a tipping point from one region—in which there is an initial point of equilibrium—to another region—in which there is a new and different point of equilibrium. That, for them, is what 'structural change' means. ZPMG are thus exclusively—or nearly exclusively—concerned with the first of the three causes of structural change which I have been discussing in this article.

Having set out their focus on tipping-points, ZPMG spend much time describing reasons why there might be a potential for tipping points in any actual economy. Here I simply list the causes of tipping points which they analyse. These are: reinforcing technical feedbacks; aggregation, clustering, and early mover advantage; processes of innovation that induce other further innovations; shifts in expectation and in the perception of risk; and cross-sectional feedbacks. These ideas are brought together in Table 1.

The question then arises: when there is a potential for a tipping point arising—for any one of the above reasons—what kind of process should govern the use of models to guide the decision-making of economic policy-makers. ZPMG argue that the conventional model simulation approach is not useful since, by and large, this approach uses models which ignore the potential for tipping points. I made this point earlier in this article when I was discussing the use of integrated assessment models (IAMs) to analyse policy in relation to climate change issues.

In a conventional model-simulation approach the effects of a change in external circumstances—a shock—are investigated using a particular model. This model is based, as far as possible, on past data. Then, subsequently, that same model is used to investigate the appropriate policy response to these effects, in order to remove, as much as possible, the deleterious consequences of the shock.

This conventional procedure is illustrated by ZPMG in Figure 2 which appears early in their paper. The left-hand side of the figure illustrates how this conventional approach uses a model to determine the appropriate response of a policy instrument—say taxes and/or a subsidy—to any shock which has happened. The figure suggests that—in the absence of any allowance for potential tipping points—the marginal cost of achieving any particular improvement through the use of this policy instrument will rise steadily as the level of policy intervention increases.[42] In such circumstances, the conventional model-based approach recommends using a policy instrument up until the point at which the marginal cost of using the policy instrument rises to the point where it is just equal to the marginal benefit of doing this. These marginal costs and benefits are computed on the same model which has been used to calculate the deleterious effects of the shock in the first place. Cost benefit analysis—or CBA—is central to these calculations. For obvious reasons, ZPMG call this conventional approach 'marginalist' or 'incrementalist'.

But when there is a potential for a tipping point of the kind illustrated in Table 1 of the paper, this conventional approach will not be appropriate, and an alternative approach will be necessary. That is precisely because the tipping point indicates that the model needs to change: the marginal costs and benefits, as calculated by the original model, will no longer be relevant. The right-hand side of Figure 2 illustrates these different circumstances. The figure suggests that—once an allowance for a potential tipping point has been made—the marginal cost of achieving any particular improvement through the use of a particular policy instrument may stop rising and—indeed—may begin to fall. The figure also suggests that, beyond such a point, it might be possible for the policy intervention to be gradually withdrawn. That is because, in some circumstances, the private sector might find it profitable to bring about the necessary improvement without any tax or subsidy. That is to say, the private sector might itself begin to take the economy in the right direction without any further policy assistance.

Of course this will not always be the case. A policy effort designed to get the economy to a tipping point will be enough when the new zero-carbon technology turns out to have lower costs than a carbon-intensive technology. That may very well be the case for most electricity generation and cars—although even with electricity, places that are particular favoured with endowments

[42] The vertical axis on the figure is not labelled, but this is a way in which the figure may be read.

for high-carbon-emissions power generation may do well in a future with no carbon price (e.g. Latrobe Valley coal in the State of Victoria in Australia, which is highly carbon-intensive and actually dirt cheap). But there is nothing in science, engineering, or economics that says that there are no economic activities that are lower cost through a carbon-intensive than a zero carbon (or more generally zero-emissions) way. Indeed, perhaps there are many. For these, we will need to have regulation or carbon pricing, including to encourage carbon-absorbing activities, forever.

Later in the paper ZPMG provide a more detailed comparison of the different kinds of models which underpin these two very different approaches to policy-making in the presence of the structural changes which are being caused by carbon emissions and global warming.

