Journal of Environmental Economics and Management 89 (2018) 1–17

Contents lists available at ScienceDirect

Journal of Environmental Economics and Management

journal homepage: www.elsevier.com/locate/jeem

Rebound, directed technological change, and aggregate demand for energy ☆

Rob Hart

Department of Economics, Swedish University of Agricultural Sciences, Box 7013, 750 07 Uppsala, Sweden

ARTICLE INFO

Article history:
Received 20 November 2015
Revised 17 December 2017
Accepted 10 March 2018
Available online 26 March 2018

JEL codes:
O41
Q43

Keywords:
Energy
Technology
Directed technological change
Structural change
Rebound

ABSTRACT

We analyse the long-run role of energy in aggregate production. The factor share of energy has been remarkably constant, despite the relative decline in the price of energy. We analyse possible explanations for this observation, ruling out the idea that endogenous directed technological change has led to a failure of energy-augmenting technology to keep pace with labour-augmenting technology. Instead we propose a model in which a combination of income and substitution effects has driven both shifts in consumption patterns towards existing energy-intensive goods and the emergence of new such goods.

1. Introduction

Over the last 150 years the increase in global aggregate energy efficiency has been modest, even though the energy efficiency of individual production processes (such as the generation of motive power or artificial light from primary energy inputs) increased greatly over the same period. This seems paradoxical: why doesn't an increase in the efficiency of individual production processes result in an increase in aggregate efficiency? In this paper we argue that the explanation for the paradox is a shift in consumption patterns over time towards energy-intensive goods, including the endogenous introduction of new product varieties which are intrinsically energy-intensive.

Our model is relevant to three fields which have rarely been connected in the literature: the explanation of the high long-run aggregate elasticity of demand for energy; growth and structural change; and the rebound effect. We now discuss these in turn. Regarding aggregate energy demand, we discuss two approaches in the literature, neither of which allows for changes in consumption patterns between final goods differing in energy intensity: the first is the putty–clay model of Atkeson and Kehoe (1999), and the second is the approach based on directed technological change (henceforth DTC).[1] In Atkeson and Kehoe (1999)

☆ Thanks to the Swedish Energy Agency (project 37758-1) and the Swedish EPA (project *Economic instruments and consumption*) for financial support. Thanks also to several anonymous referees, and to seminar participants at the Tinbergen Institute and the EAERE conference in Helsinki, for input. Thanks also to Sjak Smulders, Gerard van der Meijden and Peter Mulder for valuable discussions.

E-mail address: rob.hart@slu.se.

[1] See section 5 for further analysis. A third approach is the putty–putty model of Pindyck and Rotemberg (1983); for a discussion of this approach see Atkeson and Kehoe (1999).

https://doi.org/10.1016/j.jeem.2018.03.002
0095-0696/

long-run aggregate production is simply a Cobb–Douglas function of labour, capital, and energy, but capital is putty–clay and an increase in energy efficiency requires replacement of the capital stock. There is therefore a lag between an unexpected change in the energy price and the resultant change in energy efficiency. This approach cannot solve our paradox since the explanation of the failure of aggregate energy efficiency to rise is that product-level energy efficiency has failed to rise. Recall that the paradox is that product-level increases in energy efficiency are not reflected in the aggregate. In the standard approach to modelling DTC (see for instance Acemoglu, 2002), investment in knowledge augmenting different factors is in proportion to the shares of those factors.[2] Since energy is complementary to labour–capital, an increase in energy price leads to an increase in energy share, hence energy-augmenting knowledge is boosted, pushing the energy share back down. Hart (2013) shows that if the factor-augmenting knowledge stocks are produced independently of one another then this buffering mechanism is perfect in the sense that the long-run factor shares of complementary inputs are independent of the long-run quantities of these inputs; we could say that the 'long-run production function' is Cobb–Douglas. In the light of this result it is not surprising that the seminal paper in the modelling of DTC and energy demand—Smulders and de Nooij (2003)—implicitly assumes independent knowledge stocks, thus yielding the long-run fixed-share property. Many subsequent authors—including Gerlagh (2008), Fischer and Newell (2008), Gans (2012), and Hassler et al. (2012)—make the same assumption and thus their models have the same property.[3] Again this approach fails to solve the paradox, because again its explanation for the failure of aggregate energy efficiency to rise is the failure of product-level efficiency to rise.

If product-level energy efficiency has risen while aggregate energy efficiency has not, the only possible explanation is a shift in consumption patterns towards energy-intensive goods. It has been well known since Engel (1857) that economic growth goes hand-in-hand with systematic shifts in patterns of consumption, driven by income effects: as income increases, the share of necessities such as food declines while luxury goods increase their share.[4] But luxury is a relative concept, and Matsuyama (2002) argues that as productivity improves, households constantly expand the range of goods they consume, as more and more goods become affordable. He models this process using a household utility function with lexicographic preferences, i.e. households expand their consumption from one good to the next irrespective of relative prices.[5] If these goods—introduced successively—are successively more energy-intensive then this process could explain the data.

Shifts in consumption patterns towards energy-intensive goods may also be driven by substitution effects. Return to the period 1870–1970. Since the price of primary energy failed to rise during this period, while energy efficiency rose substantially, *ceteris paribus* we would expect the relative price of energy-intensive products to decline. This could induce substitution towards such products. A related process is studied by Acemoglu and Guerrieri (2008), who model substitution between labour and capital with the aim of explaining both the constant capital share and structural change. They posit two sectors with fixed—but different—capital shares, and show that if the elasticity of substitution between the sectors is less than one (in their calibration it is approximately 0.5) then capital deepening will cause relatively higher output from the capital-intensive sector, but a lower share of income to that sector. The net effect of these shifts is to leave the capital share roughly constant.

Many authors study structural change and energy use in an international context, by comparing the ratio of energy input to economic output across sectors: see for instance Mulder and de Groot (2012) and Mulder (2015). Mulder and de Groot (2012) find that energy efficiency is a more important driver of energy demand than structural change, and Mulder (2015) links the changes to changing global production patterns with international trade. This literature is of limited relevance for our analysis because structural change is defined as shifts in production between sectors, of which examples are (from Mulder and de Groot) 'Food and beverages' and 'Motor vehicles, trailers, and semi-trailers'. It follows that these data do not capture shifts towards more energy-intensive goods *within the same sector*, such as a shift towards heavier and more powerful cars. Our data show that such shifts may be very important for aggregate energy demand.

Recall that our central claim is that consumption patterns have shifted over time towards energy-intensive goods (including completely new goods), at the same time as energy productivity of individual production processes has increased. We focus on two possible mechanisms which could have driven this shift: the first possibility is that the shift was caused by the increase in energy productivity, implying that there is a substitution effect as energy-intensive goods become relatively cheaper (cf. Acemoglu and Guerrieri), and the second is that increasing labour productivity—and hence rising incomes—caused the shift, as richer consumers chose more energy-intensive goods (i.e. an income effect, as in Matsuyama). It is important for energy policy to know the strengths of these mechanisms, because if increases in energy productivity cause substitution to energy-intensive goods then the increases do not yield the hoped-for reduction in energy consumption: we have rebound.[6]

To define rebound, assume an economy in which total primary energy use is R, and focus on production of good i using (among other inputs) augmented primary energy flow $A_{ri}R_i$. Rebound is present when an increase in efficiency A_{ri} by a factor x

[2] This result is demonstrated by Hart (2013).

[3] In the case of Smulders and de Nooij (2003) the fixed-share result seems paradoxical, as the authors set out to explain why the energy share has *declined* in recent decades in a selection of countries. The explanation is that the authors start off their simulated economy away from the long-run balanced growth path (b.g.p.), and the decline in share occurs on the transition path.

[4] See Houthakker (1957) for a discussion of Engel's law.

[5] Assume rising income. Good 1 is food, and good 2 is not consumed at all until income is sufficient to satiate the desire for food. At this point, consumption of good 2 begins; when desire for that good is satiated, consumption of good 3 begins. Etc.

[6] An alternative explanation is that some factor other than increases in energy productivity has made energy-intensive goods more attractive; for instance, it could be that the productivity of labour and capital inputs has risen more rapidly in energy-intensive sectors than in other sectors, hence driving down the relative price of these goods. We return to this alternative in Section 4.2.

leads to a reduction of R by less than $R_i(1 - 1/x)$, where the baseline reduction may be described as that arising from a naive or 'engineering' analysis. Given this definition then even if preferences are Leontief (implying no substitution effects) there will always be some rebound given that an increase in energy efficiency of any good leads to a reduction in the price of that good, and therefore an increase in real income. If this extra income is respent on goods of average energy intensity—i.e. around 6 percent—then we have 6 percent rebound. The potential for more substantial rebound, however, depends primarily on a combination of two things: that the good in question is energy-intensive, and that the efficiency increase causes a significant shift in consumption patterns towards that good.[7]

It is widely agreed that the rebound literature suffers from a lack of coherence (see for instance van den Bergh, 2011; Turner, 2013; Broberg et al., 2015; Gillingham et al., 2016), which Turner attributes to a lack of solid theoretical foundations. Ideally a model could be developed in which total production would be divided into a range of alternative goods of varying energy intensity, and preferences and technology defined for each good. Then given input productivities and endowments we could derive analytical solutions for the allocation of resources, and for the effects of (for instance) exogenous increases in the efficiency of certain inputs or processes. Furthermore—and crucially—the model structure would be rooted in empirical reality. How does existing work match up to this ideal? We look briefly at three approaches.

Borenstein (2015) (Section 4) analyses rebound for a specific service provided by an appliance and an energy-intensive input in an economy where there is a budget constraint and remaining expenditure is spread over a good (or set of goods) of average energy intensity. The case-by-case analysis is enlightening, but there is no model of preferences and technology defined across all the goods in the economy. Chan and Gillingham (2015) assume (in their basic model) a representative consumer who spends her income on two alternative fuels and an energy-free aggregate good. Here the allocation does depend explicitly on preferences and the technology, but the empirical grounding required in our ideal model is lacking: assuming that goods are either 100 percent or 0 percent energy is highly problematic empirically, when Borenstein (2015) shows that—even in an energy-intensive sector such as transport—fuel accounts for a relatively modest proportion of marginal costs, while on the other hand no goods are completely energy-free in their production. Finally, Broberg et al. (2015) build and validate a CGE model of the economy in which differentiated goods are produced based on a series of nested CES functions of labour, capital, materials, and energy, and the same goods are demanded by consumers based on nested CES utility functions. Here the results are based on numerical models; furthermore, by construction income effects are ruled out in these models, since households are assumed to have Gorman preferences to allow straightforward aggregation.[8] This handicaps such models when applied to analysis of long-run growth and energy use, but is not so crucial when they are applied to rebound, where substitution effects are most important.

