Directed technical change as a response to natural-resource scarcity

John Hassler*, Per Krusell†, and Conny Olovsson‡

May 2, 2016

Abstract

How do markets economize on scarce natural resources? In this paper we emphasize technological change aimed at saving on the scarce resource. We develop a neoclassical macroeconomic theory that is quantitatively oriented and that views technical change as directed: it can be used to save on different inputs. At a point in time, the elasticity between inputs—in our application a capital-labor composite and fossil energy—is given by a production function with fixed parameters, but because the future values of these parameters can be changed with R&D efforts today, the long-run elasticity between the inputs is higher than it is in the short run. We demonstrate how the theory can be used to robustly derive predictions for the long-run cost share accruing to the scarce resource as well as for its rate of depletion. In an application, we look at postwar U.S. data, estimate the short-run elasticity between inputs using an aggregate CES production function, and also estimate the implied input-saving technology series. From these technology series, we can gauge what the historical tradeoff has been in the choice between allocating R&D to save on one or the other input. The implied parameter estimates are then used in our aggregate model to make long-run predictions, which indicate a marked increase in the share of costs going to fossil energy.

*Hassler: Institute for International Economic Studies (IIES) and CEPR. Support from Mistra-Swecia is gratefully acknowledged.

†Krusell: IIES, CEPR, and NBER. Support from the European Research Council and Mistra-Swecia is gratefully acknowledged.

‡Olovsson: Sveriges Riksbank. The present paper circulated in an earlier version with a different name ("Energy-Saving Technical Change").

1 Introduction

What is the future of our dependence on natural resources in finite supply? How will consumption growth be affected by scarcity? We develop theory to answer these questions and apply it quantitatively for the case of fossil fuel-based energy as an input into production. Market-oriented analysis would point to price rises as the economy's main signal and incentive to save on the scarce resource, but in this paper we focus on a different channel. Our particular angle is on technological change and how it—in the form of new techniques and products allowing us to save on inputs—is directed endogenously toward inputs that become scarce. Fossil fuel is a prime example and differs from some natural resources in that it cannot be recycled: either energy use has to fall, or alternative energy sources need to be used. All of our quantitative analysis is thus an application to fossil fuel, but we also argue that the approach we take in this paper can be applied more broadly to study technology's role in combatting scarcity.

How potent is technology, however? A key idea in the paper is that we can use historical data, from the perspective of a structural macroeconomic growth model, to assess how effectively energy saving has responded to price movements in the past. This information allows us to parameterize the structural model so that it can be used to predict how potent energy saving is likely to be in the future. Thus, it allows us to directly address the sustainability issue: what are the effects of the resource scarcity on economic growth and welfare, and how concentrated will factor income be, if one views the income accruing to the owners of fossil fuel as one of the factors?

Energy saving can be accomplished both by reducing energy waste and by shifting toward less energy-intensive products. The structural model used in this paper is aggregate in nature and thus melds these two together. In particular, we formulate an aggregate production function, aimed at describing the U.S. economy, with capital, labor, and fossil energy as inputs. We focus on a function that is not necessarily Cobb-Douglas and that therefore allows us to identify separate input-augmenting technology series. We thus look at two input aggregates: a capital-labor composite and energy. The model then contains another layer where the two corresponding (input-augmenting) technology series are subject to choice, along the lines of Acemoglu (2002). This mechanism allows us to capture the natural notion that there is very low short-run substitutability between energy and other inputs, once the technology factors at a point in time have been chosen, but significantly higher substitutability over longer periods when these factors are endogenous.[1] We then show how the long-run energy share, along with consumption growth, will be determined in the model.

With parameter values in hand and an estimate of the amount of fossil fuel left, one can use the model to compute the future paths for technologies, output, and welfare. In the selection of parameter values, there are at least two important challenges, however. One is to determine the shape of the "ex-post" aggregate production function, i.e., the input elasticities conditional on given values of the factor-augmenting technology levels. This shape is important per se but in our context it plays a special role. Given prices and quantities of the different inputs and of output, the shape of the ex-post production function, namely, affects our measures of how the key factor-augmenting technology levels move over time. The other challenge is to characterize the "technology technology": the production possibility frontier for the factor-augmenting technology levels or, put differently, how given the factor-augmenting technology levels today different future paths of these variables can be chosen. We address both questions based on past U.S. data.

We find, first, that an aggregate production function with a unitary elasticity between capital and labor and a near-zero elasticity between the capital-labor composite and fossil energy fits the data quite well and that functions with higher elasticities cannot. Given this overall ex-post shape, we then note that the implied energy-saving technology trend took off very sharply after the oil-price shocks in the 1970s, after having been dormant for decades. The capital/labor-saving technology series, meanwhile, looks very much like the standard aggregate TFP series, thus mimicking the well-known productivity slowdown episode but otherwise featuring steady growth.

We also note that the energy-saving technology trend displays a clear negative medium-run correlation with capital/labor-saving technical change. This finding is nice because it suggests, precisely, that technical change—in the form of saving on different inputs—is endogenous and, as labeled in the literature (see Hicks (1932), Kennedy (1964), Dandrakis and Phelps (1966), Acemoglu (2002), and others), "directed".

Turning to the calibration of the parameters governing the direction of technological

[1] Our approach for modeling ex-ante/ex-post distinctions between input elasticities is, we believe, novel relative to the earlier literature on the topic, which has tended to look at vintage structures. See Atkeson and Kehoe (1999) and the interesting recent study in Abrell, Rausch, and Schwerin (2016).

improvements, it turns out that the theory has a rather direct link between these parameters and observables. We prove, in particular, for a general constant-returns production function, that the key transformation elasticity between the growth rates of the two kinds of factor-augmenting technologies in the long run of the model must equal the relative cost shares of the two factors. I.e., a one-percent decrease in the growth rate of capital/labor-augmenting technical change allows an increase in the growth rate of energy-augmenting technical change by an amount that equals the ratio of the cost share of capital and labor to that of energy in production. Our finding that the growth rates in the two technology trends have varied historically allows us to identify the key transformation elasticity from past data.

Given our estimated parameters, we predict the long-run energy share to be high: no less than 14%. Thus, absent innovation into new sources of energy, the future appears to bring a much higher energy dependence than we have today. There is no collapse of economic growth, however: the implied long-run growth rate of consumption is only somewhat lower than in the past—about 1.56% per year. Thus, although the economy will show high dependence on energy it will still generate high consumer welfare.

Our findings that there is active, directed technical change is perhaps not surprising given a variety of studies using disaggregated analysis; see, e.g., Popp (2002) for energy-saving and Aghion et al. (2012) for the application to "clean" and "dirty" technologies in the case of autos. Our specific contribution here is to formulate and detect directed research efforts on the aggregate level. It should be mentioned in this context that if we were to base our parameter calibration for the relevant transformation elasticity between capital/labor- and energy-augmenting technology on microeconomic studies (as opposed to macroeconomic data), the long-run share of energy would be higher (and economic growth lower). Another relevant reference for our present study is the literature on directed technical change toward high- vs. low-skilled labor (or products intensive in these respective inputs), where beginning with Katz and Murphy's (1991) paper an argument was put forth that there has been skill-biased technical change since the late 1970s. However, no attempt in that context has been made, as far as we are aware, to look at medium-run correlations or to otherwise calibrate the relevant research technologies.

From a modeling perspective, our aggregate setup is standard and, even, rather restrictive, since we do not consider any more broadly generalized production technologies than a nested CES function. However, the data we look at does not seem to call for a generalization, and an advantage of our tight parametrization is that the implied model of directed technical change remains highly tractable. As for our way of modeling the research sector, we look at a planning problem, thus implicitly assuming that whatever research spillovers exist are properly internalized by government policy. This assumption is a natural starting point, though an important extension of the present work would be to study a decentralized setting for technology innovation. Such a setting would feature not only spillovers but also monopoly power, as in typical growth models with endogenous technological change. Scarcity by itself may not call for policy intervention but because we argue that endogenous technological change is an important response to scarcity, policy analysis becomes a natural tool in making sure markets respond appropriately to scarcity.

It should be emphasized that the model we look at has a number of implications not found in more standard growth models. First, it generates stationary income shares despite the fact that we assume no substitutability in the short run. Second, it can produce "peak oil", i.e., a period of increasing fossil fuel use.[2] The latter is observed in data but is difficult to produce in more standard models of finite natural resources; see, for example, the literature following the oil-price shocks of the 1970s, e.g., Dasgupta and Heal (1974), Solow (1974), and Stiglitz (1974, 1979). Relatedly, Jones's (2002) textbook on economic growth has a chapter on non-renewable resources with quantitative observations related to those we make here. Recently, a growing concern for the climate consequences of the emission of fossil CO_2 into the atmosphere has stimulated research into in the supply of, but also demand for, fossil fuels as well as alternatives; see, e.g., Acemoglu et al. (2011) or Hémous (2013). The recent literature, as well as the present paper, differ from the earlier contributions to a large extent because of the focus on endogenous technical change, making use of the theoretical advances from the endogenous-growth literature.[3]

We begin the analysis by developing a model of directed energy-saving technical change in Section 2. We then show analytically in Section 2.1 that the long-run energy share of income exclusively depends on how costly it is to enhance energy efficiency in terms of lost capital/labor efficiency. We also briefly discuss alternative energy sources in this section before we turn our attention to the main application in the paper: fossil fuel. Section 3 thus analyzes production functions with fossil fuel as an input, whereafter we carry out our

[2] Peak oil can also refer to the life cycle of oil production at a given well.

[3] See, e.g., Aghion and Howitt (1992).

estimation in Section 4. Section 5 computes the long-run energy share and growth rate and shows that the model can generate peak oil. Section 6, finally, concludes.

2 The model

In this section, we set up a model with the overall aim of evaluating what our future energy dependence will look like, and to quantify what this energy dependency will imply for the overall growth rate of the economy. Intermediate means to these ends include: (i) finding out what a reasonable aggregate production function that takes energy as an input looks like and (ii) measuring how input-saving technologies respond to incentives, such as input shortages. We will then use the model and our findings to project forward.

Our analysis of directed technical change also allows us to address the issue of "peak oil". Clearly, oil use has increased over time and since this resource is in finite supply it must peak at some point and then fall. A challenge here is that standard models do not predict a peak-oil pattern: they predict *falling* oil use from the beginning of time. By standard models, here, we refer to settings relying on the classic work on nonrenewable resources in Dasgupta and Heal (1974) and with high substitutability between oil and capital/labor. As will become clear, a lower elasticity of substitution between oil and capital/labor is a potentially important factor behind this phenomenon; in particular it is important to include physical capital in the analysis. We will look at these issues more carefully after laying out the basic model.[4]

In order to keep the model transparent and tractable, we make some simplifying assumptions. For the theorem below characterizing the long-run share of energy costs and the rate at which the resource in question—which we think of as oil—is extracted asymptotically, the main simplifying assumption is that oil is costless to extract. It is not in practice, but its marginal cost is, for most available deposits of oil, significantly below the market price, thus implying a large rent (the "Hotelling rent"). As for preferences and technology assumptions made to obtain our first result, we otherwise allow a rather general setting, thus admitting any preferences consistent with balanced growth (the period utility flow is a power function

[4] A number of explanations for increasing oil use, such as new discoveries, imperfect foresight, and falling marginal costs of production, have been proposed. The mechanism we use here should of course be viewed as complementary with these.

of consumption) and the production function is constant returns to scale. When we work out specific cases the later sections, including when we compute transition dynamics, we assume logarithmic utility, a Leontief (or Cobb-Douglas) production function, and full depreciation of capital. These assumptions are most often made in quantitative macroeconomic settings, except for the assumption on depreciation, but full depreciation is not that unreasonable if the time period is 10 years or more, and that is the time frame we have in mind here. Throughout, we also focus on the planning solution and sidestep any issues coming from suboptimal policy with regard to R&D externalities, monopoly power due to patents, as well as the market power in the energy sector; these assumptions are not essential to our main analysis but much simplifies the exposition. We also ignore the (global) climate externality, which might still be of concern in this context even though the damages to the U.S. economy have been estimated to be fairly small overall. So whereas these are important issues, especially from a policy perspective, and we are studying them in related work, for our focus here they should not be a primary concern and could also straightforwardly be included in the analysis.[5]

The representative consumer derives utility from a discounted sum of a power function of consumption at different dates. This is also the objective function of the planner, i.e.,

$$\sum_{t=0}^{\infty} \beta^t \frac{c_t^{1-\sigma} - 1}{1-\sigma}. \tag{1}$$

Output, y, is produced with capital, k, labor, l, and fossil energy, e, as inputs. Specifically, the production function combines these inputs according to the following function

$$y_t \equiv F\left(A_t k_t^{\alpha} l_t^{1-\alpha}, A_t^e e_t\right),$$

where F is homogenous of degree one. The first argument in the production function is a Cobb-Douglas composite of capital and labor (in efficiency units), which ensures that the relative shares of capital and labor inherits their properties from the usual Cobb-Douglas form used in growth studies. The second argument is fossil energy (in efficiency units). A and A^e are measures of input-saving: the level of the capital/labor-augmenting technology and the fossil energy-augmenting technology, respectively.

[5]See Nordhaus and Boyer (2000) and Golosov et al. (2014).

The period resource constraint is given by

$$c_t + k_{t+1} = k_t + (1 - \delta)k_t, \tag{2}$$

as in standard one-good models.

Energy comes from a fossil-fuel source which we think of as oil. Oil is a finite resource, which implies that the planner must respect the following constraint:

$$\sum_{t=0}^{\infty} e_t \leq R_0, \tag{3}$$

where R_t is the remaining stock of oil in ground in the beginning of time t; we could equivalently write $R_{t+1} = R_t - e_t$ for all t.[6] We furthermore assume that oil is costless to extract. These assumptions require comments. First, whereas it is not controversial that there is a finite amount of oil, one could easily argue that R_t is endogenous to some extent, and perhaps stochastic (from the perspective of earlier periods). Second, and relatedly, oil is available from different sources, each associated with a specific (non-zero) unit cost of extraction, so the zero-cost assumption we entertain is also not perfect (e.g., although the marginal cost of most oil in Saudi Arabia is close enough to zero, it is not close to zero in the North Sea). Moreover, the extraction costs can also be affected by $R\&D$ and may be stochastic (shale gas and tar sand are examples of recent innovations of this nature). A full quantitative treatment of oil hence needs a richer structure. It also needs the inclusion of other energy sources, fossil as well as non-fossil. Our simplifying assumptions allow us to uncover mechanisms rather clearly and to characterize long-run outcomes without resorting to numerical analysis. However, for all the reasons above, it would clearly be worthwhile to go beyond the simple assumptions we use here and look at their robustness, a task that we judged to be outside the scope of the present analysis.

