Available online at www.sciencedirect.com
ScienceDirect
Energy Economics 28 (2006) 539 – 562
Energy Economics
www.elsevier.com/locate/eneco

Representing induced technological change in models for climate policy analysis

Ian Sue Wing *

Kennedy School of Government, Harvard University, Center for Energy & Environmental Studies and Dept. of Geography & Environment, Boston University and Joint Program on the Science & Policy of Global Change, MIT

Available online 10 July 2006

Abstract

Induced technological change (ITC), whereby the relative price effects of reducing greenhouse gas emissions stimulate innovation that mitigates the cost of abatement, is both tantalizing to decision makers and challenging to represent in the computational economic and engineering models used to analyze climate change policy. This overview reconciles the divergent views of technology and technological change within different types of models, elucidates the theoretical underpinnings of ITC, introduces the reader to the techniques of their practical implementation, and evaluates the implications for models' results.

JEL classification: O3; Q2; Q4

1. Introduction

This paper is a brief overview of the subtleties and challenges involved in endogenizing technological progress in computational models for climate policy analysis. Its focus is on induced technological change (ITC), which in the climate context is the expansion of substitution possibilities for greenhouse gas (GHG) intensive inputs to production—principally fossil fuels—facilitated by inventive responses to the price changes induced by policies to mitigate global warming.

The future trajectory of technology is perhaps the most important factor influencing the cost of mitigating climate change. Concerns over the adverse economic consequences of policies to cut carbon dioxide (CO_2) emissions are motivated by the lack of carbon-free energy supply options which will be able to substitute for fossil fuels in the foreseeable future, and the limited possibilities for using other inputs to substitutes for the global economy's demand for energy.

* Rm. 141, 675 Commonwealth Ave., Boston, MA 02215. Tel.: +1 617 353 4751; fax: +1 617 353 5986.
E-mail address: isw@bu.edu.

0140-9883/$ - see front matter
doi:10.1016/j.eneco.2006.05.009

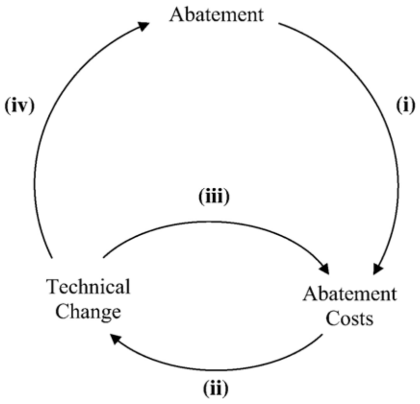

Fig. 1. Endogenous technological change and GHG mitigation: the big picture.

Innovations which facilitate substitution on either the demand or supply side are therefore seen as the saving grace which can moderate the cost of climate policies.

Computational investigations of the costs of reducing GHG emissions policies often treat technological progress as exogenous and invariant to the effects of abatement policies. In an intertemporal optimization framework, autonomous technological advance favors a “wait-and-see” strategy of continuing to emit at high levels while allowing innovation to lower future abatement costs (e.g., Wigley et al., 1996; Manne and Richels, 1997). By contrast, ITC implies that such beneficial innovation may be stimulated by near-term actions such as subsidies for the research and promotion of carbon-free energy supplies or energy efficiency improvements, or simply aggressive programs of abatement, thus favoring an “act now” strategy (e.g., Grubb et al., 1995).

At issue are analysts’ assumptions about the response of innovation to policy-induced input price changes and the consequent shift in substitution possibilities. Since Hicks (1932) initial articulation of the ITC hypothesis, theoretical and empirical elaboration of the mechanisms by which prices affect the rate and direction of innovation have until recently proved elusive.[1] But in spite of this caveat, the idea that forcing producers to bear the costs of cutting pollution induces technological change that both alleviates the costs of abatement and increases profits still captures the imagination of environmental advocates and scholars alike.[2]

Fig. 1 captures the essence of the issue. While the direct effects of emission taxes or quotas on welfare (i) have been well studied (see, e.g. Weyant (ed.), 1999), we still have only a limited

[1] Hicks (1932, p. 124): “a change in the relative prices of factors of production is itself a spur to invention, and to invention of a particular kind—directed to economizing the use of a factor which has become relatively expensive”. Attempts by Kennedy (1964), von Weizsacker (1965) and Ahmad (1966) to develop this hypothesis into a theory led to a sizeable literature (for surveys, see Binswanger and Ruttan, 1978; Thirtle and Ruttan, 1987) which after much argument and criticism (e.g. Samuelson, 1965; Nordhaus, 1973a) faded from view in the 1970s. The topic has seen a revival thanks to Acemoglu (2002) model of directed technical change, which continues to spawn applications in the environmental policy arena (e.g., Smulders and de Nooij, 2003).

[2] This is the so-called “Porter hypothesis”—see, e.g., Ashford et al. (1985), Ashford (1994) and Porter and van der Linde (1995)—as well as Palmer et al. (1995) for criticisms. In a reconciliation of these positions, Popp (2005) demonstrates that when the outcome of pollution-saving R&D is uncertain, a tax on pollution which reduces a firm’s expected profit will induce innovation which has a high probability of lowering profits even further but which nonetheless has a low probability of increases profits above their pre-tax level.

understanding of how such measures may induce technological change (ii) and what the economic and environmental consequences of this may be (iii). Thus, to gain further insight into the influence of ITC on the optimal program of mitigation (iv) requires an elaboration of the mechanisms which drive the feedback loop (ii)–(iii).

To represent this feedback loop in climate policy simulations, exogenous technological change must be replaced by a formulation which renders both the rate and the direction of innovation endogenous. It is necessary for the analyst to designate which variables represent the inputs to and the outputs of innovation, develop a reduced-form structural representation of the transformation between the former and the latter, and express the result with models using specific algebraic functions and numerical parameterizations. Building on previous reviews,[3] this paper sheds light on how these architectural options are constrained by the different structural representations of technology in the two main classes of climate policy simulations: "bottom-up" engineering models and "top-down" macroeconomic models.[4] The result is a concise roadmap to the practical methods of implementing the basic theoretical elements of technological change, which points to promising directions for future investigations of the nexus between climate policy and ITC.

The remainder of the paper is organized into four sections. Section 2 lays the groundwork for the succeeding discussions by introducing a simple conceptual framework for understanding the different representations of technology in climate policy models. Section 3 examines engineering and economic conceptions of technological change, reconciling the former's description of "micro-scale technological change" with the latter's description of substitution. Section 4 is the meat of the paper, which compares and contrasts the various methods of representing technological change within climate policy models. Section 5 provides a summary and conclusions.

2. What is "Technology"? insights from production theory

Throughout the paper we employ the analytical device of a hypothetical producer, who in each time period $t = \{0,\ldots,T\}$ generates output, $Y(t)$, according to a production function, Q, denominated over two inputs: a quantity $C(t)$ of a clean good (e.g., a composite of capital, labor and non-energy intermediate inputs) and a quantity $D(t)$ of a dirty good (e.g., fossil fuels) which generates pollution:

$$Y(t) = Q[C(t), D(t)]. \tag{1}$$

In the neoclassical model of production, Q represents the envelope of feasible techniques of combining quantities of C and D to produce the level of output Y. Each technique may be thought

---

[3] e.g., Azar and Dowlatabadi (1999), Grubler et al. (1999), Weyant and Olavson (1999), Grubb et al. (2002), Goulder (2004), and especially Loeschel (2002, 2004).

[4] Bottom-up models (e.g., MARKAL—Löschel, 2004) refer to primal activity analysis simulations that solve for the levels of capacity of energy transformation and conversion technologies that minimize the cost of fulfilling demands for energy services that are either specified according to demand curves or derived from a simple aggregate macroeconomic model (e.g. Manne et al., 1995). They contain a detailed description of the energy system in the form of a technology matrix that represents the level of capacity of individual processes for transforming different primary energy resources into a range of energy carriers, and for converting these commodities into energy services that satisfy final demands. Top-down models refer to macroeconomic simulations which come in two flavors: primal aggregate Ramsey growth models with an environmental sector (e.g., DICE and RICE—Nordhaus and Boyer, 2000), and primal–dual CGE models (e.g., EPPA—Paltsev et al., 2005) that solve for the set of commodity and factor prices, and the levels of industry activity and household income that clear all markets in the economy, given aggregate factor endowments, households' consumption technologies (specified by their utility functions) and industries' transformation technologies (specified by their production functions).

of as a particular production process for combining clean and dirty inputs in a fixed proportion given by the benchmark unit demands $\bar{C}$ and $\bar{D}$, respectively. Thus, for technique $m \in M$,[5]

$$Y_m(t) = \min[C_m(t)/\bar{C}_m, D_m(t)/\bar{D}_m]. \tag{2}$$

If $O(t) \subset M$ is the subset of techniques operated at time t, then the overall levels of output and demands for inputs in Eq. (1) are given by the corresponding sums over all active techniques:

$$Y(t) = \sum_{m \in O(t)} Y_m(t), \;\; C(t) = \sum_{m \in O(t)} C_m(t) \text{ and } D(t) = \sum_{m \in O(t)} D_m(t).$$

In the case where the set M completely specifies all possible means of producing output—at t or any other period—then Q is all-encompassing and immutable. In this world the production function is an essentially static concept, with the only margin of adjustment being substitution, whereby the producer is able to switch among different known and available techniques of production in response to a change in the relative price of inputs. The magnitude of this shift is determined the shares of the two inputs in production, their relative prices, and the elasticity of substitution, σ, which determines the curvature of Q.[6]

Fig. 2(A) is the textbook illustration of this process. The feasible production set, shown in gray, defines the combinations of C and D which can generate a unit of Y. The heavy arc defining the boundary of this set is the unit isoquant, which corresponds to the level curve of Q, or the locus of production techniques which utilize the smallest amounts of C and D per unit output. Techniques such as A_0 in the interior of the production set use larger quantities of C and/or D, and are therefore comparatively inefficient. Thus, of the set of available techniques, $A(t) \subset M$, the isoquant is the smooth approximation to the sequence of linear combinations of the efficient subset, A_1–A_4.

In each time period a profit-maximizing producer can generate a unit of output by employing one or more frontier techniques, which combines C and D in the fixed proportion $\bar{D}_m / \bar{C}_m$, shown by the slope of the dashed rays running through them.[7] In competitive equilibrium the producer chooses the quantity of each input so as to equalize its marginal physical productivity and its market price. In unit input space, the point where this occurs corresponds to the combination of techniques at which the slope of the relative price line is tangent to the unit isoquant.[8] Thus, if at time t the dirty input is relatively inexpensive, a situation represented by the price line PP, then the producer will employ technique (2), so that $O(t) = A_2$. If at $t+1$ a pollution tax makes D relatively dearer, the relative price line rotates counter-clockwise to $P'P'$ and the producer shifts to A_3.