Our contribution is to put the rebound mechanism into a long-run general equilibrium context with non-homothetic preferences so that patterns of consumption evolve over time due to both substitution effects (changes in relative prices) and increasing income. And by positing a continuum of goods with variable energy intensity, and allowing for the endogenous introduction of new varieties, we go far beyond existing papers such as Chan and Gillingham (2015) in terms of the model's ability to match stylized facts about sectoral patterns of energy consumption; furthermore this structure gives rise to novel, intuitive results.[9] For instance, in our model economy the rebound effects of increased energy efficiency of individual products are equally likely to be negative as they are to be positive: when products of lower-than-average energy intensity become more energy efficient, substitution towards these products causes an additional reduction in energy use compared to the baseline. On the other hand, when products are both very energy-intensive and on the cusp of affordability (consider supersonic passenger flight, or space tourism) then increases in energy efficiency may have very large positive rebound effects, and even backfire, as consumption moves towards these products from less energy-intensive alternatives. The rebound effect from an increase in overall energy efficiency in our model is 50 percent, which corresponds well to Broberg et al. (2015), who find effects ranging from 30 to 70 percent depending on the parameterization.

When we calibrate the model it does a reasonable job of accounting for patterns of long-run primary energy demand. However, given the very simple pattern in the data the success of the calibrated model is perhaps not a very startling achievement; furthermore, the model is simple and leaves out many factors which may (through further research) be shown to be crucial to explaining patterns in energy demand. Perhaps the most important simplification is that we assume—as mentioned above and discussed in Section 4.2—that increases in labour-augmenting knowledge are uniform across products. If the productivity of labour (or labour–capital) increases faster in energy-intensive sectors than in labour-intensive sectors this would be a powerful factor pushing the relative price of energy-intensive products down and hence driving substitution towards their consumption.

The remainder of the paper is organized as follows. In Section 2 we present data. In Section 3 we present and solve the model of expanding variety. In Section 4 we discuss the implications of the expanding variety model, and in Section 5 we compare the model to alternatives, the Atkeson–Kehoe model and models based on DTC. Section 6 concludes.

[7] There are other possibilities. The respend might be placed disproportionately on energy-intensive goods, but we agree with Borenstein (2015) that this is implausible in the absence of evidence. Another possible type of effect is the 'economywide' effect, an example of which is that if the demand curve for primary energy shifts left by x MJ/year, the price of primary energy may fall and therefore primary energy use may fall by less than x. Finally, an increase in energy-augmenting knowledge in production may lead producers to substitute energy inputs for labour–capital, as analysed by Lemoine (2016).

[8] See for instance Acemoglu (2009) chapter 5 for a textbook discussion of the issue of aggregation in macroeconomic models.

[9] Nevertheless, it should be pointed out that the production structure is highly stylized when compared to (for instance) CGE models such as Broberg et al. (2015).

2. The data

We discuss data regarding four separate subjects: long-run data on production, energy use, and energy price; biased technological change; variations in energy intensity across sectors; and structural change. And we make four claims: the long-run energy share is approximately constant or slowly declining; energy-augmenting knowledge grows at least as fast as labour-augmenting knowledge; there is a spread of energy intensities across products; and there is structural change in consumption patterns towards energy-intensive goods.

2.1. The long-run energy share

Fig. 1—showing data on production, energy use, and price—suggests that although the short-run elasticity of substitution between energy and the other inputs to production is small, the long-run elasticity is close to 1, implying that the long-run energy share of GDP is constant. Furthermore, both Griffin and Gregory (1976) and Pindyck (1979) use cross-section data across countries and estimate an own-price elasticity of energy use of 0.8, and Griffin and Gregory (1976) explicitly consider the Cobb–Douglas function and cannot reject it. If we assume that price differences between countries are persistent this supports the idea of (close to) unit long-run elasticity.

More recent work—such as Smulders and de Nooij (2003) and Csereklyei et al. (2016)—tends to focus on energy intensity rather than the energy share. Nevertheless, Csereklyei et al. find limited evidence that the long-run cost share of energy has declined slowly. Our data suggests that the energy share has been more constant, because we exclude non-market energy and energy from fodder. If we include fodder then—as shown by an anonymous referee—expenditure and quantity are raised significantly prior to 1880, reducing the growth rate of expenditure up to 1880. After 1880 the share of fodder inputs is small and the results are scarcely affected. The inclusion of fodder would raise many questions and difficulties, including how to separate fodder for food production from fodder for generation of motive power, and how to price fodder inputs. On the other hand, we do not include energy from nuclear power and renewables either: including nuclear would cause energy expenditure to increase significantly subsequent to 1960.

2.2. Bias of technological change

Now we consider evidence regarding the bias in technological change, i.e. the historical growth rate of energy-augmenting knowledge compared to labour-augmenting knowledge. Recall from the introduction that we distinguish between changes in aggregate energy efficiency and changes in the energy efficiency of production of individual products; here we are interested in

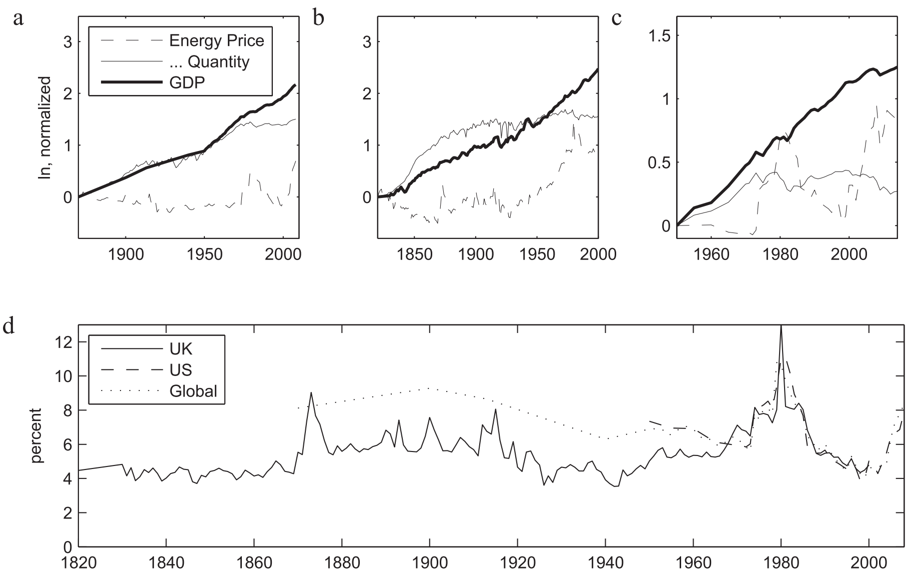

Fig. 1. Primary energy from combustion. Price and quantity per capita compared to GDP per capita: (a) global; (b) U.K.; (c) U.S. And (d) the factor shares in the three cases. Global product data from Maddison (2010). Energy: Coal, oil, natural gas, and biofuel. Global fossil quantity data from Boden et al. (2012); UK data from Warde (2007). Oil price data from BP (2012). Coal and gas price data from Fouquet (2011); note that these data are only for average prices in England; we make the (heroic) assumption that weighted average global prices are similar. Biofuel quantity data from Maddison (2003). Biofuel price data from Fouquet (2011); again, we assume that the data are representative for global prices, and we extrapolate from the end of Fouquet's series to the present assuming constant prices. Combustion of biofuels in the U.K. is negligible over the entire period. U.S. data from the EIA, downloaded 4 January 2017: quantity data from http://www.eia.gov/totalenergy/data/monthly/pdf/sec1.pdf table 1.3, coal price data https://www.eia.gov/coal/data.php (prices, coal rank), gas price data https://www.eia.gov/opendata/qb.php?category=461217&sdid=NG.N9190US3.A, and petroleum price data http://www.eia.gov/totalenergy/data/annual/showtext.php?t=ptb0524.

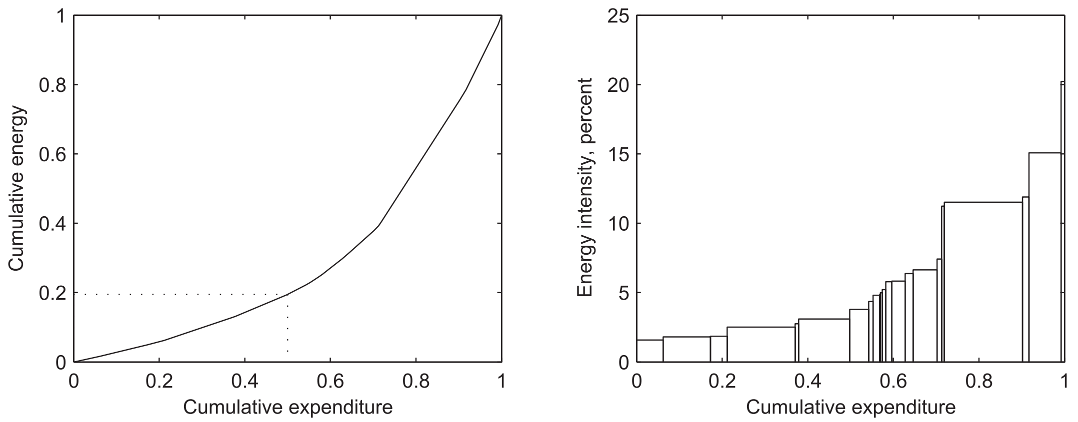

Fig. 2. Cumulative energy use and energy intensity plotted against cumulative expenditure when consumption products are sorted in order of increasing energy intensity. All the axes are normalized. Regarding energy intensity, we only have data on relative intensities, and we normalize to give an average intensity of 6 percent. Data from Mayer and Flachmann (2011). The products—in order of increasing energy intensity—are Education services; Health services; Health services and social work; Other services; Cultural and sport services; Retail and wholesale trade; Hotel and restaurant services; Office and electrical machinery; Paper and publishing; Water transport; Auxiliary transport services; Other land transport; Furniture, jewellery, musical instruments etc.; Other products; Textiles and furs; Food and tobacco; Agricultural products; Transport via railways; Habitation; Chemical products, rubber, and plastic; Motor transport; Air transport.

the latter. The two may differ greatly due to structural change: if products differ in energy intensity, then increases in energy efficiency for individual products may be masked in the aggregate data by shifts towards products of higher energy intensity. We conclude that there is strong evidence that growth in energy-augmenting knowledge has been at least as fast as growth in labour-augmenting knowledge, probably faster.

