Structural Transformation and the Oil Price[1]

Radoslaw (Radek) Stefanski
Laval University and OxCarre

September 22, 2013

Abstract

What part of the high oil price can be explained by structural transformation in the developing world? Will continued structural transformation in these countries result in a permanently higher oil price? To address these issues I identify an inverted-U shaped relationship in the data between aggregate oil intensity and the extent of structural transformation: countries in the middle stages of transition spend the highest fraction of their income on oil. I construct and calibrate a multi-sector, multi-country, general equilibrium growth model that accounts for this fact by generating an endogenously falling aggregate elasticity of substitution between oil and non-oil inputs. The model is used to measure and isolate the impact of changing sectoral composition in the developing world on global oil demand and the oil price in the OECD. I find that structural transformation in non-OECD countries accounts for up to 53% of the oil price increase in the OECD between 1970 and 2010. However, the impact of structural transformation is temporary. Continued structural transformation induces falling oil intensity and an easing of the upward pressure on the oil price. Since a standard one-sector growth model misses this non-linearity, to understand the impact of growth on the oil price, it is necessary to take a more disaggregated view than is standard in macroeconomics.

[1] I would like to thank Timothy Kehoe and Fabrizio Perri for their continued guidance and support during this project. I would also like to thank Cristina Arellano, Patricia Justino, Joe Kaboski, Narayana Kocherlakota, Rick van der Ploeg, Miguel Ricaurte and Tony Venables as well as members of the Minnesota International Trade and Development Workshop for helpful comments and suggestions. I thank Hossein Samiei and Kevin Cheng for initial discussion. I have also benefited from comments of seminar participants at the Society for Economic Dynamics (Istanbul), European Meeting of the Econometric Society (Barcelona), Summer Meeting of the Econometric Society (Pittsburgh), Minneapolis Federal Reserve Bank Seminar Series (Minneapolis), La Pietra-Mondragone Workshop (Florence), Washington University Graduate Student Conference (St. Louis), International Economic Meetings (Warsaw), the Workshop for Young Economists (Guanajuato), Tsinghua University Macro Workshop (Beijing), the Delhi School of Economics Seminar Series (Delhi) and the University of Amsterdam Seminar Series (Amsterdam). All errors are my own. Contact: radek.stefanski@ecn.ulaval.ca

1 Introduction

The average real oil price between 1970-2010 was 40.46 USD, approximately 2.79 times higher than the oil price in the forty one years preceding 1970. At the same time, world employment share in agriculture fell from 56% in 1970 to 34% in 2010 (ILO, 2003), largely driven by the spectacular industrialization of China and India.[2] While this type of correlation is appealing, it is not indicative of the role of structural transformation in developing nations on the oil price. In fact many other factors (such as the oil shocks in the 1970s) have undoubtedly also played a large role in driving the oil price over the period and one should be cautious about drawing conclusions from correlations. Nonetheless, the above observation prompts several questions. What part of the higher oil price has been driven by structural transformation in developing countries? How would the oil price have evolved if developing nations had not began the industrialization process? Finally, will continued structural transformation in these countries contribute to a permanently higher oil price?

My answer to the above questions is based on a hitherto unknown inverted-U shaped relationship in a panel of cross-country data between aggregate oil intensity and the extent of structural transformation. I find that countries in the middle stages of transition have the highest oil intensity: they spend a higher fraction of their income on oil than countries at the beginning or end of transition. If developing nations follow this pattern, their demand for oil is likely to increase initially as they industrialize and their economy becomes more oil intensive. This process should exert an upward pressure on the oil price. However, as their industrialization process comes to a close and their oil intensity begins to drop, their demand for oil should ease and the pressure on the oil price should also decline. The large size of the developing world suggests that this process may have a significant impact on the price of oil.

A decomposition of aggregate oil intensity data reveals why countries exhibit high oil intensity in the middle stages of industrialization. First, sector-specific oil intensities do not remain constant with structural transformation. In agriculture, oil intensity increases as structural transformation progresses, but in industry and services, it falls. When agriculture dominates the economy in the early stages of structural transformation, the rising oil intensity of the agricultural sector drives rising aggregate oil intensity. When non-agricultural sectors begin to dominate the economy at later stages of structural transformation, the falling intensity in industry and services drives falling aggregate oil intensity. Second, independent of the stage of structural transformation, oil intensity in agriculture and services tends to be low, whilst the

[2] Agricultural employment share in China/India declined from 78% of the labor force in 1970 to 45% in 2010. Given China and India's share in the world's total labor force is approximately 40% over the period, if Chinese/Indian agricultural employment share had remained at 78%, world employment share in agriculture would only have fallen to 47%. Thus 60% of the decline in world agricultural employment was driven by China and India.

oil intensity of industry tends to be high. The shift in the composition of an economy from one dominated by a low intensity sector (agriculture) to a high intensity sector (industry) and then back to a low intensity sector (services), will also contribute to an aggregate oil intensity curve shaped like an inverted-U.

To measure and isolate the impact of industrialization on the oil price, I construct and calibrate a multi-sector, multi-country growth model of structural transformation and compare it with the outcomes from a standard one-sector model. This allows me to disentangle the effect of an industrialization from other drivers of the oil price - such as changes in GDP, population, energy efficiency or oil reserve growth rates. The multi-sector model is similar to Echevarria (1997), Duarte and Restuccia (2010), Gollin et al. (2002) and Dekle and Vandenbroucke (2011) but allows for international trade and includes oil as an intermediate input. It is designed to replicate the process of structural transformation and changing sectoral oil intensities observed in the data. Structural transformation is driven by two standard channels: income effects arising from non-homothetic preferences as in Kongsamut et al. (2001) and substitution effects due to unbalanced productivity growth across sectors as in Ngai and Pissarides (2007). Changing sectoral intensities are generated by income effects and elasticities of substitution between oil and non-oil inputs in production that are different from 1.

Since different sectors may potentially vary with respect to elasticities of substitution between oil and non-oil inputs, the changing composition of an economy will affect the resulting aggregate elasticity. The multi-sector framework thus naturally generates an endogenously changing elasticity between oil and non-oil inputs. If, as the data suggest, agriculture is assumed to have a high enough elasticity and non-agriculture a low enough elasticity, a change in the structure of the economy induces aggregate elasticity of substitution to fall from above to below one, generating an inverted-U shaped aggregate oil intensity and contributing to a hump-shaped oil price path. By contrast, in a one-sector model with a standard CES production technology, aggregate elasticity of substitution between oil and non-oil inputs remains constant and aggregate oil intensity is (log) linear.

I find that structural transformation in developing countries - here taken to be all non-OECD countries - accounts for up to 53% of the increase in the oil price in the OECD over the 1970-2010 period. If emerging economies had not structurally transformed at the speed they did, the oil price in the OECD in 2010 would be 33% lower. Furthermore, the upward price pressure from structural transformation in the non-OECD should continue for the coming years. Importantly, this is *not* a permanent effect and should pass as industrialization in developing nations comes to a close. In the long run, structural transformation in non-OECD nations can actually contribute to an oil price that rises at a slower rate than industrialization had not occurred. This prediction is in stark contrast to the outcome of a standard one-sector growth model. In general, these types of models cannot replicate an inverted-U aggregate oil intensity and consequently miss this important non-linearity in the evolution of the oil price. As such, the qualitative takeaway from this paper is that in order to understand the impact of growth on the oil price, it is necessary to take a more disaggregated view than is standard in macroeconomics.

These results are important to both importers and exporters of natural resources. Long lasting changes in the oil price influence the value of oil windfalls in resource exporting countries and in turn impact government revenues, real exchange rates, GDP growth rates and welfare in those countries through changes in resource rents, Dutch Disease or the various channels of the "resource curse" (see, for example, van der Ploeg (2010)). Sustained changes in the oil price can also have a large impact on welfare or GDP growth in oil importing countries, since oil is a crucial input in production. See, for instance, the literature on oil price shocks and their impact on the macro-economy such as Hamilton (1983), Mork (1989) or Blanchard and Gali (2010). The aim of this paper, however, is not to measure the impact of changing oil prices on countries, but rather to investigate the source, and quantify the magnitude, of one particular transmission mechanism through which the large structural transformation in the developing world affects economies in the developed world.

Section 2 establishes the facts with respect to oil markets, industrialization in the developing world and the shape of aggregate oil intensity. Sections 3 and 4 describe and calibrate the model. Sections 5 - 7 present the quantitative predictions of the model, compare them with the data and perform two counterfactual experiments to gauge the impact of structural transformation on the oil price. Finally, section 8 examines the role of the supply side of the oil market on the oil price, whilst section 9 concludes.

2 Facts

I document three sets of facts. First, I show that the long run oil price after 1970 is higher than in the years preceding 1970. I also show the extent of the large structural transformation in non-OECD countries as well as the significant rise in their share of world oil consumption. Second, I provide a link between industrialization and oil demand: I demonstrate the existence of an inverted-U aggregate oil intensity curve as a general feature of structural transformation. Finally, I argue that this relationship, in turn, arises as a consequence of two further features of a structural transformation: the changing size and oil intensity of different sectors.