The foregoing description highlights a deep duality and important semantic distinctions between the conceptions of technology in bottom-up and top-down models. The former typically represent the individual As, treating each as “a technology”, i.e., an activity or process of a particular type. The shift from A_2 to A_3 is “technology substitution”, referred to as “microscale technological change” by

[5] Bottom-up studies often refer to these as “linear” technologies because they are specified in terms of their cost functions, i.e., the dual of Eq. (2).

[6] The elasticity of substitution can be thought of as the percentage change in the relative quantities of the inputs to production induced by a one-percent change in relative input prices, with the price and quantity of output held constant. Mathematically, $\sigma = -\left(\frac{\mathrm{d}(C/D)}{C/D} \div \frac{\mathrm{d}(p_c/p_D)}{p_c/p_D}\right)$.

[7] e.g., the gently-sloped ray through technique A_4 indicates that it employs inputs of the clean input relatively intensively, while the steeply-sloped ray through A_1 indicates that it is relatively pollution-intensive.

[8] Using the fact that under conditions of perfect competition and constant returns to scale the marginal product of each input is equal to its price ($\partial Q / \partial C = p_C$ and $\partial Q / \partial D = p_D$), the total derivative of Eq. (1) can be rearranged to yield $\mathrm{d}D / \mathrm{d}C = -p_C / p_D$, which is the marginal rate of technical substitution (i.e., the slope of the isoquant) at the point of tangency in D–C space.

Grubler et al. (1999). By contrast, macroeconomic models rarely represent discrete activities but instead consider "the technology" to be the entire envelope Q, whose discrete elements all exist prior to the relative price change, but may not all have been operated. The shift from A_2 to A_3 is therefore regarded as substitution. To understand the subtleties of this distinction it is necessary to clarify what is meant by technological change.

3. Technological change and its inducement

3.1. Economic vs. engineering conceptions of technological change

Technological progress is a change in the character of productive activity. This change can be radical, creating fundamentally new products or processes, or incremental, improving the performance or efficiency of ones which already exist. Technological change encompasses the processes of invention, innovation or development, and diffusion or adoption. Invention is fueled by individual human creativity and the state of scientific knowledge, while development relies on engineering know-how to scale-up working prototypes into commercially useful, routinized production processes, whose widespread adoption by firms depends on both ruling prices and the level of producers' technical knowledge.

The narrower economic conception of technological progress is the change in the character of production, with prices held constant, which enables more output to be produced using the same quantities of inputs, or, symmetrically, which allows the same level of output to be produced from smaller quantities of inputs. In the present context, innovation alters the recipe for combining C and D to make Y, a process which constitutes an evolution in the shape of the production function. The implication is that Q is no longer the envelope of all possible production techniques, only those feasible at time t. Radical technological change is synonymous with the appearance in subsequent periods of new techniques, while incremental technical change corresponds to increases over time in the efficiency with which known techniques transform inputs into output. Technical progress occurs if either of these processes shifts the unit isoquant inward toward the origin.

Fig. 2(B) provides an illustration. First, imagine that at t, prices are constant at PP and the production function is given by Q. Then radical technological change creates the new technique A_5, which is capable of generating a unit of output using the same amount of D as A_2, but with one-third less of input C. The result is an expansion of the feasible production set and a shift in the unit isoquant toward the origin to Q', which is a new transformation frontier given by the sequence of linear combinations of $\{A_1, A_5, A_4\}$. A_2 and A_3 are eclipsed, and join A_0 in the interior of the production set. The key outcome is the shift in the point of tangency between the relative price line and the isoquant, with the producer shifting from A_2 to a combination of techniques A_1 and A_5, using relatively less of the dirty input in the process.

A similar outcome arises if instead of radical innovation there is incremental improvement in the efficiency of techniques A_2 and A_3, represented by their inward shifts to A_2' and A_3'. The necessary condition here is the *differential* movement of each technique along its ray toward the frontier, with the largest improvements being concentrated in the most pollution-intensive activities. Thus, with incremental technological change the sequence of techniques which define Q and Q' are identical, but the change is in the *character* of its constituent activities. In this case the producer also adopts a cleaner technique of production, shifting from A_2 to a combination of A_2' and A_3'.

The C-using and D-saving shift of the production frontier in Fig. 2(B) is an example of biased technical change, which occurs when an innovation saves relatively more of one input to production than another. The strength of this effect is known as the bias of technical progress,

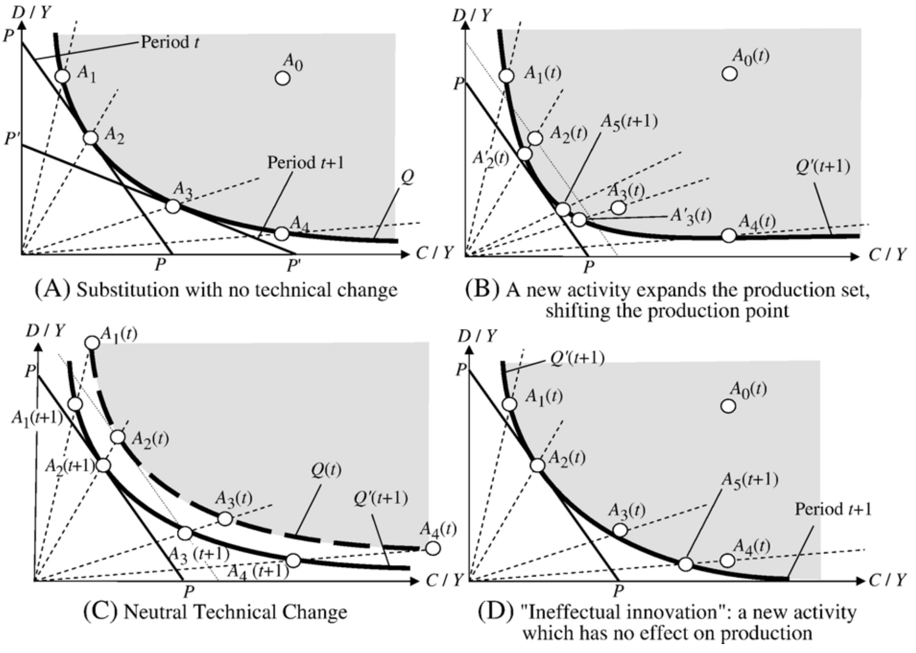

Fig. 2. Substitution versus technical progress.

which is the rate of change in the shares of the inputs to production when the prices of inputs and output are held constant (Binswanger and Ruttan, 1978). Fig. 2(C) illustrates the alternative case of neutral technical progress, in which the rates of reduction in input demand are identical, and shift the entire isoquant toward the origin without changing its shape. The upshot is that with unchanged relative prices the producer sticks with technique A_2, but is able to create a unit of output from proportionately fewer units of all inputs.

But while technical progress implies an expansion of the subset of efficient techniques and/or a change in their efficiencies, the converse is not necessarily true. Obvious examples are a small increase in the efficiency of A_0 or the appearance of a new technique within the grey area, both of which are improvements which have no effect on the position of the unit isoquant. It is less obvious, however, that a shift in the frontier does not automatically imply progress. This is illustrated by the pathological case of "ineffectual innovation" in Fig. 2(D). Despite the efficiency improvement from A_5 to A'_5, this technique is sufficiently distant from the pre-existing production point A_2 in normalized input space that the producer continues to be use the latter. However, this taxonomic convention is complicated by price changes. If prices shift from PP to $P'P'$ in Fig. 2 (B), then the existence of a highly efficient technique such as A'_5 facilitates greater substitution of C for D than would be the case in Fig. 2(A). Only then can it be said that progress has actually taken place.

3.2. Technological change vs. substitution

The foregoing examples shed light on the relationship between substitution and technical change. Once we relax the assumption of constant prices, technological change in the broadest sense of the term can be logically thought of as a two-step process of innovation followed by

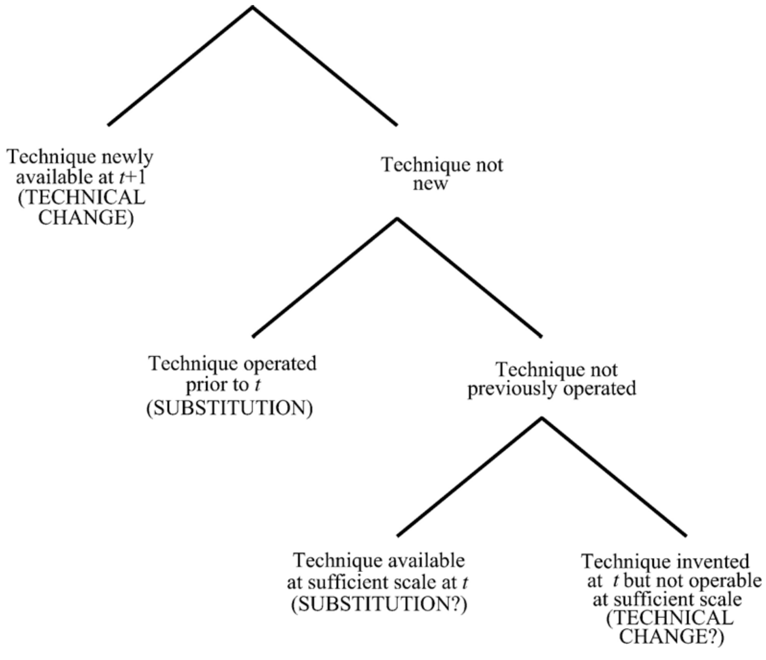

Fig. 3. Substitution vs. technical change: a conceptual classification.

substitution toward newly-available techniques. The general caution is that maintaining a neat conceptual distinction between substitution and technical change is not a straightforward task.[9]

To classify an observed shift in technique it is necessary to discern whether the production process which ends up being used pre-dates the price change. Several cases are possible in Fig. 2 (A), a taxonomy for which is given in Fig. 3. If A_3 does not exist in period t then technical progress has clearly taken place. If it does exist, the question is whether A_3 has been operated prior to t, or is a still "new" in the sense of being available but unused. In the first instance, once A_3 has not experienced intervening efficiency improvements, then the shift is clearly substitution. If A_3 has never been used but has been available at sufficient scale, then whether or not substitution has occurred is ambiguous. But in any case this circumstance seems implausible, as it is unlikely that producers will build up sizeable capacity in unprofitable technologies only to leave them idle in expectation of uncertain future price changes. The more likely alternative is where A_3 has been *invented* but is not a fully-fledged production process at t. Then the movement from A_3 to A_2 must involve development, diffusion, and scale-up of operations, all of which are components of technological change.