We analyse energy-efficiency in three cases: U.S. automobiles; light production in the U.K.; and generation of motive power in general. Our evidence on the U.S. automobile industry is taken from Knittel (2011), who estimates that—for a vehicle of fixed characteristics in terms of weight and engine power—fuel economy would have increased by 58 percent over the period 1980–2006 due to technological change, i.e. by 1.8 percent per year. This rate is slightly greater than the average growth rate of U.S. labour productivity over the same period, which was 1.7 percent per year.[10]

Our evidence on artificial light is based on Fouquet and Pearson (2006). Light is a convenient product category for analysis since light is a consumption good which is rather homogeneous and unchanging over very long timescales, and the energy efficiency of its production is easily measured. Fouquet and Pearson conclude that the efficiency of light production in the U.K. (measured by lumen produced per watt of energy used) increased 1000-fold from 1800 to 2000; the productivity of labour in the U.K. over the same period rose by a factor of around 12 (based on Maddison, 2010).

Finally we turn to the production of motive power from fossil fuels, a very large sector. Motive power is typically an intermediate good rather than a final good, nevertheless increases in the efficiency with which energy inputs are used to generate motive power are very likely to be reflected in the overall efficiency with which energy is used to generate final goods, as long as the final goods are homogeneous and do not change over time. In the 19th century motive power was largely generated by steam engines, while over the last 100 years we must consider electric power generation and the internal combustion engine. Regarding steam engines, sources such as Hills (1993) suggest that their efficiency in generating power from coal inputs increased steadily from their invention in the early 1700s up to 1900, and by a factor of around 20 over the entire period; this growth in efficiency is again more rapid than the growth in labour productivity over the same period. Subsequently, the efficiency of coal-fired power stations has continued to increase but at a declining rate; see for instance Yeh and Rubin (2007) for detailed evidence.

2.3. Energy intensities across consumption categories

In Fig. 2 we show variations in energy intensity across consumption categories. In Fig. 2(a) we see that if we divide consumption into two equal parts, one energy-intensive the other not, then the low-energy-intensity consumption accounts for just under 20 percent of energy consumption. In 2(b) we see the energy intensity and expenditure share of different consumption categories: different types of services—of low energy-intensity—account for more than half of expenditure, while the two major energy-intensive categories are habitation and motor transport, and the final category (with highest intensity but only a small expenditure share) is air transport.

[10] We calculate the latter figure using data from stats.oecd.org, 'Labour productivity growth in the total economy'.

2.4. Structural change

We have established that aggregate energy efficiency has grown rather slowly (especially prior to 1974), while the energy-efficiency of production of specific energy-intensive goods has tended to increase more rapidly. Logically, this implies that there must have been simultaneous shifts in consumption patterns towards energy-intensive goods. Here we consider direct evidence for such shifts using historical and more recent data. We know—recall Engel (1857) and Matsuyama (2002)—that economic growth is accompanied by systematic changes in consumption patterns, as households expand the range of goods they consume. But are the new goods more energy intensive?

We consider evidence from first U.K. lighting, and then U.S. transport. Fouquet and Pearson (2006) find that per capita consumption of artificial light in the U.K. rose by a factor of 7000 between 1800 and 2000. This factor should be compared to the approximately 12-fold increase in per capita GDP over the same period; without shifts in consumption patterns, consumption of all products should have risen by this factor over the period. Given the 1000-fold increase in efficiency noted earlier, if light consumption had tracked GDP growth then energy use per capita for lighting would have fallen by a factor of approximately 80, but instead it increased by a factor of 7.

Regarding transport—which accounts for around 30 percent of U.S. primary energy consumption—we make three observations. Firstly, in the period 1960–1999 there was a steady increase in passenger-miles per capita, roughly tracking GDP growth, and an almost equally rapid increase in energy use per capita in the sector. This is illustrated in Fig. 3(a); note that passenger-miles per unit of energy rose by just 0.6 percent per year over the period, whereas GDP and passenger-miles per capita rose by around 2.1 percent per year. These data taken in isolation would be broadly consistent with balanced growth (no structural change) and a very slow rate of increase in energy efficiency. But our second and third observations paint a different picture. The second observation is that—at the same time as the steady increase in passenger-miles—there was a substantial shift in travel patterns away from cars, and towards heavier vehicles and air travel which are more energy-intensive. In

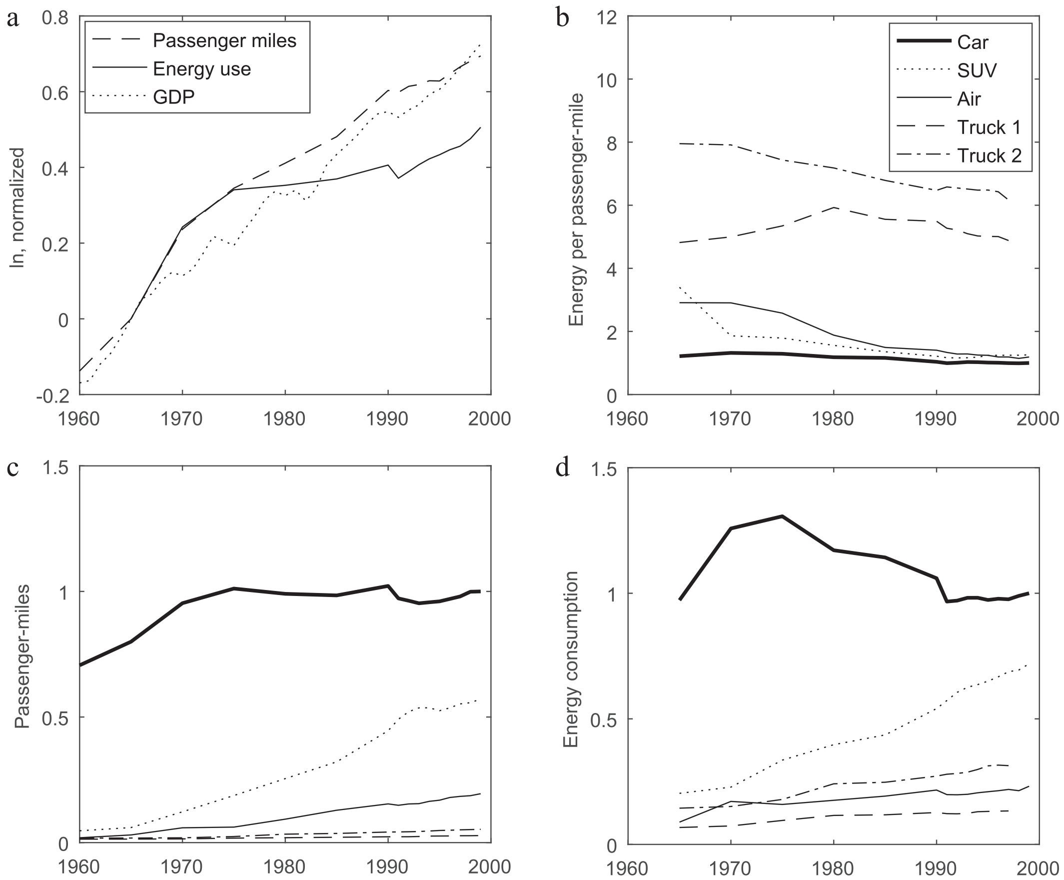

Fig. 3. Aggregate data for passenger-miles and energy consumption in the U.S. for private vehicles and air travel (combined): (a) Growth in total passenger–miles and energy use, compared to GDP; (b) Changes in energy per passenger–mile for different transport categories; (c) Changes in passenger-miles by category; (d) Total energy consumption by category. All data from U.S. Department of Transportation, National Transportation Statistics, except GDP data from Maddison (2010). Note that we have renamed the DOT categories: our *SUV*; *Truck 1*; and *Truck 2* are (respectively) *Other 2-axle 4-tire vehicles*; *Truck, single-unit 2-axle 6-tire or more*; and *Truck, combination* (i.e. an articulated vehicle).

Fig. 3(c) we see how car travel grows steeply up to 1970 and then flatlines, whereas the other categories—especially SUVs and air travel—grow rapidly throughout the period, starting from almost nothing. The relative energy-intensities (per passenger-mile) are illustrated in 3(b), where we see that cars are the least energy-intensive category, hence the shift out of car travel drives up total energy consumption. Our third observation is that there have been major changes within transport categories, changes which have typically also contributed to increasing fuel consumption. Between 1970 and 1997 (the period for which we have consistent data) average fuel efficiencies of cars, SUVs, artics, and trucks—measured as passenger-miles per unit of fuel—increased by 1.0, 1.5, 0.9, and 0.1 percent per year respectively (Fig. 3(b)), much less that the 1.8 percent per year increase in the efficiency of new vehicles (vehicle-miles per unit of fuel) noted by Knittel (2011) for the period 1980–2006. The difference is surely due to the increases in weight and power of vehicles within each category, as documented by Knittel. On the other hand, the efficiency of air travel increased by 3.4 percent per year over the same period: here the increase in aircraft efficiency has been complemented by an increase in load factor (i.e. actual passengers per seat) of 1.1 percent per year over the same period, as documented by Lee et al. (2001), leaving a 2.3 percent increase to be explained by aircraft efficiency.[11]

3. A model of expanding variety, structural change and energy demand

In this section we develop our model of structural change and energy demand. Our aim is to build a model capturing the above four characteristics of the data: the long-run energy share should be approximately constant or slowly declining; energy-augmenting knowledge should grow at least as fast as labour-augmenting knowledge; there should be a spread of energy intensities across products; and there should be structural change in consumption patterns towards energy-intensive goods, including the endogenous introduction of completely new such goods.