ZPMG argue that conventional models assume the following, explicitly or implicitly. Resource availability is fully understood as well as is possible, based on knowledge of past data. Resources are allocated in the best possible way, based on this knowledge. The capability to increase these resource endowments, and their productivity, is taken as given. Of course the models used in this conventional approach do allow for the fact that resources are being devoted to innovation, and to the development of new ideas, in the way which was captured by early versions of endogenous growth theory.[43] But there are no tipping points in these models of the kind which are identified in Table 1.

By contrast, ZPMG look for an alternative kind of process which recognises the following features. The amount of resources available at any time is recognized to be uncertain, but there is a capacity to supply them if that is necessary. The economy innovates to create resources, and to improve productivity, according to what it anticipates it will need. All innovation has uncertain rewards, requires investment. If the circumstances are right, in the way described in Table 1, the investment will be carried out. What is being discussed here is the kind of directed technical progress which Dan Susskind examines in his paper described in section III above. The winners of this year's Nobel Prize in Economics—Philippe Aghion, Peter Howitt and Joel Mokyr—focus with great clarity on when these self-perpetuating circumstances become sustainable.[44]

Of course scenario analysis is the kind modelling process that is needed to guide policy in these circumstances. That is because only scenario analysis can help one to understand why the phenomena depicted in Table 1 can—in fact—lead to the emergence of new and sustainable equilibria. The purpose of the scenario analysis will be to describe how different the economy might look, and in what ways, when it has crossed a particular tipping point. For example, ZPMG write as follows in one of their entries in Table 1:

> Combinatorial technologies and network and coordination effects mean there are greater advantages of moving in tandem with others, with spillovers, shared benefits, and more potential to generate new ideas.

Scenario analysis can help one to comprehend whether collective moves like this might actually occur. By contrast, cost benefit analysis (CBA), using a particular model in the way described above, may well turn out to examine only *part* of a self-fulfilling process of this kind—and so may well declare it to be undesirable. A policy intervention which targets only part of such an outcome might well not appear to be welfare improving. And—in fact—such an intervention might fail to stimulate any move at all.

Although ZPMG highlight the advantages of scenario analysis, they are nevertheless cautious in their proposals. They argue that economic models, when constructed in a suitable manner, are valuable in providing insights. However, such models also have fundamental limitations when making predictions in the context of the reinforcing feedbacks and increasing-returns features, as described in the ZPMG paper. Nevertheless, the authors argue that models are best used to inform risks. They suggest that a variety of models, complemented by a range of qualitative and non-modelling analytical approaches, with different strengths and weaknesses, can play an important role in helping to identify tipping points. Scenario analysis using such models can steer policies in the appropriate direction, even if the models used are suggestive rather than tightly specified or precisely calibrated.

43 See Romer (1990)
44 See Royal Swedish Academy of Sciences (2025)

(ii) Criteria that models should satisfy in the presence of complexity and uncertainty

The second paper in Part VII of this issue is by Cameron Hepburn, Matthew Ives, Sam Loni, Penny Mealy, Pete Barbrook-Johnson, Leon Clarke, Doyne Farmer, Nicholas Stern, and Joseph Stiglitz. The paper examines the most effective way that economic models and frameworks can help guide one particular set of policy choices, namely the transition to net zero policy. (Hepburn *et al*., 2025, this issue, henceforth HILMBCFSS). The emphasis of this paper is on the need to say something useful about what kinds of models can be used, despite the complexity and uncertainty involved in this transition process. The authors deliberately focus on a subset of the key economic policy questions involved in the transition to net zero, namely the questions that are of most interest to finance ministries, based on a recent survey.

Reaching net-zero emissions is a process which will involve deep, and far-reaching, uncertainty. That is because it will involve a structural transformation of the global economy which will also. be deep and far-reaching. HILMBCFSS focus on three aspects of this uncertainty as follows.