The unit isoquant is thus a snapshot of the techniques amongst which the producer is assumed to be able to shift frictionlessly over a certain time-frame. The length of this interval (which in climate policy models is on the order of 5–10 years) distinguishes substitution—which is

[9] This point harks back to an old debate on the possibility of econometrically identifying both the elasticity of substitution and the bias of technological change given prices and quantities of inputs and output (Nerlove, 1967; Diamond et al., 1978). This question has been rendered moot by the use of flexible cost and production functions such as the translog or generalized Leontief, in which bias of technical progress is given by the coefficient on the time-trend in an input's intensity or share, controlling statistically for the effects of the prices and/or quantities of inputs and the scale of production.

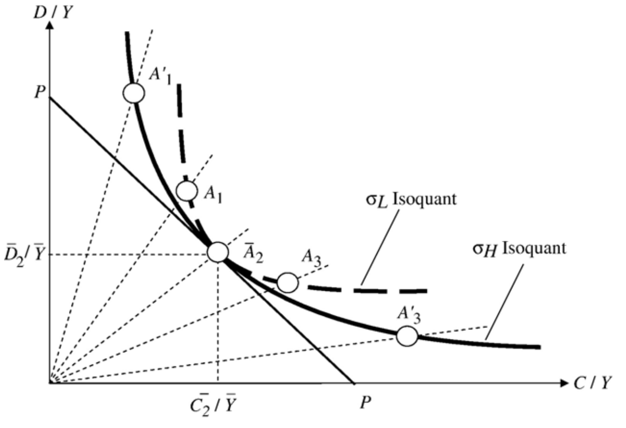

Fig. 4. The elasticity of substitution as an efficiency parameter.

considered to be a contemporaneous phenomenon, from technological change—which, depending on how radical a new technology is, may take several decades. Judgments about the scope of substitution possibilities in the short run are implicit in the value of the elasticity of substitution, which is perhaps the most important exogenous parameter employed in the numerical calibration of top-down models.

Fig. 4 portrays a hypothetical situation where benchmark data indicate that $\bar{Y}$ units of output are produced from $\bar{C}_3$ units of the clean input and $\bar{D}_3$ units of the dirty input at relative input prices *PP*. The missing piece of necessary information is the elasticity of substitution, whose value defines the curvature of the isoquant through $\bar{A}_2$ which is tangent to *PP*.[10] While $\bar{A}_2$ is a known benchmark, there is often a dearth of information about the range of alternative activities to which the producer can switch in the short run, implying that σ must be assumed. For given shifts in relative prices, a low elasticity σ_L reflects the availability of alternatives such as A_1 and A_3 whose input proportions are similar to $\bar{A}_2$, while a high elasticity σ_H reflects the availability of the radically different (and relatively efficient) alternatives A_1' and A_3'.

If a given expenditure on research allows technological breakthroughs appear with some constant flow probability, then the more time the producer has to adjust to a given price change, the greater the chance of that an innovation comes into being which is capable of transforming currently unprofitable techniques with radically different input proportions into feasible production alternatives at the new ruling prices. Thus, the larger the value of σ, the greater the potential for conflating long-run substitution with elements of technological change.

An important corollary is that with stable prices, the availability of new clean production techniques (i.e., ones close to the *C* axis) is likely to have little contemporaneous impact on production. This point is suggested by Fig. 5, which plots historical data on average U.S. fossil fuel prices from DOE/EIA, trends in patenting of energy supply and demand technologies from

[10] Models' predictive validity rests on their ability to replicate real-world observations. Calibration refers to the numerical procedures which solve for the values of the coefficients in a model's equations which enable the model to reproduce the data on which it is empirically benchmarked. Because there are frequently fewer observations than unknown parameters, the calibration problem is often under-determined, and can only be solved if the modeler exogenously specifies the values of key parameters (Jorgenson, 1984; McKitrick, 1998; Dawkins et al., 2001). Where econometric estimates are not available, the latter are often based on assumptions and judgment.

Popp (2001, 2002), indices of the effects of substitution and embodied technical progress on industries' energy intensity, and an index of aggregate energy intensity net of structural change from Sue Wing and Eckaus (2004). Despite the surge in energy technology patenting in response to the first energy price shock (epoch I), it took many more years of high prices (epoch II) before aggregate energy intensity began to decline. This lag is suggestive of the time and resources

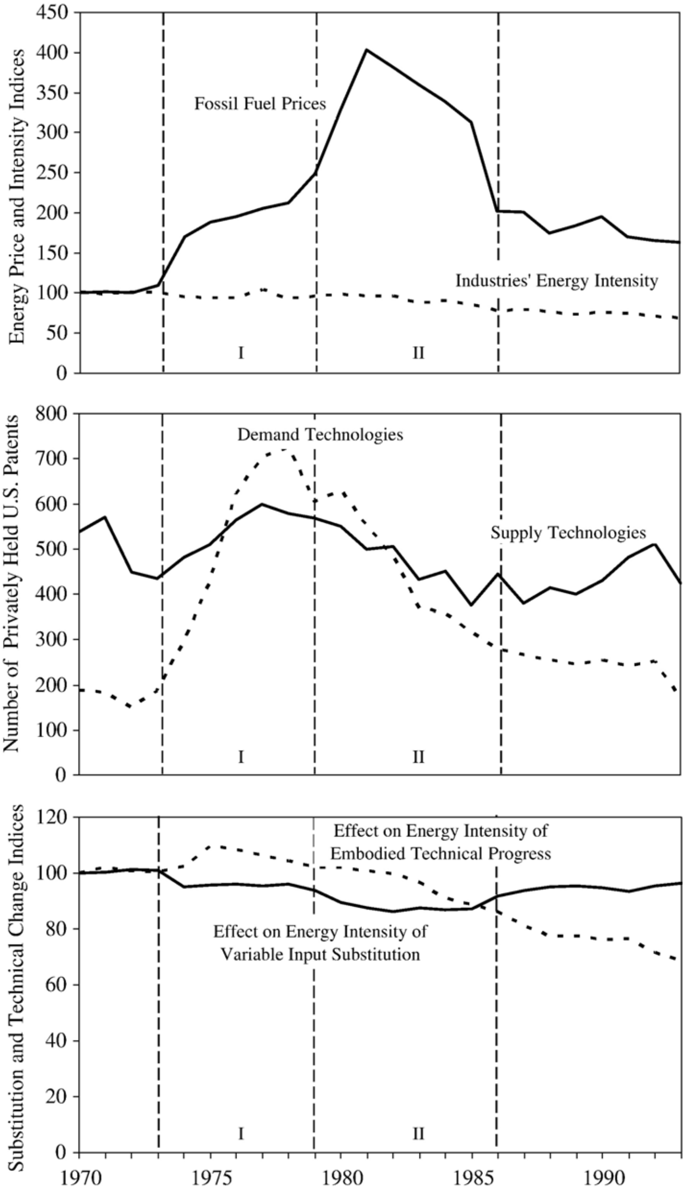

Fig. 5. The impact of energy prices on technological change and energy intensity. Sources: DOE/EIA (various years); Popp (2001, 2002); Sue Wing and Eckaus (2004).

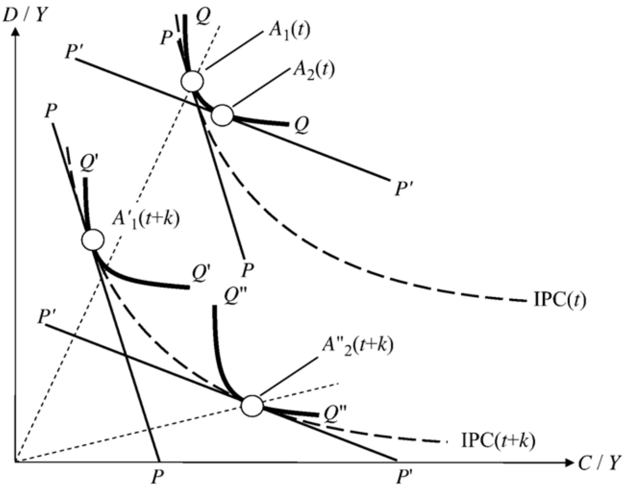

Fig. 6. The Ahmad–Binswanger–Ruttan model of ITC.

necessary for producers to move between invention and substitution. In particular, Sue Wing and Eckaus find that while energy intensity was temporarily diminished by substitution during this period, the former's sustained decline (even after the collapse of energy prices) was due to embodied technical change—namely, the capitalization of energy-saving innovations into industries' stocks of equipment, machinery and especially information technology capital.

The implication is that we should be pessimistic about the likelihood of success of subsidies to bring high-efficiency or low-carbon technologies "down off the shelf" and into commercial use without complementary measures to increase the price of inputs which embody GHGs (IPCC 2001: §7.3.4.1). This raises the question of how prices affect both patenting (i.e., invention) and the intervening processes of innovation and diffusion, which we turn to next.

3.3. Understanding inducement

In what is perhaps the clearest theoretical articulation of ITC, Ahmad (1966) and Binswanger and Ruttan (1978) treat the isoquant in Figs. 2 and 3 as only one of many potential families of techniques for producing output under different relative price regimes. Each of these alternative production functions represents a draw from a superset of production possibilities which is determined by the state of technological knowledge in each period. The key determinant of the alternative which is actually realized is the *innovation possibility curve* (IPC), which represents the efficient frontier of this superset in input space, and defines the envelope of unit isoquants which the producer is capable of developing with a given research and development expenditure.

Fig. 6 provides an illustration. The IPC in each period is given by the dashed lines and the heavy curves show the production functions which are actually realized. The latter are the ones that share a common point of tangency with the relative price line and the IPC, so that the marginal rates of technical substitution along the isoquant and the technology menu both equal the relative price ratio. Prices therefore not only induce the producer to select the particular technique which lies at the point of tangency between the relative price line and the unit isoquant, they also induce the selection of the family of activities which lies at the point of tangency between the relative price line and the IPC.