3.1. The model

The model is parsimonious. We begin with the production side, then turn to the consumption side. There is an infinite continuum of goods which can potentially be made, indexed by i. The goods are made by perfectly competitive firms using labour and energy, so price is equal to marginal cost. At a given time the production function of the representative firm for good i is

$$y_i = \min\{A_l l_i, (A_r/i) r_i\},$$

where A_l and A_r are productivity indices for labour and energy, and l_i and r_i are quantities of labour and energy (prices w_l and w_r). So we have a Leontief production function in which labour productivity is the same across all products, whereas energy productivity is declining in i. There are constant returns to scale in labour and energy, so payments to labour and energy are equal to revenue. Define total labour and energy use as L and R respectively, so

$$L = \int_0^\infty l_i \mathrm{d}i \quad \text{and} \quad R = \int_0^\infty r_i \mathrm{d}i.$$

The representative consumer's utility is a CES function of the quantities of the goods consumed by that agent—where each good has the same weight—but adjusted by a constant factor $\overline{y}$ which ensures that no good is essential even though the goods are complements ($\varepsilon > 0$):

$$u = \left\{ \int_0^\infty [(y_i/L + \overline{y})^{-\varepsilon} - \overline{y}^{-\varepsilon}] \mathrm{d}i \right\}^{-1/\varepsilon}.$$

The dynamics of the model are straightforward, in that we let A_l and A_r grow exogenously, and furthermore the ratio of the energy price to the wage w_r/w_l also evolves exogenously; we return to this when we calibrate the model.

3.2. Solution

To solve the model we need an expression for l_i as a function of i alone. To get there, we obtain two expressions for p_i/p_0, the price of good i relative to the price of good 0, first from the production function (price = marginal cost), and then from the utility function (relative price = relative marginal utility). We use these expressions to eliminate p_i and obtain an expression for l_i as a function of i and l_0. We then know that the integral of l_i across all the goods produced must be equal to L, which allows us to solve the problem.

To obtain the first expression for p_i/p_0 define total production costs for good i as C_i, and note that it follows from the Leontief production function that

[11] An additional factor may be increased density of seats (for instance, more seats on the same airplanes).

$$y_i = A_l l_i,$$

and $$r_i = (A_l/A_r) i l_i,$$

so $$C_i = w_l y_i/A_l + w_r (y_i/A_r) i,$$

where w_l and w_r are the input prices. Since p_i is equal to marginal cost we have

$$p_i = w_l/A_l + w_r i/A_r. \tag{1}$$

Finally, normalize (without loss of generality) such that $w_l = A_l$, implying that $p_0 = 1$ and

$$p_i/p_0 = 1 + w_r i/A_r. \tag{2}$$

To obtain the second expression for p_i/p_0 note that the cost of a marginal increase in utility must be equal for all i, hence

$$\frac{p_i}{p_0} = \frac{\partial u/\partial y_i}{\partial u/\partial y_0} = \left(\frac{y_0/L + \bar{y}}{y_i/L + \bar{y}}\right)^{1+\varepsilon} = \left(\frac{w_l l_0/L + \bar{y}}{w_l l_i/L + \bar{y}}\right)^{1+\varepsilon}.$$

Eliminating p_i/p_0 we have

$$l_i = \left(l_0 + \frac{\bar{y}L}{w_l}\right)\left(1 + \frac{w_r}{A_r} i\right)^{-1/(1+\varepsilon)} - \frac{\bar{y}L}{w_l}. \tag{3}$$

So if we were to plot l_i against i for given parameters, prices, and productivities, and an arbitrarily chosen value of l_0, we would obtain a monotonically declining curve, reaching $l_i = 0$ at some (finite) value of i, and this value (denoted N) is a decreasing function of l_0. To solve the model we find the unique values of N and l_0 which ensure that the area under the curve is equal to L, total labour. That is, l_0 is chosen such that

$$\int_0^N l_i \mathrm{d}i = L,$$

where N (a function of l_0) is the value of i at which $l_i = 0$.

The solution to this problem is straightforward but somewhat long-winded, hence we leave the working to Appendix A. Note however that for tractability we assume $\varepsilon = 1$. The result is that $l_0/L = [\bar{y} w_r/(w_l A_r)]^{1/2}$, and hence

$$N = \frac{A_r}{w_r}\left[\left(1 + \hat{A}\right)^2 - 1\right]; \tag{4}$$

$$y_0 = L\bar{y}\hat{A}; \tag{5}$$

$$y_i = \frac{y_0 + L\bar{y}}{[1 + (w_r/A_r)i]^{1/2}} - L\bar{y}; \tag{6}$$

$$r_i = y_i i/A_r; \tag{7}$$

and $$\frac{w_r R}{w_l L} = \frac{2}{3}\hat{A}\left(1 + \frac{\hat{A}^2}{4}\right), \tag{8}$$

where $$\hat{A} = \left(\frac{w_l w_r/A_r}{\bar{y}}\right)^{1/2}. \tag{9}$$

So the factor share of energy is increasing in $\hat{A}$, and $\hat{A}$ is increasing in w_l (the wage) and w_r/A_r (the energy price in efficiency units). Furthermore, since we know that the factor share of primary energy is small—around 6 percent—this tells us that $\hat{A}$ is small (around 0.06) and hence $\hat{A}^2/4 \ll 1$, so we can approximate (after defining $C = 2/(3\bar{y}^{1/2})$, a constant)

$$\frac{w_r R}{w_l L} \approx C\left(w_l w_r/A_r\right)^{1/2}, \tag{10}$$

and, using $w_l = A_l$, $$\frac{R}{L} \approx C \cdot A_l\left(\frac{1}{w_r}\frac{A_l}{A_r}\right)^{1/2}. \tag{11}$$

So if the energy price is constant while the levels of labour-augmenting and energy-augmenting technology grow at equal rates then the energy share is constant, while primary energy consumption per capita tracks GDP growth. On the other hand, since the elasticity of the energy share to the energy price is 0.5, an increasing energy price leads to a slow increase in the energy share and a slow decline in energy use per capita. Finally, if energy efficiency A_r grows faster then labour efficiency A_l this reduces both the energy share and energy use.

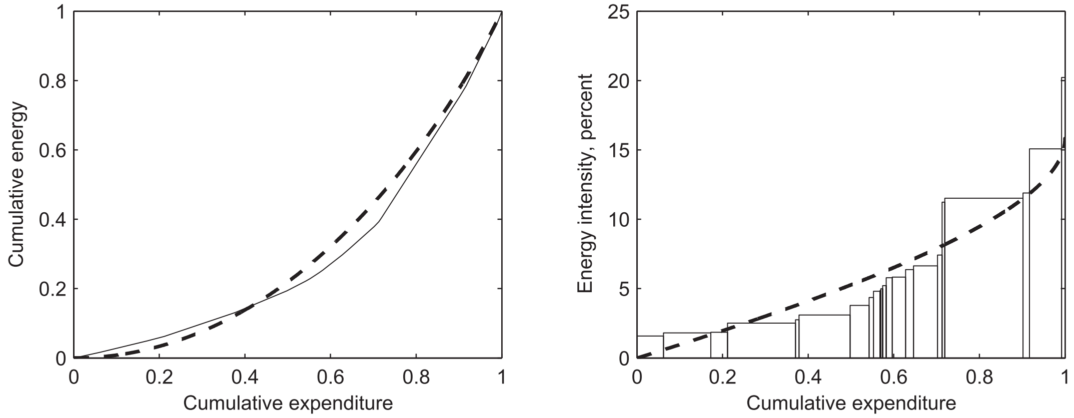

Fig. 4. Cumulative energy use and energy intensity plotted against cumulative expenditure when consumption products are sorted in order of increasing energy intensity: the model economy (dashed lines) and data compared. Data as in Fig. 2.

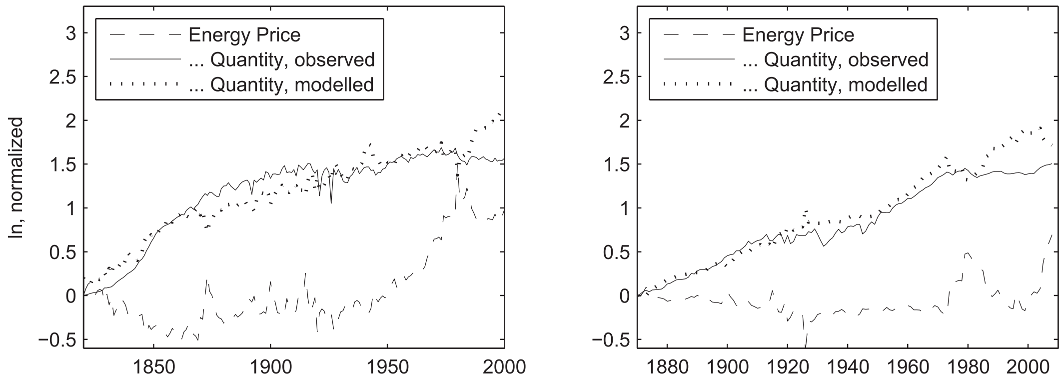

Fig. 5. Long-run growth in observed energy consumption compared to modelled energy consumption: (a) U.K.; (b) Globally. Data as in Fig. 1.

3.3. Parameterization

We parameterize the model to fit the data of Figs. 1 and 2, beginning with Fig. 2. Here we begin by setting the relative share of primary energy (i.e. $w_rR/(w_lL)$) to 6 percent. Given a fixed relative share S we can show using the above equations that

$$\frac{r_i}{R/N} = \frac{i}{N}\frac{[(1+\widehat{A})^2-1]^2}{\widehat{A}^2}\left[\frac{1+\widehat{A}}{\{1+[(1+\widehat{A})^2-1]i/N\}^{1/2}}-1\right]\frac{1}{S}.$$

Since $\widehat{A}$ is also a function of S this shows that—for given S—total energy use for good i relative to the average of total energy use across all the goods is a function of i/N alone. So for given S the pattern of energy use across the range of goods in production is independent of the other variables (technology levels A_l and A_r, the labour force L, and the energy price w_r). We plot this pattern in Fig. 4, where we compare the pattern derived from the model to the pattern found in the data previously shown (Fig. 2). Note that the model does a reasonable job of matching the observed spread of expenditure across goods of differing energy intensity.