> First, the transition to a net zero economy will lead to profound structural alterations or breaks in several different sectors, not least energy, agriculture, transport, and much of heavy industry. These transformations must overcome path dependencies and deliver a shift from one system to another. Very large coordination challenges will emerge, in which policy-makers in different countries are actively intervening to gain comparative advantage in key technologies.… Second, in most sectors, the transition will be characterised by strategic complementarities between multiple agents, technologies and sectors. In particular, the returns to deploying one technology often increase as usage of a complementary technology expands, especially with the right coordination policy in place.… Third, non-linear dynamics will be a central feature of the net-zero transition. Decarbonising the global economy will involve the sort of tipping points, technology cost thresholds, investment frictions, regulatory ceilings or floors that inevitably create non-linear dynamics. In such circumstances, traditional methods that rely on linearization or small perturbations may fail. Good policy guidance will therefore often require models that can account for such complexities. Reliance on static or comparative-statics between optimal outcomes thus risks overlooking these crucial feedback loops.

In the face of this uncertainty, HILMBCFSS have set themselves the task of considering the most effective way in which economic models and frameworks can help guide policy. To do this they take the reader through a convincingly organised three-step procedure.

First, they identify five sets of questions which will need to be dealt with. Here is their list: the working of prices and markets (including, of course, any pricing of carbon and also including a large number of regulatory issues); the relative competitive positions which different countries will face; the nature of technological uncertainties; the implications for the labour market and for skills; and finally nature of the fiscal position—including taxes and budgeting—which countries will face themselves in.

Second, the authors suggest five general criteria that models and frameworks should meet. The models should be parsimonious; will need to be tractable, must provide insight, have to be empirically consistent; and—crucially—must be suitable for the task in hand. The authors discuss each of these criteria in some detail and provide some guidance on how to select the right model for whatever happens to be the task in hand. Armed with their five criteria, the authors reach an important, but obvious, conclusion. There will be many useful models. There will be no single model which encompasses all of the other models. The reason for this is that will be many different kinds of things to be done: their five sets of questions, which I have just enumerated, are so very different from each other that each of them will require its own kind of modelling support, depending on the question which is being analysed.

Third, the authors point out that there already exists a very wide variety of models. They first categorise models along two dimensions—from small to large and from theoretical to empirical. Then, armed with this taxonomy, they produce a list of the kinds of models which can already be identified, and briefly describe the features of each of them. This list includes: IAMs (of an optimising and a non-optimising kind); energy system models; CGE models; DSGE models; sector-specific models; agent-based models; system-dynamic models; and, lastly, hybrid models (which bring together more than one of these different kinds of models into one computational system). For each of these kinds of models, Table 4 provides a brief description of its capabilities.

This three-step argument builds up systematically to Table 5, which is the outstandingly useful feature of this paper. The Table provides a careful and comprehensive taxonomy, running over seven pages, of what the available models are and what they do. Each of the models listed in the Table is examined in a way which is illuminated by the above three-step argument. Thus for each model we learn what question it was designed to answer, what kind of model it actually is, and how it fares in relation to the five criterial which I have listed above.

All policy makers working in the climate-change field will be enormously grateful to the authors of this paper for the taxonomy which they have provided.

VIII. Conclusion

In section II of this article, I presented a taxonomy of reasons for structural change in macroeconomic models. First, there can be tipping points, both in single-sector and in multisectoral versions of a model. Second, there can a 'leading sector' which grows more rapidly than other sectors, for reasons which the model explains. Finally, there can be changes in external circumstances, or in the policy environment, which render the existing macroeconomic model inadequate and require it to be changed.

In section III I used this taxonomy to help me describe models of structural change to the growth process that are published in this issue of the *Oxford Review of Economic Policy*. I have suggested that the causes of structural change in these models can be found in one, or more than one, of the reasons described above. In section IV, I discussed papers which show how structural change can also arise in the short term, for similar kinds of reasons, as a result of interactions between households and firms in the private sector and macroeconomic policy-makers located in the central bank. In section V, I examined two papers that can help us to understand how structural change can happen not just in single economies but in the world economic system as a whole. Section VI provided an analysis of a paper which shows how chaotic behaviour can lead to structural change. And finally, in section VII, I discussed two papers which demonstrate the way in which the models of structural change set out in this *OxREP* issue can be of use to macroeconomic policy-makers.