2.1 The Oil Price and Emerging Economies

Oil Price The curve labeled "Raw Data" in Figure 1, shows the 1900-2012 average annual oil price in 2005 US dollars deflated by the OECD consumer price index (see Appendix 10.1 for

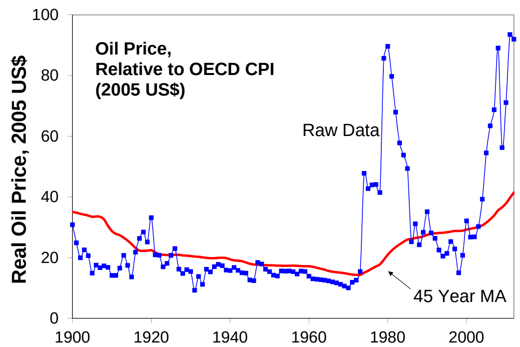

Figure 1: Real oil price and its trend (Relative to the US CPI).

construction details). The price pre-1970 declines at a relatively slow and stable rate, whilst the period after 1970 is characterized by a large increase in both the volatility and the long run level of the oil price. These facts have also been documented by Dvir and Rogoff (2009). The focus of this paper is *not* the increase in volatility, but the high overall level of the oil price after 1970. Quantifying changes in the long run level of the oil price is tricky, exactly because of its high volatility. However, by any reasonable measure, the level of the oil price after 1970 is higher than before. For example, the average oil price in the 43 years before and after 1970 increased by 195% - from 14.56 to 42.90 USD. The median price over these two periods jumped from 15.29 to 35.13 USD - a 129% increase. Even in 1998, when the oil price fell to its lowest level in recent history, it was nonetheless 49% higher than in 1970. To focus on the long run trend, in this paper I pick one particular measure of the long run oil price - a 45-year moving average of the raw data. Naturally, there is enormous volatility in the data that stems from a multitude of sources and is not captured by this trend. I make no attempt to account for all these changes in the oil price. Instead, I focus on the simpler task of isolating the impact of developing nation's industrialization on the long run level of the price of oil.

Structural Transformation in Non-OECD Nations I divide the countries of the world into two regions: the OECD or developed nations and the non-OECD or developing/emerging nations.[3] Figure 2(a) shows how the employment share in agriculture in non-OECD countries has fallen from 67% in 1970 to 40% by 2010. Meanwhile, the share of employment in industry

[3] In this paper I consider the world to consist of 178 countries. Here, and in the rest of the paper, due to the available data I take the OECD to be: Australia, Austria, Belgium, Canada, Denmark, Finland, France, Germany, Greece, Ireland, Italy, Japan, Mexico, Netherlands, New Zealand, Norway, Portugal, Spain, Sweden, Switzerland, Turkey, the United Kingdom and the United States. The non-OECD group consists of the remaining 155 countries. Due to a lack of data, the only countries left out of this group (besides micro- and island-states), are the Democratic Republic of the Congo, Libya, North Korea, Saudi Arabia, the United Arab Emirates and Yemen.)

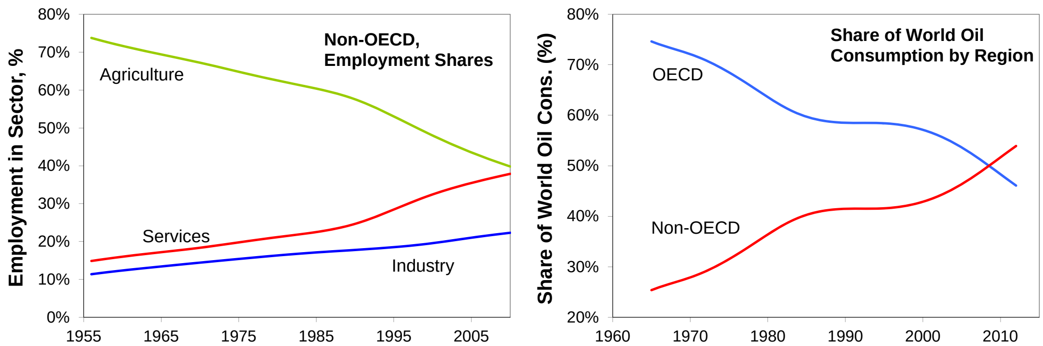

(a) Employment shares in Non-OECD countries (1956-2010)

(b) Changing world oil shares (1970-2007)

Figure 2: Structural transformation in non-OECD countries and oil consumption by region.

has risen from approximately 14% to 22% in industry and from 18% to 38% in services. See Appendix 11 for construction details and data sources. Notice that China and India are the key drivers of this structural transformation. According to official data, more than 57% of the decline in agricultural employment in the non-OECD between 1970-2010 has been driven by a declining agricultural employment share in China and India alone.[4] This industrialization process has also had a large impact on oil demand. Figure 2(b) shows that the share in world oil consumption of non-OECD countries rose by approximately 24 percentage points between 1970-2010 (BP, 2013). Over 50% of the increase in the non-OECD's oil consumption share stems from an increase of China and India's oil consumption.

2.2 The Demand Channel of Structural Transformation

Aggregate Oil Intensity Next, I describe the channel through which structural transformation influences oil demand. In the baseline experiment, I index the progress of a country along a structural transformation by its share of employment in agriculture: countries with high shares of employment in agriculture are relatively structurally undeveloped, whereas countries that have lower agriculture shares are more structurally developed. The share of GDP spent on oil - the aggregate oil intensity - varies with the progress of a structural transformation. Countries at the beginning and end of a structural transformation spend the lowest share of their income on oil, whilst countries in the "middle" of a structural transformation spend the highest share. This fact is shown in Figure 3. The smooth, continuous line plots decile averages

[4] This is even probably an underestimate as Brandt and Zhu (2010) and Rawski and Mead (1998) argue that official Chinese data significantly underreports the outflow of labor from agriculture. In either case, the last forty years have seen the largest inter-sectoral movements of labor in human history.

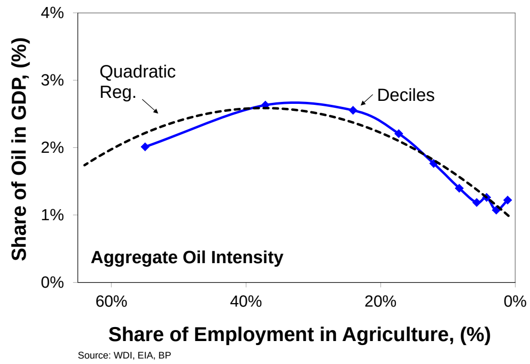

Figure 3: Share of oil in GDP vs. share of employment in agriculture for a panel of the 100 largest countries by population for the years, 1980-2011. Line indicates the decile averages of the data.

of aggregate oil intensity versus the share of employment in agriculture for a panel of the 100 largest countries (for the years 1980-2011). The pooled data is first sorted according to employment share in agriculture, then divided into ten groups. Finally, the average employment share in agriculture and the average oil intensity of each group is calculated and shown in the above graph. To demonstrate this fact more formally, I run four additional quadratic regressions. The first relates aggregate oil intensity to the share of employment in agriculture and its square. The results are shown in column (1) of Table 1 and plotted as the dashed line in Figure 3. The remaining regressions control for time- and country-fixed effects. In each of the regressions, the coefficients are highly significant and an inverted-U shape exists, as can be seen from the negative sign of the squared term. This demonstrates the existence of the hump-shaped pattern both across countries and within countries over time whilst controlling for exogenous drivers of intensity. For data and construction details of the above, see Appendix 10.2.[5]

2.3 Sources of the Hump Shape

Aggregate oil intensity, N, is the sum of oil intensities of individual sectors, n_i, weighted by their share in GDP, s_i. To see this, let O and Y be aggregate oil consumption and value added respectively, p_O be the oil price and i be an index over all sectors then,

$$N \equiv \frac{p_O O}{Y} = \sum_i \left(\frac{p_O O_i}{Y_i} \frac{Y_i}{Y} \right) = \sum_i n_i s_i, \tag{1}$$

To understand the evolution of aggregate oil intensity, it is thus necessary to understand how the size and the oil intensity of specific sectors change over a structural transformation.

[5] An online appendix on my website also contains additional robustness checks that demonstrate similar results using alternative yardsticks of structural transformation and alternative measures of oil intensity. I also show examples of hump-shaped oil intensities in particular countries over time.

COEFFICIENT	(1) Agg. Oil Int.	(2) Agg. Oil Int.	(3) Agg. Oil Int.	(4) Agg. Oil Int.
agrShare	0.0901***	0.0917***	0.0417***	0.0468***
	(0.0043)	(0.0041)	(0.0039)	(0.0038)
agrShareSq	-0.1203***	-0.1211***	-0.0698***	-0.0712***
	(0.0073)	(0.0071)	(0.0064)	(0.0061)
Time FE	no	yes	no	yes
Country FE	no	no	yes	yes
Observations	2102	2102	2102	2102
R^2	0.2190	0.2782	0.8210	0.8401
Argmax	0.374	0.379	0.299	0.329

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 1: Column (1) shows the quadratic regression of aggregate oil intensity versus share of employment in agriculture for a panel of 100 countries over the period 1980-2011. Columns (2)-(4) add time- and country-fixed effects (Coefficients for time dummies and constant terms not shown). Argmax refers to the share of employment in agriculture at which implied quadratic functions reach a peak.