Turning to the time-series data, we now have a little more flexibility because we can test different assumptions about the growth of A_r. We have data on GDP and population, and we calculate A_l (labour productivity) as the ratio of the two. Given our labour-productivity series, we let energy productivity grow 10 percent faster each year (so if labour productivity is observed to grow by 2.2 percent in a given year, energy productivity is assumed to grow by 2.42 percent that year). The weighted price of primary energy is calculated as in Fig. 1. Finally, we calibrate initial energy productivity to deliver a relative energy share of 6 percent in 1970.[12]

The results are shown in Fig. 5. Again we see that the model does a reasonable job in matching the data, although given the very simple pattern in the data this is perhaps not a very startling achievement. The main failure of the model (apart from not matching short-run fluctuations) is that it predicts a steep increase in energy consumption following the price declines after

[12] Note that this implies separate normalizations for the UK and global data.

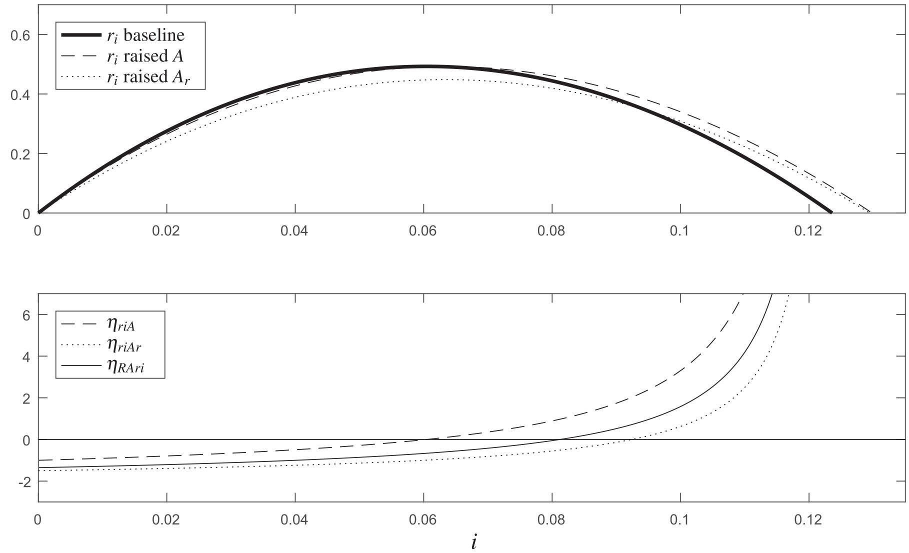

Fig. 6. Energy consumption for individual products r_i as a function of i. The top panel shows the baseline curve and the effect of 10 percent increases in A and A_r respectively. The lower panel shows the elasticity of energy consumption with respect to overall technological progress, η_{riA}, energy-augmenting technological progress, η_{riAr}, and product-specific energy-augmenting technological progress, η_{RAri}.

1980, an increase not found in the data. We discuss possible reasons for this divergence in the next section.

4. Implications of the expanding variety model

4.1. Rebound and second-best energy policy in the model economy

In order to better understand the implications of the expanding variety model we need a clearer picture of the structure of the economy and the nature of structural change, which is given by Fig. 6. In Fig. 6 we have i—i.e. the index of different goods—on the x-axis, and on the y-axis we show energy use r_i (upper panel) and three different elasticities (lower panel). The expressions for these elasticities are derived in Appendix A. The first thing to note from the figure is that most energy use comes from products in the middle of the energy-intensity range, which is to be expected since production of the most energy-intensive goods approaches zero.

Assume now that a regulator has policy tools (such as technology subsidies) with which to selectively boost specific productivity levels A_{ri} and A_{li}, or the overall productivity levels A_r and A_l. First assume that the regulator induces a general increase in productivity (A_l and A_r both increase by the same factor, while relative prices w_r/w_l are constant), and define the elasticity of r_i w.r.t. the productivity increase as η_{riA}. From (11) we know that the elasticity of total energy use to the TFP increase is 0.5, but now we want to study individual products. The top panel of Fig. 6 shows that the r_i curve is flattened and shifted to the right as the variety of goods produced increases. This leads to a fall in the total energy use from low-intensity goods, and steep rises in energy use from high-intensity goods. The lower panel of Fig. 6 shows that η_{riA} is -1 when $i = 0$, and approaches infinity when $i \to N$.[13] The reason for the latter result is that when i is close to N the goods are on the cusp of affordability, hence their price elasticity of demand is extremely high.

Secondly, assume instead a general increase in A_r alone, holding prices constant, and define the elasticity of r_i w.r.t. the increase in A_r alone as η_{riAr}. Again we can use (11) to find the elasticity of total energy use to such a change, which is -0.5; this implies that the rebound effect from an increase in overall energy efficiency in our model is 50 percent, which corresponds well to Broberg et al. (2015), who find effects ranging from 30 to 70 percent. Turning to individual products, again we see from Fig. 6 that the r_i curve is flattened and shifted to the right. The difference is that η_{riAr} is now negative for the majority of the products, because the rise in A_r has a strong direct downward effect on energy use. However, it also causes a relatively strong substitution effect towards energy-intensive products, hence the rise in total energy use for products of very high energy intensity.

Thirdly we investigate the effect on total energy consumption of an increase in A_{ri}, the energy-efficiency of a specific product. The elasticity of R w.r.t. A_{ri} is zero since product i is one of an infinite number and any change in just one product will have a negligible effect on total energy consumption. Therefore we define the 'elasticity' η_{RAri} as $dR/r_i/(dA_{ri}/A_{ri})$, hence it shows the total change in energy consumption as a percentage of r_i, divided by the percentage change in A_{ri}. The curve is broadly similar to the curve for η_{riAr}. A rise in the energy-efficiency of a medium-intensity product yields a reduction in energy use per unit of

[13] Recall that the expressions for the elasticities are derived in Appendix A.

that product made, and also causes substitution towards that product. But because the product is of medium energy-intensity this substitution has little effect on total energy, hence the total effect of the rise in A_{ri} is negative. For a high-intensity product, on the other hand, the model indicates backfire: an increase in A_{ri} leads to an increase in total energy consumption R.

Finally we consider the effect of a rise in w_r, the energy price. This is not illustrated, because the effect is almost a mirror-image of the effect of the rise in A: consumption is pushed back towards low-intensity products, leading to increases in total energy use associated with these products, at the same time as energy use in producing high-intensity products declines dramatically.

The most interesting results—and most relevant to the discussion of the rebound effect—are perhaps those for η_{RAri}, relating to the effect on total energy consumption of product-specific improvements in energy efficiency. They show that the effect is very strongly dependent on the energy-intensity of the product in question: for products of low energy-intensity the rebound effect is *negative*, that is (recalling the discussion in the introduction), an increase in A_{ri} by a factor x leads to a reduction of R by *more* than $R_i(1-1/x)$. The reason is that substitution by consumers towards products of lower-than-average energy intensity makes a positive contribution towards reducing total energy demand. However, for products of the highest energy intensity we have instead *backfire*, i.e. increases in A_{ri} lead to increases in total energy consumption. This suggests that policies to reduce energy-intensity of production might best be aimed at products of medium intensity: those of low intensity have very low energy consumption anyway, so the effects are likely to be small, whereas rebound effects are large for those of high intensity. Products of medium intensity account for a large part of energy consumption, and rebound effects are likely to be small or even negative.

4.2. Weaknesses and possible extensions

The model matches the long-run data fairly well. Furthermore, it is consistent with the data of Fig. 2, showing the spread of production across energy intensities; the parameterization is broadly consistent with evidence regarding the bias of technological change, i.e. that there has been some bias towards the growth of energy-augmenting technology; and the model includes both income and substitution effects in its explanation for the long-run failure of the energy share to decline. Finally, the expansion of consumption into new energy-intensive sectors predicted by the model is clearly relevant: the two most energy-intensive sectors in Fig. 2 are motor transport and air transport. These sectors rose from scratch during the 20th century as a result of a combination of income growth and progress in energy-augmenting technology. Nevertheless, the model has a number of weaknesses that limit our ability to use it to draw conclusions about real economies, of which we discuss the following: the poor match of the calibrated model to recent data; the high price-elasticity of demand for goods and hence the unrealistically powerful rebound effects for energy-intensive goods; the lack of flexibility in treating income effects; the failure to account for international trade; the treatment of technological change as exogenous; and the simplified nature of the utility and production functions.

The model predicts a rapid increase in global consumption of combustibles between 1980 and 2000, more rapid than the observed increase. There are several possible explanations for this, of which we mention three. The first possible explanation is that our data includes only combustibles, which is a good approximation for total primary energy up to the early 1970s, but not so good after this time since installed nuclear capacity grew steeply throughout the '70s and early '80s, and without the growth in nuclear energy growth in combustible inputs would have been significantly greater. The second explanation is that the energy crisis of 1974 led to a series of energy-related policies which amount to implicit rises in the price of energy paid by firms and consumers, even though these increases have not been captured in our energy-price data. Preliminary analysis suggests that when energy taxes are included the fit of the model to the data is improved. The third explanation is that the price shocks of the 1970s had a profound and long-lived effect on investment decisions and the direction of technological change, due in turn to a change in long-run price expectations: pre-1974 the dominant expectation may have been the prices would continue their slow decline, whereas post-1974 it seems reasonable that expectations about future prices may have been revised dramatically upwards. Such a change could lead to an increase in the growth rate of A_r relative to A_l, thus holding energy consumption down.

The model works because demand for individual goods is highly price-elastic, which allows consumers to substitute quite freely between goods of low and high energy-intensity, depending on their relative prices. The other side of the coin is that income effects—while present in the model—are limited in power. Due to the need for tractability we set $\varepsilon = 1$ and therefore we cannot calibrate the model to raise the strength of income effects on demand for energy-intensive products; neither can we calibrate consumers' willingness to pay for expanding variety. If we could lessen the substitutability between goods and strengthen income effects we could potentially continue to fit the historical data, while weakening the rebound effects predicted by the model. More generally, the balance between income and substitution effects is crucial to the predictions and policy conclusions of the model. In the model this balance is stable, by construction. However, it is possible that the income-elasticity of demand for energy-intensive products declines with rising income, and is very low or even negative when incomes are very high (i.e. in rich industrialized nations). This is a very optimistic scenario as it suggests that energy demand will tail off endogenously without the need for price rises or technology policy. This vision of a clean, post-industrial global economy is tempting, but little convincing evidence has been presented that it will materialize without powerful policy interventions.