In the introduction to this article, I suggested that a better understanding of the macroeconomics of structural change might have a beneficial influence on the way in which macroeconomics is practised as a discipline. Such an understanding might also have a beneficial influence on the way it is taught to graduate students. I say this for the following reason. In most core macroeconomics principles courses, growth theory is presented through a process which involves the gradual, logical development of the canonical, representative-agent single-sector Ramsey model of the economy, in the way which I described in the introduction. The objective of doing this is to present an all-encompassing abstract model of what, in principle, any macroeconomic growth process might actually look like in the longer term. However, this is a model with no structural change. Then, in a subsequent part of the core macroeconomics course, students are introduced to a similar logical development of the representative-agent, single-sector New Keynesian model of the behaviour of the economy in the short term, showing the way in which stabilization policy can and should be carried out in the presence of nominal rigidities. This is also a model with no structural change. Finally, at the end of the course, these two pieces of analysis are brought together in what I have called the extended canonical model; students are shown how these two sets of ideas can be brought together in a single overarching model of both the short term and the longer term. And again, this abstract picture of the economy contains no structural change.

There is an alternative. One might instead present macroeconomic theory as a set of particular responses to evolving historical circumstances and the structural changes which these circumstances bring about. Macroeconomic theory would then be presented as a set of exercises designed to help provide an understanding of those changing circumstances.[45] This is a key feature, for ex-

[45] Mark Blaug makes a similar point when he identifies two different ways to present ideas in the history of economic thought (Blaug, 1968, p. xi). One approach describes the development of economic theory as emerging from a need,

ample, of the outstanding Ohlin Lectures on international economics given by Roland Findlay in 1991 (Findlay, 1995). It is what Lewis did when he invented his theory of economic growth. And—in an earlier generation—this is what Keynes did when he invented the model in his *General Theory.* This is the kind of re-orienting of the teaching which the authors of the CORE project have successfully achieved, in relation to the teaching of economics to first-year undergraduate students.[46]

The approach which I have adopted in this paper suggests that, for graduate students, there may be a way of doing both of these two things at once. I have shown how, in a number of different ways, models of structural change can be bolted onto the canonical model of the growth process. I have also shown how models of structural change can be bolted onto its extended version, when analysing macroeconomic policy-making in the shorter term. I have myself found this dual approach rewarding, both for me as a teacher and for my students as well.

References

Anderson, K. (1987), 'On Why Agriculture Declines With Economic Growth', *Agricultural Economics*, 1(3), 195–207. https://doi.org/10.1111/j.1574-0862.1987.tb00020.x

Benzell, S.G., Kotlikoff, L. J., Kazakova, M., Lagarda, G., Nesterova, K., Yifan Ye, V., and Zubarev, A. (2025), 'The Future of Global Economic Power', *Oxford Review of Economic Policy*, 41(2), 519–54.

Bernanke, B. (2024), *Forecasting for Monetary Policy Making and Communication at the Bank of England*, available at https://www.bankofengland.co.uk/independent-evaluation-office/forecasting-for-monetary-policy-making-and-communication-at-the-bank-of-england-a-review/forecasting-for-monetary-policy-making-and-communication-at-the-bank-of-england-a-review

Blake, A., Dennis, R., Kirsanova, T., and Yates, A. (2025), 'Multiple Equilibria in the Absence of Commitment', *Oxford Review of Economic Policy*, 41(2), 404–25.

Blanchard, O. J. (1985), 'Debt, Deficits, and Finite Horizons', *Journal of Political Economy*, 93(2), 223–47. https://doi.org/10.1086/261297

— Fischer, S. (1989), *Lectures on Macroeconomics*, Cambridge, MA: MIT Press.

Blaug, M. (1968), *Economic Theory in Retrospect*, London: Heinemann.

Brigden, J. B. *et al*., (1929), *The Australian Tariff: An Economic Enquiry*, Melbourne University Press.