Given the state of technical knowledge in period *t*, relative prices *PP* determine both the production function *QQ* and the active technique A_1. Over a span of *k* periods, the increase in knowledge due to the secular progress of scientific discovery makes innovation more efficient, causing the IPC to move toward the origin in a neutral fashion. The production function thus shifts inward until it is jointly tangent with the new IPC and the relative price schedule. The upshot is that even with constant prices, biased technical progress can occur as the production function moves to $Q'Q'$. If relative prices change to $P'P'$ the producer's instantaneous response is to shift to technique A_2, which gives rise to a slight reduction in the demand for the dirty input.[11] But over time, prices serve to direct the process of technical change, causing the entire production function to shift to $Q''Q''$ and the producer's shift to the vastly different technique A_2'' on this frontier, which employs a far smaller input of *D*. The shift from *QQ* to $Q''Q''$ thereby represents price-induced *C*-using and *D*-saving technical progress.[12]

Fig. 6 is also suggestive of the importance of σ. By pinning down the curvature of *QQ*, it determines which realization of the production function will be jointly tangent to the IPC and the relative price line. The elasticity of substitution is therefore a de facto technology parameter (Jacoby et al., this volume), which leads immediately to the question of whether there exist similar parameters which determine the position and curvature of the IPC. It turns out that the answer is no. Although the Ahmad–Binswanger–Ruttan framework provides a short-cut to the *outcome* of ITC, the IPC's shape is determined by the *mechanisms* through which technological change proceeds.

Recall that the IPC's distance from the origin is indicative of the overall efficiency of production, which in turn is determined by the economy's stock of skills, experience, ideas and blueprints. Collectively these can be viewed as an intangible asset—"knowledge capital"—whose accumulation is driven by education and the additions to economically-useful human understanding as a result of production experience or investments in research. The IPC's position and curvature reflect the potential influence of prices on both the precursors and the outcome of the accumulation process, namely, the inducement of R&D and the influence of intangible capital on the character of the set *A*.

To understand why, it is necessary to develop a process-based elaboration of the feedback loop in Fig. 1. An example is illustrated schematically in Fig. 7. Besides inducing contemporaneous substitution among tangible inputs, shifts in p_D/p_C also stimulate inventive activity in the form of R&D (A) through a mechanism which still is not well understood. But although the need to mitigate rising unit costs creates the demand for innovation, the supply is constrained by the firm's current revenues, which determine the pool of resources available for research (B). The key implication is that quantity of R&D is determined by the equilibration of the forces (A) and (B). Innovation can be thought of as the process by which R&D is transformed into new knowledge (C). Together, these factors are responsible for endogenously shifting the IPC toward the origin.

Because it is the outcome of the accumulation process, technological change occurs only with a lag. In the analogue of diffusion or adoption, the envelope of substitution possibilities shifts in

[11] A technical detail is that once the production function at *t* is realized given prices and the state of knowledge, the remainder of the IPC disappears. Therefore, in Fig. 6 the price change does not cause the contemporaneous movement of *QQ* along IPC(*t*). Rather, the prices $P'P'$ determine a new point of tangency between the production function and the IPC at $t+1$.

[12] Jones (2005) uses the concept of the IPC to build up production functions from stochastic microfoundations.

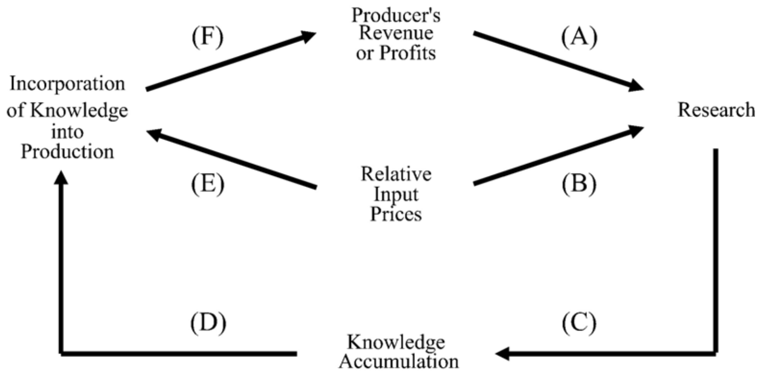

Fig. 7. A process-based model of ITC.

response to the increased supply of knowledge (D) and the effect of relative prices on the incorporation of knowledge into the production function (E).[13] These processes enhance productivity and profits (F) either by increasing the efficiency of production or facilitating the substitution of knowledge for relatively costly tangible inputs. By specifying the differential improvements in the efficiency with which C and D are employed, these factors implicitly define the curvature of the IPC.

In the present context it is useful to think of knowledge not as monolithic but as heterogeneous in character, differentiated in terms of the ideas, blueprints, etc. that improve the producer's efficiency of employing C versus that of using D. As we shall see, the implication is that all six of

[13] The deeper point here is the essential dependence of the input-saving consequences of an invention on relative prices, notwithstanding the inventor's intention to save on one input or another. Samuelson (1965: 355) captures the essence of the issue: "For the most part, labor saving innovation has a spurious attractiveness to economists because of a fortuitous verbal muddle. When writers list inventions, they find it easy to list labor-saving ones and exceedingly difficult to list capital-saving ones. (Cannan is much-quoted for his brilliance in being able to think up wireless as a capital-saving invention, the syllable 'less' apparently being a guarantee that it does in fact save capital!) That this is all fallacious becomes apparent when one examines a mathematical production function and tries to decide in advance whether a particular described invention changes the partial derivatives of the marginal-productivity imputation one way or another. Thus, consider a locomotive. It is big and heavy. So the literary mind thinks that it must correspond to a capital-using invention and hence to a labor-saving invention. Or think of a complex Rube Goldberg-like modern contraption. It is intricate and round about. So it must be regarded as labor-saving and capital-using. And yet there is not the slightest pretext for such inference. In the steady state, when human labor is organized through time with locomotives rather than without them, there is no way to tell in advance whether the relative share of labor in comparison with property has gone up or down in the steady state with production at all stages vertically integrated. We have the unfortunate tendency to use labor as the denominator in making productivity statements. Any invention, whether capital or labor saving, just by virtue of its definition as an invention rather than a disimprovement will, other things being equal, result in more output with the same labor or the same output with less labor. That could be said with any factor substituted for labor. But we know how difficult it is with a changing technology to get commensurable non-labor factors to put in the denominator of a productivity comparison. So we tend to concentrate on labor, and then we fall for the pun, or play on words, which infers a labor-saving invention whenever there is an invention. Thus, consider a simple case where output acts as if it were produced by a Cobb–Douglas [C–D] function with coefficients 3/4 and 1/4. Now let the locomotive, or the wheel, or fire, or the calculus be invented. Now can one have the least idea whether the function is merely increased in scale as against being twisted one way or the other in terms of its C–D coefficients? And when considers embodied technical change, and changing elasticities of substitution, as one must be prepared to do, how far from intuitive the problem becomes".

the foregoing factors play a role in determining the shape of the IPC. Armed with this understanding, we are now equipped to examine the methods by which technological change is represented in computational models for climate policy analysis.

4. Technological change in climate policy simulations

4.1. Exogenous input-augmenting technological change

In climate policy, the bulk of modeling analyses treat technological change as exogenous. The most common method of representing technical change is through parametric improvements in efficiency which alter the shape of the envelope of techniques. The typical device consists of augmentation coefficients, α, applied to each input, whose values grow over time:

$$Y(t) = Q[\alpha_C(t)C(t), \alpha_D(t)D(t)]. \tag{3}$$

Neutral efficiency improvement occurs when $\alpha_D(t)=\alpha_C(t)=\alpha(t)$. Provided the production function has "nice" properties,[14] Eq. (3) can be re-written $Y(t)=\alpha(t)\ Q[C(t), D(t)]$, where α now defines the overall rate of technical progress.

Biased technical progress occurs when $\alpha_D(t)\neq\alpha_C(t)$. In climate policy models this is used extensively to represent the empirically-observed phenomenon of expanding output accompanied by declining energy- and emissions intensity. If D denotes carbon-rich fossil fuels, the direction of technical progress is determined by the autonomous energy efficiency improvement (AEEI) parameter, which defines the degree to which the neutral rate of technological progress overstates the augmentation of fossil fuels, α'_D:

$$Y(t) = \alpha(t)[C(t), \alpha'_D(t)D(t)]. \tag{4}$$

Here, $\mathrm{AEEI}(t)=\dot{\alpha}'_D(t)/\alpha'_D(t)$, with $-\dot{\alpha}(t)/\alpha(t)<\mathrm{AEEI}\leq 0$.[15] In bottom-up simulations, the analogous procedure is to stipulate technology-specific neutral productivity parameters, α_m, whose values increase over time

$$Y_m(t) = \alpha_m(t)\min[C_m(t)/\bar{C}_m, D_m(t)/\bar{D}_m]\cdot \tag{5}$$

The rate of efficiency improvement is determined by demand decoupling factor (DDFs): $\mathrm{DDF}_m(t)=\dot{\alpha}(t)/\alpha_m(t)$. For biased technical progress synonymous with the AEEI the DDFs must generally increase with the ratio $\bar{C}_m/\bar{D}_m$, implying that improvements are concentrated in activities which are relatively clean.

[14] Technically, if $Q[\cdot]$ is homogeneous of degree one in its arguments (reflecting constancy of returns to scale in production and perfect competition in input and output markets) then α represents the rate of total input productivity, analogous to the Solow residual. The result is the shift in the unit isoquant illustrated in Fig. 2(C).

[15] A dot over a variable denotes a time-derivative. The first documented use of the AEEI is Edmonds and Reilly (1985), who cite the historical reduction in the energy intensity of GDP with increasing economic development as justification for a declining coefficient on energy input. They create a simulation model with an exogenously increasing index of energy-saving technology (the equivalent of α_D above) whose inverse is used to attenuate price-determined demands for fuels. This trick is routinely employed in climate policy models, whose AEEI values tend to cluster around 1%.

The limitations of this approach are by now well known.[16] First, the rate and direction of technical change are both pre-specified by the modeler and invariant to the effects of climate policy, with the result that substitution is the only mechanism through which climate mitigation measures can affect input demands. Second, technical progress is subsumed under the rubric of incremental improvement, leaving no room for radical technological change. We take up this issue next.