A further weakness of the model is the need to assume a representative consumer, implying that a product on the cusp of affordability for one consumer is on the cusp of affordability for all consumers. This causes the demand elasticities for such products to be seriously overestimated: in reality the good may be inferior for the richest consumers, and completely out of reach for the poorest. A useful extension to the model would therefore be to disaggregate consumers into groups at different income levels and treat their demand for goods separately, as is done in much of the literature on rebound effects for individual products. In such a model it is possible that income effects would dominate, and (in the most optimistic scenario) that in the very long run energy consumption will level off or decline as the richest consumers substitute away from energy-intensive goods. A possible approach would be to assume a utility function consistent with PIGL preferences, as Boppart (2014) does in his model of structural change between goods and services.

Not only do we assume a representative consumer, we also assume a single country, i.e. we do not include international trade. International trade is primarily relevant to the manufacturing and agricultural sectors, and Mulder (2015) shows that changing trade patterns are indeed relevant to explaining country-level trends in energy intensity, as countries increasingly specialize in particular sectors. Thus it would be valuable to build an extended model of expanding variety and energy use in an open economy, which could capture shifts in energy-intensive manufacturing industries out of OECD countries and in to countries such as China, tending to reduce domestic energy share in the OECD. However, Mulder also shows that the manufacturing sector typically accounts for less than 30 percent of current energy use in OECD economies. So the incorporation of trade and international specialization into the model—while highly desirable—should not fundamentally change its properties.

We treat the bias of technological change as exogenous, but there is plenty of evidence—such as that put forward by Newell et al. (1999)— to show that energy-augmenting knowledge responds positively to increased energy prices, an effect that could be captured by incorporating a model of directed technological change (DTC). From equation (10), the elasticity of the energy share to the energy price is (when A_r is exogenous) 0.5. So if the energy price rises it leads to an increase in the energy share, and—given DTC—this should lead to upward pressure on energy-augmenting knowledge A_r, tending to pull the share back down. Such a model could be calibrated to give a similar overall rate of energy-augmenting technological progress as we assume exogenously; the difference would be that this rate would vary endogenously as a function of the energy price.

Finally, the model includes a range of products varying in energy intensity, with technological change driving increases in energy- and labour-augmenting knowledge which are uniform across products. While this is a step forward compared to models including just two or a handful of products (such as Chan and Gillingham, 2015), it remains extremely simplified when compared to the real economy in which quality and input-intensity vary in multiple dimensions, and rates of technological change vary (endogenously) from one product to another. This extra complexity leads to a range of caveats to the results, of which we highlight two. Firstly, if the price elasticities of demand for different goods vary, much of the intuition from the model is robust—for instance that the elasticity of demand for a good on the cusp of affordability is large, and that if such a good is highly energy-intensive then increases in energy efficiency of that good may lead to backfire—but the quantitative results will depend on the elasticity of demand for the good in question, hence the model results should be interpreted with caution; the intuition should act as a guide for further empirical research in order to support effective policy. Secondly, if technological change affects the prices of different goods in more complex ways than are allowed for in our model this opens up for alternative explanations for the shift in consumption patterns towards energy-intensive goods. One such explanation could be that the efficiency of labour inputs in production of these goods has grown faster than the corresponding efficiency in the production of other goods; recall that in our model A_l is assumed to grow at the same rate in all sectors. In a more general setting with capital, further possibilities arise; for instance, it could be that energy-intensive goods also tend to be capital-intensive, and increases in the efficiency of capital inputs in production has grown more rapidly than labour. It would be a major advance on the current paper to build and calibrate a model in which such mechanisms were included, ideally with endogenous technological change at the sectoral or product level.

5. Comparison of the expanding variety model to alternatives

In this section we return to the alternative models of aggregate energy demand discussed briefly in the introduction, Atkeson–Kehoe and models based on directed technological change with a single sector. In both cases we conclude that the mechanisms of these models—in which short run inelasticity is turned to long-run elasticity through the gradual reshaping of stocks of capital (Atkeson–Kehoe) or knowledge (DTC)—are almost certainly relevant, but do not seem to be the most important mechanisms in explaining the long-run data.

5.1. The Atkeson–Kehoe model

How well does the model of Atkeson and Kehoe (1999) account for the data of Section 2? In this model we have a final-good sector with the production function

$$Y = (A_l L)^{1-\alpha} Z^{\alpha},$$

where Z is capital services.[14] Capital services are produced using inputs of capital and energy, where capital is putty–clay and depreciates at a strictly positive rate. For a given quantity of homogeneous capital, k^q (where q is a superscript denoting quality or type of capital), the production function of capital services is Leontief:

$$z^q = \min\{A_k^q k^q, A_r^q r^q\}.$$

[14] Note that I adapt Atkeson and Kehoe's notation to fit with that used previously in this paper.

Here A_k^q and A_r^q are the respective factor productivities associated with the capital of type q, and r^q is the quantity of resource inputs used together with the capital k^q. Total capital services Z are simply the integral across q of z^q; in other words, capital services produced using different types of capital are perfectly substitutable. The choices of A_k^q and A_r^q are restricted as follows:

$$(A_r^q)^{\omega}(A_k^q)^{1-\omega} = A_r^{\omega} A_k^{1-\omega},$$

where A_r and A_k are given, and ω is a parameter between 0 and 1.

The upshot of the model is that, effectively, the long-run production function for capital services is Cobb–Douglas, whereas the short-run production function is Leontief. To understand what I mean by this statement, assume that the relative price of energy to capital is permanently fixed at some level W^*. Then the shares of energy and capital will be ω and $1-\omega$ respectively, irrespective of W^*. Now assume that W^* doubles, suddenly and unexpectedly, and then remains fixed (with certainty) at the new level. Then the relative quantities of energy and capital will be unchanged in the short run (hence the energy share will double). However, in the long run the capital stock will be replaced and the shares will return to their original levels.

Turning to energy policy, in the model economy the (aggregate) bias of technological change has no observable effects in the long run, and efforts specifically targeted at boosting energy-augmenting knowledge A_r will not—even if they are successful—reduce aggregate energy use R. For a fixed energy price w_r the long-run production function is

$$Y = (A_l L)^{1-\alpha}[(A_r R)^{\omega}(A_k K)^{1-\omega}]^{\alpha}$$

Assume that energy is supplied at an exogenous price w_r. The result then implies that (given perfect markets) $R = \alpha\omega Y/w_r = (\alpha\omega A K^{\alpha(1-\omega)} L^{1-\alpha}/w_r)^{1/(1-\alpha\omega)}$, where A is total factor productivity (TFP), $A_l^{1-\alpha}(A_r^{\omega}A_k^{1-\omega})^{\alpha}$. So any boost to TFP will increase energy use, irrespective of whether it is through an increase in A_l or A_r.

To compare the Atkeson–Kehoe model to ours we should begin by extending the former by adding exogenous increases in A_l and A_r, as in our model; such an extended model is broadly consistent with the data in Fig. 1. However, it does not explain why there is such a broad menu of energy-intensities available to firms investing in new capital, and furthermore—crucially—it implies that all final goods at a given time are perfect substitutes and have the same energy intensity, which is contradicted by the data in Fig. 2. So our model captures structural change, which is ruled out by construction in Atkeson–Kehoe. However, their model allows for the substitution of energy for capital on the production side, which we rule out by construction since our production function is Leontief in labour and energy, and does not include capital. Ideally we would set up our model with a more flexible production function such as a nested CES including labour and capital–energy, and estimate the separate roles of substitution between production inputs for given technology, biased technological change, and structural change.

5.2. Single-sector models with DTC

Atekson and Kehoe deliver the contrast between short-run inelasticity and long-run elasticity through the need to replace inflexible capital assets; an alternative is to build a model in which it takes time to develop technologies appropriate to relative input prices, i.e. a model of DTC. In Section 4.2 we discussed adding DTC to a model of structural change. Here we discuss whether DTC can account for the data in a single-sector model. We conclude that although it is straightforward to build a single-sector DTC model which yields the constant-energy-share property—as we see from the existing literature—this success comes only at the expense of predicting that energy-augmenting knowledge fails to rise when energy prices are constant, which is contradicted by the evidence above (Section 2.2). Furthermore, such a model falls foul of a range of other evidence, as well as failing to account for the spread of energy intensities across products, and structural change. Note that we do not claim that DTC is unimportant, we simply claim that DTC in a single-sector model cannot explain the data. Recall the discussion in Section 4.2, where we note that ideally we would incorporate a realistic model of DTC (consistent with the results of Newell et al., 1999 and other related work) into an extended model of structural change.

The fundamental idea behind single-sector DTC models of energy demand is that high long-run elasticity of substitution comes as a result of the difference between the physical inputs (labour, primary energy) and the augmented inputs, which are the physical inputs multiplied by their respective levels of factor-augmenting technology. The augmented inputs are poor substitutes, but the physical inputs may be highly substitutable in the long run if the relative levels of factor-augmenting technology can be changed through directed investment by firms. Assuming a CES production function for the representative firm in equilibrium, we have (using the same notation as above)

$$y = [(A_r r)^{-\rho} + (A_l l)^{-\rho}]^{-1/\rho}, \tag{12}$$

where ρ is large and positive, delivering the low short-run elasticity of substitution $1/(1+\rho)$. Given perfect competition it follows straightforwardly that to deliver unit long-run elasticity of substitution (implying constant factor shares) we require that

$$\frac{w_r/A_r}{w_l/A_l}$$

must be constant, i.e. the relative prices of energy and labour *in efficiency units* must be constant. This implies that if the price of energy falls relative to the price of labour (as we observe in the data), there must be a countervailing—and equal—fall in the level of energy-augmenting knowledge relative to labour-augmenting. Intriguingly, this exact result follows if we assume that the knowledge stocks grow purely as a function of directed investment in that stock, with the growth rate being independent of the current level of knowledge[15]:

$$A_{rt} = A_{rt-1}(\zeta_r z_{rt})^{\phi}; \tag{13}$$

$$A_{lt} = A_{lt-1}(\zeta_l z_{lt})^{\phi}. \tag{14}$$

Here the ζs are parameters, the zs are investment levels, and ϕ is a parameter.