The growth rates of the technology trends are assumed given by

$$A_{t+1}/A_t \equiv g_{A,t} = f(n_t) \tag{4}$$

$$A^e_{t+1}/A^e_t \equiv g_{A^e,t} = f_e(1 - n_t), \tag{5}$$

[6]We of course impose that e and R be non-negative at all times; these constraints will not bind for any cases of interest.

where we interpret n as the share of a fixed amount of R&D resources that is allocated to enhancing the efficiency of the capital/labor bundle. The functions f and f_e are increasing. Of course, an increase in A (A^e) is equivalent to a decrease in the input requirement coefficient for the capital/labor (energy) bundle. By changing n, the planner can direct technical change to either of the two. When A grows at a different rate than A^e, the requirement of energy relative to the requirement of the capital/labor bundle changes. Thus, there is factor substitutability in the long run. We also restrict labor supply so that $l_t = 1$ for all t. This can be thought of as "full employment".

2.1 The long-run energy share of income

Our overall aim is to evaluate what our future energy dependence will look like. We consider a balanced growth path where c, k, and y all grow at the rate g, where e grows at a constant rate g_e, and the saving rate and n are both constant which implies that g_A and g_{A^e} also are constant. A natural measure to look at then is energy's long-run share of income. Proposition 1 below shows analytically that this share exclusively depends on the R&D tradeoff. In particular, it does not depend directly on the elasticity of substitution between capital/labor and energy, or on the stock of fossil energy. The proposition also establishes at what rate energy use falls to zero.

Proposition 1 *On a balanced growth path (BGP), the following features must hold:*

1. *The two arguments of the aggregate production function, $A_t k_t^{\alpha}$ and $A_t^e e_t$, both have to grow at the rate of output g.*

2. *Energy use must fall at rate $\beta g^{1-\sigma}$.*

3. *R&D effort n and the consumption growth rate g are determined by $f_e(1-n)\beta = f(n)^{\frac{\sigma}{1-\alpha}} = g^{\sigma}$.*

4. *Energy's share of income is exclusively determined by how costly it is to enhance energy efficiency in terms of lost capital/labor efficiency. Specifically, the long-run energy share is implicitly given by equation (6):*

$$\frac{1 - e^{share}}{e^{share}} = -\frac{\partial g_{A^e}/g_{A^e}}{\partial g_A/g_A}. \tag{6}$$

Proof. To derive (6), start by maximizing (1) subject to constraints (2), (3), (4), and (5). Let $\lambda_t \beta^t$, κ, $\mu_t \beta^t$, and $\mu_t^E \beta^t$ denote the multipliers on these four constraints. The first-order conditions with respect to c_t, k_{t+1}, e_t, n_t, A_{t+1}, and A_{t+1}^e are then given by

$$\lambda_t = c_t^{-\sigma} \tag{7}$$

$$\lambda_t = \beta \lambda_{t+1} F_1 (t+1) \left(\alpha A_{t+1} k_{t+1}^{\alpha - 1} + 1 - \delta \right) \tag{8}$$

$$\kappa = \beta^t \lambda_t F_2 (t) A_t^e \tag{9}$$

$$\frac{\mu_t}{\mu_t^e} = \frac{A_t^e f_e' (1 - n_t)}{A_t f' (n_t)} \tag{10}$$

$$\mu_t = \beta \left[\lambda_{t+1} F_1 (t+1) k_{t+1}^{\alpha} + \mu_{t+1} f(n_{t+1}) \right] \tag{11}$$

$$\mu_t^e = \beta \left[\lambda_{t+1} F_2 (t+1) e_{t+1} + \mu_{t+1}^e f_e (1 - n_{t+1}) \right], \tag{12}$$

where $F_i (t+1)$ refers to the derivative of F with respect to its ith argument in period $t+1$.

On a BGP, the resource constraint dictates that c, k, and y must all grow at the rate g. Both the arguments then have to grow at the rate of output, since the production function is homogeneous of degree one (stated feature 1). It follows from Euler's theorem that both $F_1(t)$ and $F_2(t)$ have to be constant on the BGP. From (7), it then follows that the multiplier λ must grow at rate $g^{-\sigma}$. Using this fact in combination with (11) reveals that the balanced-growth rate for μ is $g^{\alpha - \sigma}$. Equation (9) shows that the balanced growth rate for A^e is $g_{A^e} = g^{\sigma} / \beta$, which implies $g_e = \beta g^{1-\sigma}$: feature 2. Equation (12) can then be used to infer that μ^e will grow at rate $\beta g^{1-2\sigma}$ on the BGP. The first feature stated in the proposition implies $f(n) g^{\alpha} = f_e(1-n) g_e = g$. This equation and the determination of g_e delivers stated feature 3.

Combining (11) and (12) and dividing through by $F(t+1)$ delivers the following expression

$$\frac{F_2 (t+1) A_{t+1}^e e_{t+1}}{F (t+1)} = \frac{F_1 (t+1) A_{t+1} k_{t+1}^{\alpha}}{F (t+1)} \frac{\mu_t^e / \beta - \mu_{t+1}^e f_e (1 - n_{t+1})}{\mu_t / \beta - \mu_{t+1} f(n_{t+1})} \frac{A_{t+1}^e}{A_{t+1}}. \tag{13}$$

Using (10) in the above equation gives

$$\frac{1 - e_{t+1}^{share}}{e_{t+1}^{share}} = \frac{\frac{\mu_t}{\mu_{t+1} \beta f(n_{t+1})} - 1}{\frac{\mu_t^e}{\mu_{t+1}^e \beta f(1 - n_{t+1})} - 1} \frac{f_e' (1 - n_{t+1}) / f_e (1 - n_{t+1})}{f' (n_{t+1}) / f (n_{t+1})}. \tag{14}$$

Inserting the balanced growth rates for μ and μ^e into (14) reveals that the first ratio on the right-hand side of this equation on a BGP equals one. Finally, substituting the fact that $-\frac{\partial g_{A^e,t}/g_{A^e,t}}{\partial g_{A,t}/g_{A,t}} = \frac{f_e(1-n_t)/f_e(1-n_t)}{f'(n_t)/f(n_t)}$ into (14), and recognizing that e^{share} and n are both constant on a BGP delivers (6). ■

Proposition 1 reveals that the steady-state income share of energy is fully determined by technological constraints implied by the R&D functions. These, in turn, depend on the BGP value of n: n determines the relevant local curvature $\frac{\partial g_{A^e}/g_{A^e}}{\partial g_A/g_A}$. The value of n, moreover, only depends on the features of f and f_e and α and σ. In particular, it does not depend on the amount of oil left in the ground (but does rely on oil use going to zero asymptotically at rate $\beta g^{1-\sigma}$). Key for this result is the constant-returns assumption, which hardwires that the two arguments of the production function need to grow at the same rate on any balanced growth path. This is also a reason why the elasticity between inputs is not relevant in the proposition or in the determination of n and the energy share. Of course, for a balanced growth path to exist, which the proposition assumes, the elasticity cannot be too high or else there will be specialization. See Acemoglu (2009) for a discussion of such cases.

Notice that the characterization of the long-run energy share is proved using the endogenously determined growth rate of energy on the BGP. The result that in the long run, and no matter what production function is used so long that it has constant returns to scale, energy use must fall at rate $\beta g^{1-\sigma}$ was quite a surprise to us. Dasgupta and Heal (1974) characterized the case with exogenous and constant technology parameters and arrived at the same result. But it is also clear that this result cannot hold when the technologies grow at arbitrary exogenous rates g_A and g_{A^e}, because then $g_A g^{\alpha} = g_{A^e} g_e = g$ on a balanced growth path.

In order to quantify the long-run energy share, we thus need to compute the tradeoff $\frac{\partial g_{A^e,t}/g_{A^e,t}}{\partial g_{A,t}/g_{A,t}}$. Specifically, the evolutions of the two technology trends A and A^e need to be separately identified. In order to do that, a few more assumptions about the production function are needed. This is the purpose of Section 3. Before we look more closely at the production function, however, let us for a moment contrast the model entertained here with settings with different energy sources and with cases where there are research spillovers.

2.2 Alternative sources of energy and research spillovers

We now briefly point to two possible generalizations, the first involving alternative energy sources and the second research spillovers.

2.2.1 Alternative energy sources

The general production function posited here—$F\left(Ak^{\alpha}l^{1-\alpha}, A^{e}e\right)$, where F has constant returns to scale—can be thought of more generally than from our limited-resource example. In particular, "e" can be any source of energy, or it could be a function of multiple energy sources. So what if we consider an alternative to oil: what are the implications then for the future energy share?

Because we focus on balanced growth paths where *both* the capital-labor composite and energy are actively used and command constant (positive) income shares, we can immediately focus on the following equations:

$$g = g_A g^{\alpha} = g_{A^E} g_e.$$

The second of these equations state that the two arguments of the production function, due to constant returns to scale, have to grow at the same rate; the first equation says that this rate also has to equal the growth rate of output. From the first equation, we can deduce that $g = g_A^{\frac{1}{1-\alpha}}$. Using the assumptions on the research technologies, we then have that

$$f(n)^{\frac{1}{1-\alpha}} = f_e(1-n)g_e$$

must hold. That is, given a growth rate for energy, this equation determines how the research input (n) must be allocated. Above, we saw that in the oil case, g_e is endogenous and given by a simple function of g, σ, and β. Suppose, instead, that an alternative energy source were considered, and let us look at some different possibilities. First, consider "solar power", and let us treat solar power—when fully developed in a cost-effective way—as providing a constant energy flow (fundamentally given by the amount of sunlight reaching earth per unit of time). I.e., for solar power $g_e = 1$ in the long run. Similarly, wind power and power generated by ocean waves arguably involve $g_e = 1$ in the long run.[7] When it comes to other

[7]Over a foreseeable future, when these technologies are being developed, one may also consider $g_e > 1$,

resources in finite supply, such as coal or nuclear power, their long-run values must be the same as that derived for oil.[8]

Recall that the long-run energy share in our setting will be pinned down by (decreasing in) the balanced-growth value of $-\frac{\partial g_{A^e}/g_{A^e}}{\partial g_A/g_A}$, a result that holds regardless of the energy source. As already pointed out, this expression only depends on n, since we can write $g_{A^e} = f_e(1 - f^{-1}(g_A))$. How, then, is n (and, thus, the energy share) affected by the type of energy source considered? In this model, we will consider assumptions on f and f_e such that $g_{A^e} = f_e(1 - f^{-1}(g_A))$ describes a concave function in the positive orthant: a "technological possibility frontier for growth rates". Thus, a higher n (or g_A) implies a higher derivative $\frac{\partial g_{A^e}}{\partial g_A}$ in absolute value as well as a higher g_A/g_{A^e}, and hence a higher $-\frac{\partial g_{A^e}/g_{A^e}}{\partial g_A/g_A}$. In sum, the higher is n, the lower must the long-run energy share be.

Putting this insight together with that above, we conclude that a higher g_e, which obtains to the extent the energy source is not based on a resource in finite supply, must attract R&D away from energy-saving, so n goes up and hence the energy share falls. The extent to which it falls depends on the global properties of f and f_e. This is how the "finiteness" matters for long-run income shares in this model. Of course, there can be long transition periods, and indeed the case we look at below displays a long transition to the steady-state value of the energy share.

2.2.2 Research spillovers

What if research efforts into saving on energy also helps in saving on the capital-labor composite, and vice versa? This model allows these extensions rather straightforwardly. As an example, consider the following concrete formulation:

$$g_A^{\delta_1} g_{A^e}^{\delta_2} = f(n)$$

and

$$g_{A^e}^{\delta_3} g_A^{\delta_4} = f_e(1 - n),$$

though it would then not be literally as a long-run value.

[8] Here as well, these energy sources may involve $g_e > 1$ for a long period of time, as the relevant technologies experience efficiency gains. Arguably, new types of nuclear power, like thorium breeders and fusion power reactors, may have the potential to imply $g_e > 1$ for a period at least as long as the time over which we have seen increasing fossil fuel use.

where the δs are such that overall concavity in the mapping from g_A to g_{A^e} is maintained. Hence, in its reduced form, this formulation could be analyzed just like our benchmark case (which corresponds to $\delta_1 = \delta_3 = 1$ and $\delta_2 = \delta_4 = 0$), only with altered f and f_e functions to obtain the appropriate mapping from g_A to g_{A^e}. In particular, in our empirical assessment of this mapping in Section 4, nothing would need to be changed. Of course, any analysis of a market equilibrium and of the effects of research policy in such contexts would depend greatly on the presence of different kinds of spillovers.

Finally, it may be instructive to relate our present findings to the discussion in Acemoglu's (2009) textbook on growth, in particular its chapter 15 on directed technical change. First, Acemoglu looks at a linear f and a linear f_e and derives a relation between the relative income shares of the two inputs and the rate of technological transformation between the growth rates of the two input-saving technology series. This is similar to what we do, though our use of this relation is focused on how one can think about the determination of the long-run share as a function of the technological transformation (which we treat as non-linear and which is potentially observable in historical data). Second, Acemoglu's chapter 15 considers two kinds of cases: one where the two inputs into the equivalent of our F is each in fixed supply (skilled and unskilled labor) and one where one input is growing (capital) and the other one is fixed (labor). The latter case is somewhat reminiscent of short discussion of the case with solar energy, which in the long run will be like a fixed input ($g_e = 1$). In the Acemoglu setting of this sort, only labor-augmenting technology growth is possible in the long run, and as a corollary there can be no research spillovers (since if there are such spillovers there will be capital-augmenting technology growth too). How, then, can it be that our setting allows $g_A > 1$? The reason is that we consider a capital-labor composite as the other input and not capital alone. Suppose we only had capital: suppose α would equal one. Then g_A would need to be 1 from the equations above, in line with Acemoglu's result. But an input in fixed supply (labor, here) precisely allows $g_A > 1$ and given that there can be long-run growth in A, moreover, spillovers can also be allowed. This is why our setting, though very similar to that in Acemoglu's text, in some ways allows very different qualitative conclusions.[9]

[9] The result that g_A can exceed one is not due to our capital-labor composite being Cobb-Douglas; it would go through for any constant-returns composite of capital and labor so long as the relevant input-saving here concerns labor: A would need to be labor-augmenting, and the key equation above would then read $f(n) = f_e(1-n)g_e$, as in a case with two fixed stocks—now labor and solar power.