Changing Sector Size The process of systematic change in sectoral size with development has been widely documented in the literature. It is characterized by shares of employment and value added that are falling in agriculture, rising in services, and initially rising and later falling in industry.[6] Figure 4 shows this typical pattern for employment shares in the United States (1860-2004) as an example. Maddison (1982) presents evidence for this process for 16 industrialized countries since 1820-1973. Echevarria (1997) provides examples of this pattern holding in cross-section. Duarte and Restuccia (2010) construct a panel of 29 countries for the period 1956-2000 and document a similar pattern of structural transformation (and its influence on aggregate productivity) in each of the countries over time. Finally, Buera and Kaboski (2008) document similar trends in value added shares for 30 countries from 1820 to 2001.

Changing sectoral oil intensities Next, I demonstrate that another feature of structural transformation is a systematic change in the oil intensity of individual sectors (agriculture, industry and services). I run the following regression:

$$SectOilShare^s_{i,t} = \beta_0 + \beta_1 agrSh_{i,t} + \sum_{i=1}^{T-1} D_{i,t} + \varepsilon_{i,t}, \tag{2}$$

which relates the oil intensity of a sector s in country i at time t, $SectOilShare^s_{i,t}$, to how far countries are in the process of structural transformation measured as the employment share in

[6] Here, and in the rest of the paper unless noted otherwise, I divide sectors according to the standard ISIC III classification. See Appendix 11 for details.

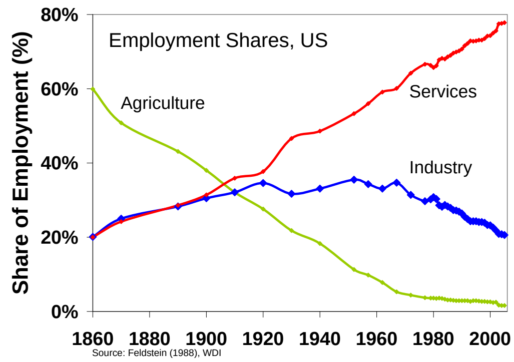

Figure 4: Employment share by sector in the US, 1860-2004.

agriculture, $agrSh_{i,t}$. Since I use panel data I include time dummies, $D_{i,t}$, to control for time-fixed effects. The data itself comes from Input-Output tables from the OECD (2006), for OECD countries as well as Argentina, Brazil, China, Israel, India, Indonesia, Russia and South Africa for 1970-2000. See Appendix 10.3 for construction details and summary statistics. Figure 5 shows the regression results and the corresponding regression lines (extended to the full domain of agricultural employment shares for illustrative purposes). The regressions are significant at the one percent level and provide a good fit of the data. As a country structurally develops (i.e. as its share of employment in agriculture falls), sectoral oil intensity in agriculture *increases* whilst sectoral oil intensity in industry and service *falls*. This pattern could be explained as a movement away from traditional towards more modern, energy intensive agriculture and by improvements in oil use efficiency in non-agriculture.

So, why the inverted-U aggregate oil intensity? This particular pattern of changing structure and oil intensity can result in an inverted-U aggregate oil intensity curve. Consider Figure 5(b). In the early stages of structural transformation, two factors contribute to rising oil intensity. First, the economy is shifting from predominantly oil unintensive agriculture towards oil intensive industry and services. Second, oil intensity of the largest sector - agriculture - is rising. Both of these developments contribute to rising aggregate oil intensity. In the late stages of structural transformation however, there are also two factors contributing to falling oil intensity. First, the economy shifts from oil intensive industry to (relatively) oil unintensive services. Second, the oil intensities of the largest sectors - industry and services - are falling. If oil intensity in agriculture rises slowly enough and oil intensity in industry and services falls fast enough, aggregate oil intensity can fall.

Notice, however, that an inverted-U is not - by any means - inevitable in the above setup. If

COEF.	(1) Oil Int. Agr.	(2) Oil Int. Ind.	(3) Oil Int. Ser.
agrSh.	-0.0492***	0.0599***	0.0750***
	(0.0176)	(0.0205)	(0.0092)
Constant	0.0431***	0.0460***	0.0221***
	(0.0124)	(0.0145)	(0.0065)
Obs.	104	104	104
R^2	0.380	0.283	0.503

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

(a) Regressions results.

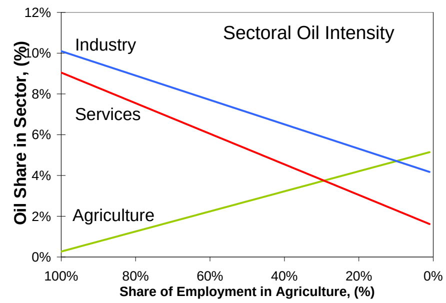

(b) Regression lines. Extended from 0%-100% employment for illustrative purposes.

Figure 5: Changes in sectoral oil intensity over structural transformation.

in the late stages of structural transformation oil intensity in agriculture rises quickly enough, or oil intensity in non-agriculture does not fall fast enough, aggregate oil intensity may not fall. To a large extent the existence of an inverted-U aggregate oil intensity hinges on underlying parameters of the economy.

3 The Model

The model is constructed to capture a shift of labor across sectors and changing sector-specific oil intensities which together can result in an aggregate oil intensity that first rises, then falls as countries structurally transform. The model is then used to isolate the effect of rising oil demand caused by structural transformation in non-OECD countries on the oil price.

3.1 The Economic Environment

There are three regions in the model - the non-OECD (C), the OECD (D) and an Oil Producer (O). C and D are qualitatively identical: they produce agriculture (A), industry (I) and service (S) goods which they then trade with each other and with O. The motive for trade is a taste for variety and the need for oil as a productive input. Quantitatively, C and D differ along the following dimensions: (1) initial levels of sector-specific TFP, B_s^i; (2) sector-specific TFP growth rates, g_s^i; (3) size of their labor force, L^i; and (4) labor force growth rates, g_N^i, where $s = A, I, S$ and $i = C, D$. Differences in TFP will result in countries being at different stages of their structural transformation. Country O is assumed to: (1) only produce oil, (2) to be the only oil producer and (3) to have a small labor force. The model is taken to be a sequence of static problems that vary over time, through exogenous changes in sectoral TFP.

Consumers' problems At each point in time t, the representative consumer in each country $i = C, D, O$ allocates income across sector $s = A, I, S$ intermediate goods according to:

$$\max_{\{A^i_j, I^i_j, S^i_j\}_{j=C,D}} \left(\alpha_A (c_A[A^i_C, A^i_D] - \bar{A})^{\frac{\rho-1}{\rho}} + \alpha_I c_I [I^i_C, I^i_D]^{\frac{\rho-1}{\rho}} + \alpha_S (c_S[S^i_C, S^i_D] + \bar{S})^{\frac{\rho-1}{\rho}}\right)^{\frac{\rho}{\rho-1}} \quad (3)$$

$$\text{s.t.} \sum_{j=C,D} \left(p^j_A A^i_j + p^j_I I^i_j + p^j_S S^i_j\right) = \begin{cases} w^i & \text{if } i = C, D \\ \frac{p^O O^i}{L^i_t} & \text{if } i = O \end{cases}.$$

In the above, $\alpha_s \in (0, 1)$ is the utility weight on sector $s = A, I, S$ so that $\sum_s \alpha_s = 1$, $\bar{A} > 0$ is a subsistence level in agriculture, $\bar{S} > 0$ can be interpreted as a constant level of production of service goods in the home, whilst ρ is the elasticity of substitution between goods. Consumers in $i = C, D$ have wage income w^i, whilst consumers in $i = O$ have income from oil sales, $\frac{p^O O^O}{L^O}$, where p^O is the price of oil and L^O is country O's labor force. Consumers in country i then choose how much of each type of good $s = A, I, S$ from country $j = C, D$ to consume, s^i_j, at price p^j_s. Goods from the same sector but different countries are then bundled together to produce final sectoral goods using the Armington aggregator, $c_s[C, D] = \left(\nu^i_s C^{\frac{\gamma-1}{\gamma}} + (1 - \nu^i_s) D^{\frac{\gamma-1}{\gamma}}\right)^{\frac{\gamma}{\gamma-1}}$, where $\nu^i_s \in (0, 1)$ is country i's preference weight on country C's good. I assume that consumers in $i = C, D$ place the same weight on their home goods, $\nu^i_s = 1 - \nu^{i'}_s$ and consumers in $i = O$ value consumption goods from C and D equally, $\nu^O_s = 0.5$.

Firms' problems At each point in time t, in countries $i = C, D$ and sectors $s = A, I, S$, firms choose how much oil to buy and labor to hire in order to maximize profits:

$$\max p^i_{s,t} (g^i_s)^t B^i_s F_s[O^i_{s,t}, L^i_{s,t}, t] - p_{o,t} O^i_{s,t} - w^i_t L^i_{s,t}, \quad (4)$$

where, $F_s[L, O, t] = (\eta_s (g^t_E O)^{\frac{\sigma_s-1}{\sigma_s}} + (1 - \eta_s) L^{\frac{\sigma_s-1}{\sigma_s}})^{\frac{\sigma_s}{\sigma_s-1}}$. Firms across countries differ in their TFP levels - both the initial sector-specific TFP, B^i_s, and the sector-specific TFP growth factors, g^i_s, potentially vary across countries, $i = C, D$. The oil-use parameter, η_s, and the elasticity of substitution between oil and labor, σ_s, however, remain constant across countries but can vary across sectors. Finally, I assume that there exists energy-specific technical progress, g_E, which is common across countries and sectors.