4.2. Backstop technologies: Semi-endogenous technological change

To represent radical technological change, we need a way to model the appearance of wholly new production techniques. In top-down simulations this is often an alternative production function, Q', which switches on in future periods in response to rising prices—a "backstop" technology.[17] These are often techniques which are forecast to become available after a certain future date, $\bar{t}_b$, and only actually produce output once their unit costs of production, p', become competitive with that of the conventional technology, p:

$$Y(t) = Q'[C(t), D(t)] \quad \text{if} \qquad p'(t) \le p(t) \text{ and } t \ge \bar{t}_b. \tag{6}$$

In bottom-up models the expansion of the set of operable technologies is explicitly represented, most often by designating the timing with which individual activities become available. The notation above can be used to capture this procedure for backstop technique b:

$$Y_b(t) = \min[C_b(t)/\bar{C}_b, D_b(t)/\bar{D}_b], \tag{7}$$

where $b \in A(t)$ if $t \ge \bar{t}_b$.

Eqs. (6) and (7) are referred to as semi-endogenous because the timing of the backstop's penetration is determined by the values of other variables for which the model solved. Nevertheless, given the impossibility of representing the creation of new techniques whose characteristics are not known at the start of the simulation and somehow emerge endogenously, the modeler must still make predictions about the key parameters $\bar{t}_b$, $\bar{C}_b$ and $\bar{D}_b$. The result is a sequence of available techniques $\{\bar{A}(0),\ldots,\bar{A}(T)\}$ which is inevitably exogenous, but which can have huge impacts on long-run mitigation costs. Because backstops are idealized representations of generic unprofitable or speculative energy technologies, it is inevitable that their parameters will be based in large part on engineering judgment. Consequently, the more radical the technology or the farther in the future it appears, the more uncertain its attributes.

Another issue is the nature of the transition between Q and Q'. It is plausible for these techniques to coexist for a number of periods, with the latter gaining market share at the expense of the former. In top-down models, as soon as $p' \le p$ the backstop usually takes over the entire market in what is known as "flip-flop" or "bang-bang" behavior. This dynamic stems from the

[16] In multi-sectoral economic models the AEEI is a short-hand approximation for not only energy-saving innovation, but also the shift in the composition of the economy toward activities that demand smaller quantities of fossil fuels, environmental policies restricting the use of fossil fuels, and the removal of "market barriers" to the diffusion of more energy-efficient technologies. Sue Wing and Eckaus (2005) assess the implications of this for the role of technological change in projecting baseline emissions.

[17] Coined by Nordhaus (1973b), the term backstop refers to a production process which becomes available only at high input prices, but can generate an unlimited quantity of output at constant marginal cost (a classic example is the plutonium breeder reactor). Within models, the practical effect is to cap the long-run rise in fossil-fuel prices due to their depletion or restrictions on their supply mandated by GHG emission reduction policies.

perfect substitutability of the outputs of the two production functions, which modelers remedy by imposing ad-hoc bounds on the backstop penetration rate, specifying the outputs of Q and Q' as imperfect substitutes (Popp, 2004b; Sue Wing, in press) or modeling the clean input as being in imperfectly elastic supply. With respect to the latter, Sue Wing (in press) focuses on capital as a component of C, modeling that input as activity-specific and imperfectly malleable. An additional device which is employed to regulate the penetration of the backstop is a fixed factor, Z_B, which replaces the dirty input:

$$Y(t) = Q'[C(t), Z_B(t)] \tag{8}$$

Z_B is a proxy for non-price "market barriers" such as technical standards or the lack of complementary infrastructure whose evolution proceeds exogenously (McFarland et al., 2004; Jacoby et al., this volume), or for the scarcity of technical know-how necessary to operate Q competitively—which may be alleviated by investment in R&D (Popp, 2004b).

Similar issues arise in bottom-up simulations, as Section 4.4 elaborates. But even though formulations such as Eq. (7) are a mainstay of energy system models, their impact on the simulation results is difficult to judge because of the multiplicity of activities represented therein, each of whose attributes—especially $\overline{t}_b$—tend to be poorly documented.[18] Furthermore, the results reflect the trajectory of activities which are operated, $O(t)$, from which it is not possible to reconstruct the evolution of the envelope of available techniques, $A(t)$. All this points to the need for bottom-up studies which explicitly track the time paths of the sets O and A, and apportion the changes in emissions and system costs between shifts of the former (i.e., substitution) and the expansion of the latter (i.e., technological change).

4.3. Price-induced input augmentation

The most direct way of appending inducement to the framework of Eq. (4) is to formalize the Ahmad–Binswanger–Ruttan model by specifying the augmentation coefficients as functions of the relative prices of the inputs. What is required is a relation, ψ, which is the analogue of the IPC, and whose function is to translate shifts in relative prices into changes in the values of the augmentation coefficients:

$$\dot{\alpha}_i(t) = \psi[p_C(t), p_D(t)] \qquad i = \{C, D\}. \tag{9}$$

Formulations of this kind are rarely used.[19] The reasons are the lack of empirical information about the shape of ψ, the difficulty of specifying a function which can simultaneously capture both complementarities and tradeoffs among the augmentation of different inputs (especially when the number of inputs is large), and the challenge of representing the potential for small price changes to stimulate technological change while simultaneously allowing large increases to stifle innovation.

[18] For example, IIASA's Carbon Dioxide Technology (CO_2DB) Database contains detailed technical, economic and environmental characteristics of 3000-odd energy technologies, but the basis of these data in the peer reviewed literature is unclear.

[19] Dowlatabadi (1998) and Dowlatabadi and Oravetz (1996, in press) specify ψ as a function of per capita GDP and current and lagged energy prices. Jakeman et al. (2004) model the aggregate amount of innovation as exogenous but endogenously allocate the resulting pool of "technological change" among industries according to their unit costs.

4.4. Learning by doing: a critical assessment

Perhaps the most popular method of endogenizing technological change is learning by doing (LBD), which is based on Arrow's (1962) model of unit cost decline driven by producers' accumulation of knowledge via their experience with the use of particular techniques of production. This approach is favored by bottom-up models, which employ cumulative capacity—or production, as the utilization decision is not explicitly represented—as the proxy for experience.[20] In a given period, each activity's demand decoupling factor is specified as an increasing function of the history of its output:

$$\alpha_m(t) = \psi[V_m(t); \lambda_m], \tag{10}$$

where $V_m(t) = \int_0^t Y_m(s)\mathrm{d}s$ is an index of cumulative production at t, and ψ is an experience function (ψ, $\psi' \geq 0$, $\psi'' < 0$). The key parameter is the learning rate, λ_m, which determines the increase in productivity, or symmetrically the reduction in unit costs, as a result from a doubling of V. A typical formulation is $\psi = V_m^{\zeta m}$, in which the learning exponent $\zeta_m = \log(1+\lambda_m)/\log 2$.[21]

Under this approach ITC is the stimulation by climate policies of aggressive near-term investment in high-cost, low-carbon energy technologies, which then enjoy cost reductions due to LBD (e.g., Grubler and Messner, 1998: Fig. 4). The upshot is uniformly greater mitigation effort early in the simulation horizon, and often large reductions in overall policy costs relative to the no-learning baseline. This outcome is the result of perfect-foresight models' balancing the discounted future outlays from switching away from conventional energy technologies against the discounted future savings in energy supply costs due to learning in alternative technologies. Subsidizing capacity addition in initially-unprofitable alternative technologies permits the accumulation of experience with their operation, inducing declines in their unit costs which enable them to more quickly compete with conventional technologies, and yield the societal benefits of the subsequent stream of savings on energy costs. Models' implicit subsidization of the initial unit cost differential (especially between fossil fuels and renewables) is thus rationalized as a "learning investment" in generating these future benefits (Wene, 2000).

Although widely used, the LBD approach suffers from a number of important limitations. First and foremost, introducing LBD makes the problem of minimizing total energy system cost non-convex, giving rise to multiple equilibria and attendant instability of models' numerical solution.[22] The positive feedback of cost reductions on output and further experience gains in Eq. (10) predisposes learning technologies to become dominant only a few periods after their introduction, which can cause intertemporal bottom-up models to exhibit implausible market share dynamics. As with backstop technologies, the common remedy is to include penetration constrains on technologies which enjoy LBD (e.g., upper bounds on capacity or investment rates), but their effect is to render the trajectories of cost reductions exogenous (Loulou et al., 2004: 71–

[20] "The first question is that of choosing the economic variable which represents 'experience'. The examples given above suggest the possibility of using cumulative output (the total of output from the beginning of time) as an index of experience, but this does not seem entirely satisfactory. If the rate of output is constant then the stimulus to learning presented would appear to be constant, and the learning that does take place is a gradual approach to equilibrium behavior". (Arrow, 1962: 155–156).

[21] Grubb et al. (1995) and Goulder and Matthai (2000) employ a different formulation in which cost reductions due to learning are a function of cumulative *abatement*. In top-down framework of Eq. (4) this is equivalent to an acceleration of the AEEI with reduction in the use of the dirty input: AEEI(t) = $-\psi[D(t)]$.

[22] Manne and Barreto (2004) demonstrate multiplicity of equilibria in a simple three-technology model with LBD.

72). However, the main problems with these antidotes are their tenuous theoretical and empirical bases, and particularly their tendency to be poorly documented.[23] In large bottom-up models, the sheer number of activities—each with a different learning rate, whose cost reductions may even be amplified by inter-technology spillovers due to “learning clusters” (e.g., Gritsevskyi and Nakicenovic, 2000)[24]—only exacerbates this lack of transparency.

A second issue is the critical assumption shared by virtually all LBD studies that carbon-free technologies experience the most rapid learning and cost reductions, while their conventional counterparts enjoy little or no improvement. Renewables then often grow to satisfy a large share of total energy demand (Mattson and Wene, 1997; van der Zwaan et al., 2002), sometimes even without the imposition of climate mitigation measures (Chakravorty et al., 1997). Not only is this outcome quite speculative given our limited understanding of the association between unit cost reductions and the diffusion of new technologies, it is virtually a pre-determined the outcome of the simulation.[25] Not only is the potential for activities to learn fundamentally exogenous, even modest learning potentials predestine them to rapidly increase their market share once they become operational. The simulated producer’s role is then relegated to allocating capacity accordingly.

The third shortcoming is the phenomenological, heuristic character of LBD itself. That cost reductions flow automatically from capacity additions in new technologies is not just mechanistic, it implies that innovation is a costless by-product of clean production, rather than the outcome of deliberate, costly investments in research. Thus, apart from shedding light on the production subsidies justified by the influence of learning potentials on the timing and technological characteristics of responses to emission or concentration targets,[26] the LBD approach is unsatisfying because it yields little insight into the forces which drive ITC, even as it holds out the possibility of a technological “free lunch”.