Although the theory here is elegant, this elegance is beguiling, at least for the case of energy and labour. On the one hand, models of this type have been developed and applied in a series of papers[16]; on the other hand, the theory is found wanting when confronted by empirical evidence. We know from work such as Newell et al. (1999) that firms respond to higher energy prices by investing more in energy-augmenting technology, which in turn delivers higher energy-augmenting knowledge. However, to deliver unit long-run elasticity of substitution between labour and energy in a single-sector model we require a much more specific result, that an x percent fall in long-run energy prices relative to wages yields—in the long run—an x perfect fall in energy-augmenting knowledge relative to labour-augmenting knowledge. This is contradicted by the evidence we presented in Section 2.2, which suggests that energy-augmenting knowledge grew at least as fast as labour-augmenting knowledge during the long period (up to 1974) when energy prices fell steeply relative to the wage.

The problem lies in the specification of the production function for energy-augmenting knowledge. In the single-sector model, we must specify such that energy-augmenting knowledge grows independently of other knowledge (e.g. labour-augmenting knowledge) in order to deliver the unit-elasticity result, as shown by Hart (2013). This does not stand up to scrutiny, for several reasons, of which we mention three. (1) It implies that the research input must be pure labour (no lab equipment), otherwise we create a link between the stocks. This is because the price of lab equipment is constant, and in a growing economy more of it will be used in all research sectors, thus growth in overall productivity (A_l in equation (14)) will lead to an increase in the growth rate of A_r; when research workers are the only input this effect does not exist because growth drives up the price of labour rather than the quantity supplied. (2) In a calibrated model, research labour inputs should rise over time due to population growth, giving us accelerating growth (scale effects); such an acceleration in growth is not observed. We can neutralize scale effects if we assume decreasing returns to knowledge accumulation, but again this generates a link between the knowledge stocks, as a smaller stock tends to grow faster. (3) We know in practice that energy-augmenting knowledge builds not only on existing such knowledge, but also on other knowledge such as labour-augmenting knowledge. There is not much evidence in the literature about the strength of these links; however, Acemoglu (2002)—based on data from Trajtenberg et al. (1992) —argues that there are strong links between knowledge accumulation in different sectors. The assumption of independent knowledge stocks, on the other hand, implies that technologies which allow us to make better use of raw energy inputs (the steam engine, the steam turbine for the generation of electricity, the internal combustion engine) are developed completely independently of other technological advances in the economy. This seems to be an indefensible idea: such technologies are developed hand-in-hand with advances in mathematics, physics, engineering etc., advances which are also relevant to augmenting inputs of labour–capital.

6. Conclusions

We focus on the failure of the energy share to decline despite the long-run decline in the price of energy relative to labour, and claim that a shift in consumption patterns over time towards goods of high energy intensity must be an important part of the explanation, since careful scrutiny shows that alternative explanations—such as the idea that directed technological change has led to slow growth in energy-augmenting knowledge —are insufficient. We propose a novel model in which the shift in consumption patterns consists not just of increasing consumption of existing energy-intensive goods, but also the production and consumption of completely new such goods. The switch is driven by a combination of income and substitution effects.

In the model economy, policy-induced rises in the price of energy will reduce energy consumption, as will policy-induced increases in the growth rate of energy-augmenting knowledge. However, technology policy is more effective if it can be directed towards goods which lie towards the lower end of the distribution of energy intensities. The reason is that an increase in the energy-efficiency of such goods causes their price to decline (albeit weakly), inducing consumers to substitute towards consumption of these goods. The resulting drop in consumption of energy-intensive goods leads to a 'reverse-rebound' effect: an increase in energy-augmenting knowledge in production of good i, A_{ri}, by a factor x leads to a reduction of total energy consumption R by *more* than $R_i(1 - 1/x)$. On the other hand, somewhat paradoxically, increases in the energy-efficiency of the most energy-intensive goods (such as air transport) are much more likely to lead to rebound or even backfire, i.e. an increase in total energy consumption. Because these goods are assumed to be on the cusp of affordability, their price elasticity of demand is extremely high.

Our analysis shows that a shift towards energy-intensive consumption, including new energy-intensive goods, has occurred over the industrial period, and is to some extent continuing even in the most advanced economies. Our model proposes an

[15] This result is demonstrated by Hart (2013), the intuition being that firms invest in proportion to factor costs, so if the cost of one factor rises, investment in factor-augmenting knowledge drives the cost back down. If knowledge stocks grow independently then equilibrium is always attained at the same level of relative costs.

[16] As mentioned in the introduction; see for instance Gerlagh (2008), Fischer and Newell (2008), Gans (2012), and Hassler et al. (2012).

explanation for this shift—a combination of substitution and income effects—but we present little evidence of the veracity of the model. For concrete policy recommendations the model should be backed up by microeconometric evidence on the causes of the shifting consumption patterns in relevant sectors, especially for high-energy products on the cusp of affordability. In order to accurately measure income effects such studies need to build on household data. Preliminary work on air travel in Sweden with a double-hurdle model (first choose whether to be an air traveller, then choose how much to fly) shows that the decision to become an air traveller is strongly dependent on income, suggesting that there is an important overall role for income effects.

The model predicts that if the energy price tracks the wage in the future, this will brake the growth in energy consumption but not stop it. Such an increase in w_r could be cancelled out if energy efficiency A_r stops rising. This is bound to happen in some sectors, such as lighting and motive power, where the laws of physics impose strict limits on what is achievable, limits which we are already approaching. This points to the need for new models of directed technological change which base the innovation possibilities frontier on evidence rather than assumption (*cf.* Hart (2013) and Nordhaus (1973)). In the most pessimistic scenario—with long-run growth but a slowdown in growth of energy efficiency—the model would predict that the relatively stable global energy consumption since 1974 may be only a temporary phenomenon, with consumption set to rise again in the future absent price rises or policy interventions.

Appendix A. The model of increasing variety

Appendix A.1. Solving for N

Recall that we have equation (3),

$$l_i = \left(l_0 + \frac{\bar{y}L}{w_l}\right)\left(1 + \frac{w_r}{A_r}i\right)^{-1/(1+\varepsilon)} - \frac{\bar{y}L}{w_l},$$

and that the highest value of i at which production takes places is denoted N. Set up the following integral:

$$\int_0^N l_i \mathrm{d}i = L = \int_0^N \frac{y_0 + \bar{y}L}{A_l}\left(1 + \frac{w_r}{A_r}i\right)^{-1/(1+\varepsilon)} - \frac{\bar{y}L}{A_l}\mathrm{d}i.$$

Solve to yield

$$A_l L = \frac{1+\varepsilon}{\varepsilon}\frac{y_0 + \bar{y}L}{w_r/A_r}\left[\left(1 + \frac{w_r}{A_r}i\right)^{\varepsilon/(1+\varepsilon)}\right]_0^N - \bar{y}LN$$

$$= \frac{1+\varepsilon}{\varepsilon}\frac{y_0 + \bar{y}L}{w_r/A_r}\left[\left(\frac{w_r}{A_r}N + 1\right)^{\varepsilon/(1+\varepsilon)} - 1\right] - \bar{y}LN.$$

Here we have an equation in two unknowns, y_0 and N. To solve it we need another equation, which we obtain by setting $l_i = 0$ when $i = N$. This gives us

$$\frac{w_r}{A_r}N + 1 = \left(\frac{y_0/L + \bar{y}}{\bar{y}}\right)^{1+\varepsilon}.$$

Eliminate $y_0 + \bar{y}$ from these two equations to obtain

$$A_l + \bar{y}N = \frac{1+\varepsilon}{\varepsilon}\frac{\bar{y}}{w_r/A_r}\left(\frac{w_r}{A_r}N + 1\right)^{1/(1+\varepsilon)}\left[\left(\frac{w_r}{A_r}N + 1\right)^{\varepsilon/(1+\varepsilon)} - 1\right].$$

For tractability we now assume $\varepsilon = 1$. Multiply through by $w_r/(A_r\bar{y})$ to obtain

$$\frac{A_l w_r}{A_r \bar{y}} + \frac{w_r}{A_r}N = 2\left(\frac{w_r}{A_r}N + 1\right)^{1/2}\left[\left(\frac{w_r}{A_r}N + 1\right)^{1/2} - 1\right].$$

For clarity denote $(w_r/A_r)N + 1 = x^2$, $A_l(w_r/A_r)/\bar{y} = \widehat{A}^2$, so

$$\widehat{A}^2 + x^2 - 1 = 2x(x-1)$$

and

$$x^2 - 2x + 1 - \widehat{A}^2 = 0.$$

Solve to show that[17]

$$\left[(w_r/A_r)N + 1\right]^{1/2} = 1 + \left(\frac{A_l w_r/A_r}{\bar{y}}\right)^{1/2}.$$

[17] We must take the positive root since N must be positive.