Oil production I assume a simple reduced form version of oil supply, $O_t = g^t_O (\widetilde{p}^O_t)^\varepsilon B_O$. In this expression, g_O is a growth factor that captures technological progress in the ability to locate, extract or process oil as well as the rate of extraction, whilst $\widetilde{p}^O_t$ is the price of oil relative to a Laspeyres consumption index of the oil producer, $\widetilde{p}^O_t$.[7] Over time oil output can

[7] So that $\widetilde{p}^O_t = \frac{p^O_t}{p^{CPI}_{O,t}}$ where $p^{CPI}_{O,t} = \frac{\sum_{j=C,D} \sum_{s=A,I,S} p^j_{A,t} s^O_{j,t}}{\sum_{j=C,D} \sum_{s=A,I,S} p^j_{A,0} s^O_{j,0}}$.

change due to the exogenous evolution in g_O but it is also responsive to changes in the oil price where ε is the price elasticity of oil supply. In Appendix 11.3 I show how the above production function emerges naturally from a system of equations describing the evolution of oil reserves and extraction choices.

Market Clearing Finally, goods, labor and oil markets clear so that for $s = A, I, S$:

$$\sum_{j=C,D,O} \bar{L}^j s^j_{i,t} = (g^i_s)^t B^i_s F_s[L^i_{s,t}, O^i_{s,t}], \ \sum_s L^i_{s,t} = (g^i_N)^t \bar{L}^i, \ O_t = \sum_{i=C,D} \sum_s O^i_{s,t}. \tag{5}$$

3.2 Discussion of Model

Structural Transformation Structural transformation in the model is driven by two channels. First, there are income effects arising from non-homothetic preferences in agriculture and services such as in Kongsamut et al. (2001) or Duarte and Restuccia (2010). I assume that there exists a subsistence level of agricultural consumption, $\bar{A}$, and a term capturing (in reduced-form) home-production in services, $\bar{S}$. At low levels of TFP, a higher proportion of the labor force must be devoted to agriculture in order to produce the required output. As TFP in agricultural grows, the subsistence output can be produced with fewer workers, releasing them to other sectors. In the same way, at low levels of service sector TFP, workers satiate their demand for service sector goods via ‘home-production’, but as service sector productivity rises, service sector demand shifts towards the marketplace and a higher proportion of the labor force must be devoted to services. Hence, these preferences imply that the income elasticity of service goods is greater than one and the income elasticity of agricultural goods is less than one. Second, there are substitution effects caused by unbalanced sectoral TFP growth and a non-unitary elasticity in preferences between goods (Ngai and Pissarides, 2007). If TFP growth rates are such that $g_A, g_I > g_S$, then setting a low elasticity, $\rho < 1$, results in labor moving away from agriculture and industry towards services. Intuitively, with a low elasticity, consumers enjoy goods in relatively fixed proportions. With unbalanced TFP growth, the only way to maintain fixed proportions in consumption is for labor to move from the faster to the slower growing sectors. Conversely, if $g_S > g_A, g_I$, an elasticity greater than one is needed to induce labor to move towards services. Notice that the model allows for both income and substitution channels to drive structural transformation since - like Duarte and Restuccia (2010) - I find that a single channel alone is not capable of reproducing the extent of the reallocation towards the service sector.

Sectoral Intensities Changing sector-specific oil intensity is captured by a non-unitary oil-labor elasticity of substitution in production. To see this, I write sectoral oil intensity as:

$$\frac{p_O O_s}{V_s} = \frac{p_O O_s}{w L_s} = \frac{p_s F_{s,O} O_s}{p_s F_{s,L} L_s} = \frac{\eta_s}{1-\eta_s}\left(\frac{g_E^t O_s}{L_s}\right)^{\frac{\sigma_s - 1}{\sigma_s}} \tag{6}$$

The first equality follows from the fact that oil is imported so that sectoral value added is simply the sectoral wage. The second equality follows from the profit maximization problem of firms, where $F_{s,L}$ and $F_{s,O}$ are derivatives of F with respect to labor and oil. If the technology-augmented oil-labor ratio rises over time, setting $\sigma_A > 1$ and $\sigma_I, \sigma_S < 1$ allows me to capture rising oil intensity in agriculture and falling oil intensity in industry and services. If the technology-augmented oil-labor ratio decreases, setting $\sigma_A < 1$ and $\sigma_I, \sigma_S > 1$ captures the corresponding oil intensity trends. Notice that in the data the world oil-labor ratio over the 1970-2010 period decreased by 14.8%, or 0.4% a year. It is highly plausible that energy-saving technological progress has increased by more than this, which would imply a rising technology-augmented oil-labor ratio in the data. This assertion is later confirmed in the calibration.

Oil Production The primary focus of this paper is the impact of changing demand on the oil price. As such, I abstract from many supply-side issues such as uncertainty, imperfect competition and storage.[8] Furthermore, I abstract from the impact of non-renewability of oil on its price, a connection first highlighted by Hotelling (1931). In a survey of the theoretical and empirical literature, Krautkraemer (1998) concludes very decisively that the exhaustibility channel seems not to play a first order role in explaining the behavior of the oil price.[9] Whilst the above assumptions are primarily made to simplify the analysis, they are nonetheless unlikely to drive the results. I am interested in the difference in the oil price between scenarios where non-OECD countries do and do not industrialize. In as far as the above supply side assumptions affect both scenarios to a similar extent, the *difference* in the oil price between scenarios will remain largely unaffected. Finally, notice that I have also experimented with a setup that explicitly modeled oil output as emerging from a production function where agents in country O chose the proportion of their labor to devote to the oil sector. This, however, added substantially to expositional complexity and had very little impact on the results. As such, I have decided to adopt a reduced-form version of oil supply.[10]

[8] See Dvir and Rogoff (2009) and the references therein for a discussion where these issues take center stage.

[9] As is pointed out by Dvir and Rogoff (2009), world oil reserves have been increasing in recent decades in spite of an ever rising production. As such "it may well be that technological advances in oil exploration and utilization will be enough to satisfy demand in the foreseeable future."

[10] The results with the more generalized oil production function are available on request.

Parameter	Non-OECD	OECD	Oil Prod.	Target
$B^i_{A,1970}$	0.216	1	–	Empl. Share Agr 1970; Normalization
$B^i_{I,1970}$	0.006	1	–	Empl. Share Ind 1970; Normalization
$B^i_{S,1970}$	0.252	1	–	Rel. Aggregate Prod. 1970; Normalization
g^i_A	1.011	1.032	–	Empl. Share Agr. 2005; Prod. growth in A
g^i_I	1.136	1.018	–	Empl. Share Ind. 2005; Prod. growth in I
g^i_S	0.989	1.008	–	Rel. Aggregate Prod. 2005; Prod. growth in S
g_N	1.021	1.013	1.013	Labor force growth, 1970-2005
$\bar{L}^i$	3.590	1	0.001	Size of Labor force, 1970
g_E	1.025	1.025	-	Change in Real Oil Price, 1970-2005
g_O	–	–	1.008	World oil output growth, 1970-2005
ε	–	–	0.33	Krichene (2006); Dahl and Duggan (1996)

Table 2: TFP, TFP growth rates and labor force parameter values and targets in the model.

4 Calibration

The model is calibrated to OECD and non-OECD data over the 1970-2005 period. The calibration strategy involves selecting parameters to match features of structural transformation, aggregate productivity and oil consumption in both regions. Throughout I assume that a period in the model is one year.

OECD Productivity and Labor Force Parameters I normalize the size of the labor force in D to 1 in 1970. Using data from Heston et al. (2012) I choose g_N^D, g_N^C and $\bar{L}^c$ to match the levels and the growth of the labor force in the OECD and the non-OECD over the 1970-2005 period. For simplicity, I assume that the labor force in O is small at 0.1% of the D's labor force and that it grows at the same rate as in D. I then normalize 1970 productivity levels in each sector of the OECD, so that $B_s^D = 1$ for $s = A, I, S$. Then, using data from Heston et al. (2012), UN (2008) and OECD.Stat (2013) as well as the labor force data above, I calculate g_s^D for $s = A, I, S$ to match the annualized growth rates of sectoral labor productivity found in the data. For details, see Appendix 11 and the online Appendix. The calibration results are given in Table 2.

Preference Parameters I follow Duarte and Restuccia (2010) in calibrating parameters $\bar{A}$, α_A, $\bar{S}$, α_I and ρ. Given the time path of agricultural productivity, parameters $\bar{A}$ and α_A are chosen to match the 1970 and 2005 employment share in agriculture in the OECD. Given the relative time path of labor productivity in industry to services and ρ, parameters $\bar{S}$ and α_I are chosen to match employment shares in the OECD in 1970 and 2005. Data on the proportion of workers employed in each sector comes from OECD.Stat (2013). Since ρ determines how much relative productivity growth between sectors is needed to produce a given reallocation of labor across sectors, different values of ρ result in different patterns of aggregate productivity growth. Hence, following Duarte and Restuccia (2010), I select ρ to match the annualized growth rate

Parameter	Values	Target
$\bar{A}$	0.118	Empl. in Agr. in D, 1970
α_A	10^{-9}	Empl. in Agr. in D, 2005
$\bar{S}$	0.051	Empl. in Ind. in D, 1970
α_I	0.161	Empl. in Ind. in D, 2005
ρ	0.679	Agg. Labor Prod. Growth in D, 1970-2005
σ_A	2.236	Oil cons. in Agr. in D, 1970
σ_I	0.722	Griffin and Gregory (1976)
σ_S	0.750	Oil cons. in Agr. in D, 1970
η_A	0.013	Oil cons. in Agr. in C, 1970
η_I	0.023	Oil cons. in Ind. in C, 1970
η_S	0.012	Oil cons. in Ser. in C, 1970
ν_A^i	0.924	Trade share in Agr. in C, 2005
ν_I^i	0.512	Trade share in Ind. in C, 2005
ν_S^i	0.833	Trade share in Ser. in C, 2005
γ	1.429	Change in oil cons. share in C, 1970-2005

Table 3: Preference, production and trade parameter values and targets in a multi-country model.

of aggregate labor productivity in the OECD between 1970 and 2005 calculated as 1.5% from Heston et al. (2012). For details, see Appendix 11. The calibration results are given in Table 3.