[23] Consistent with Arrow’s intuition (cf. footnote 18), technologies do not learn forever but achieve a steady state once they gain sizeable market share (Nakicenovic, 2002), which implies an upper bound on α_m given by its cost reduction potential—the difference between the cost of a unit of capacity of m at $t=0$ and some ultimate “floor” cost at $t=\infty$. The problem is that the latter is unknown and is often arbitrarily assumed. According to Grubler et al. (1999): “[If one] assumes that the cost of the incremental technology falls at an exogenously determined rate [then] the result is typical for simple optimization models: a new technology diffuses rapidly and widely at the moment it becomes cheaper than alternatives. Indeed, large-scale technology optimization models, which are widely used to assess the costs of abating various environmental problems, display similar ‘flip-flop’ behavior. Published runs typically do not illustrate such behavior only because additional constraints, such as restrictions on the rate and pattern of technological diffusion, tuned according to the modeler’s sense of plausibility, are widely used to make the outputs appear more realistic. Like sausage, the final product is evidently tasty, but the method of producing it is best left shrouded in mystery”.

[24] A cluster is a subset of available technologies x each of whose unit costs depends on the *sum* of the cumulative capacities of all the technologies in the cluster. Thus for any $m \in x$, $V_m(t) = \int_0^t \sum_{m \in X} Y_m(s) \mathrm{d}s$.

[25] Chakravorty et al’s result seems far-fetched, given renewable energy technologies’ low capacity factors and the abundance of low-cost hydrocarbon resources (e.g., Gregory and Rogner, 1998). The use of a technology’s cumulative capacity as a sufficient statistic for the impact of LBD conflates the accumulation of experience with economies of scale and the embodiment of technological progress in the firm’s capital equipment. An additional complication is that time series of the costs of new energy technologies are barely long enough to statistically distinguish these factors. Isoard and Soria (2001) is one of the few studies to rigorously document such estimates—and their shortcomings—in the peer-reviewed literature (see, e.g., McDonald and Schrattenholzer, 2001: Table 2). The consequent dearth of information on learning rates in energy technologies has led modelers to adopt progress ratios drawn from a range of non-energy manufactured goods (e.g., Dutton and Thomas, 1984).

[26] See, e.g., van der Zwaan et al. (2002), Gerlagh and van der Zwaan (2003) and Gerlagh et al. (2004). These studies’ strength is their disaggregation of energy production into only two activities (generic fossil and non-fossil technologies), which makes their analysis transparent and simplifies interpretation of their simulation results.

Attempts to address this limitation have focused on the complementarities between costly R&D and cost reductions due to learning.[27] Kouvaritakis et al. (2000) incorporate this effect by including cumulative research, Ω_m, as an additional argument in the experience function:

$$\psi = V_m(t)^{\zeta_m}\Omega_m(t)^{\mu_m}, \tag{11}$$

where $\Omega_m(t) = \int_0^t R_m(s)\mathrm{d}s$, μ_m indicates the strength of the effect of R&D on learning, and m-specific R&D, R_m, is modeled as an additional control variable. In so far as this formulation amplifies the cost-reducing effect of cumulative production it is likely to exacerbate the non-convexity described above. Moreover, the assumption that renewables enjoy larger values of μ than fossil fuels can bias the simulation toward a carbon-free future to an even greater extent. Nevertheless, Eq. (11) is important because it acknowledges that the reconfiguration of the production function is a direct outcome of the application of new knowledge—for which experience, and in turn, cumulative capacity, are merely proxies, and moreover, which is costly to accumulate.[28] It is therefore a gateway to the explicit representation of the processes of technological change.

4.5. Process-based models of ITC: The stock of knowledge approach

This approach is the algebraic analogue of Fig. 7, but with knowledge represented as a differentiated asset. Its centerpiece is the stock of knowledge, H, whose accumulation is determined by obsolescence (which is captured by the rate of depreciation parameter δ) and investments in research, R, following the perpetual inventory framework developed in the new economic growth literature:

$$\dot{H}_i(t) = \psi[R_i(t), H_i(t)] - \delta H_i(t) \qquad i = \{C, D\}. \tag{12}$$

Here, ψ is a transformation function which represents the innovation process. A generic formulation is:

$$\psi = \beta_i R_i(t)^{\theta_i} H_i(t)^{\xi_i}\{\sum_i R_i(t)\}^{\omega_i}, \tag{13}$$

in which β denotes the efficiency of innovation, $\theta < 1$ indicates contemporaneous diminishing returns to R&D, and $\xi < 1$ indicates diminishing returns to intertemporal knowledge spillovers.[29] If $\omega \neq 0$ there is an additional *contemporaneous* spillover effect due to complementarities among different lines of research, analogous to the learning clusters in LBD models described in footnote 24.

The stock of knowledge approach is the almost exclusive preserve of top-down models, in which ITC encompasses three distinct processes: the price effects of climate policy's stimulation of additional dirty-input saving research, the consequent rise in the rate of accumulation of D-saving knowledge and expansion of substitution possibilities for D, and the resulting increase in

[27] See, e.g., Lieberman (1984); Cohen and Levinthal (1989).

[28] Arrow himself clearly notes this in the introduction to his seminal paper (Arrow, 1962: 155–156): "[The] empirical studies of Abramovitz and Solow [...] do not directly contradict the neo-classical view of the production function as an expression of technological knowledge. *All that has to be added is the obvious fact that knowledge is growing in time.* [...] *The acquisition of knowledge is usually what is termed 'learning'*, and [...] I do not think that the picture of technical change as a vast and prolonged process of learning [...] is in any way a far-fetched analogy". The italics are my own.

[29] See Rivera-Batiz and Romer (1991), Jones (1995), Jones and Williams (2000) and Popp (2004a). The parameter θ proxies for duplication externalities in research ("stepping on toes"), while ξ indicates the intensity of creative destruction ("standing on shoulders"). The necessary condition for an intertemporally convex optimization problem is $\theta+\xi+\omega \leq 1$.

the substitution of H_D for D as a result of relative prices.[30] ITC's overall impact tends to be modest, and turns on reductions the welfare cost of mitigation through the inducement of near-term emission-saving R&D, as opposed to early abatement. Thus, as suggested by the discussion in Section 3.3, Eq. (12) is the intertemporal conduit for the real action in models of ITC, namely, the influence of relative prices on R&D—which can be thought of as an "accumulation effect", and the joint impact of relative prices and the expansion of the supply knowledge on production costs—which can be thought of as a "substitution effect".

Regarding the latter, the present framework provides an opportunity to endogenize technological change in the models of earlier sections. In the input augmentation model in Eq. (3) this is accomplished by specifying each coefficient as a function of its respective knowledge stock: $\alpha_i(t)=f_i[H_i(t)]$, with f, $f'>0$, $f''<0$. Alternatively, C- and D-saving knowledge may be specified as additional inputs to the production correspondence:

$$Y(t) = Q[g_C\{C(t), H_C(t)\}; g_D\{D(t), H_D(t)\}], \tag{14}$$

where the g_is are nested production functions which define the substitutability between each type of knowledge and its associated input.[31] These methods capture incremental technological change, but radical change can be accommodated as well. The backstop model in Eq. (8) may be endogenized by specifying a stock of backstop-augmenting knowledge, H_B, as an intangible fixed factor (e.g., $Z_B=H_B$).[32]

The foregoing devices explicitly model the mechanism via which prices and knowledge supplies affect the rate and direction of technical change. As discussed in Section 3.3, at each point in time the IPC is implicitly defined by two factors: the difference in the size of the knowledge stocks, and either the knowledge elasticity of augmentation (the curvature of f), or the elasticity of substitution between each kind of knowledge and its associated tangible input (the curvature of g). But although these have the potential to tremendously influence mitigation costs (Sue Wing, 2003), the dearth of empirical estimates on the ease of incorporating knowledge into production precludes empirical validation of models' calculations of the benefits of ITC.[33]

Turning now to the accumulation effect, the signal advantage of the stock of knowledge approach is that the engine of technological change, i.e., research, is fundamentally endogenous. In perfect-foresight models, R&D is a control variable whose trajectory is implicitly determined by the

[30] The main reason appears to be the lack of disaggregate data on R&D at the level of individual energy technologies, which militates against the calibration of initial knowledge stocks. This is a problem even in the case of Eq. (12). Although time series of R&D in published statistics may be accumulated into "R&D capital", these data are often not disaggregate enough to permit stocks of C- and D-augmenting knowledge to be constructed. The practical consequence is that in climate policy simulations the starting value for H is arrived at by a variety of ad-hoc means, or calibrated along with other unknown parameters. For different approaches see Goulder and Schneider (1999), Nordhaus (2002), Sue Wing (2003) and Popp (2004a).

[31] e.g., Popp (2002, 2004a,b,c), in which g_D as a CES composite of carbon-energy and energy-saving knowledge.

[32] This formulation, due to Popp (2004b), is $Y(t)=Q[C(t); g\{D(t), B(t)\}]$, where B is a backstop input to production which is imperfectly substitute for the dirty good (fossil energy), and H_B increases elasticity of backstop supply according to $p_B=\eta[H_B]$, $\eta'<0$, $\eta''>0$. Popp finds that the rate of accumulation of H_B exerts a much larger effect on the long-run penetration of B than the curvature of η.

[33] Popp (2001) is the sole study which estimates the elasticities of unit energy demand with respect to the stocks of energy-saving knowledge in different industries. Comparable estimates of the elasticity of substitution between knowledge and tangible inputs do not exist, even at the level of aggregate economies. The cost implications can be understood by considering the case where the gs are Leontief—here the trajectories of H_C and H_D completely determine the difference ITC makes for the C–D substitution response to a given relative price change.

optimality condition which equates its marginal cost with the present value of its future input cost savings (Popp, 2004a). In recursive–dynamic simulations, the lack of intertemporal linkages forces modelers to specify R&D supply and demand using functions of current-period variables which are for the most part ad-hoc. Sue Wing (2003) models the latter conceptually through the device of an *inducement function*, ρ. This is a reduced-form representation of the effect on the equilibrium in R&D markets of input prices (which create the demand for innovation) and the producer's sales (a proxy for the revenue which determines the resources for the supply of R&D):

$$R_i(t) = \rho[p_i(t), Y(t)]. \tag{15}$$

Despite its shortcomings, this expression captures the balance of influences (A) and (B) in Fig. 7. This may be seen by assuming a downward-sloping demand curve for the output good and differentiating with respect to p_i using the chain-rule:

$$\frac{\partial R_i}{\partial p_i} = \underset{(+)}{\frac{\partial \rho}{\partial p_i}} + \underset{(+)}{\frac{\partial \rho}{\partial Y}}\,\underset{(-)}{\frac{\partial Y}{\partial p}}\,\underset{(+)}{\frac{\partial p}{\partial p_i}}. \tag{16}$$

The result illustrates that net impact of an increase in p_i on R&D is the sum of the effects shown by the two terms in this expression: the direct effect of Hicksian induced i-saving innovation, and the indirect adverse effect of p_i on unit costs, p, and the consequent reduction in sales, revenue, and the pool of resources available to finance research. Thus, if the demand for output is sufficiently elastic, with a large enough increase in p_D the indirect effect may dominate, *reducing* the time-path of R_D relative to its the counterfactual no-policy baseline.