Butlin, N. G., and Gregory, R. G. (1989), 'Trevor Winchester Swan 1918–1989', *Economic Record*, 65, 369–77. https://doi.org/10.1111/j.1475-4932.1989.tb00689.x

Calvo, G. A. (1983), 'Staggered Prices in a Utility-maximizing Framework', *Journal of Monetary Economics*, 12(3), 383–98. https://doi.org/10.1016/0304-3932(83)90060-0

Carlin, W., and Soskice, D. (2018), 'Stagnant Productivity and Low Unemployment: Stuck in a Keynesian Equilibrium', *Oxford Review of Economic Policy*, 34(1–2), 169–94. https://doi.org/10.1093/oxrep/grx060

Christiano, L. J., Eichenbaum, M., and Evans, C. L. (2005), 'Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy', *Journal of Political Economy*, 113(1), 1–45. https://doi.org/10.1086/426038

Devarajan, S., Lewis, J., and Robinson, S. (1993), 'External Shocks, Purchasing Power Parity and the Equilibrium Real Exchange Rate', *The World Bank Economic Review*, 1, 45–63. https://doi.org/10.1093/wber/7.1.45

Farmer, J. D. (2025), 'Quantitative Agent-based Models: A Promising Alternative for Macroeconomics', *Oxford Review of Economic Policy*, 41(2), 571–90.

Findlay, R. (1995), *Factor Proportions, Trade, and Growth*, Cambridge, MA: MIT Press.

Galí, J. (2015), *Monetary Policy, Inflation, and the Business Cycle*, 2nd edn, Princeton, NJ: Princeton University Press.

Garnaut, R. (2025), 'The Blair Black Hole in Global Climate Policy', *Oxford Review of Economic Policy*, 41(2), 395–403.

Gollin, D., and Rogerson, R. (2025), 'Structural Change and Macro Development: Beyond the One-sector Growth Model', *Oxford Review of Economic Policy*, 41(2), 291–312.

Henckel, T., Menzies, G., and Zizzo, D. (2025), 'The Straw that Breaks the Camel's Back: Inferential Expectations and Sudden Belief Changes', *Oxford Review of Economic Policy*, 41(2), 426–51.

Henderson, D. W., and McKibbin, W. (1993), 'A Comparison of Some Basic Monetary Policy Regimes for Open Economies: Implications of Different Degrees of Instrument Adjustment and Wage Persistence', *Carnegie-Rochester Conference Series on Public Policy*, 39, 221–317. https://doi.org/10.1016/0167-2231(93)90011-K

internal to the models in question, to develop and extend the model for analytical reasons. The other describes it as emerging from a need to analyse a changing set of historical events, external to the models.

[46] CORE stands for 'Curriculum Open-access Resources in Economics'. See https://www.core-econ.org/

Hepburn, C., Ives, M., Loni, S., Mealy, P., Barbrook-Johnson, P., Clarke, L., Farmer, J. D., Stern, N., and Stiglitz, J. E. (2025), 'Economic Models and Frameworks to Guide Climate Policy', *Oxford Review of Economic Policy*, 41(2), 616–52.

Hicks, J. (1965), *Capital and Growth*, Oxford: Oxford University Press.

Hirano, T., and Stiglitz, J. E. (2025*a*), 'Overlapping Generations Models, Multiplicity of Steady States and Momentary Equilibria, and Economic Fluctuations', *Oxford Review of Economic Policy*, 41(2), 260–90.

— (2025*b*), 'Henry George, Land Speculation, and Economic Growth and Transformation', *Oxford Review of Economic Policy*, 41(2), 326–57.

Hodge, A., Jakab, Z., Lindé, J., and Nguyen, V. (2025), 'Did Expansionary Fiscal and Monetary Policies Cause the Inflation Surge?', *Oxford Review of Economic Policy*, 41(2), 484–518.

Intergovernmental Panel on Climate Change (2022), *Climate Change 2022: Impacts, Adaptation and Vulnerability: IPPC Sixth Assessment Report*, available at https://www.ipcc.ch/report/ar6/wg2/

Jones, R. (1965), 'The Structure of Simple General Equilibrium Models', *Journal of Political Economy*, 73(6), 557–72. https://doi.org/10.1086/259084

Keynes, J. M. (1936), *The General Theory of Employment, Interest, and Money*, London: Palgrave Macmillan.

Laxton, D., Igityan, H., and Mkhatrishvili, S. (2025), 'Monetary Policy Credibility, Avoiding Dark Corners, and Risk Management: A Response to Ben Bernanke's Review of Monetary Policy-making at the Bank of England', *Oxford Review of Economic Policy*, 41(2), 452–83.