Non-OECD Productivity Parameters Next, I calibrate the productivity parameters of the non-OECD. Available evidence indicates that sectoral TFP levels differ systematically with development - see for example Caselli (2005), Restuccia et al. (2008) or Kuralbayeva and Stefanski (2013). To overcome this problem, I follow Duarte and Restuccia (2010) and use model-based inference to derive sectoral productivity time paths for the non-OECD.[11] Given the vector of preference parameters $(\bar{A}, \alpha_A, \bar{S}, \alpha_I, \rho)$ and the OECD levels and growth-rates of sectoral productivity, I choose B_A^C and B_I^C as well as g_A^C and g_I^C to respectively reproduce the 1970 and 2005 agricultural and industrial employment shares in non-OECD countries calculated using data from ILO (2003), Heston et al. (2012) and OECD.Stat (2013). I choose B_S^C and g_S^C to match the ratio of aggregate labor productivity between non-OECD and OECD countries in 1970 and 2005 from Heston et al. (2012). For details, see Appendix 11. The results are shown in Table 2. The calibration indicates that most of the catch-up in the developing world over the period under consideration occurred through the rapid productivity growth of the industrial sector from a low initial level. Whilst there was also productivity growth in agriculture in the non-OECD, the calibration implies that productivity in the service sector declined slightly over the entire period.

[11] Another way to proceed would be to use PPP-adjusted sectoral value added data and calibrate these parameters in a similar fashion to those of the OECD. The PPP-adjustment data, however, is difficult - if not impossible - to obtain for most non-OECD countries - especially for the service sector.

Oil Production and Oil Efficiency The price elasticity of oil supply, ε, is set to 0.33 - the mid-point between the highest and lowest estimate in the literature.[12] Section 8 shows that results are robust to different choices of ε. I normalize B_O to 1 and choose g_O so that the model matches world oil output growth over the 1970-2005 period of approximately 1.1% calculated from BP (2013). Finally, I set $g_E = 1.0245$ to match the observed increase in the real oil price in the OECD over the 1970-2005 period. The model is thus *calibrated* to replicate the increase in the (long run) oil price post 1970 and can then be used to measure the contribution of industrialization in developing countries towards the increase. Importantly however, the estimate of g_E is eminently plausible. For example, Hassler et al. (2011) find this value to be 2.24% per year after 1973. The calibration results are given in Table 2.

Trade Parameters The home bias parameters are chosen to match the proportion of foreign consumption in a given sector in the non-OECD in 2005 using data from OECD (2006). In particular, I restrict $1 - \nu_s^i$ to be the same in C and in D and I choose the sectoral weights of foreign goods in D, $1 - \nu_s$, to match the non-OECD's sectoral import share in the data. See Appendix 11 for details. The home bias parameters in O are set to 0.5 - the oil producer shows no preference for C's or D's good. I choose the elasticity of substitution between C's and D's goods, γ, to roughly match the change in the share in world oil consumption of the non-OECD and the OECD over the 1970-2005 period. This parameter impacts oil consumption by driving substitution between local and foreign goods. Since oil needs to be imported to produce goods locally, γ influences oil consumption over time. The calibration results are given in Table 3.

Oil Intensity Parameters Oil parameters, σ_s and η_s are chosen to match sectoral oil consumption in non-OECD and OECD countries in 1970. From a firm's first order conditions we find:

$$\frac{O_{s,t}^i}{L_{s,t}^i} = \left(\frac{\eta_s}{1-\eta_s}\right)^{\sigma_s} \left(\frac{w_t^i}{p_t^O}\right)^{\sigma_s} \left(g_E^{\sigma_s - 1}\right)^t. \tag{7}$$

The oil price is common across regions, but countries with a higher wage will have a higher sectoral oil consumption. The sector-specific elasticity of substitution between oil and labor, σ_s, determines *how* much more oil per worker richer countries will use. The η_s term, on the other hand, determines how much more oil is used by a particular sector within a country. The calibration proceeds as follows. Since total world oil supply in the first period is exogenous,[13] only five of the six oil parameters are needed to match oil consumption in each sector at a point

[12] Estimates of long run price elasticity of oil supply tend to be quite low. In non-OPEC countries they range from 0.08 reported by Krichene (2006), 0.29 reported by Alhajji and Huettner (2000) to 0.15-0.58 reported by Gately (2004) or 0.58 reported by Dahl and Duggan (1996) in the US. For OPEC countries, these elasticities are likely to be even lower.

[13] In the initial period $\widetilde{p}_0^O = 1$ since p_t^O is normalized to one and so $O_0 = 1$.

in time. As such, the elasticity of substitution between oil and labor is set to the mid-range of the values estimated by Berndt and Wood (1975) and Griffin and Gregory (1976), $\sigma_I = 0.72$. Due to a lack of data on oil consumption by sector in 1970 for all countries in the sample, I use the cross-sectional properties of sectoral oil intensity at different stages of structural transformation to infer sectoral oil consumption (see Appendix 11 for details). The results from the third step of the calibration are given in Table 3. The calibration implies that industry has the highest oil-use ($\eta_I = 0.023$), whilst services ($\eta_S = 0.012$) and agriculture ($\eta_A = 0.013$) have roughly the same oil-use. It also implies that oil and labor are (gross) substitutes in agriculture with elasticity $\sigma_A = 2.24$ and (gross) complements in industry and services with elasticities $\sigma_I = 0.72$ and $\sigma_S = 0.75$.

Discussion of the Calibration Like Duarte and Restuccia (2010) I find that given the observed time path of sectoral labor productivity in the OECD, the observed process of labor reallocation between industry and services could not be achieved in the model without an income elasticity greater than one in services. Similarly to these authors setting $\bar{S} = 0$, all else equal, implies that the model generates lower employment shares in industry than in the data and lower growth rate of aggregate productivity. On the other hand, setting $\rho = 1$, ceteris paribus, the model implies slightly more labor reallocation to services than the data and somewhat lower annualized growth rate of labor productivity.

Next, in the calibration of oil parameters, I use the variation in cross-country wage to oil price ratios, rather than variation across time, to pin down key oil parameters. I do this for two reasons. First, there is only limited variation in wage to oil price ratios over time. For example, in 1994, the wage-oil price ratio was approximately 19 times higher in the OECD than in the non-OECD (WDI, BP). On the other hand in OECD time-series data, between 1970-2010, the highest wage-oil price ratio was only 6.66 times higher than the lowest. Furthermore, much of this variation came from the relatively short but very sharp oil shocks of the 1970s, which may not be relevant for the long term substitutions examined in this paper.

Second, it is unclear whether points on a long-run cost function are, in fact, being observed with annual time-series data. It is quite plausible that the ease of substitution between labor and oil depends on the type of capital that is in place. A time-series approach reflects (relatively) short-run adjustment to price changes, with a fixed technological character of the capital stock. A cross-sectional approach reflects a capital stock whose technological character has had time to adjust to different energy prices. In the short and medium run, substitution between labor and energy may be limited due to the type of capital that is in place. In the long run, it may be easier to substitute between labor and energy due to the technological changes inherent in the capital. Although my model does not explicitly include capital, calibrating to cross-sectional