The foregoing arguments suggest that ITC should not be thought of as a process by which climate policy increases R&D *in general*. Resources for R&D are supplied inelastically (Goolsbee, 1998), which means that more rapid accumulation of knowledge in one area of the economy is likely to reduce the rate of accumulation in other areas (Goulder and Schneider, 1999; Popp, 2004a,c). The consequent "crowding out" effect is an additional precursor of the IPC which is rooted in the accumulation process, and implies the existence of a fundamental tradeoff between R_D and R_C, or more starkly, between "economy-growing" and "emission-reducing" innovation.[34] This result, which stands in contrast to the LBD approach, seems to be unaffected by whether knowledge is modeled as homogeneous (Sue Wing, 2003) or heterogeneous and accumulating differentially (Goulder and Schneider, 1999): in both of these multi-sectoral simulations, a carbon tax can induce progressive *reductions* in R&D relative to its baseline level, not only within individual industries but also at the aggregate level.[35]

[34] e.g., in the sense of the knowledge stocks which drive productivity (α) and pollution augmentation (α_D^j) in Eq. (4). Not surprisingly, the effect of ITC on the simulation results is sensitive to the specification of crowding out. Nordhaus (2002) assumes complete crowding out whereby each additional dollar of R&D displaces an equal value of investment in physical capital, and finds that knowledge accumulation has very little effect on the economic impacts of climate policy or the timing of mitigation measures. By contrast, Buonanno et al. (2003) adopt the less plausible assumption of crowding *in* whereby a dollar of investment in new knowledge generates both neutral productivity growth and reductions in abatement costs. Paradoxically, ITC has the effect of *increasing* the total dollar cost of stabilizing GHG emissions in their model. Popp (2004a) models each dollar of R&D as crowding out half as much investment, but because of the higher social return to R&D (e.g., Mansfield, 1996; Jones and Williams, 1998) the effective reduction in investment is four times as large. Popp finds that a GHG emission limit significantly increases energy-saving R&D relative to the counterfactual, but does not report the corresponding reduction in the trajectory of the physical capital stock.

[35] The reductions in R&D (and, in Goulder and Schneider's model, industries' knowledge stocks) are concentrated in fossil-fuel production sectors where the incidence of the tax is highest.

Fortuitously, the impact of crowding out appears to be small. It is also partially alleviated by assuming the existence of spillovers, or by specifying R&D subsidy policies which correct distortions in the market for innovation.[36] Focusing on the first element, the main caveat is that assumptions about the locus of spillovers matter crucially for model behavior. Modeling *C*-augmenting knowledge as benefiting disproportionately from spillovers ($\omega_C > \omega_D$) reduces the costs and increases the attractiveness of near-term abatement, because the spillover effect compensates for the diminished accumulation of *C*-saving knowledge which as a result of crowding out by the inducement of *D*-saving R&D. Similar reasoning implies that costs are amplified in the reverse case (Goulder and Schneider, 1999). The upshot is that, as with the substitution effect, the importance of spillovers for knowledge accumulation is belied by a lack of empirical estimates.

We close on this note by indicating that a consistent theme underlying the discussion above is the challenge posed by the essential unobservability of knowledge itself. The lack of empirically-validated proxies for *C*-and *D*-augmenting knowledge diminishes models' transparency and the comparability of their results (e.g., footnote 30), and inhibits evaluation of the robustness of their insights into the influence of ITC. While the focus here has been on structural formulations of the inducement process, it is only the development of data along the lines of, e.g., Popp (2002), and estimation of key parameters which govern the incorporation of knowledge in production (e.g., Bernstein and Nadiri, 1997; Popp, 2001) which provide the kind of empirical guidance that can enable the stock of knowledge approach to fulfill its great potential.

5. Summary and conclusions

This paper has given an overview of the methods of representing technological change in the computational simulations for climate policy analysis. A simple conceptual framework was developed to compare and contrast the different representations of technology in bottom-up and top-down climate policy simulations. The framework was then employed to reconcile the bottom-up engineering conception of "microscale technological change" with the top-down macroeconomic conception of substitution, and to deductively formulate a process-based representation of ITC along the lines of the Ahmad–Binswanger–Ruttan model. Finally, the structural foundations of the methods of representing technological change within climate policy models were described.

While the learning by doing approach to modeling ITC generally results in significantly lower near-term emissions and dramatic reductions in the cost of emission reductions, the alternative stock-of knowledge approach tends to produce more modest cost savings and has little impact on the timing of abatement. Though it may be tempting to interpret this dichotomy in terms of the debate over adoption of "act-now" versus "wait-and-see" climate policies, it is important to realize that *both* of the approaches to modeling ITC appear to indicate early action. The essential difference is qualitative: in the LBD approach, ITC is synonymous with faster micro-scale shifts to low-carbon production activities which induce cost reductions themselves—reflecting the guiding assumption of *complementarity* between abatement and the overall efficiency of production. By contrast, ITC in

[36] Spillovers can be purely intertemporal in nature—the sole option for models with a single stock of knowledge (e.g., Popp, 2004a), or be contemporaneous in models with multiple intangible stocks—occurring across industries (Goulder and Schneider), regions (Buanonno et al) or types of knowledge (Eq. (13)). Popp (2002) uses patent data to estimate the strength of intertemporal spillovers in energy technologies. Although the productivity impacts of inter-spillovers have been studied (Bernstein and Nadiri, 1997), the influence of potential complementarities among related lines of energy research have not been empirically investigated. Distortions include pre-existing taxes on the inputs to R&D (Goulder and Schneider, 1999; Sue Wing, 2003) and the under-investment in R&D relative to the social optimum as a consequence of imperfect appropriability (Popp, 2004c).

the stock of knowledge approach is associated with faster accumulation of emission-saving knowledge—which, despite reducing emissions and improving productivity, incurs the opportunity costs of foregone "economy-growing" R&D-reflecting the guiding assumption of *substitutability* between abatement and aggregate efficiency.

This divergence, as well as the advantages and deficiencies of both methods examined in the paper, suggest the critical importance of investigating the empirical content of the assumptions which underlie the representation of ITC in climate policy models. Although it is by no means a simple undertaking, such a program of research is the modeling community's quickest route toward much-needed validation and reconciliation of its disparate efforts to address the climate problem.

Acknowledgements

This paper has benefited enormously from the discussions with Henry Jacoby, John Reilly, Karen Fisher-Vanden and David Popp. Errors and omissions are the sole responsibility of the author. This work was supported by the U.S. Department of Energy Office of Science (BER) Grant No. DE-FG02-02ER63484.

References

Acemoglu, D., 2002. Directed technical change. Review of Economic Studies 69, 781–809.
Ahmad, S., 1966. On the theory of induced invention. Economic Journal 76 (302), 344–357.
Arrow, K.J., 1962. The economic implications of learning by doing. Review of Economic Studies XXIX, 155–173.
Ashford, N.A., 1994. An Innovation-Based Strategy for the Environment. In: Finkel, A.M., Goulding, D. (Eds.), Worst Things First? The Debate Over Risk-Based National Environmental Priorities. Resources for the Future, Washington DC, pp. 275–314.
Ashford, N.A., Ayers, C., Stone, R.F., 1985. Using regulation to change the market for innovation. Harvard Environmental Law Review 9, 419–466.
Azar, C., Dowlatabadi, H., 1999. A review of technical change in assessment of climate policy. Annual Review of Energy and the Environment 24, 513–544.
Bernstein, J.I., Nadiri M.I., 1997.Product Demand, Production Cost, and Interindustry R&D Spillovers, mimeo, New York University, <http://www.econ.nyu.edu/user/nadiri/pub98.PDF>.
Binswanger, H.P., Ruttan, V.W., 1978. Induced Innovation: Technology, Institutions and Development. Johns Hopkins University Press, Baltimore.
Buonanno, P., Carraro, C., Galeotti, M., 2003. Endogenous induced technical change and the costs of Kyoto. Resource and Energy Economics 25, 11–34.
Chakravorty, U., Roumasset, J., Tse, K., 1997. Endogenous substitution among energy resources and global warming. Journal of Political Economy 105, 1201–1234.
Cohen, W.M., Levinthal, D.A., 1989. Innovation and learning: the two faces of R&D. Economic Journal 99, 569–596.
Dawkins, C., Srinivasan, T.N., Whalley, J., 2001. Calibration. In: Heckman, J.J., Leamer, E. (Eds.), Handbook of Econometrics, vol. 5. Elsevier, Amsterdam, pp. 3653–3703.
Diamond, P., McFadden, D., Rodriguez, M., 1978. Measurement of the Elasticity of Factor Substitution and Bias of Technical Change. In: Fuss, M., McFadden, D. (Eds.), Production Economics: A Dual Approach to Theory and Applications, vol. II. North-Holland, Amsterdam, pp. 125–147.
Dowlatabadi, H., 1998. Sensitivity of climate change mitigation estimates to assumptions about technical change. Energy Economics 20, 473–493.
Dowlatabadi, H., Oravetz, M.A., 1996. Understanding Trends in Energy Intensity: A Simple Model of Technical Change, mimeo, Center for Integrated Study of the Human Dimensions of Global Change, Carnegie Mellon University.
Dowlatabadi, H., Oravetz, M.A., in press. U.S. Long-Term Energy Intensity: Backcast and Projection, Energy Policy.
Dutton, J.M., Thomas, A., 1984. Treating progress functions as a managerial opportunity. Academy of Management Review 9, 235–247.
Edmonds, J., Reilly, J., 1985. Global Energy: Assessing the Future. Oxford University Press, New York.