Lewis, A. (1954), 'Economic Development with Unlimited Supplies of Labour', *The Manchester School*, 22(2), 139. https://doi.org/10.1111/j.1467-9957.1954.tb00021.x

— (1955), *The Theory of Economic Growth*, London: Allen & Unwin.

Liu, W. L., McKibbin, W., Shuetrim, G., and Wilcoxen, P. (2025), 'Modelling Structural Change in Global Macroeconomic Models', *Oxford Review of Economic Policy*, 41(2), 555–70.

McKibbin, W., and Vines, D. (2023), 'Longer-term Structural Transitions and Short-term Macroeconomic Adjustment: Quantitative Implications for the Global Financial System', *Oxford Review of Economic Policy*, 39(2), 245–66. https://doi.org/10.1093/oxrep/grad004

— Wilcoxen, P. (2013), 'A Global Approach to Energy and the Environment: The G-cubed Model', ch. 15 in *Handbook of CGE Modeling*, North-Holland, 995–1068.

Menzies, G., and Vines, D. (2008), 'The Transfer Problem and Real Exchange Rate Overshooting in Financial Crises: The Role of the Debt Servicing Multiplier', *Review of International Economics*, 16(4), 709–27. https://doi.org/10.1111/j.1467-9396.2008.00765.x

Michl, T., and Rowthorn, R. (2024), 'Optimal Inflation Targeting with Anchoring', *Review of Keynesian Economics*, 12(2), 220–38. https://doi.org/10.4337/roke.2024.02.06

Molana, H., and Vines, D. (1989), 'North–South Growth and the Terms of Trade: A Model on Kaldorian Lines', *The Economic Journal*, 99(396), 443–53. https://doi.org/10.2307/2234035

Nordhaus, W. (1992), 'The "Dice" Model: Background and Structure of a Dynamic Integrated Climate-Economy Model of the Economics of Global Warming', Cowles Foundation Discussion Paper No. 1009.

— (2018), 'Evolution of Modelling of the Economics of Global Warming: Changes in the DICE Model, 1992–2017', *Climatic Change*, 148(4), 623–40. https://doi.org/10.1007/s10584-018-2218-y

Pasinetti, L. (1960), 'A Mathematical Formulation of the Ricardian System', *The Review of Economic Studies*, 27(2), 78–98. https://doi.org/10.2307/2296129

Pitchford, J. (2002), 'Trevor Swan's 1956 Economic Growth "Seminar" and Notes on Growth', *Economic Record*, 78(243), 381–87. https://doi.org/10.1111/1475-4932.00065

Raworth, K. (2017), *Doughnut Economics: Seven Ways to Think Like a Twenty-first Century Economist*, New York: Random House.

REDDAWAY., W. B. (1937), 'Some Effects of the Australian Tariff', *Economic Record*, 13(1–2), 22–30. https://doi.org/10.1111/j.1475-4932.1937.tb02812.x

Ricardo, D. (1817), *On the Principles of Political Economy and Taxation*, London: John Murray.

Romer, P. (1990), 'Endogenous Technological Change', *Journal of Political Economy*, 98(5, Part 2), S71–S102. https://doi.org/10.1086/261725

Rowthorn, R. (1977), 'Conflict, Inflation and Money', *Cambridge Journal of Economics*, 1(3), 215–39.

— (2024), 'The Conflict Theory of Inflation Revisited', *Review of Political Economy*, 36(4), 1302–13. https://doi.org/10.1080/09538259.2024.2332297

Royal Swedish Academy of Sciences (2025), 'Scientific Background to the Sveriges Riksbank Prize in Economic Sciences in Memory of Alfred Nobel 2025', *Statement by the Committee for the Prize in Economic Sciences in Memory of Alfred Nobel*, available at https://www.nobelprize.org/uploads/2025/10/advanced-economicsciencesprize2025.pdf

Smets, F., and Wouters, R. (2007), 'Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach', *American Economic Review*, 97(3), 586–606. https://doi.org/10.1257/aer.97.3.586

Smith, A. (1776), *An Inquiry into the Nature and Causes of the Wealth of Nations*, London: W. Strahan.