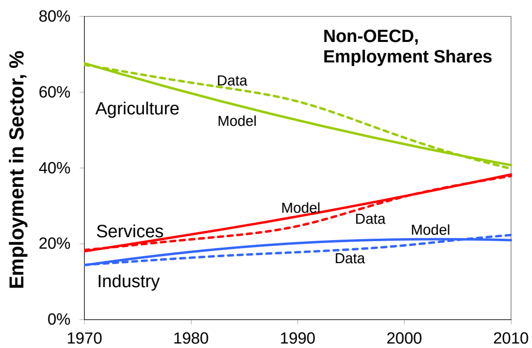

(a) Non-OECD, employment shares

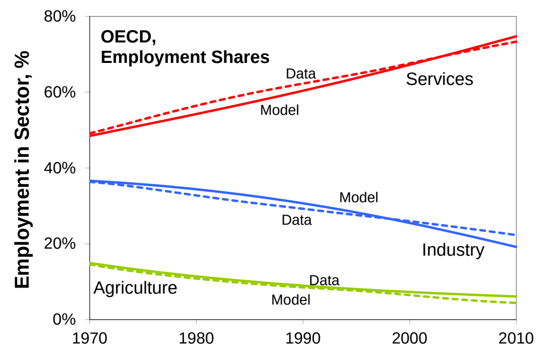

(b) OECD, employment shares

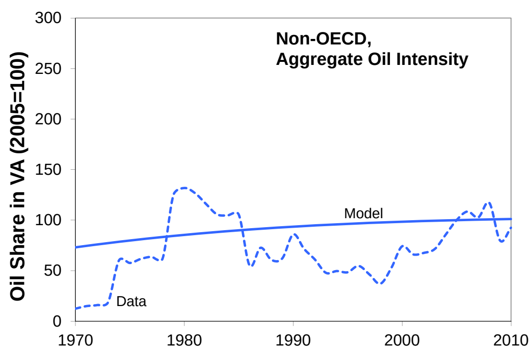

(c) Non-OECD, aggregate oil intensities

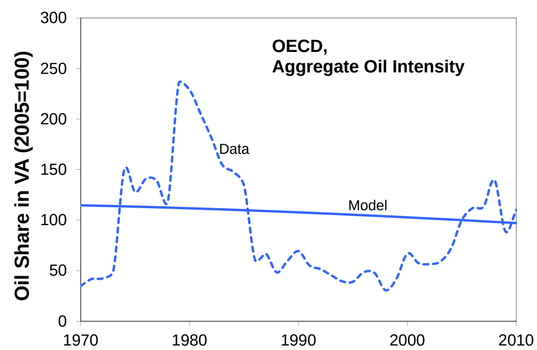

(d) OECD, aggregate oil intensities

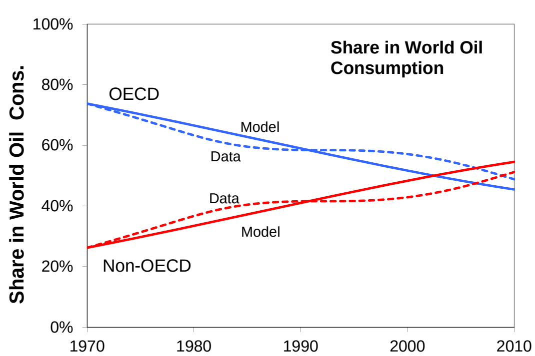

(e) Oil consumption shares

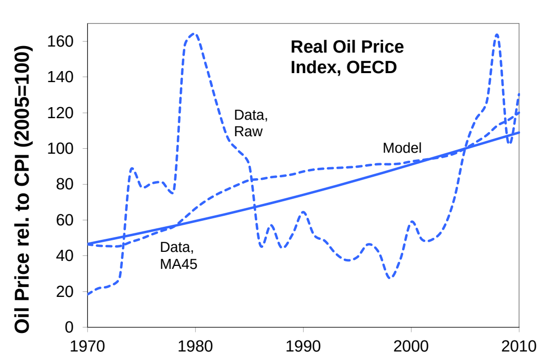

(f) OECD oil price

Figure 6: Simulation and data.

data captures the long run substitution possibilities between labor and oil when faced with implicit improvements in the capital stock.

In Appendix 11, I perform two exercises to check the robustness of the above results. First, I re-estimate these elasticities using equation (7) directly and find similar values to those above. Second, I compare my results to the literature and again find that the magnitudes found above are representative of those found by other authors.

5 Baseline Simulation

This section examines the results from the baseline simulation and compares them to the data. The top two panels of Figure 6 show the changing shares of employment over time. In the non-OECD, employment shares fall in agriculture and rise in industry and services. In the OECD, a similar pattern emerges for agriculture and services. Since the OECD has higher TFP levels, it is further along the structural transformation path and its share of employment in industry falls. The model is calibrated to match the trend of employment shares in both regions over the period and hence matches the data well. Panels (c) and (d) show aggregate oil intensity in both regions. Since non-OECD countries are in the earlier stages of structural transformation, the model predicts the initial part of the hump-shape - a rising aggregate oil intensity. Since the OECD is in the closing stages of its structural transformation, we see intensity peak early on and later decline. The model captures some of the increase of aggregate oil intensity in the non-OECD and the relative decline of intensity in the OECD - but not the high volatility. Figure 7 confirms this by showing the ratio of intensity across countries which eliminates volatility common to both countries. The model does better here, although it under-predicts the relative increase in intensity. Finally, panel (e) shows the oil shares for both regions and panel (f) shows the model and oil price data. Since the model is calibrated to replicate these values, it matches the data well. Appendix 12 further discusses robustness and fit.

6 Hump-shaped Oil Intensity

This section examines the mechanism driving aggregate oil intensity in the model. For illustrative purposes, I extend the simulation to 2100.[14] As I argued in section 2.3, the inverted-U aggregate oil intensity for non-OECD countries labeled ‘baseline’ in Figure 8(a) is driven by two factors. First, the economy shifts from predominantly oil-unintensive agriculture towards oil-intensive industry and then to oil-unintensive services. Second, oil intensity of agriculture is rising and

[14] To do this, I assume that future sectoral and oil productivity growth rates in the three regions stay at their 1970-2005 levels.

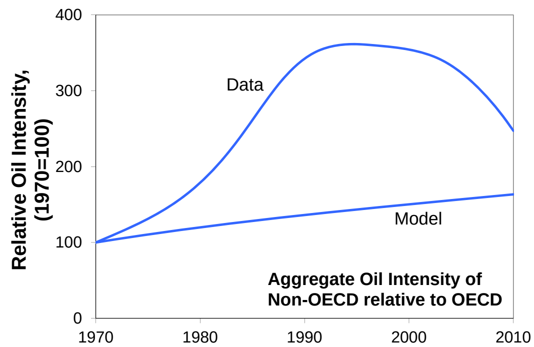

Figure 7: Simulation and data for aggregate oil intensity of Non-OECD countries relative to the OECD.

that of non-agriculture is falling. These facts are demonstrated in Figure 8(b). When the economy is dominated by agriculture, this contributes to a rising aggregate oil intensity. When the economy becomes dominated by non-agriculture, this contributes to a falling aggregate oil intensity. In the model, these two channels are captured by two sets of parameters. First, the pattern of oil-use across sectors found in section 4, $\eta_I > \eta_A, \eta_S$, contributes to the overall high level of oil intensity in industry. Second, the pattern in sector-specific elasticities of substitution between oil and labor, $\sigma_A > 1 > \sigma_I, \sigma_S$ generates rising oil intensity in agriculture and declining intensity in industry and services through equation (6).

To see which of these forces are key in generating the hump-shaped aggregate intensity, I perform three counterfactuals. First, I set the oil-use parameter in all sectors equal to the oil use parameter in the service sector, $\eta_s = 0.012$. Second, I replace all sectoral oil-labor elasticity parameters by the oil-labor elasticity in the service sector, $\sigma_s = 0.75$. Finally, I take a version of the model, where both the oil-use and sectoral oil-labor elasticity parameters are replaced by the corresponding service sector parameters so that $\eta_s = 0.012$ and $\sigma_s = 0.75$. Doing so allows me to examine the importance of different elasticities and oil-use parameters across sectors in driving aggregate oil intensity. Figure 8(a) shows the implication of these counterfactuals for aggregate oil intensity. It is evident that the key driver of the rising part of the oil-intensity hump-shape is the difference in elasticities across sectors. On the other hand, different oil-use parameters contribute to a greater decline in oil intensity in the second part of the oil-intensity hump-shape. In a model where both channels are switched off, oil intensity declines monotonically.

Finally, I consider how changes in relative income affect relative oil demand across countries and how this drives income elasticities of oil and hump-shaped aggregate oil intensity patterns.

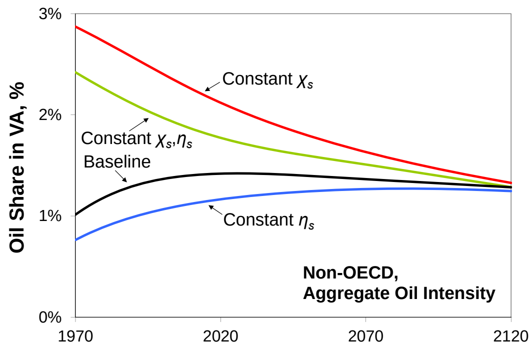

(a) Aggregate oil intensity, non-OECD

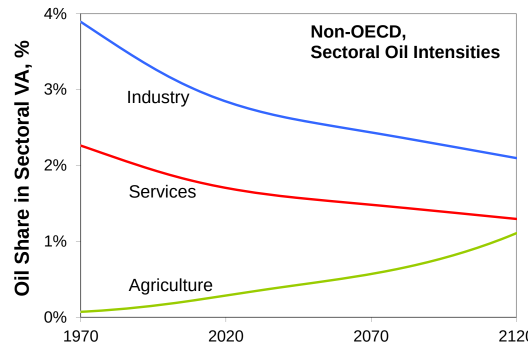

(b) Non-OECD, sectoral oil intensities

Figure 8: Sources of the hump-shaped oil intensity in the non-OECD and their impact on the oil price in in the OECD. Versions of the model where either η_s, χ_s or both parameters are constant.