3 The aggregate production function

We assume that the production function features a constant elasticity of substitution between a capital/labor composite and fossil energy:

$$y_t \equiv F\left(A_t k_t^{\alpha} l_t^{1-\alpha}, A_t^e e_t\right) = \left[(1-\gamma)\left(A_t k_t^{\alpha}\right)^{\frac{\varepsilon-1}{\varepsilon}} + \gamma\left(A_t^e e_t\right)^{\frac{\varepsilon-1}{\varepsilon}}\right]^{\frac{\varepsilon}{\varepsilon-1}}, \tag{15}$$

where ε is the elasticity of substitution between capital/labor and fossil energy. γ is a share parameter.[10] Note that when $\varepsilon = \infty$, the Cobb-Douglas composite and fossil energy are perfect substitutes, when $\varepsilon = 1$, the production function collapses to being Cobb-Douglas in all input arguments; and when $\varepsilon = 0$ the Cobb-Douglas composite and energy are perfect complements, implying a Leontief function in the capital-labor composite and energy.

As pointed out in Section 2, the model allows the short-run elasticity of substitution to be different from the long-run elasticity. Specifically, the elasticity in the short run is determined by the parameter ε, whereas significantly higher substitutability is possible over longer periods through the channel of endogenous technical change, i.e., through changes in A and A^e.[11]

Even though the production function takes a specific form, it should be pointed out that the specific nesting of capital, labor and energy in (15) is not important for our results. In fact, all the results below still hold with alternative specifications where the elasticity of substitution between capital and energy is allowed to differ from the elasticity of substitution between labor and energy.[12]

Another potential concern is that the function (15) abstract from non-fossil sources of energy such as nuclear power and renewable energy. However, substitution of other energy sources for fossil fuel is part of the process whereby the economy responds to changing

[10] A similar production function is considered by Stern and Kander (2012).

[11] Early empirical investigations into the elasticity of substitution across inputs are Hudson and Jorgenson (1974) and Berndt and Wood (1975) who both use time-series data to estimate the substitutability of energy with other inputs. Both find energy to be substitutable with labor and complementary to capital. Griffin and Gregory (1976) instead use pooled international data and find capital and energy to be substitutes. They argue that there data set better captures the long-run relationships. None of these early studies cover the time of the oil price shocks.

[12] All the results are, for instance, similar with the alternative specification $y_t = \left(\left[(1-\gamma)\left[A_t k_t\right]^{\frac{\varepsilon-1}{\varepsilon}} + \gamma\left[A_t^e e_t\right]^{\frac{\varepsilon-1}{\varepsilon}}\right]^{\frac{\varepsilon}{\varepsilon-1}}\right)^{\alpha} l_t^{1-\alpha}$.

fossil-fuel prices, a process we want to estimate.[13]

Below, we solve the model for two special cases: Cobb-Douglas ($\varepsilon = 1$) and Leontief production (i.e., $\varepsilon = 0$). Throughout the analysis, we also specialize the utility function to the logarithmic case, since that case allows us to go a bit further in terms of closed-form expressions and since this degree of curvature is rather standard in much of the quantitatively oriented macroeconomic research of this kind. The Cobb-Douglas function with a unitary elasticity of substitution between all inputs has been a focal point in the neoclassical growth literature, as well as in business-cycle analysis, primarily because it fits the data on capital and labor shares rather well, without structural or other changes in the parameters of the function. More importantly, however, it was also proposed in influential contributions to the literature on non-renewable resources; see, e.g., Dasgupta and Heal (1980). Since then it has also been employed heavily, including in Nordhaus's DICE and RICE models of climate-economy interactions (see, e.g., Nordhaus and Boyer, 2000). The benefit of the Cobb-Douglas function is that it allows for an analytical solution.

The Leontief production function, in contrast, allows for no substitution possibilities in the short run. Energy's share of income does not have to be constant, and even though energy use falls asymptotically at rate β, it admits different short-run dynamics. As we will see, this case can be solved almost fully analytically. Intermediate cases for ε are not as easy to characterize but as will become clear, the data will speak in favor of almost Leontief.

3.1 Cobb-Douglas production

When $\varepsilon = 1$, the production function is a Cobb-Douglas function in all inputs:

$$F\left(A_t, k_t, l_t, e_t\right) = A_t^{1-\gamma}\left(A_t^e\right)^{\gamma}\left(k_t^{\alpha}\right)^{1-\gamma} e_t^{\gamma}.$$

The above assumptions make it possible to fully solve the model in closed form. Specifically, the Euler equation (8) can be used to show that the saving rate is always constant:

$$k_{t+1} = \beta\alpha\left(1-\gamma\right) y_t. \tag{16}$$

[13]This also appears as a reasonable abstraction, because fossil fuel is and has been the dominant source of energy throughout the whole sample period. According to the Energy Information Administration, fossil fuels constituted 91 percent of the total energy consumption in 1949 and are around 85 percent in 2008 (see table 1.3 in the Annual Energy Review, 2008). Looking at energy's share of income is thus our first task.

Similarly, combining (8), (9), and (16) reveals that the extraction rate rate for oil is constant over time, i.e.,

$$\frac{e_{t+1}}{e_t} = \beta. \tag{17}$$

With Cobb-Douglas production, oil use should, thus, *always* fall at the rate of utility discount. The depletion rate is completely independent of technological progress and of the amount of capital available in a given time period.

An immediate implication of the result that oil use must fall at a constant rate from time 0 is that it seems hard for the Cobb-Douglas model to account for peak oil, i.e., for a rising path of oil use (and, by necessity since the resource is in finite supply, later on a decreasing path). Could it be, however, that this result merely a function of the assumption on utility, namely, logarithmic curvature? Strictly speaking, yes. However, for a moment consider a general period utility function $u(c_t)$. Then it is straightforward to show that $e_{t+1}/e_t = \alpha/s_t$, where $s_t \equiv k_{t+1}/y_t$ is the saving rate (which under logarithmic utility becomes constant and equal to $\alpha\beta$).[14] Thus, to the extent the saving rate displays only minor movements, the result that oil use falls at a constant rate—thus not admitting peak oil—will hold approximately. Moreover, the rate of saving in the U.S. has been remarkably stable over a long period of time. Hence, to the extent one can find a utility function for which the saving rate is temporarily much below its long-run value, thus allowing increasing oil use initially, such a utility function will violate the aggregate data on saving.

To solve for the optimal level of n, we note that equations (11) and (12) are two difference equations in μ_t and μ_t^e, respectively. With Cobb-Douglas production, they read

$$\mu_t = \beta \frac{1}{C_{t+1}} (1-\gamma) \frac{y_{t+1}}{A_{t+1}} + \mu_{t+1} \beta f(n_{t+1}) \quad \text{and} \tag{18}$$

$$\mu_t^e = \beta \frac{1}{c_{t+1}} \gamma \frac{y_{t+1}}{A_{t+1}^{\epsilon}} + \mu_{t+1}^e \beta f_e(1-n_{t+1}). \tag{19}$$

[14]This result holds if depreciation is at 100%, which we take as a rather innocuous assumption as we can treat a model period as comprising 20 years or more.

Substituting $c_t = (1 - \beta\alpha(1-\gamma))\, y_t$ into (18) and (19) and iterating forward delivers

$$\mu_t = \frac{\beta}{1-\beta\alpha(1-\gamma)} \frac{1}{1-\beta} \frac{1}{A_t f(n_t)} \quad \text{and} \tag{20}$$

$$\mu_t^e = \frac{\beta\gamma}{1-\beta\alpha(1-\gamma)} \frac{1}{1-\beta} \frac{1}{A_t^e f_e(1-n_t)}. \tag{21}$$

Inserting (20) and (21) into the first-order condition with respect to n, (10), we obtain

$$\frac{f_e(1-n_t)}{f(n_t)} \frac{f'(n_t)}{f_e'(1-n_t)} = \frac{\gamma}{1-\gamma}. \tag{22}$$

Equation (22) shows that n only depends on the functional forms of the R&D functions. Since these functions are constant over time, n is also constant over time, and it therefore cannot depend on the levels of A, A^ϵ, or k.

Finally, all income shares are constant with a Cobb-Douglas production function (and perfectly competitive input markets). In the model outlined here, energy's share of income should thus equal γ in all time periods.

We are now ready to sum up the predictions from the Cobb-Douglas production function. They state that the saving rate, the extraction rate of oil, the research share devoted to improving A^e over time, as well as energy's share of income, are all constant over time. Specifically, these shares/rates are independent of all macroeconomic state variables. Two of these predictions can be immediately confronted with the data. Figure 1 shows the evolution of energy's share of income, the fossil-fuel price, as well as energy use in the United States.[15]

As can be seen, energy's share of income is highly correlated with its price and it is not constant. Specifically, the share starts out around three percent in 1949 and then decreases somewhat up to the first oil price shock when it increases dramatically. The share then falls drastically between 1981 and the second half of the nineties and then finally increases again. The share does not seem to have an obvious long-run trend, implying that the possibility that the unitary elasticity is a good approximation for the very long run cannot be excluded. For the medium term, however, the data seems hard to square at least with an exact Cobb-

[15] Energy's share of output is defined as ep/q, where e is fossil-energy use, p is the fossil-fuel price in chained (2005) dollars, and q is GDP+net export of fossil fuel in chained (2005) dollars. The data on e and p are both taken from the U.S. Energy Information Administration. The data is described in more detail in Section 4.

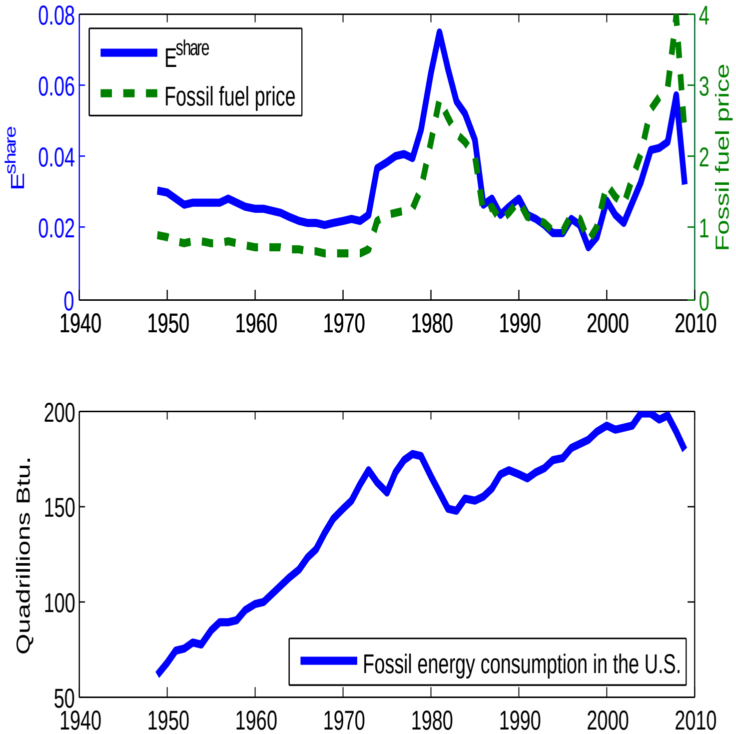

Figure 1: Top graph: fossil energy share (scale to the left) and its price (scale to the right). Bottom graph: fossil energy consumption in the United States.

Douglas production function.

The bottom graph shows that fossil energy use has been increasing throughout the post-war period in the United States. As we discussed above, the Cobb-Douglas assumption rather robustly predicts that energy use must always be exponentially falling. An explanation for the positive trend could potentially be new discoveries of fossil energy. For the Cobb-Douglas model to be able to account for the data, however, one would then have to assume that people make systematic prediction errors and over a long period of time: essentially, this explanation requires there to be a new discovery in each period, and that each such discovery is a complete surprise.

Another potential explanation could come from decreasing extraction costs. We abstract from such costs here, but if they are falling over time, it makes sense to postpone more extraction to the future, since extraction is more profitable then. There are surely improvements in extraction and exploration technologies. Unfortunately, there is relatively limited empirical evidence on the issue of whether ongoing such improvements could explain the ongoing upward trend in aggregate oil use.[16] An important study is Cuddington and Moss (2001), which provides some evidence suggesting that decreasing exploration costs have largely offset increasing scarcity of natural gas in the U.S. over the period 1967–1990. For oil, however, they find the impact of better exploration technologies to have been "more modest". In any case, our proposed mechanism should be viewed as complementary to those just discussed.

3.2 Leontief production

Let us now instead consider a Leontief production function, i.e., where $F\left(A_t k_t^{\alpha}, A_t^e e_t\right) = \min\left\{A_t k_t^{\alpha}, A_t^e e_t\right\}$. We focus on interior solutions such that capital is fully utilized. This requires initial conditions where capital is not too large, in which case it could be optimal to let some capital be idle for some time. In a deterministic model with full depreciation and forward-looking behavior, less than full utilization can only occur in the first period.[17] Due to solutions being interior, we replace the Leontief production function by the equality

$$A_t k_t^{\alpha} = A_t^e e_t \tag{23}$$

and let the planner maximize

$$\sum_{t=0}^{\infty} \beta^t \log(A_t k_t^{\alpha} - k_{t+1})$$

subject to condition (23) for all t, which will be referred to as the *Leontief condition*, $A_{t+1} = A_t f(n_t)$ and $A_{t+1}^e = A_t^e f_e(1 - n_t)$ for all t, and $\sum_{t=0}^{\infty} e_t = R_0$, by choice of $\left\{k_{t+1}, A_{t+1}, A_{t+1}^e, n_t, e_t\right\}_{t=0}^{\infty}$.