Solow, R. M. (1956), 'A Contribution to the Theory of Economic Growth', *The Quarterly Journal of Economics*, 70(1), 65–94. https://doi.org/10.2307/1884513

Stiglitz, J. E. (1994), 'Endogenous Growth and Cycles', in Y. Shionoya, and M. Perlman (eds), *Innovation in Technology, Industries, and Institutions*, Ann Arbor, MI: University of Michigan Press, 121–56. This paper was originally presented as the Schumpeter Lecture at the fourth biennial meeting of the International Schumpeter Society, held in Kyoto in August 1992 and appeared as NBER Working Paper No. 4286.

Susskind, D. (2024), *Growth: A History and a Reckoning*, New York: Allen Lane. https://doi.org/10.4159/9780674297050

— (2025), 'Directed Technological Change: A History and a Critical Agenda', *Oxford Review of Economic Policy*, 41(2), 313–25.

Swan, P. (2022), *Trevor Winchester Swan: Life and Contribution to Economic Theory and Policy*, vol. 1, London: Palgrave Macmillan.

Swan, T. (1956), 'Economic Growth and Capital Accumulation', *Economic Record*, 32(2), 334–61. https://doi.org/10.1111/j.1475-4932.1956.tb00434.x

— (1957), 'Population Growth and Economic Development', Presidential Address, Section G, ANZAAS, Dunedin, New Zealand; notes on address taken by H. W. Arndt, and published in *Trevor Winchester Swan: Contributions to Economic Theory and Policy*, vol. 2, 62–3, London: Palgrave Macmillan.

Taylor, J. B. (1993), 'Discretion versus Policy Rules in Practice', *Carnegie-Rochester Conference Series on Public Policy*, 39, 195–214. https://doi.org/10.1016/0167-2231(93)90009-L

van der Ploeg, F., and Venables, A. J. (2025*a*), 'Green Transitions: Complementarities, Multiple Equilibria, and Tipping Points', *Oxford Review of Economic Policy*, 41(2), 377–94.

— (2025*b*), 'Positive Tipping Points and Transitional Dynamics: Policies for the Green Transition', CEPR Discussion Paper No. 520764, Paris and London, https://cepr.org/publications/dp20764

Vines, D., and Luk, P. (2025), 'The Simple Extended New Keynesian DSGE Model: A Constructive Toy-model Response to the Criticisms of Joe Stiglitz', paper prepared for the Festschrift Volume to be presented to Joseph Stiglitz.

— Wills, S. (2018), 'The Rebuilding Macroeconomic Theory Project: An Analytical Assessment', *Oxford Review of Economic Policy*, 34(1–2), 1–42. https://doi.org/10.1093/oxrep/grx062

— — (2020), 'The Rebuilding Macroeconomic Theory Project Part II: Multiple Equilibria, Toy Models, and Policy Models in a New Macroeconomic Paradigm', *Oxford Review of Economic Policy*, 36(3), 427–97. https://doi.org/10.1093/oxrep/graa066

— Luk, P., Wills, S., and Zeitlin, A. (2008), 'Dual Economies', in S. Durlauf, and L. Blume (eds), *The New Palgrave Dictionary of Economics*, 2nd edn, London: Palgrave Macmillan.

Warr, P., and Yusuf, A. A. (2025*a*), 'What Causes Structural Change?', *Oxford Review of Economic Policy*, 41(2), 358–76.

— (2025*b*), 'Structural Change and Income Inequality: Evidence from Thailand', *Working Papers in Trade and Development, No. 2025/06*, Canberra: Australian National University, https://crawford.anu.edu.au/acde/content-centre/research/structural-change-and-income-inequality-evidence-thailand

Weil, P. (1989), 'Overlapping Families of Infinitely-lived Agents', *Journal of Public Economics*, 38(2), 183–98. https://doi.org/10.1016/0047-2727(89)90024-8

Yaari, M. E. (1965), 'Uncertain Lifetime, Life Insurance, and the Theory of the Consumer', *The Review of Economic Studies*, 32(2), 137–50. https://doi.org/10.2307/2296058

Zenghelis, D., Pollitt, H., Mercure, J.-F., and Geels, F. (2025), 'Understanding and Modelling Structural Economic Change as a Dynamic Resource Creation Process—An Application to Low-carbon Transitions', *Oxford Review of Economic Policy*, 41(2), 591–615.