Taking the ratio of equation 7 between both regions gives the following expression:

$$\frac{O^C_{s,t}/L^C_{s,t}}{O^D_{s,t}/L^D_{s,t}} = \left(\frac{w^C_t}{w^D_t}\right)^{\sigma_s}, \tag{8}$$

where σ_s is the (relative) income elasticity of (relative) sectoral per-capita oil consumption. In a model where η_s and σ_s are the same across sectors, the above expression holds at the aggregate level so that a one percent increase in relative income results in a $\sigma = \sigma_s$ percent increase in relative per-capita oil consumption. In a model with sectoral differences in oil-use and elasticity parameters, the impact of a one percent increase in relative income can vary over time. Figure 9(a), shows the evolution of relative per-capita oil consumption with relative wages of the non-OECD in the baseline and the counterfactuals. In a model where oil-use and oil-labor elasticity parameters are constant across sectors, relative oil demand increases (log) linearly with relative wages. When these parameters differ across sectors however, there is a strongly non-linear relationship. Figure 9(b) plots the corresponding (relative) aggregate income-elasticity of oil at different points in time for the different counterfactuals. In the version of the model where σ_s and η_s are constant across sectors, the relative income elasticity is constant and equal to 0.75. However, when oil-use and elasticity parameters differ across sectors, the aggregate elasticity declines from above to below 1. Denoting this elasticity by σ, it is easy to show that $\sigma - 1$ is the (relative) income elasticity of (relative) oil intensity.[15] When $\sigma > 1$, rising relative wages translate into rising relative oil intensities. When $\sigma < 1$, the relationship is negative. Since the multi-sector model generates elasticities that are generally declining from above to below 1 over time, this gives rise to the hump-shaped aggregate intensity. Finally, notice from Figure 9(b)

[15] To see this, notice that the oil intensity of country i is $N^i_t \equiv O^i_t/(w^i_t L^C_t)$, so that $\frac{\partial \log(N^C_t/N^D_t)}{\partial \log(w^C_t/w^D_t)} = \sigma - 1$.

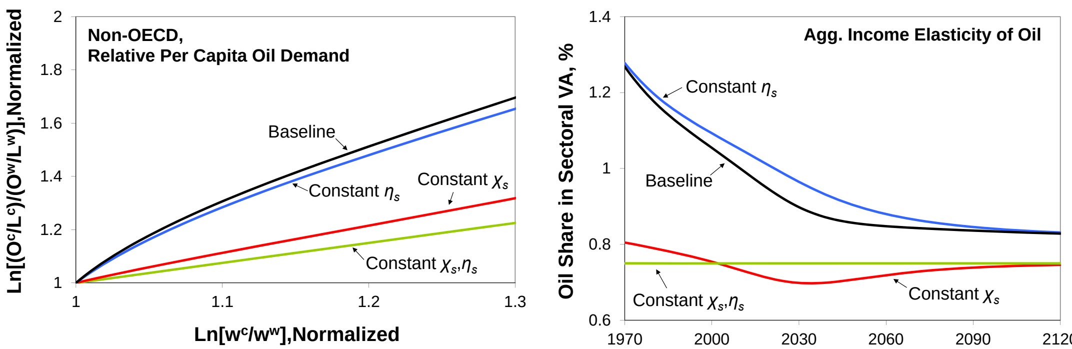

(a) Relative oil demand, non-OECD

(b) Income elasticity of oil, non-OECD

Figure 9: Sources of the hump-shaped oil intensity in the non-OECD, relative oil demand and implied (relative) income elasticity of oil demand.

that the large decline in elasticity stems predominantly from differences in sectoral elasticities of substitution rather than from variation in sectoral oil weights.

7 Prices

The previous section demonstrated the importance of the multi-sector framework for the formation of hump-shaped aggregate oil intensity. In this section, I perform two counterfactuals that 1) gauge the effect of growth and structural transformation in the developing world on the oil price in the OECD and 2) highlight the importance of the multi-sector framework (in contrast to the more standard one-sector framework) for modeling oil prices.

No Growth In the first counterfactual, I switch off productivity growth in non-OECD countries in all sectors and compare the resulting price to the one obtained in the baseline model. The result is shown in Figure 10(a), and labeled 'No growth'. This experiment allows me to measure the total effect that growth and structural transformation in developing countries have had on the oil price. If countries within the non-OECD had not undergone structural transformation at the speed that they did, oil prices would have increased, but the increase would have been substantially smaller. Comparing the baseline and no-growth scenarios in Figure 10(b), the developing country effect results in oil prices in the model that are 56% higher at their peak in 2024 than they would be without the non-OECD effect. Furthermore, according to the model, the long run oil price in the OECD in 2010 is 49% (or roughly 12.44 USD) higher than if non-OECD countries had not structurally transformed at the speed that they did. Since the long run

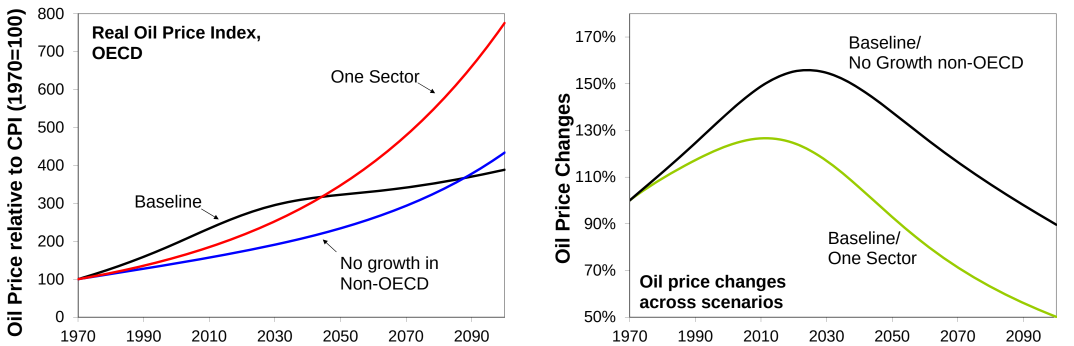

(a) Real oil price index, OECD

(b) Baseline oil price change relative to counterfactuals

Figure 10: Counterfactual experiments, prices.

oil price increased from approximately 14.65 to 37.88 USD between 1970 and 2010, growth and structural transformation in non-OECD countries can account for 53% of the observed increase in the oil price.

The Role of Structural Transformation In the second exercise, I compare the results of the multi-sector model to those of a corresponding one sector model.[16] In this manner I turn off the impact of structural transformation but keep all other growth effects. The resulting price index is shown in Figure 10(a) (labeled 'One Sector') and increases at a constant rate. Omitting structural transformation thus misses a crucial non-linearity in prices. This is highlighted in Figure 10(b) which shows the change in oil prices in the baseline relative to the one sector world. Notice that prices increased 27% more by 2011 at their peak in the baseline than in the one-sector model.[17] Importantly however, continued structural transformation contributes to an oil price that eventually increases at a lower rate than a standard one-sector model would predict and gives rise to the hump shape visible in the figure. The source of this non-linearity is the hump-shaped oil intensity curve, generated - in the case of the multi-sector model - by different oil-labor elasticities of substitution and oil-use parameters across sectors as well as structural transformation. By comparing the two curves in Figure 10(b), it becomes evident

[16] To obtain a one-sector version of the model, I first replace all sector-specific parameters with the corresponding parameters of the service sector as in section 6 so that $\sigma_s = 0.750$, $\eta_s = 0.012$ and $\nu_s = 0.833$. Second, in each country I restrict productivity growth rates and levels to be equal across sectors in each sector and choose these parameters to reproduce the observed paths of aggregate labor productivity of each region over the 1970-2005 period so that $g_s^C = 1.021$, $B_{s,1970}^C = 0.142$, $g_s^D = 1.014$ and $B_{s,1970}^D = 1$. Finally, I recalibrate $g_O = 1.004$ to maintain observed world oil output growth of 1.1% a year over the 1970-2005 period. All other parameters remain the same.

[17] Since the models are different in the initial period, they can no longer be viewed as normalized price *levels*. However, it is still correct to compare the price dynamic over time between the two scenarios.

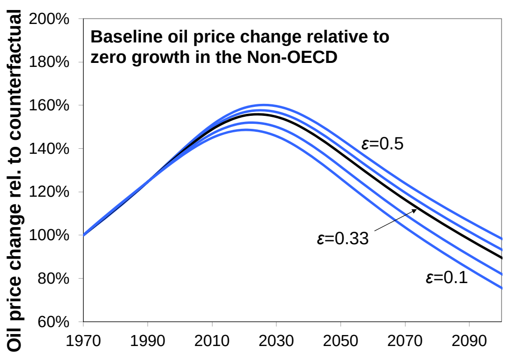

(a) Baseline oil price change relative to the no-growth in the non-OECD.

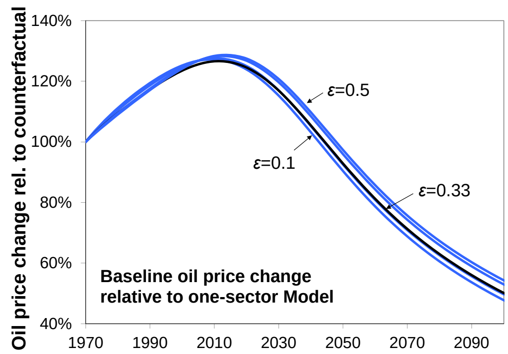

(b) Baseline oil price change relative to one-sector model.

Figure 11: Impact of varying price elasticity of oil supply on OECD price of oil.

that 54% of the increase in the relative price of oil observed between 1970 and 2010 in the multi-sector model was driven by structural transformation effects whilst the rest was driven by growth factors. The lesson here is that to understand the impact of growth on the oil price, it is crucial to take a more disaggregated view than is standard in macroeconomics.