Let $\beta^t \lambda_t$ denote the multiplier on the Leontief condition (23). As we are formally solving a special case to that considered in Proposition 1, we will refer to some of the first-order conditions there. In particular, the Euler equation (8) and the first-order condition with

[16] The literature on the relation between theory and data for prices of non-renewables is much larger; see, e.g., Krautkraemer (1998) and the more recent Cuddington and Nulle (2014).

[17] In a model with shocks, one can imagine recurring periods of less than full capital utilization.

respect to energy (9) then read

$$\frac{1}{(1-s_t)A_t k_t^\alpha} = \beta\left(\frac{1}{(1-s_{t+1})A_{t+1}k_{t+1}^\alpha} - \lambda_{t+1}\right)\alpha A_{t+1}k_{t+1}^{\alpha-1} \tag{24}$$

and

$$\beta^t \lambda_t A_t^e = \kappa.$$

Here, s_t is the saving rate out of output.[18] Notice that the Euler equation has a new term: any additional unit used in the future requires an additional energy expense, as represented by the appearance of λ_{t+1}, the multiplier on the Leontief constraint.

Furthermore, equations (11) and (12) under Leontief production are now given by

$$\mu_t = \beta\left[\left(\frac{1}{(1-s_{t+1})A_{t+1}k_{t+1}^\alpha} - \lambda_{t+1}\right)k_{t+1}^\alpha + \mu_{t+1}f(n_{t+1})\right] \tag{25}$$

and

$$\mu_t^e = \beta\left[\lambda_{t+1}e_{t+1} + \mu_{t+1}^e f_e(1-n_{t+1})\right]. \tag{26}$$

The first order condition with respect to n_t is unchanged and still given by (10).

3.2.1 The behavior of saving, conditional on energy use

The second first-order condition above can be used to solve for λ_t in a useful way:

$$\lambda_t = \kappa\frac{1}{A_t k_t^\alpha}\hat{e}_t,$$

where we have defined $\hat{e}_t \equiv \beta^{-t}e_t$. If this expression is inserted into (24) we obtain

$$\frac{s_t}{1-s_t} = \alpha\beta\left(\frac{1}{1-s_{t+1}} - \kappa\hat{e}_{t+1}\right). \tag{27}$$

Using $\hat{s}_t \equiv s_t/(1-s_t)$ and the algebraic identity $\frac{1}{1-s} = 1 + \frac{s}{1-s}$, we can rewrite this equation as

$$\hat{s}_t = \alpha\beta\left(1 + \hat{s}_{t+1} - \kappa\hat{e}_{t+1}\right).$$

[18] Of course, in a recursive formulation, optimal savings depend on the vector of state variables. Here, we solve the model sequentially and present the solution as a time series.

Solving this forward yields

$$\hat{s}_t = \frac{\alpha\beta}{1-\alpha\beta} - \kappa \sum_{k=0}^{\infty} (\alpha\beta)^{k+1} \hat{e}_{t+k+1}. \tag{28}$$

Thus, given a sequence $\{\hat{e}_t\}$ and a value of κ, this equation uniquely, and in closed form, delivers the full sequence of saving rates. We can see from this equation that the more energy is used in the future (in relative terms), the lower is current $\hat{s}$, implying a lower current saving rate. The intuitive reason for this is that more capital requires more fossil fuel and/or more energy efficiency. This limits the value of accumulating capital and more so, the more scarce is fossil fuel. If fossil fuel were not scarce, we would have $\kappa = 0$ and $\hat{s} = \frac{\alpha\beta}{1-\alpha\beta} \Rightarrow s = \alpha\beta$; in that case, the model would thus look just like the Cobb-Douglas model (and the textbook Solow model).

3.2.2 Optimal energy use and R&D

From (24), we have that

$$\beta \left(\frac{1}{(1-s_{t+1})A_{t+1}k_{t+1}^{\alpha}} - \lambda_{t+1} \right) k_{t+1}^{\alpha} = \frac{1}{\alpha A_{t+1}} \frac{k_{t+1}}{(1-s_t)A_t k_t^{\alpha}} = \frac{s_t}{\alpha(1-s_t)A_{t+1}}.$$

Inserted into (25), we obtain

$$\mu_t = \frac{s_t}{\alpha(1-s_t)A_{t+1}} + \beta\mu_{t+1} f(n_{t+1}).$$

Multiplying this equation by A_{t+1} and defining $\hat{\mu}_t \equiv \mu_t A_{t+1}$, we obtain

$$\hat{\mu}_t = \frac{s_t}{\alpha(1-s_t)} + \beta\hat{\mu}_{t+1} = \frac{\hat{s}_t}{\alpha} + \beta\hat{\mu}_{t+1}. \tag{29}$$

Given a sequence of saving rates, this equation can be solved for the $\hat{\mu}_t$ sequence. Iterating forward on (29) one obtains another discounted sum:

$$\hat{\mu}_t = \frac{1}{\alpha} \sum_{k=0}^{\infty} \beta^k \hat{s}_{t+k}.$$

Recalling that $\hat{\mu}_t$ is the current marginal value of capital-augmenting technology, it is intuitive that the more is saved for the future, the higher is the value of $\hat{\mu}_t$. Using equation (28), we can write this directly in terms of the $\hat{e}$ sequence. By direct substitution and slight simplification we obtain

$$\hat{\mu}_t = \frac{\beta}{(1-\beta)(1-\alpha\beta)} - \frac{\kappa}{\alpha}\sum_{k=0}^{\infty}\beta^k \sum_{j=0}^{\infty}(\alpha\beta)^{j+1}\hat{e}_{t+1+k+j}. \tag{30}$$

Similarly, from the expression for λ_t above, we can write

$$\lambda_{t+1}e_{t+1} = \kappa\frac{e_{t+1}}{A_{t+1}k_{t+1}^{\alpha}}\hat{e}_{t+1} = \kappa\frac{\hat{e}_{t+1}}{A_{t+1}^{e}},$$

where the last equality follows from the Leontief assumption. Inserted into (26), we obtain

$$\mu_t^e = \beta\left[\kappa\frac{\hat{e}_{t+1}}{A_{t+1}^{e}} + \mu_{e,t+1}f_E(1-n_{t+1})\right].$$

Multiplying this equation by A_{t+1}^e and defining $\hat{\mu}_t^e \equiv \mu_t^e A_{t+1}^e$, we obtain

$$\hat{\mu}_t^e = \beta\left[\kappa\hat{e}_{t+1} + \hat{\mu}_{t+1}^e\right].$$

Given a sequence of normalized fossil fuel uses, this equation can be solved for the $\hat{\mu}_t^e$ sequence. It can also be solved forward:

$$\hat{\mu}_t^e = \kappa\sum_{k=0}^{\infty}\beta^{k+1}\hat{e}_{t+k+1}. \tag{31}$$

In parallel to the case of capital-augmenting technology, the more energy is used in the future and the larger is fossil fuel scarcity, the higher is the marginal value of energy-enhancing technology.

Finally, rewriting the first-order condition for n_t above in terms of the redefined multipliers, we arrive at

$$\hat{\mu}_t\frac{f'(n_t)}{f(n_t)} = \hat{\mu}_t^e\frac{f_e'(1-n_t)}{f_e(1-n_t)}.$$

From this equation we can solve for n_t uniquely for all t, given the values of the multipliers, so long as f is (strictly) increasing and (strictly) concave.[19] Using the expressions above for the multipliers, we obtain

$$\frac{f_e'(1-n_t)f(n_t)}{f_e(1-n_t)f'(n_t)} = \frac{\frac{\beta}{\kappa(1-\beta)(1-\alpha\beta)} - \frac{1}{\alpha}\sum_{k=0}^{\infty}\beta^k\sum_{j=0}^{\infty}(\alpha\beta)^{j+1}\hat{e}_{t+1+k+j}}{\sum_{k=0}^{\infty}\beta^{k+1}\hat{e}_{t+k+1}}. \tag{32}$$

This equation enables us to solve for n_t directly as a function of the future values of $\hat{e}$. The more energy is used in the future, the lower is current n, i.e., the more R&D labor is allocated to energy-saving now. It is straightforward to solve this equation numerically for transition dynamics; we discuss details below. However, one can characterize long-run growth features more exactly.

3.2.3 Balanced growth with Leontief production

On a balanced path, the saving rate is constant and so, from the Euler equation, $\hat{e}_t$ is constant. Recalling the definition $\hat{e}_t \equiv \beta^{-t}e_t$, this means that $e_t = (1-\beta)R_t$ and $R_t = \beta^t R_0$, implying $\hat{e} = (1-\beta)R_0$. The saving rate s_t on the balanced growth path is constant and $\hat{s}_t = \frac{s_t}{1-s_t} = \frac{\alpha\beta(1-\kappa\hat{e})}{1-\alpha\beta}$.

The solution for n is given by feature 3 of Proposition 1 (which holds for any CRS production function): $f_e(1-n)\beta = f(n)^{\frac{1}{1-\alpha}}$. On a balanced growth path, equation (32) becomes

$$\frac{(f_e)'(1-n)f(n)}{f_e(1-n)f'(n)} = \frac{1-\kappa\hat{e}}{\kappa\hat{e}(1-\alpha\beta)}. \tag{33}$$

Given our solution for n, we can solve for the utility value of the spending on energy, $\kappa\hat{e}$, from equation (33), thus also delivering the per-unit value κ as $\hat{e} = (1-\beta)R_0$. Notice that any increase in R_0 just lowers the κ by the same percentage amount on a balanced growth path. We then also obtain the saving rate, as specified above, since it depends on $\kappa\hat{e}$. Of course, a balanced path from time 0 also requires that the initial vector (k_0, A_0, A_0^e, R_0) satisfy

$$A_0 k_0^{\alpha} = A_0^e(1-\beta)R_0;$$

if and only if this equation is met will the economy grow in a balanced way all through time.

The model implies a simple relation between the energy share and the rate of saving. To

[19] Either the monotonicity or the concavity need to be strict.

derive it, first find an expression for the energy share: divide the marginal utility value of an energy unit, κ, by the marginal utility of consumption at t and then multiply by by $\frac{e_t}{y_t}$ to obtain

$$e_t^{share} = \frac{\kappa}{\beta^t u'(c_t)} \frac{e_t}{y_t} = \kappa \hat{e}_t (1 - s_t),$$

where we have used that $u'(c_t) = 1/c_t$ and $c_t = (1 - s_t)Y_t$.

Second, insert the implied expression for $\kappa\hat{e}$ as a function of the energy share in the stationary version of the Euler equation derived in the beginning of this section. We then obtain that along a balanced path it must be that

$$e^{share} = 1 - \frac{s}{\alpha\beta}.$$

Put differently, we have that $s = \alpha\beta(1 - e^{share})$. That is, in this model, the rate of saving is smaller than if energy were not needed in production (it is $\alpha\beta$ with Cobb-Douglas production, logarithmic utility and full depreciation of capital).

3.2.4 Transitional dynamics

The long-run growth patterns of the present model involve a non-trivial determination of the rate of growth of technologies and of consumption as well as the long-run share of energy in production costs. What about the transition dynamics? They are interesting in that, as we shall illustrate later with a numerical example based on calibration of the model's parameters to data, they allow "peak oil", i.e., an initially rising path of oil use, which we saw was not possible as an outcome under Cobb-Douglas production and which is more broadly somewhat of a modeling challenge in this literature. Before moving to the model calibration and its use, however, let us make some remarks on the nature of transition and what kinds of initial conditions can generate peak oil.

Let us first point out that whereas it is not possible to characterize the full transition paths of the model in closed form, these are very easy to solve for numerically, since for a given sequence for $\{\hat{e}_t\}_{t=0}^{\infty}$, all other variables can be solved in closed form. The numerical algorithm thus iterates over this particular sequence. In the appendix, we describe how this is accomplished.

So what are the state variables in this model? They are four: k, A, A^e, and R, i.e., physical capital, the two kinds of technology, and the remaining amount of oil. Despite the high dimensionality (relative to those in standard growth models), it is relatively simple to think about transition dynamics here. To do this, let us first recall how the dynamics work without oil, i.e., in the same model except without the Leontief constraint, and then consider the full model.

In the model without oil, all research is devoted to increasing A and its gross growth rate will hence equal $f(1)$. This model, which only has two state variables, k and A, delivers a balanced growth rate g for output, capital, and consumption that will satisfy $g = f(1)^{\frac{1}{1-\alpha}}$. For any two initial condition (k_0, A_0) such that $gk_0 = sA_0k_0^\alpha$, where s is the saving rate, the model features balanced growth at all dates: with $k_{t+1} = gk_t$ and $A_{t+1} = f(1)A_t$, the equation $k_{t+1} = sA_tk_t^\alpha$ will then always be satisfied.[20] Hence, given any A_0, there is a unique value for k_0 leading to balanced growth. It is also straightforward to see, given the closed-form nature of the system, that any higher (lower) value for k_0 will lead to monotone convergence of $k_t^{1-\alpha}/A_t$ to s/g from above (below).

With the full model, one first need to think about whether the Leontief condition will hold at all times. We will devote all our attention below to the case in which it does hold at all times, but let us briefly consider other possible cases first. So can it ever be that $A_tk_t^\alpha < A_t^ee_t$? No, because then e_t could be lowered and higher utility obtained: energy could then be used at some other point in time when it would increase production and consumption. But could $A_tk_t^\alpha > A_t^ee_t$? For $t = 0$ clearly yes, since R_0 is bounded: for a low enough A_0^e, this inequality would then have to hold. Could such an inequality hold at later points in time too? It is straightforward to show that it could not, since such choices would be dominated by saving less in physical capital and/or by lowering n in the period prior, thereby freeing up resources for consumption and/or for raising A^e. The possibility that an inequality will obtain from time zero will hereby be abstracted from.[21]

Restricting matters to the case where the Leontief constraint is binding, on the balanced path, the growth rates for A and for k are in the same relation to each other as in the model without oil. However, these growth rates are lower now since some research efforts must be devoted to making A^E grow, implying that A must grow less fast. Moreover, the

[20] In this model, $s = \alpha\beta$ on an off the balanced path, but that is not relevant for the argument.