We thus have a solution for N and hence y_0 and hence y_i:

$$w_r/A_rN = \left[1 + \left(\frac{A_l w_r/A_r}{\bar{y}}\right)^{1/2}\right]^2 - 1;$$

$$\frac{y_0}{L\bar{y}} = \left(\frac{A_l w_r/A_r}{\bar{y}}\right)^{1/2};$$

$$\frac{y_i}{L\bar{y}} = \frac{y_0/(L\bar{y}) + 1}{(1 + w_r/A_r i)^{1/2}} - 1.$$

The next step is to find R, total energy use. We know that

$$r_i = y_i i/A_r,$$

hence

$$R = \frac{L\bar{y}}{A_r}\int_0^N \left[\frac{y_0/(L\bar{y}) + 1}{(1 + w_r/A_r i)^{1/2}} - 1\right] i\,\mathrm{d}i.$$

Solve the integral and insert the solution for N above to yield

$$\frac{R}{L} = \frac{2}{3}\frac{\bar{y}}{A_r}\left(\frac{A_l}{\bar{y}}\right)^2\left(\frac{A_l w_r/A_r}{\bar{y}}\right)^{-1/2}\left(1 + \frac{1}{4}\frac{A_l w_r/A_r}{\bar{y}}\right).$$

Finally, use the fact that $w_r/A_r = w_r/A_r$ to write

$$\frac{w_r R}{w_l L} = \frac{2}{3}\widehat{A}\left(1 + \frac{\widehat{A}^2}{4}\right),$$

where

$$\widehat{A} = \left(\frac{A_l}{A_r\bar{y}}w_r\right)^{1/2}.$$

Appendix A.2. Elasticities illustrated in Fig. 6

Now we turn to the elasticities shown in Fig. 6. To find η_{riA} write

$$r_i = \left[\frac{y_0 + L\bar{y}}{(1 + w_r i/A_r)^{1/2}} - L\bar{y}\right]\frac{i}{A_r}, \tag{A.1}$$

and note that y_0 is invariant in A_l/A_r. Thus we can find η_{riA} by differentiating the above expression for r_i w.r.t. A_r and multiplying by A_r/r_i, yielding

$$\eta_{riA} = -1 + \frac{1}{2}\frac{r_i A_r/i + L\bar{y}}{r_i A_r/i}\frac{w_r i/A_r}{1 + w_r i/A_r}.$$

Turning to η_{riAr}, note that y_0 is decreasing in A_r so we can use equation A.1 (and the expression for y_0 above) to find η_{riAr}, since $\eta_{riAr} = \eta_{riA} + \partial r_i/\partial y_0 \cdot \partial y_0/\partial A_r \cdot A_r/r_i$, hence

$$\eta_{riAr} = \eta_{riA} - \frac{1}{2}y_0\frac{i/(A_r r_i)}{(1 + w_r i/A_r)^{1/2}}.$$

Now η_{RAri}. First define L_i^* as labour remaining after the allocation to sector i, and define the number of sectors as $M \times N$. (Since we have a continuum of sectors measure N, we will let M approach infinity later on.) Then

$$\frac{\mathrm{d}R}{\mathrm{d}A_{ri}} = \frac{\partial R}{\partial L_i^*}\frac{\partial L_i^*}{\partial l_i}\frac{\partial l_i}{\partial A_{ri}} + \frac{\partial R}{\partial r_i}\left(\frac{\partial r_i}{\partial A_{ri}} + \frac{\partial r_i}{\partial l_i}\frac{\partial l_i}{\partial A_{ri}}\right).$$

We now discuss the seven differentials in order. The first is simply R/L. The second is equal to $-1/M$. The third follows from equation (3) and the solution for y_0, noting that y_0 is unaffected by a change in A_{ri} when $M \to \infty$:

$$\frac{1}{2}\left(l_i + \frac{L\bar{y}}{A_l}\right)\frac{w_r i/A_r}{1 + w_r i/A_r}\frac{1}{A_r}.$$

The fourth is equal to $1/M$. The fifth is $-r_i/A_{ri}$, the sixth is r_i/l_i, and the final differential is the same as the third. Finally, to find η_{RAri} multiply through by $(A_{ri}/R) \cdot [R/(r_i/M)]$ to obtain

$$\frac{A_{ri}}{r_i}\frac{\mathrm{d}R}{\mathrm{d}A_{ri}} = -1 + \frac{1}{2}\left(\frac{r_i}{l_i} - \frac{R}{L}\right)\left(l_i + \frac{L\bar{y}}{A_l}\right)\frac{w_r i/A_r}{1 + w_r i/A_r}\frac{1}{r_i}.$$

Finally, the elasticity of R w.r.t. w_r is calculated in an analogous way to the above, yielding

$$\eta_{riwr} = \frac{1}{2}\frac{i}{A_r r_i}\frac{y_0 + L\bar{y}w_r i/A_r}{(1 + w_r i/A_r)^{3/2}}.$$

References

Acemoglu, D., 2002. Directed technical change. Rev. Econ. Stud. 69, 781–809.
Acemoglu, D., 2009. Introduction to Modern Economic Growth. Princeton, New Jersey.
Acemoglu, D., Guerrieri, V., Jun. 2008. Capital deepening and nonbalanced economic growth. J. Polit. Econ. 116 (3), 467–498.
Atkeson, A., Kehoe, P.J., Sep. 1999. Models of energy use: putty-putty versus putty-clay. Am. Econ. Rev. 89 (4), 1028–1043.
Boden, T.A., Marland, G., Andres, R.J., 2012. Global, Regional, and National Fossil-fuel CO_2 Emissions. Tech. rep., CDIAC, Oak Ridge National Laboratory, U.S. Department of Energy, Oak Ridge, Tenn., U.S.A.
Boppart, T., 2014. Structural change and the Kaldor facts in a growth model with relative price effects and non-Gorman preferences. Econometrica 82 (6), 2167–2196.
Borenstein, S., 2015. A microeconomic framework for evaluating energy efficiency rebound and some implications. Energy J. 36 (1), 1–21.
BP, 2012. Statistical Review of World Energy. Annually published data, British petroleum.
Broberg, T., Berg, C., Samakovlis, E., Aug. 2015. The economy-wide rebound effect from improved energy efficiency in Swedish industries—a general equilibrium analysis. Energy Policy 83, 26–37.
Chan, N.W., Gillingham, K., 2015. The microeconomic theory of the rebound effect and its welfare implications. Journal of the Association of Environmental and Resource Economists 2 (1), 133–159.
Csereklyei, Z., Rubio Varas, M. d. M., Stern, D., 2016. Energy and economic growth: the stylized facts. Energy J. 37, 223–255.
Engel, E., 1857. Die productions- und consumptionsverhaeltnisse des koenigsreichs sachsen. In: Zeitschrift des Statistischen Bureaus des Königlich Sächsischen Ministeriums des Inneren 8–9.
Fischer, C., Newell, R.G., Mar. 2008. Environmental and technology policies for climate mitigation. J. Environ. Econ. Manag. 55 (2), 142–162.
Fouquet, R., 2011. Divergences in long-run trends in the prices of energy and energy services. Rev. Environ. Econ. Policy 5 (2), 196–218.
Fouquet, R., Pearson, P.J.G., 2006. Seven centuries of energy services: the price and use of light in the United Kingdom (1300–2000). Energy J. 27, 139–177.
Gans, J.S., Nov. 2012. Innovation and climate change policy. Am. Econ. J. Econ. Policy 4 (4), 125–145.
Gerlagh, R., 2008. A climate-change policy induced shift from innovations in carbon-energy production to carbon-energy savings. Energy Econ. 30, 425–448.
Gillingham, K., Rapson, D., Wagner, G., 2016. The rebound effect and energy efficiency policy. Rev. Environ. Econ. Policy 10 (1), 68–88.
Griffin, J.M., Gregory, P.R., 1976. Intercountry translog model of energy substitution responses. Am. Econ. Rev. 66 (5), 845–857.
Hart, R., 2013. Directed technological change and factor shares. Econ. Lett. 119, 77–80.
Hassler, J., Krusell, P., Olovsson, C., 2012. Energy-saving Technical Change. NBER working paper series (18456).
Hills, R.L., 1993. Power from Steam: a History of the Stationary Steam Engine. Cambridge University Press.
Houthakker, H.S., 1957. An international comparison of household expenditure patterns, commemorating the centenary of Engel's law. Econometrica 25 (4), 532–551.
Knittel, C.R., Dec. 2011. Automobiles on steroids: product attribute trade-offs and technological progress in the automobile sector. Am. Econ. Rev. 101 (7), 3368–3399.
Lee, J.J., Lukachko, S.P., Waitz, I.A., Schafer, A., et al., 2001. Historical and future trends in aircraft performance, cost, and emissions. Annu. Rev. Energy Environ. 26 (1), 167–200.
Lemoine, D., 2016. General Equilibrium Rebound from Energy Efficiency Policies. University of Arizona. Working paper 14-02.
Maddison, A., 2003. Growth accounts, technological change, and the role of energy in Western growth. In: Economia e Energia, secc.XIII-XVIII. Istituto Internazionale di Storia Economica “F. Datini” Prato, Florence.
Maddison, A., 2010. Historical Statistics of the World Economy: 1–2008 ad. Tech. rep., Groningen growth and development centre.
Matsuyama, K., 2002. The rise of mass consumption societies. J. Polit. Econ. 110, 1035–1070.
Mayer, H., Flachmann, C., 2011. Extended Input-output Model for Energy and Greenhouse Gases. Tech. rep., Federal Statistical Office, Germany.
Mulder, P., 2015. International specialization, sector structure and the evolution of manufacturing energy intensity in OECD countries. Energy J. 36, 111–136.
Mulder, P., de Groot, H.L., 2012. Structural change and convergence of energy intensity across OECD countries, 1970–2005. Energy Econ. 34 (6), 1910–1921.
Newell, R.G., Jaffe, A., Stavins, R., August 1999. The induced innovation hypothesis and energy-saving technological change. Q. J. Econ. 941–974.
Nordhaus, W.D., 1973. Some skeptical thoughts on theory of induced innovation. Q. J. Econ. 87 (2), 208–219.
Pindyck, R.S., 1979. Interfuel substitution and the industrial demand for energy: an international comparison. Rev. Econ. Stat. 61 (2), 169–179.
Pindyck, R.S., Rotemberg, J.J., 1983. Dynamic factor demands and the effects of energy price shocks. Am. Econ. Rev. 1066–1079.
Smulders, S., de Nooij, M., 2003. The impact of energy conservation on technology and economic growth. Resour. Energy Econ. 25, 59–79.
Trajtenberg, M., Henserson, R., Jaffe, A., 1992. Ivory Tower versus Corporate Lab: an Empirical Study of Basic Research and Appropriability. NBER working paper 4146.
Turner, K., 2013. “Rebound” effects from increased energy efficiency: a time to pause and reflect. Energy J. 34 (4), 25–42.
van den Bergh, J.C.J.M., Jan. 2011. Energy conservation more effective with rebound policy. Environ. Resour. Econ. 48 (1), 43–58.
Warde, P., 2007. Energy Consumption in England and Wales, 1560–2000. CNR–ISSM.
Yeh, S., Rubin, E.S., 2007. A centurial history of technological change and learning curves for pulverized coal-fired utility boilers. Energy 32, 1996–2005.