Gerlagh, R., van der Zwaan, B.C.C., 2003. Gross world product and consumption in a global warming model with endogenous technological change. Resource and Energy Economics 25, 35–57.
Gerlagh, R., van der Zwaan, B.C.C., Hofkes, M.W., Klaassen, G., 2004. Impacts of CO_2 taxes in an economy with niche markets and learning-by-doing. Environmental and Resource Economics 24, 367–394.
Goolsbee, A., 1998. Does government R&D policy mainly benefit scientists and engineers? American Economic Review 88 (2), 298–302.
Goulder, L.H., 2004. Induced Technological Change and Climate Policy, Pew Center on Global Climate Change Report, Washington DC.
Goulder, L.H., Matthai, K., 2000. Optimal CO_2 abatement in the presence of induced technological change. Journal of Environmental Economics and Management 39, 1–38.
Goulder, L.H., Schneider, S.H., 1999. Induced technological change and the attractiveness of CO_2 abatement policies. Resource and Energy Economics 21, 211–253.
Gregory, K., Rogner, H.-H., 1998. Energy resources and conversion technologies for the 21st century. Mitigation and Adaptation Strategies for Global Change 3, 171–230.
Gritsevskyi, A., Nakicenovic, N., 2000. Modeling uncertainty of induced technological change. Energy Policy 28, 907–921.
Grubb, M., Chapuis, T., Duong, M.H., 1995. The economics of changing course: implications of adaptability and inertia for optimal climate policy. Energy Policy 23, 417–431.
Grubb, M., Köhler, J., Anderson, D., 2002. Induced technical change in energy and environmental modeling: analytic approaches and policy implications. Annual Review of Energy and the Environment 27, 271–308.
Grubler, A., Messner, S., 1998. Technological change and the timing of mitigation measures. Energy Economics 20, 495–512.
Grübler, A., Nakićenović, N., Victor, D.G., 1999. Modeling technological change: implications for the global environment. Annual Review of Energy and the Environment 24, 545–569.
Hicks, J., 1932. The Theory of Wages. Macmillan, London.
Isoard, S., Soria, A., 2001. Technical change dynamics: evidence from the emerging renewable energy technologies. Energy Economics 23, 619–636.
Jacoby, H.D., Reilly, J.M., McFarland, J.R., this volume. Technology and Technical Change in the MIT EPPA Model.
Jakeman, G., Hanslow, K., Hinchy, M., Fisher, B.S., Woffenden, K., 2004. Induced innovations and climate change policy. Energy Economics 26, 937–960.
Jones, C.I., 1995. R&D-based models of economic growth. Journal of Political Economy 103, 759–784.
Jones, C.I., 2005. The shape of production functions and the direction of technical change. Quarterly Journal of Economics 120, 517–549.
Jones, C.I., Williams, J.C., 1998. Measuring the social return to R&D. Quarterly Journal of Economics 113, 1119–1135.
Jones, C.I., Williams, J.C., 2000. Too much of a good thing? The economics of investment in R&D. Journal of Economic Growth 5, 65–85.
Jorgenson, D.W., 1984. Econometric Methods for Applied General Equilibrium Analysis. In: Scarf, H.E., Shoven, J.B. (Eds.), Applied General Equilibrium Analysis. Cambridge University, Cambridge, pp. 139–203.
Kennedy, C., 1964. Induced bias in innovation and the theory of distribution. Economic Journal 74, 541–547.
Kouvaritakis, N., Soria, A., Isoard, S., Thonet, C., 2000. Endogenous learning in world post-Kyoto scenarios: application of the POLES model under adaptive expectations. International Journal of Global Energy Issues 14, 222–248.
Lieberman, M.B., 1984. The learning curve and pricing in the chemical processing industries. RAND Journal of Economics 15 (2), 213–228.
Löschel, A., 2002. Technological change in economic models of environmental policy: a survey. Ecological Economics 43, 105–126.
Löschel, A., 2004. Technologies, energy consumption and environmental policy in energy–economy models. International Journal of Energy Technology and Policy 2 (3), 250–261.
Loulou, R., Goldstein, G., Noble, K., 2004. Documentation for the MARKAL Family of Models,<http://www.etsap.org/documentation.asp>.
Manne, A.S., Barreto, L., 2004. Learn-by-doing and carbon dioxide abatement. Energy Economics 26, 621–633.
Manne, A.S., Richels, R.G., 1997. On stabilizing CO_2 concentrations: cost-effective emission reduction strategies. Environmental Modeling and Assessment 2, 251–265.
Manne, A.S., Richels, R.G., Mendelsohn, R., 1995. MERGE: a model for evaluating regional and global effects of GHG reduction policies. Energy Policy 23 (1), 17–34.
Mansfield, E., 1996. Microeconomic policy and technological change. In: Fuhrer, J.C., Sneddon Little, J. (Eds.), Technology and Growth: Federal Reserve Bank of Boston Conference Series No. 40, pp. 183–200.
Mattson, N., Wene, C.-O., 1997. Assessing new energy technologies using an energy system model with endogenized experience curves. International Journal of Energy Research 21, 385–393.

McDonald, A., Schrattenholzer, L., 2001. Learning rates for energy technologies. Energy Policy 29, 255–261.
McFarland, J.R., Reilly, J.M., Herzog, H.J., 2004. Representing energy technologies in top-down economic models using bottom-up information. Energy Economics 26, 685–707.
McKitrick, R.R., 1998. The econometric critique of applied general equilibrium modelling: the role of functional forms. Economic Modeling 15, 543–573.
Nakicenovic, N., 2002. Technological Change and Diffusion as a Learning Process. In: Grübler, A., Nakicenovic, N., Nordhaus, W.D. (Eds.), Technological Change and the Environment. RFF Press, Washington DC, pp. 160–181.
Nerlove, M., 1967. Recent Empirical Studies of the CES and Related Production Functions. In: Brown, M. (Ed.), The Theory and Empirical Analysis of Production. National Bureau of Economic Research, New York, pp. 55–122.
Nordhaus, W.D., 1973a. Some skeptical thoughts on the theory of induced innovation. Quarterly Journal of Economics 87, 208–219.
Nordhaus, W.D., 1973b. The allocation of energy resources. Brookings Papers on Economic Activity 3, 529–570.
Nordhaus, W.D., 2002. Modeling induced innovation in climate-change policy. In: Grubler, A., Nakicenovic, N., Nordhaus, W.D. (Eds.), Technological Change and the Environment. Resources for the Future, Washington, DC, pp. 182–209.
Nordhaus, W.D., Boyer, J., 2000. Warming the World: Economic Models of Global Warming. MIT Press, Cambridge, MA.
Palmer, K., Oates, W.E., Portney, P.R., 1995. Tightening environmental standards: the benefit–cost or the no–cost paradigm? Journal of Economic Perspectives 9 (4), 119–132.
Paltsev, S., Reilly, J.M., Jacoby, H.D., Eckaus, R.S., McFarland, J., Sarofim, M., Asadoorian, M., Babiker, M., 2005. The MIT Emissions Prediction and Policy Analysis (EPPA) Model: Version 4. MIT Joint Program on the Science and Policy of Global Change Report, vol. 125.
Popp, D., 2001. The effect of new technology on energy consumption. Resource and Energy Economics 23, 215–239.
Popp, D., 2002. Induced innovation and energy prices. American Economic Review 92, 160–180.
Popp, D., 2004a. ENTICE: endogenous technological change in the DICE model of global warming. Journal of Environmental Economics and Management 48, 742–768.
Popp, D., 2004b. ENTICE-BR: Backstop Technology in the ENTICE Model of Climate Change. NBER Working Paper 10285.
Popp, D., 2004c. R&D Subsidies and Climate Policy: Is there a "Free Lunch". NBER Working Paper 10880.
Popp, D., 2005. Uncertain R&D and the Porter Hypothesis. Contributions to Economic Analysis and Policy 4 (1: Art. 6), <http://www.bepress.com/bejeap/contributions/vol4/iss1/art6/>.
Porter, M.E., van der Linde, C., 1995. Toward a new conception of the environment-competitiveness relationship. Journal of Economic Perspectives 9, 97–118.
Rivera-Batiz, L.A., Romer, P.M., 1991. Economic integration and endogenous growth. Quarterly Journal of Economics 106, 531–555.
Samuelson, P.A., 1965. A theory of innovation along Kennedy–Weisacker lines. Review of Economics and Statistics 47, 343–356.
Smulders, S., de Nooij, M., 2003. The impact of energy conversation on technology and economic growth. Resource and Energy Economics 25, 59–79.
Sue Wing, I., 2003. Induced Technical Change and the Cost of Climate Policy. MIT Joint Program on the Science and Policy of Global Change Report, vol. 102.
Sue Wing, I., in press. The synthesis of bottom-up and top-down approaches to climate policy modeling: electric power technologies and the cost of limiting US CO_2 emissions. Energy Policy.
Sue Wing, I., Eckaus, R.S., 2004. Explaining Long-Run Changes in the Energy Intensity of the U.S. Economy. MIT Joint Program on the Science and Policy of Global Change Report, vol. 116.
Sue Wing, I. Eckaus, R.S., 2005. The Decline in U.S. Energy Intensity: Its Origins and Implications for Long-Run CO_2 Emission Projections, paper presented at the NCCR Climate Symposium 2005, Interlaken, Switzerland, Mar. 3–5.
Thirtle, C.G., Ruttan, V.W., 1987. The Role of Demand and Supply in the Generation and Diffusion of Technical Change. Harwood, Chur.
van der Zwaan, B.C.C., Gerlagh, R., Klaassen, G., Schrattenholzer, L., 2002. Endogenous technological change in climate change modelling. Energy Economics 24, 1–19.
von Weizsacker, C.C., 1965. Tentative notes on a two-sector model with induced technical progress. Review of Economic Studies 33, 245–251.
Wene, C.-O., 2000. Experience Curves for Energy Technology Policy. OECD/IEA, Paris.
Weyant, J. (Ed.), 1999. The Costs of the Kyoto Protocol: A Multi-Model Evaluation, Energy Journal Special Issue.
Weyant, J., Olavson, T., 1999. Issues in modeling induced technological change in energy, environment and climate policy. Environmental Modeling and Assessment 4, 67–85.
Wigley, T.M.L., Richels, R.G., Edmonds, J.A., 1996. Economic and environmental choices in the stabilization of atmospheric CO_2 concentrations. Nature 379, 240–243.