[21] The exact condition on A_0^e—as a function of the other state variables—under which the inequality will hold cannot be derived analytically as far as we are aware.

connection between k_0 and A_0 required for exact balanced growth at all dates still has to hold: $gk_0 = sA_0k_0^\alpha$, where g is lower but s is lower too, as we saw in the previous section. So what initial conditions on A^e and R are required for exact balanced growth from time 0 and on? There is no condition connecting the growth rates of R and A^e: any combination of A_0^e and R_0 satisfying $A_0k_0^\alpha = A_0^eR_0$ would be consistent with balanced growth (again, assuming that they are large enough that the Leontief condition will hold at time 0). Thus, the initial conditions under which there is balanced growth at all times are merely $gk_0 = sA_0k_0^\alpha$ and $A_0k_0^\alpha = A_0^eR_0$. In fact, we show in the appendix that one can transform the maximization problem so that the only two relevant initial conditions of the model are the values for $\tilde{k}_0 \equiv k_0/(A_0^eR_0)$ and $\tilde{A}_0 \equiv A_0/(A_0^eR_0)^{1-\alpha}$.

It is particularly interesting to discuss to what extent the paths for technology grow faster or more slowly than on their balanced paths and whether energy use is monotone over time. Through the initial relations between technologies and capital, this model has much richer implications than the Cobb-Douglas model. Suppose that A_0 (or k_0) are lower than on the balanced path. Then energy, and energy-saving technology level, are abundant, and the economy is better off accumulating A and k than increasing A^e. Since the Leontief condition will still hold, e will then grow with Ak^α until the technology levels are more balanced and A^e will start growing; at some point, e will then reach its peak and begin to fall. This scenario thus follows rather intuitively and we illustrate with a numerical example in a calibrated version of the model in Section 5.3. Of course the reverse is also possible, i.e., that A^e will start below its balanced path and then A^e will grow fast initially, whereas the growth rate of e will then be below its balanced path.

4 Measuring technical change

In order to compute the long-run energy share, and the long-run growth rate we need to compute the tradeoff $\frac{\partial g_{A^e,t}/g_{A^e,t}}{\partial g_{A,t}/g_{A,t}}$. Specifically, the evolution of the two technology trends A and A^e needs to be separately identified. In the next section, we will formally estimate these trends along with the key elasticity parameter ε. However, let us first start with how one could use a procedure similar to that in Solow (1957), provided one knew the correct production function. In his seminal paper, Solow shows how to measure technology residuals with a general production function with two key assumptions: perfect competition and constant returns to scale. We can thus use his procedure to back out an overall residual but how can the separate technology trends A and A^e be identified? Here the key insight is that with one more structural assumption than that Solow used, one can: we use the functional form of the production function.[22]

Output gross of energy expenditure is now defined as $q_t \equiv y_t$ plus net export of fossil fuel. We thus look at a broad measure of fossil fuel, even though the sub-components are somewhat heterogeneous; natural gas and oil are rather similar in terms of their production technologies but coal is different, with a higher marginal cost as a fraction of the market price and with higher estimated proved reserves. We also abstract from non-fossil sources of energy (such as water power and nuclear power). An alternative approach would be to model the energy sector in much more detail; such a study would be valuable, but is probably beyond the scope of the present paper as more detailed modeling would necessitate also bringing in more technology series—in principle, one for each energy source—and as the range of estimates of substitution possibilities between different energy sources shows large variation across studies (see, e.g., the meta study by Stern, 2012).[23] Throughout our discussion in this section, we impose $\alpha = 0.3$, as we know that it will fit the data on the relative capital/labor shares well. We spend considerable time discussing how to choose elasticity parameter ε, including both an informal approach and structural estimation, but for now assume that we know its value, along with γ. Under perfect competition in input markets, marginal products equal factor prices, so that labor's and energy's shares of income are given by

$$l_t^{share} = (1-\alpha)(1-\gamma)\left[\frac{A_t k_t^{\alpha} l_t^{1-\alpha}}{q_t}\right]^{\frac{\varepsilon-1}{\varepsilon}} \tag{34}$$

and

$$e_t^{share} = \gamma \left[\frac{A_t^e e_t}{Q_t}\right]^{\frac{\varepsilon-1}{\varepsilon}}, \tag{35}$$

respectively. Equations (34) and (35) can be rearranged and solved directly for the two

[22] A similar approach has been used in a number of other applications; see, e.g., Krusell et al. (2000) and Caselli and Coleman (2006).

[23] We did also consider an oil-only case, with very similar results to those reported in the text.

technology trends A_t and A_t^e. This delivers

$$A_t = \frac{q_t}{k_t^{\alpha} l_t^{1-\alpha}} \left[\frac{l_t^{share}}{(1-\alpha)(1-\gamma)} \right]^{\frac{\varepsilon}{\varepsilon-1}} \tag{36}$$

and

$$A_t^e = \frac{q_t}{e_t} \left[\frac{e_t^{share}}{\gamma} \right]^{\frac{\varepsilon}{\varepsilon-1}}. \tag{37}$$

Note that with ε and γ given, and with data on q_t, k_t, l_t, e_t, l_t^{share} and e_t^{share}, equations (34) and (35) give explicit expressions for the evolution of the two technologies. Clearly, the parameter γ is a mere shifter of these time series and will not play a role in the subsequent analysis. The key parameter, of course, is the elasticity ε.

As for how to identify the elasticity in the data, it turns out that it is quite informative to first look at the data informally: we vary ε, back out the resulting technology series, and compare their properties. This procedure will reveal that unless ε is chosen to be very low, the series for A^e will be extremely volatile and, in particular, not look like a technology series at all. By a technology series here we refer to one that is rather smooth and mostly non-decreasing. After our informal look at the data, then, we structurally estimate ε, an exercise that delivers a value in line with that suggested by the informal procedure.

4.1 Data

We begin by describing our data, which is annual and covers the period 1949–2009. The GDP data is denoted in chained (2005) dollars and is taken from the Bureau of Economic Analysis.[24] The net export of fossil fuel is also converted to chained (2005) dollars and is taken from the Annual Energy Review 2009, Table 3.9. The data used to compute the wage share is taken from the Bureau of Economic Analysis, and it includes data on self-employment

[24]The data is found at http://www.bea.gov/histdata/Releases/GDP_and_PI/2009/Q4/Third_March-26-2010/Section1ALL_xls.xls. For the USEIA reference, see tables 1.3 and 3.1, respectively, available at http://www.eia.gov/totalenergy/data/annual/. We use online data from the BLS on the labor force. The data is available at http://www.bls.gov/webapps/legacy/cpsatab1.htm#a1.f.1. The data on the capital stock is taken from FRED database. Labor's share of income is calculated as (compensation of employees / (compensation of employees + private surplus - proprietors' income)) and is taken from BEA, National Accounts.

income.[25]

The data on fossil energy-use e, as well as the fossil-fuel price p, are both taken from the U.S. Energy Information Administration. Specifically, we construct a composite measure for fossil energy use from U.S. consumption of oil, coal, and natural gas. Similarly, a composite fossil-fuel price is constructed from the individual prices of the three inputs, and it is converted to chained (2005) dollars.[26] The method for constructing the two composites is described in the Appendix A.1.

4.2 Comparing high and low substitution elasticities

We now look at how varying ε changes the implied technology series A and A^e. We will just look at two cases: one near Cobb-Douglas and one near Leontief, i.e., one where ε is close to 1 and one where it is close to 0. For each case, we offer interpretations of the changes in the implied series over the sample, and we will find that the Leontief case offers an interpretation that is overall reasonable and very much in line with the overall theme of our paper, whereas the interpretation of the Cobb-Douglas case, though clear in its logic, is highly implausible from a quantitative perspective. We begin with the Cobb-Douglas case.

4.2.1 A high elasticity

The Cobb-Douglas production function does not allow the two technology series to be separately identified, because in this case ($\varepsilon = 1$) the shares will be constant and unaffected by the technology series. For any non-unitary value, however, the shares will depend on these series. Here, we look at a value slightly below 1: $\varepsilon = 0.8$.[27] The evolution of the fossil energy-saving technology implied by equations (36)-(37) is displayed in Figure 2. The

[25] The alternative specification of the production function as given in footnote 12, requires annual data on capital's share of output. We computed this share as $r_t k_t / q_t = 1 - w_t l_t / q_t - p_t e_t / q_t$.

[26] Since housing services are included in GDP, we include fossil fuel used for heating residential houses in e. Transportation services provided by the private use of cars are, however, not included in GDP so the fossil-fuel use for these purposes should ideally be deducted from total fossil fuel use when e is constructed. Unfortunately, the only data on petroleum consumption that we have is an aggregate for all passenger cars, i.e., it also includes fuel for passenger cars that are used for professional services (and, thus, are measured in GDP). The correlation between fuel for passenger cars and e is higher than 0.95, so such a deduction only has a level effect. As a benchmark, we therefore do not deduct the fuel used by passenger cars, but we have verified that our estimate of ε does not increase when it is deducted from e.

[27] A value above one will share the main features reported on below—extreme volatility and non-technology like behavior of A^e—but the up-and-down turns will switch signs.

obtained series features very large jumps and dives; it increases by more than 50,000 percent between 1980 and 1998 and decreases by more than 7,600 percent between 1998 and 2009. It appears impossible to interpret it as technology.

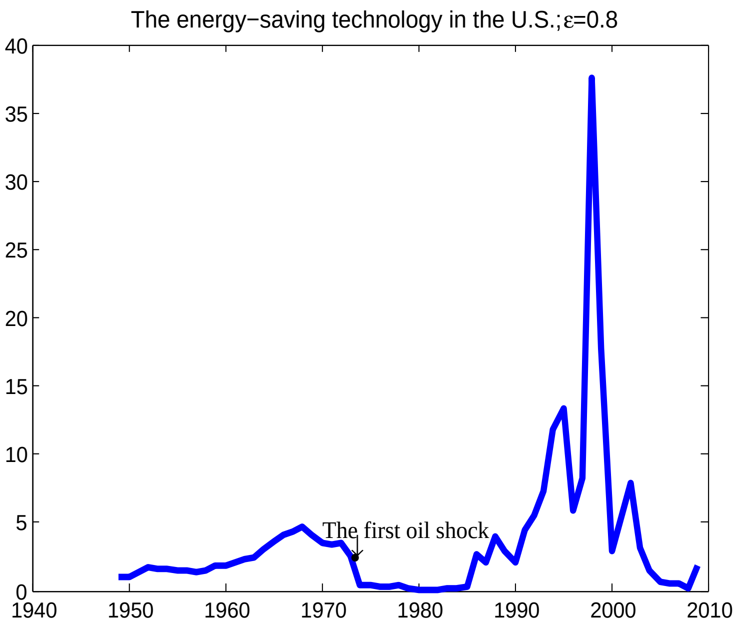

Figure 2: Energy-saving technology with an elasticity of 0.8.

To understand why we obtain such a volatile series, take the first oil shock as an example. The observed price increase in 1973 would then, with a fairly high substitution elasticity, call for much lower energy consumption in U.S. business. However, this fall in consumption is not observed in the data. According to the model, the lack of a drop in oil consumption must then have occurred because of a huge fall in the energy-specific technology value. That is, oil is now much more expensive, but it is also less efficient in producing energy services, so that the net result is a rather constant level of oil use. Similarly, when the price falls in the nineties, an increased demand for energy is expected but, again, this rise is not observed in the data. The reason, from the perspective of the model, must then be that the energy-specific technology increases sharply—by 350 percent in just one year—so that no increased oil use is necessary. The mean annual growth rate in A^e during the period with the chosen elasticity is negative (-1.42 percent) and the standard deviation is very high (62.8 percent).[28]

When thinking about the implications of the price movements for fossil fuel, note that the exact reason for the large volatility of fossil fuel prices is not important. This is key since there are two very different, and contrasting, interpretations of the large movements in oil prices. Barsky & Kilian (2004), in particular, argue that the conventional view, i.e., that events in the Middle East and changes in OPEC policy are the key drivers of oil-price changes, is incorrect. Rather, they contend, the price changes were engineered by U.S. administrative changes in price management. Be that as it may, what is important here is simply that firms actually faced the prices we use in our analysis. In conclusion, at least from the perspective of the assumptions of the theory, it is not possible to maintain as high an elasticity of substitution between capital/labor and energy as 0.8. In fact, it takes much smaller εs to make the volatile, non-technology-like features go away—only values below $\varepsilon \sim 0.05$ make the series settle down. We now turn to what those series look like for the near-Leontief case.

4.2.2 A low elasticity

The near-Leontief case, i.e., ε close to zero, is rather robust in that once the substitution elasticity is in this low range, the features of the technology series do not vary noticeably. So consider now instead a very low elasticity and set $\varepsilon = 0.02$. The implied energy-saving technology trend is presented in Figure 3. The figure shows the path for fossil energy-saving technology A^e. As is evident, we observe a smooth, increasing, and overall reasonable-looking graph for fossil energy-specific technology. The mean growth rate is 1.47 percent and the standard deviation is 2.25 percent.[29]

Moving toward interpretations of the obtained technology series, the figure also shows separate trends lines before and after the first oil-price shocks: 1949–1973 and 1973–2009. Clearly, the technology series appears to have a kink around the time of the first oil price shock. In fact, the growth rate is 0.1 percent per year up to 1973 and 2.54 percent per year after 1973. The fact that the kink occurs at the time of the first oil price shock suggests that the higher growth rate in the technology is an endogenous response to the higher oil price.

[28]The capital/labor-augmenting technology is highly volatile too when ε is high, but the effects of high substitutability are larger on the energy-saving technology.

[29]The energy-saving technology does not change much for elasticities in the interval $(0, 0.05)$.

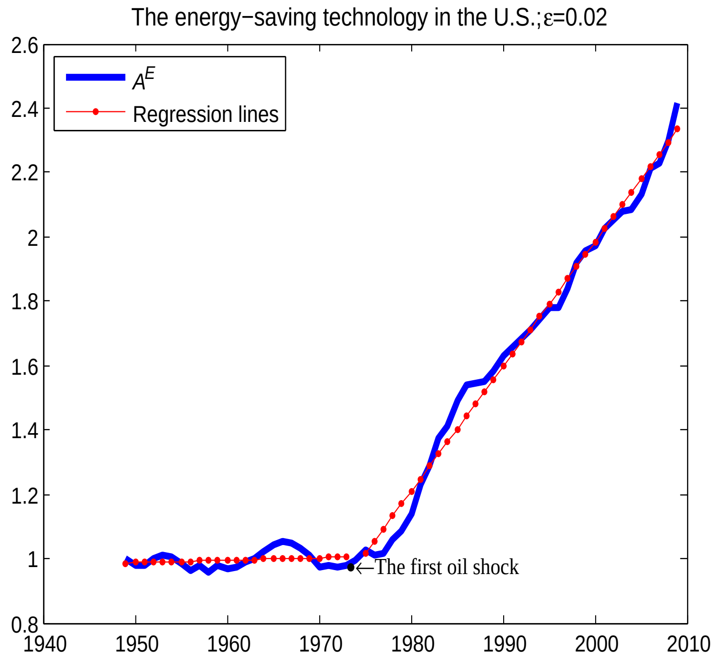

Figure 3: Energy-saving technology with an elasticity of 0.02.

A possible alternative explanation for the kink around 1973 could be structural transformation, e.g., the expansion of the service sector relative to the manufacturing sector. Specifically, if the production of services requires relatively less energy, then such a process could be mistaken for energy-saving technical change. At the same time, however, if services in fact do require relatively less energy than manufactured goods then this type of transformation could also be an endogenous response to the oil price shocks. Regardless of the direction of causation, Figure 8 in Appendix A.5 shows the energy-saving technology in the manufacturing sector, and it is qualitatively very similar to that in the aggregate; the kink is somewhat less pronounced, but the interpretation is that there appears to have been a drastic increase in energy saving within the service sector as well.

What does a low substitution elasticity imply for the evolution of the capital/labor-augmenting technology? The series for A is plotted the solid line in Figure 4, alongside the A^e series. A too is smooth and increasing graph and very much looks like the conventional total-factor productivity (TFP) series. The mean growth rate in A is 1.28 percent and the standard deviation is 1.63 percent.

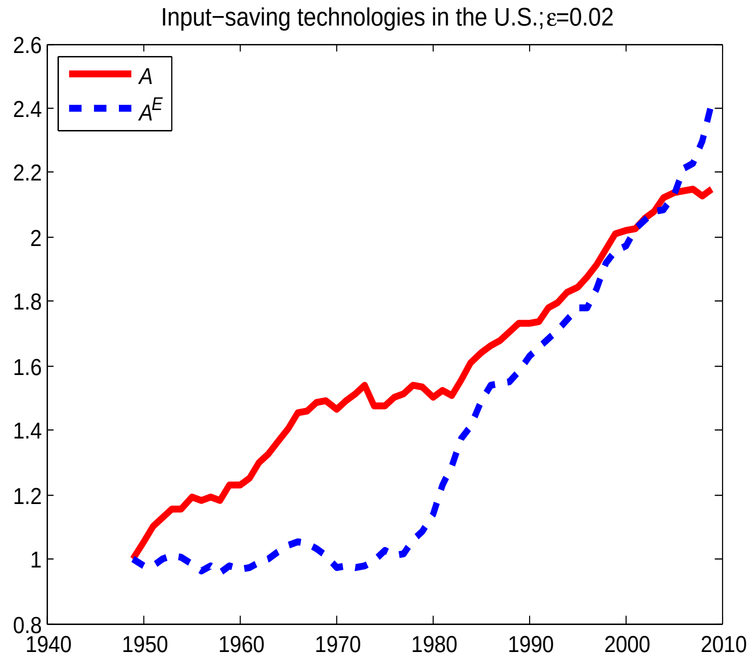

Figure 4: Energy- and capital/labor-saving technologies compared.

4.3 Estimation

We now estimate the elasticity ε structurally, together with some other parameters, using a maximum-likelihood approach. The idea behind the estimation is perhaps simplistic but, we think, informative for our purposes: we specify that the technology series are exogenous processes of a certain form and then estimate the associated parameters along with ε. The technology processes have innovation terms and the maximum likelihood procedure, of course, chooses these to be small. Hence, the key assumption behind the estimation is to find a value of ε such that the implied technology series behave smoothly, or as smoothly as the data allows. This may not be appropriate for other kinds of series but it seems reasonable precisely to require that changes in technology are not abrupt. Based on our findings above, however, we allow the average growth rates to be different before and after the oil shocks of the 1970s.

As for the particular specification of our technology processes, we also choose parsimony: we require that they be stationary in first differences and have iid, but correlated, errors.

Other formulations, allowing different moving trends or serially correlated errors, will undoubtedly change the details of the estimates we obtain, but they are unlikely to change the broad findings.

Our chosen formulation is

$$\begin{bmatrix} a_t \\ a_t^e \end{bmatrix} - \begin{bmatrix} a_{t-1} \\ a_{t-1}^e \end{bmatrix} = \begin{bmatrix} \theta_{pre}^A \chi_T + \theta_{post}^A (1-\chi_T) \\ \theta_{pre}^e \chi_T + \theta_{post}^E (1-\chi_T) \end{bmatrix} + \begin{bmatrix} \varpi_t^A \\ \varpi_t^e \end{bmatrix}, \tag{38}$$

where $a_t = \log(A_t)$, $a_t^e = \log(A_t^E)$, $\varpi_t \equiv \begin{bmatrix} \varpi_t^A \\ \varpi_t^e \end{bmatrix} \sim N(\mathbf{0}, \Sigma)$, and $\chi_T$takes the value 1 for all observations before the oil shocks (i.e., for $t < T$) and 0 after. The θs, thus, measure the average growth rates before and after the oil shocks for the two technologies.

Dividing equations (36) and (37) by their counterparts in period $t-1$ gives

$$\frac{A_t}{A_{t-1}} = \frac{q_t}{k_t^\alpha l_t^{1-\alpha}} \frac{k_{t-1}^\alpha l_{t-1}^{1-\alpha}}{q_{t-1}} \left[\frac{l_t^{share}}{l_{t-1}^{share}}\right]^{\frac{\varepsilon}{\varepsilon-1}} \tag{39}$$

and

$$\frac{A_t^e}{A_{t-1}^e} = \frac{q_t}{e_t} \frac{e_{t-1}}{q_{t-1}} \left[\frac{e_t^{share}}{e_{t-1}^{share}}\right]^{\frac{\varepsilon}{\varepsilon-1}}. \tag{40}$$

Taking logs of (39) and (40) and using (38) in these expressions allows us to write the system as

$$\mathbf{s}_t = \theta - \frac{\varepsilon}{\varepsilon - 1}\mathbf{z}_t + \varpi_t, \tag{41}$$

where

$$\mathbf{s}_t \equiv \begin{bmatrix} \log\left(\frac{q_t}{k_t^\alpha l_t^{1-\alpha}}\right) - \log\left(\frac{q_{t-1}}{k_{t-1}^\alpha l_{t-1}^{1-\alpha}}\right) \\ \log\left(\frac{q_t}{e_t}\right) - \log\left(\frac{q_{t-1}}{e_{t-1}}\right) \end{bmatrix} \quad \text{and} \quad \mathbf{z}_t \equiv \begin{bmatrix} \log l_t^{share} - \log l_{t-1}^{share} \\ \log e_t^{share} - \log e_{t-1}^{share} \end{bmatrix}.$$

The log-likelihood function is now given by

$$\begin{gathered} l(\mathbf{s}|\theta,\varepsilon,\Sigma) = \\ -\frac{N}{2}\log|\Sigma| - \frac{1}{2}\sum_{t=1}^{N}\left(\mathbf{s}_t - \left(\theta - \frac{\varepsilon}{\varepsilon-1}\mathbf{z}_t\right)\right)^T \Sigma^{-1} \left(\mathbf{s}_t - \left(\theta - \frac{\varepsilon}{\varepsilon-1}\mathbf{z}_t\right)\right) \\ + const. \end{gathered} \tag{42}$$

Maximization of (42) with respect to θ, ε, and Σ gives the estimated parameters straightforwardly.[30]

The results are presented in Table 1.

Table 1: Estimated parameters

θ^A_{pre}	θ^e_{pre}	θ^A_{post}	θ^e_{post}	ε
0.0132	0.0047	0.0118	0.0264	0.0013
(0.0029)	(0.0014)	0.0030	0.0036	(0.0055)

Standard errors in parenthesis

There are two noteworthy features here. One is that the technology trends are both positive and of very similar, and a priori reasonable, magnitude. The estimation also confirms the results above, i.e., that the capital/labor-saving technology grows faster than the energy-saving technology before the oil shocks, whereas this result is reversed after the shocks. The second, and more important, point is that the elasticity of substitution between the capital/labor composite and energy is very close to, and in fact not significantly different from, zero. Thus, the CES function that fits the annual data on shares best—where we again emphasize that the estimation procedure penalizes implied technology series that are not smooth—is essentially a Leontief function.[31] This is all in line with our informal examination of the implications of different values for ε in the previous section.

Turning to the estimated covariance matrix, it is given by

$$\Sigma = 10^{-3} * \begin{bmatrix} 0.2629 & -0.0052 \\ -0.0052 & 0.3735 \end{bmatrix}.$$

Thus, the energy-saving shocks are somewhat more volatile than are capital/labor-saving technology shocks.

Figure 5 provides some interpretation for the estimated low elasticity. It shows histograms of the estimated shocks ϖ_t as computed from (41), alongside the normal distributions with the estimated parameters. The top graphs show the results for the estimated value of ε, i.e., $\varepsilon = 0.0013$, whereas the bottom graphs show the residuals for the higher value of $\varepsilon = 0.75$.

Clearly, the variance for the ϖ^e shock in particular is an order of magnitude larger for

[30]We find the parameters through an iterative algorithm that is described in Appendix A.4.

[31]The exact choice of T is unimportant for this finding.

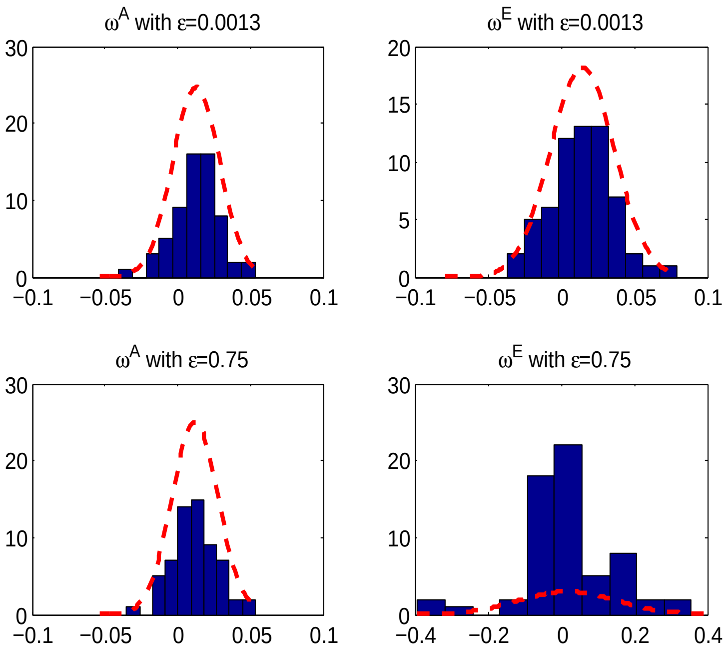

Figure 5: Histograms of the data on the shocks ϖ_t as computed from (41), alongside the normal distributions with the estimated parameters. The top graphs show the residuals for the estimated value of ε, i.e., $\varepsilon = 0.0013$, whereas the bottom graphs show the residuals for the higher value of $\varepsilon = 0.75$.

the high value of ε relative to the low value. In addition, a high elasticity seems inconsistent with Gaussian innovations. A low ε implies a low variance and, thus, a high value for the log-likelihood.

5 Implications

The short-run elasticity, ε, has now been estimated and found to be close to zero. We have also seen that a smooth evolution of energy-saving technical change seems to require a low elasticity. We now turn to analyse the relation between the two growth rates. In particular, the model assumes technical change to be directed. To what extent is this true in the data?

5.1 The relation between A and A^e

Figure 4 reveals that the two series nearly mirror each other.[32] In the beginning of the period, the capital/labor-augmenting technology series grows at a relatively fast rate, whereas the growth rate for the energy-augmenting technology is relatively slow. This goes on until around 1970, i.e., somewhat just before the first oil price shock. After 1970, the energy-augmenting technology grows at a faster rate and the growth rate for the capital/labor-augmenting technology slows down. This continues up to the mid-1980s. Hence, the much-discussed productivity slowdown coincides with a faster growth in the energy-saving technology.

Hence, the growth rates in A and A^e appear negatively correlated, as in the our model of directed technical change. This interpretation also indicates that there are substantial costs associated with improving energy efficiency, since a higher energy efficiency seems to come at the cost of lower growth of capital/labor-efficiency.

5.2 Using the model and production function to estimate the long-run energy share and growth rate

The previous sections suggest that energy-saving and capital/labor saving, captured as shift parameters in an aggregate production function, respond to incentives similar to in the model with directed technical change. Specifically, there seems to be a medium-run negative correlation between the two technology series. We are then, finally, ready to compute the long-run energy-share of income from equation 6.

Figure 6 displays the negative relation between the growth rates, by plotting pairs of average growth rates for three specific periods: the pre-oil crisis (1949–1973), the oil crisis (1974–1984), and the post-oil crisis (1984–2009).

The three points form a straight line with the following relation

$$g_A = 1.5179 - 0.1595 g_A^e. \tag{43}$$

Using $-\frac{\partial g_{A^e}/g_{A^e}}{\partial g_A/g_A} = \frac{1}{0.1595}$ in (6) reveals that the long-run energy share should be 13.76 %. This is significantly higher than the historical average (see the top graph in Figure 1). As

[32]Both series have been normalized so that the initial value is 1.

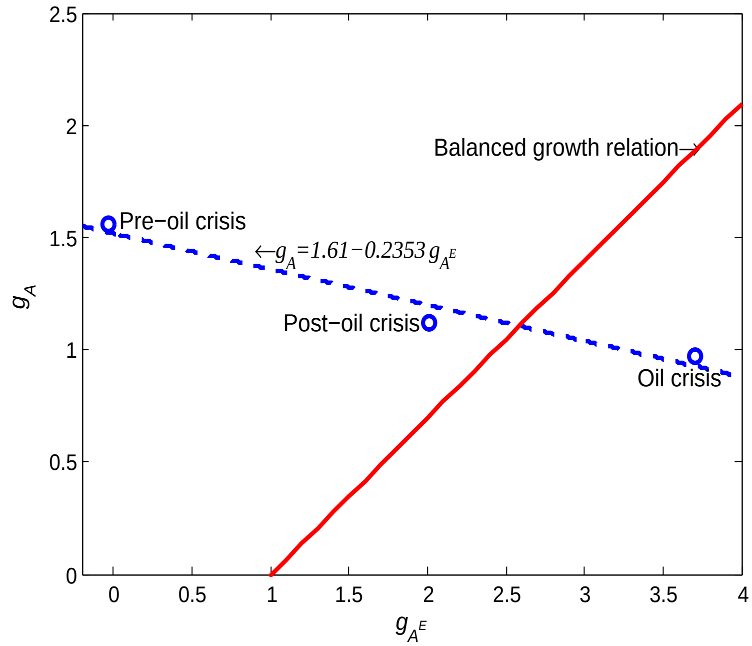

Figure 6: The negative relation is plotting pairs of average growth rates for three specific periods: the pre-oil crisis (1949–1973), the oil crisis (1974–1984), and the post-oil crisis (1985–2009). The positive relation comes from the balanced-growth requirement that the two inputs are growing together. The line with a positive slope is given by (44) with $\alpha = 0.3$ and $\beta = 0.99$.

noted above, this share is only determined by the ratio $\frac{\partial g_{A^e}/g_{A^e}}{\partial g_A/g_A}$, and to obtain an income share more in line with historically observed values, we need a higher slope than -1/0.1595: it would have to be around -1/0.05, i.e., one unit less of growth in A would deliver 20 units more of growth in g_A^e. This is only possible if our estimate based on historical data is off by more than an order of magnitude.

An alternatively method for computing $-\frac{\partial g_{A^e}/g_{A^e}}{\partial g_A/g_A}$ would be to apply an HP filter to the two series, take out the cyclical component, and then regressing the trend growth of energy-augmenting technology g_{A^e}, on the trend growth of the capital-augmenting technology g_A. This, however, results in even higher long-run energy shares. Hence, 0.14 is effectively a lower bound for the share.

Our findings adds macroeconomic support to the findings in Popp (2002), who uses patent data from 1970-1994 to estimate a long run price elasticity between energy prices and energy patents of 0.35. Even though Popp's findings have implications for the impact of factor prices on the direction R&D will take, he does not explicitly compute the tradeoff between the two growth rates for the technologies.

So how does the larger energy dependency affect the growth rate for consumption? To answer this question, we note that in addition to the equations (4)-(5) that dictates a negative relation between the growth rates for A and A^e, balanced growth also requires a positive relation between these two growth rates. As stated in Proposition 1, e falls at rate β on the BGP. Balanced growth then requires that

$$\frac{A_{t+1}}{A_t}\left(\frac{k_{t+1}}{k_t}\right)^{\alpha} = \frac{A^e_{t+1}}{A^e_t}\beta,$$

implying (since n_t has to be constant)

$$f(n)\frac{k^{\alpha}_{t+1}}{k^{\alpha}_t} = f_e(1-n)\beta.$$

A constant saving rate implies, if k is to grow at a constant rate, that k grows at the gross rate $f(n)^{\frac{1}{1-\alpha}}$. Thus, we have

$$f_e(1-n) = \beta^{-1} f(n)^{\frac{1}{1-\alpha}}. \tag{44}$$

This equation describes a positive relation between the growth rates of A and A^e: $1+g_{A^e} = \beta(1+g_A)^{\frac{1}{1-\alpha}}$. Note here that the positive relation is general and always holds on a BGP with our constant-returns-to-scale production function, i.e., it is not specific for Leontief production. Figure 1 plots the positive long-run relation betwen g_{A^e} and g_A as given by (44), where we have set $\alpha = 0.3$ and $\beta = 0.99$. The long-run equilibrium is found at the intersection of the two lines, and they intersect at $g_A = 1.11\%$ and $g_{A^e} = 2.58\%$. This implies a long-run growth rate of consumption of 1.56 percent per year.

Hence, our findings suggest that energy will earn a higher scarcity rent in the future. This, however, does not need to be too harmful for economic growth. In fact, despite oil running out, consumption will grow in the long run, albeit at a somewhat lower rate relative to during the last hundred years.[33]

[33] Table 2 in Appendix A.6 shows how the forecast for the long-run growth rate depends on α and β.

5.3 Peak oil?

In this section the aim is to calculate the transition path for our economy, and to analyze whether the model can produce peak oil. The idea here is rather simple: if the initial value of A (relative to A^e) is low enough, then (i) the economy would mainly accumulate A initially and (ii), as a consequence, due to the implied slow growth in A^e and the strong complementarity between the capital-labor composite and energy, e would have to rise. That is, an initial phase of rising oil use would be possible. To use the model for quantitative predictions, we specify the R&D functions to be

$$f(n_t) = 1 + B n_t^{\phi}, \tag{45}$$

$$f_e(1 - n_t) = 1 + B_e (1 - n_t)^{\phi}. \tag{46}$$

Inverting the first function, substituting, and taking logs yields

$$g_{A^e,t} \approx \log B_e + \phi \log \left(1 - \frac{g_{A,t} - \log B}{\phi} \right). \tag{47}$$

Our linear regression in (6) can be thought of as a linearization of equation (47). The three parameters of these functions are calibrated using the relation between the two technological growth rates found in the first part of the paper. Specifically, we use the average growth rates of the two technology series for the pre-oil crisis (1946–73), the oil crisis (1974–84), and the post-oil crisis (1985–2009). These six observations allow us to identify the three parameters of the R&D functions and the three unobserved values of n.[34] In addition we use $\alpha = 0.3$ and $\beta = 0.99^{10}$. The numerical algorithm is described in the Appendix: Section A.2.

We find that if A_0 and k_0 are set to 45 and 50 percent of their balanced-growth values, respectively, we obtain six decades of increasing fossil fuel use. During the transition, e has to grow fast to keep up with the quickly increasing capital and capital/labor-saving technology growth. The results are plotted in Figure 7 below.

The bottom-right graph shows that energy's share of income is growing steadily from

[34] For expositional clarity, the parameter B_e was, in fact, adjusted slightly upwards to make energy's long-run share of income in the model equal to the estimate in the previous section. With the specified non-linear R&D functions and without this correction, energy's long-run share would be somewhat higher.

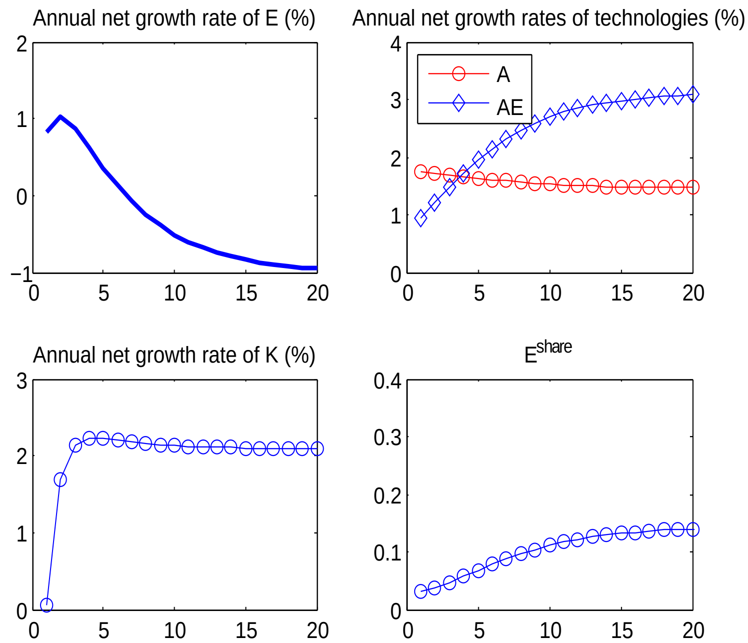

Figure 7: Transition

today's value up toward the long-run value of 0.14.

As we have seen, during the transition path the growth rate of fossil-fuel use is quite different from its balanced-growth value. In contrast, the real interest rate does not depart much from its balanced-growth level: it is not off by more than around half a percentage point at any point in time. Thus, our model produces very slow convergence in one dimension—fossil-fuel use—while it behaves as a standard neoclassical model in terms of interest rates. This is in line with data, where the secular increase in fossil-fuel use has occurred without any apparent trend in the interest rate. An intuitive explanation for why convergence is faster in the capital-labor ratio than in the ratio of the two technologies is that capital depreciates but knowledge does not. Finally, in our model, the price of oil rises steadily, as it satisfies the Hotelling formula: it rises at the real rate of interest. This prediction, which can be regarded as a weakness of the model (and indeed a large class of models), is discussed briefly in the concluding section.

6 Concluding remarks

In this paper we propose a framework for thinking about technological change as an economy's response to the finiteness of natural resources. Through endogenous technical change directed at valuable resources that become scarce, our theory naturally builds in substitutability across (input) goods that is higher ex ante than it is ex post. The theory captures a tradeoff between directing the research efforts toward different inputs and, to the extent the result of these efforts can be measured in the data, we can obtain some insights about the nature of these tradeoffs. We formulate an aggregate production function that is rich enough to allow us to measure the levels of input-saving technologies given data on output, inputs, and prices, and we use it to make quantitative use of our framework.

Along these lines, we thus estimate an aggregate production function in capital, labor, and fossil energy on historical U.S. data in order to shed light on how the economy has dealt with the scarcity of fossil fuel. The evidence we find strongly suggests that the economy actively directs its efforts at input-saving so as to economize on expensive, or scarce, inputs. As outlined, we then use this evidence to inform estimates of R&D technologies, allowing us to make projections into the future regarding energy use and sustainability. Our conclusion is that we can expect fossil fuel to demand a significantly higher share of costs in the future than now; our projection suggests roughly 14 percent. This number should be viewed as a form of upper bound, as alternative energy sources, such as the use of solar power, would bring the share down, at least if one examine such alternatives through the lens of our model. The conclusion that the energy share will be significantly higher in the future in the absence of innovation into new sources of energy will not, however, have major implications for consumption growth. Our model implies a long-run growth rate of consumption that is only somewhat lower than in the past—about 1.56% per year—so from this perspective the energy dependence looks less problematic.

An interesting by-product of our analysis is that the model is able to generate "peak oil", i.e., a rising path for fossil-fuel use (before it will eventually have to drop and go to zero as the resource is depleted). This is accomplished through the low (ex-post) elasticity of substitution between capital/labor and energy: if energy is abundant early on in the industrial development process, it means that output is instead mainly held back by low levels of capital and labor, so that input saving will be directed toward these inputs. As capital and labor are accumulated and used more efficiently, energy use will rise as well due to complementarity. Only as its scarcity will become more problematic will its use then peak and begin falling, at the same time as input saving is instead directed toward energy.

Our framework is rather aggregate and stylized in nature and any results we derive of course will suffer from not introducing more detail, such as to the energy sector, where not only fossil energy is used but also a range of other energy sources. The setup is transparent and tractable, however, and we think of it as a possible blueprint for addressing sustainability issues more broadly in economics. Thus, it is ready to be applied in different contexts and to be made richer, as we see no conceptual or computational difficulties for most extensions of interest. What we currently regard as the most troubling weakness of the model is the implications for prices: for finite resources, Hotelling's (1931) characterization applies, which is that the marginal profit per unit extracted must grow at the real rate of interest. The problem is that in the data, natural-resource prices (less marginal costs) do not seem to follow an exponential trend, thus violating Hotelling's rather robust logic. This model weakness is stark but we know of no convincing solution to this contrast between the basic theory and data. Some arguments have been suggested to make the Hotelling model fit the data better. One is the assumption that marginal cost of extracting the resource falls over time; then there would be a counteracting downward pressure on prices. If new deposits of the resource are discovered, the price will jump down—the "Hotelling rent" will fall—and hence if there are continuous new discoveries one can imagine a continuing downward pressure on prices to the extent these discoveries keep surprising the market. An interesting recent paper (Anderson et al., 2014) formulates oil production as a process with rather fixed extraction rates; in their framework, thus, production does not respond to prices. This helps explain oil prices and production in the data. Investment rates do, however, respond to prices so over a sufficiently long horizon one must still find a mechanism that avoids an exponential trend in prices (less marginal extraction costs).[35] To improve the predictions for prices, one could straightforwardly combine any of these features with the setting in this paper. More generally, a medium-term framework allowing us to jointly account for both oil prices and oil production would be valuable as "oil shocks" are regarded as important drivers of short- to medium-term macroeconomic developments; for a recent cross-country analysis, see Arezki

[35] A form of "short-sightedness" has also been proposed to explain the lack of an upward trend in oil prices; see Spiro (2014).

et al. (2016). The goal of producing such a framework is, however, beyond the scope of the present paper.

Aside from a discussion of price formation, our paper calls for several different kinds of elaboration and follow-ups. One of the most interesting aspects would be policy analysis; it is straightforward to describe a setting with monopolistic competition, patents, and decentralized R&D, so that policy analysis can be conducted. Here it becomes very important to isolate technology spillovers. We do not include such spillovers in the present formulations since a planner would only look at the "reduced form" for technologies (as we argue in Section 2.2.2). When our economies face sustainability issues, then, what is an appropriate government response? If it were not for endogenous technological change, a standard answer would be: no government intervention is needed at all, as the price mechanism should be expected to work, thus making scarce resources more expensive, inducing appropriate saving. However, we argue in this paper that a likely response is technological change, and we know that technology innovation tends to be associated with market frictions (monopoly power and externalities primarily), hence instead suggesting that government policy is indeed needed! Theory for modeling these frictions is available off-the-shelf and could be used here to work on appropriate policy responses to sustainability challenges.

As for further improvements on the setup we have developed here, one would amount to a more detailed study of the postwar data with both demand and supply shocks in the energy market. As indicated above, one would like to add richness to the energy sector. Our modeling of technical progress is also extremely simple; we chose this approach since we focus on planning problems throughout for example, we do not consider spillovers across the different kinds of factor saving, and we have not grounded our calibration in microeconomic estimates of R&D functions. It would be very valuable to take several more steps in this direction, especially when several energy sources are considered jointly. Finally, in this paper we assume that the total amount of R&D is fixed, but it could easily be turned into a choice variable, thus adding a more standard layer of endogenous growth to our setting.

References

Acemoglu, Daron (2002), "Directed Technical Change", *Review of Economic Studies* LXIX, 781–810.

Acemoglu, Daron (2009), *Introduction to Modern Economic Growth*, Princeton, NJ: Princeton University Press.

Acemoglu, Daron, Philippe Aghion, Leonardo Bursztyn, and David Hémous (2012), "The Environment and Directed Technical Change," *American Economic Review* 102(1), 131—166.

Aghion, Philippe, and Peter Howitt, (1992), "A Model of Growth Through Creative Destruction", *Econometrica* LX, 323–351.

Aghion, Philippe, Antoine Dechezlepretre, David Hémous, Ralf Martin, and John Van Reenen (2012), "Carbon taxes, Path Dependency and Directed Technical Change: Evidence from the Auto Industry", mimeo.

Anderson, Soren T., Ryan Kellogg, and Stephen W. Salant (2014), "Hotelling under Pressure", NBER Working Paper 20280.

Arezki, Rabah, Valery Ramey, and Liugang Sheng (2016), "News Shocks in Open Economies: Evidence from Giant Oil Discoveries", mimeo.

Atkeson, Andrew, and Patrick J. Kehoe (1999), "Models of Energy Use: Putty-Putty versus Putty-Clay", *American Economic Review*, 89(4), 1028–1043.

Barsky, Robert and Lutz Kilian (2004), "Oil and the Macroeconomy Since The 1970s," *Journal of Economic Perspectives* 18(4), 115–134.

Berndt, Ernst and David Wood (1975), "Technology, Prices, and the Derived Demand for Energy", *The Review of Economic Studies* 1975, 57, 259–268.

Caselli, Francesco, and John W. Coleman (2006), "The World Technology Frontier", *American Economic Review* 96:3, 499–522.

Cuddington, John T. and Diana L. Moss (2001), "Technological Change, Depletion, and the U.S. Petroleum Industry", *American Economic Review*, 91(4): 1135–1148.

Cuddington, John and Grant Nülle (2014), "Variable Long-Term Trends in Mineral Prices: The Ongoing Tug-of-War between Exploration, Depletion, and Technological Change", *Journal of International Money and Finance*, vol. 42, issue C, 224–252.

Dasgupta, Partha and Geoffrey Heal (1974), "The Optimal Depletion of Exhaustable Resources", *Review of Economic Studies*, 41.

Dasgupta, Partha and Geoffrey Heal (1974), "The Optimal Depletion of Exhaustible Resources.", *The Review of Economic Studies* 41, 3–28.

Dasgupta, Partha and Geoffrey Heal, (1980), *Economic Theory and Exhaustible Resources*, Cambridge University Press.

Drandakis, Emmanuel, and Edmund Phelps (1966), "A Model of Induced Innovation, Growth and Distribution", *Economic Journal* 76, 823–40.

Golosov, Michael, John Hassler, Per Krusell, and Aleh Tsyvinski (2014), "Optimal Taxes on Fossil Fuel in General Equilibrium", *Econometrica*, 82:1, 41-88.

Griffin, James and Paul Gregory (1976), "An Intercountry Translog Model of Energy Substitution Responses", *American Economic Review* 66(5), 845–857.

Hémous, David (2013), "Environmental Policy and Directed Technical Change in a Global Economy: The Dynamic Impact of Unilateral Environmental Policies?", mimeo.

Hicks, John (1933), *The Theory of Wages*, London: McMillan.

Hotelling, Harold (1931), "The economics of exhaustible resources", *Journal of Political Economy* 39, 137–175.

Hudson, Edward and Dale Jorgenson (1974), "U.S. Energy Policy and Economic Growth, 1975–2000", *Bell Journal of Economics* 5, 461–514.

Jones, Charles, (2002), *Introduction to Economic Growth*, Second Edition, W.W. Norton.

Kennedy, Charles (1964), "Induced Bias in Innovation and the Theory of Distribution", *Economic Journal* LXXIV, 541–547.

Kilian, Lutz (2008), "The Economic Effects of Energy Price Shocks", *Journal of Economic Literature* 46(4), 871–909.

Krautkraemer, Jeffrey A. (1998), "Nonrenewable Resource Scarcity", *Journal of Economic Literature*, Vol. 36, No. 4, 2065–2107.

Krusell, Per, Lee Ohanian, José-Víctor Ríos-Rull, and Giovanni Violante (2000), "Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis", *Econometrica* 68(5), 1029—1053.

Nordhaus, William and Joseph Boyer (2000), *Warming the World: Economic Modeling of Global Warming*, MIT Press, Cambridge, Mass.

Popp, David (2002), "Induced Innovation and Energy Prices", *American Economic Review* 92, 160–180.

Solow, Robert (1957), "Technical Change and the Aggregate Production Function", *The Review of Economics and Statistics* 39(3), 312–320.

Solow, Robert (1974), "Intergenerational Equity and Exhaustible Resources", *The Review of Economic Studies* 41, 29–45.

Spiro, Daniel (2014), "Resource prices and planning horizons", *Journal of Economic Dynamics and Control*, Vol. 48, 159–175.

Stern, David I. (2012), "Interfuel Substitution: A Meta-Analysis", *Journal of Economic Surveys*, 26, 307—331.

Stern, David I. and Astrid Kander, (2012), "The Role of Energy in the Industrial Revolution and Modern Economic Growth", *The Energy Journal* 33, 125–152.

Stiglitz, Joseph (1974), "Growth with Exhaustible Natural Resources: Efficient and Optimal Growth Paths", *The Review of Economic Studies* 41, 123–137.

Stiglitz, Joseph (1979), "A Neoclassical Analysis of the Economics of Natural Resources", in Scarcity and Growth Reconsidered, Resources for the Future, (Ed.) Kerry Smith, Chapter 2, 36–66. Baltimore: The Johns Hopkins University Press.

A Appendix

A.1 Fossil energy use

The fact that the price per ton of oil is several times higher than coal can be interpreted as representing higher general efficiency of oil. To estimate these efficiency differences, we calculate the average price per Btu for oil and natural gas relative to the coal price over the period 1949-2009. The average price per Btu of coal is on average roughly 3.5 times higher for oil and 1.6 times higher for natural gas. The assumption implies that coal is priced at marginal extraction costs and that this determines the price of oil and gas. The latter two are priced so that it compensates for the higher efficiency, providing rents to extractors due to its lower extraction cost per unit of efficiency.[36] Denoting the period-t consumption of coal, oil and gas respectively by e_t^c, e_t^o and e_t^g, the sum of all fossil fuel consumption is $e_t^c + e_t^o + e_t^g$. We then calculate the Btu equivalent fossil fuel consumption to be $e_t^c + 3.5 * e_t^o + 1.6 * e_t^g = e_t$.

We compute the fossil fuel composite price according to the following formula:

$$p = \frac{p^c}{1} * \frac{e_t^c}{p_t} + \frac{p^o}{3.5}\frac{e_t^o}{e_t} + \frac{p^g}{1.6}\frac{e_t^g}{e_t},$$

where the indices c, o and g respectively denotes coal, oil and gas.

[36] The supply of oil and gas will depend on how the supply side is organized, which we we abstract from here.

A.2 Algorithm for solving for transition dynamics

Specifically, the following algorithm has shown to work speedily and accurately.

1. Fix the parameters of the model including initial conditions R_0, A_0, k_0 and A_0^e.

2. Guess on a sequence $\{\hat{e}_t\}_{t=0}^{\infty}$. Make the sequence satisfy

   (a) $\sum_{t=0}^{\infty} \beta^t \hat{e}_t = R_0$ (it exhausts the resource); and

   (b) $\hat{e}_0 = e_0 = A_0 k_0^{\alpha} / A_0^e$ (the Leontief condition).[37]

3. Guess on κ.

4. Directly compute each s_t from (28) and each n_t from (32).

5. To verify the guess, first

   (a) construct the sequences $\{k_{t+1}, A_{t+1}, A_{t+1}^e\}_{t=0}^{\infty}$ (straightforward using the computed saving rates and R&D labor allocations); and

   (b) using the Leontief condition between inputs, find the implied updated sequence $\{\hat{e}_t\}_{t=1}^{\infty}$.

6. To update, first iterate on κ until the implied updated fossil-fuel sequence (just obtained) is resource-feasible.

7. Next, if the updated sequence $\{\hat{e}_t\}_{t=0}^{\infty}$ is not equal to the originally guessed sequence, go to step 2 and use a new sequence $\{\hat{e}_t\}_{t=0}^{\infty}$ in the following way: use the same $\hat{e}_0$ and, for the rest, take a weighted average of the earlier guess and the just obtained update, and use the updated κ as well.

[37] To implement this more specifically, select a T (like 25 or 50) and guess "freely" on a path up to T. Then for $t > T$ assume we are on a balanced path so that $\hat{e}_t = \hat{e}_T$ for all $t > T$. Thus, in computing n one can compute the infinite sums explicitly (or approximate them well also for t below but close to T.

A.3 Transforming the model

Let $\tilde{k}_t = k_t/(A_0^e R_0)$, $\tilde{A}_t = A_t/(A_0^e R_0)^{1-\alpha}$, $\tilde{A}_t^e = A_t^e/A_0^e$, and $\tilde{e}_t = e_t/R_0$. Then the former maximization problem becomes a constant plus

$$\max_{\left\{\tilde{k}_{t+1}, \tilde{A}_{t+1}, \tilde{A}_{t+1}^e, \tilde{e}_t, n_t\right\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t \log\left(\tilde{A}_t \tilde{k}_t^\alpha - \tilde{k}_{t+1}\right)$$

subject to

$$\tilde{A}_t \tilde{k}_t^\alpha = \tilde{A}_t^e \tilde{e}_t, \quad \tilde{A}_{t+1} = \tilde{A}_t f(n_t), \quad \tilde{A}_{t+1}^e = \tilde{A}_t^e f_e(1-n_t), \quad \text{and} \quad \sum_{t=0}^{\infty} \tilde{e}_t = 1,$$

where now $\tilde{A}_0^e = 1$. The solution of this maximization problem does not depend on A_0^e. In particular, the paths for $\tilde{e}$, n, and $s_t \equiv k_{t+1}/(A_t k_t^\alpha) = \tilde{k}_{t+1}/(\tilde{A}_t \tilde{k}_t^\alpha)$ will not depend on A_0^e.

A.4 Estimation

We find the parameters through the following iterative algorithm.

1. Choose a tolerance limit ζ.
2. Make a guess for ε and θ.
3. Compute Σ from the first order condition of (42) with respect to Σ.
4. Compute new values for ε and θ from the first order conditions of (42) with respect to ε and θ respectively, and compare the updated values to the previous values for ε and θ.
   - If the absolute value of the difference is smaller than ζ- stop. The solution has been found.
   - Otherwise, goto step 3 with the updated values for ε and θ.

A.5 Energy-saving in the manufacturing sector

Figure 8 plots the evolution of the level of the energy-saving technology for the manufacturing sector. The kink is less pronounced, but the energy-saving technology is clearly growing at a faster rate after the oil-shocks than before.

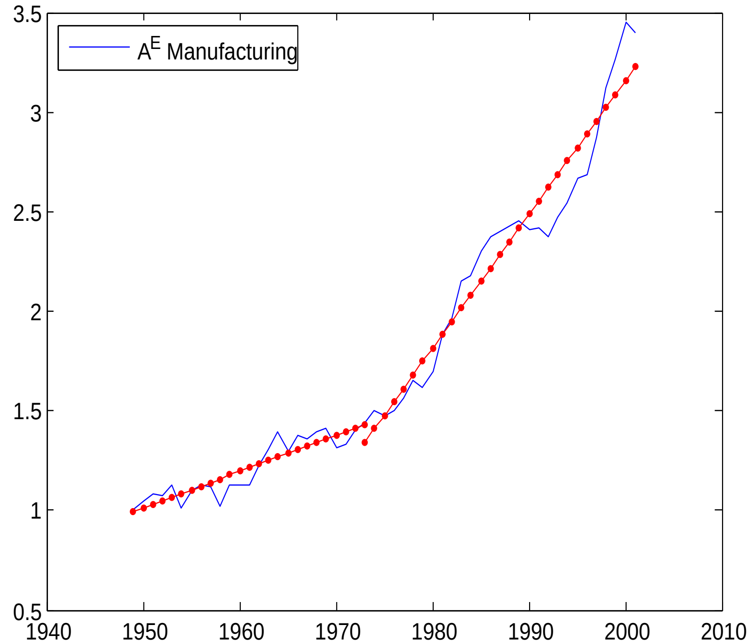

Figure 8: Energy-saving technology with an elasticity of 0.02. Data on industrial-sector energy consumption is taken from the U.S. Energy Information Administration, and the data on output in U.S. manufacturing is from the Board of Governors of the Federal Reserve System (i.e., the FRED database).

A.6 Sensitivity analysis

Section 5.2 shows that the long-run growth rate is given by the intersection of equations (43) and (44). Table 2 shows how this growth rate is affected by different parameter values.

Clearly, the growth rate is increasing in both β and α.

Table 2: Long-run growth rates for different parameter values

	$\beta = .97$	$\beta = .98$	$\alpha = 0.35$	$\alpha = 0.4$
g	1.14	1.36	1.67	1.78

The long-run growth rate as a function of the discount factor, β, and capital's share of output, α.