8 Supply Side of the Oil Market

Elasticity of Oil Supply What role does the price elasticity of oil supply, ε play in driving the results? I vary this elasticity so that $\varepsilon = 0.1; 0.2; 0.4; 0.5$ and re-calibrate the model to match the oil growth rate of 1.1% per year between 1970 and 2005. First, aggregate oil intensity in non-OECD countries is nearly indistinguishable between the scenarios. Next, I switch off productivity growth in the non-OECD and use the model to measure the additional impact of non-OECD growth and industrialization on the oil price in the OECD. The results are shown in Figure 11(a). Interestingly, a higher price elasticity of oil supply implies that growth in the non-OECD contributes even more to changes of the oil price. The contribution of growth and industrialization however are robust to the chosen elasticity between 1970-2010, and only somewhat sensitive after 2010. The impact of industrialization alone is also robust to chosen elasticities. Figure 11(b) shows changes in the oil price in the baseline model relative to the corresponding one-sector model. The differences across scenarios are small.

The Role Oil Output Growth If growth and industrialization in developing countries explain only 53% of the post-1970 higher oil price, what accounts for the rest? One obvious candidate is the slow down in world oil output growth rates. The average growth rate of world

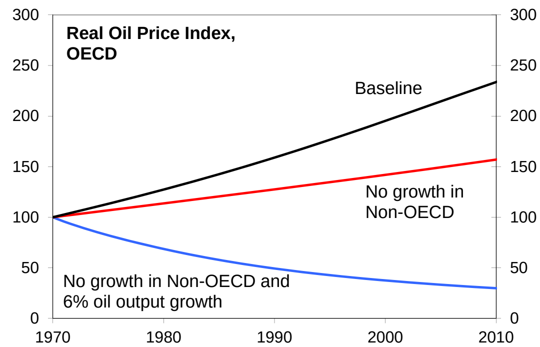

Figure 12: The role oil output growth

oil output fell from approximately 6% between 1900 and 1970 to 1.1% after 1970 (EIA). I can use the model to measure the contribution of the decline in world oil output growth rates towards the higher post-1970 oil price. In this counterfactual, I continue to assume that non-OECD countries did not grow over the 1970-2005 period. In addition however, I choose $g_O = 1.072$ so that world oil output in the model grew at 6%. The results are shown in Figure 12. Without industrialization in the non-OECD and with the higher oil output growth, the price of oil over the period would have fallen by approximately 67% . Whilst the decline in oil output growth was key in generating a higher oil price post-1970, industrialization in the developing world still played a significant role: relative to the counterfactual, 38% of the higher oil price in 2010 can be explained by the industrialization that took place in non-OECD countries.

9 Conclusion

As structural transformation progresses, aggregate oil intensity first rises and then falls - forming an inverted-U shape. This can result in oil prices that follow a similar pattern with structural transformation. As large, developing countries such as China and India enter the most oil intensive phases of their structural transformation, oil prices will rise faster than they otherwise would. The faster growth is not necessarily permanent. In the medium to long run, the pressure on the oil price will ease, as industrialization in developing countries comes to an end and oil intensity falls. Since standard growth models do not generate a hump-shaped intensity, they miss this non-linearity and can give misleading implications about the long-term oil price. To understand the impact of growth on the oil price, it is necessary to take a more disaggregated view than is standard in macroeconomics.

This paper is the first to identify an inverted-U aggregate oil intensity curve in the data, the first to build a model that theoretically justifies its existence through endogenously changing aggregate elasticities of substitution between oil and non-oil inputs and the first to consider the long term price path implications of such a curve. The main contribution of the paper however, is to take a systematic approach to a contentious topic - the impact of industrializing countries on the oil price. In particular, the model presented here predicts that as long as the structural transformation in developing nations follows past patterns, the upward pressure on oil prices from growth and industrialization will continue for many years but that it will not be permanent. The model predicts that in the more distant future, oil prices can return - or even fall below - the trend they would have been following without industrialization in the developing world as these economies become dominated by services. As such, the impact of industrialization on the oil price in developing nations such as China and India is not necessarily permanent.

10 Empirical Appendix

10.1 Oil Price

The curve labeled "Raw Data" in Figure 1, shows the 1900-2012 average annual oil price in 2005 US dollars deflated by the OECD consumer price index. The nominal oil price data comes from BP (2013). The OECD CPI that is used to deflate the nominal price is constructed as follows. From the WDI I obtain CPI indices for all OECD countries under consideration from 1970 to 2012, with the exception of Germany and the United Kingdom that only have data from 1990 and 1988 respectively. Next, from OECD (2012) I obtain weights for each of the CPI indices which are "based on the previous year's private final consumption expenditure of Households and Non-profits institution serving Households expressed in purchasing power parity terms".[18] I then calculate a weighted average of the CPI's for the period 1970-2012. Due to lack of CPI index data for a large majority of countries before 1970, I extend the resulting CPI index back in time using the implied growth rates of the US CPI obtained from BP (2013). The resultant index in normalized to 1 in 2005, and the nominal oil price in USD is then divided by this index to obtain a 'real' measure of the oil price in 2005 USD used in the paper.

10.2 Aggregate Oil Intensity

I construct a baseline panel of data that consists of the world's largest 100 countries (by population in the year 2000) over the 1980-2011 period.[19] Each data point is composed of: 1) a time period, 2) the share of employment in agriculture and 3) the aggregate oil intensity. The data for the shares of employment in agriculture comes from WDI (2013). The aggregate oil intensity of the economy is constructed by calculating the current year value of oil consumed in an economy (in dollar terms), divided by the current year GDP of the country (in dollar terms) which comes from UN (2008). The population data comes from Heston et al. (2012). Total country-specific oil consumption (in thousands of barrels a day) comes from the EIA (2013) whilst the current oil price data come from BP (2013). Summary statistics for the data are shown in Table 4(a). The baseline oil price used is the 45-year moving average. Using un-smoothed oil price data results in oil intensities that are significantly higher due to the high oil prices shock in the early 1980s. The shape of the curve however, remains unchanged if observed oil price data is used. The first four columns of Table 4(b), replicate the regressions found in Table 1 uses the full data sample (instead of only the largest 100 countries) and leaves the data un-smoothed. The data still exhibit a strong, inverted-U shape. As a robustness check, I construct an alternative

[18] The data is available at: `http://stats.oecd.org/mei/default.asp?lang=e&subject=8&country=OTO`.

[19] In the baseline, I also drop Tajikistan, Azerbaijan and Iraq from the sample as these are clear outliers with oil intensity as high as 40%. Notice, that this in no way drives the results, as I show below where I consider the full sample of the data.

Variable	Obs	Mean	Std. Dev.	Min	p5	p50	p95	Max
Agr. Emp. Share	2102	0.167	0.168	0	0.012	0.099	0.521	0.922
Agg. Oil Int. (smoothed)	2102	0.017	0.011	0.001	0.006	0.014	0.041	0.095
Agg. Oil Int.	2529	0.047	0.040	0.002	0.011	0.036	0.122	0.408
Agg. Oil Int. (2005 PPP)	2398	0.036	0.021	0.005	0.017	0.031	0.081	0.260

(a) Summary statistics for aggregate oil intensity and agricultural employment share data.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Agg. Int.	Agg. Int.	Agg. Int.	Agg. Int.	Agg. Int. PPP	Agg. Int. PPP	Agg. Int. PPP	Agg. Int. PPP
agrSh	0.209***	0.205***	0.159***	0.134***	0.053***	0.052***	0.070***	0.046***
	(0.012)	(0.011)	(0.019)	(0.015)	(0.007)	(0.007)	(0.008)	(0.007)
agrShSq	-0.278***	-0.270***	-0.238***	-0.226***	-0.093***	-0.094***	-0.082***	-0.087***
	(0.019)	(0.017)	(0.031)	(0.023)	(0.011)	(0.011)	(0.012)	(0.012)
Time FE	no	yes	no	yes	no	yes	no	yes
Ctry FE	no	no	yes	yes	no	no	yes	yes
Obs.	2529	2529	2529	2529	2398	2398	2398	2398
R^2	0.112	0.303	0.589	0.765	0.031	0.062	0.761	0.798
Argmax	0.376	0.380	0.334	0.296	0.285	0.277	0.427	0.264

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

(b) Columns 1-4 leave the oil price data un-smoothed. Columns 5-8 calculate oil intensity using the 2005 price of oil and 2005 PPP GDP.

Table 4: Analog of Table 4(b) which tests robustness of baseline empirical results.

measure of oil intensity, where the value of oil consumed in a country is measured using the 2005 USD oil price and the value of GDP is measured using 2005 PPP USD from Heston et al. (2012). This measure eliminates any potential influence of changing relative prices that are not captured by the country and time-fixed effects. Columns 5-8 of Table 4(b) shows the results. Once more the data still exhibit a strong, inverted-U shape.

10.3 Sectoral Oil Intensity

The oil intensity by sector data, is derived from Input-Output tables constructed by the OECD (2006) and is calculated by dividing the value of sectoral inputs in the category "Refined petroleum products, coke and nuclear fuel" by total sectoral value added[20] in a given country and year. The data under consideration is for Australia, Canada, Denmark, France, Germany, Italy, Japan, the Netherlands, the UK and the US for the years 1970, 1972, 1975, 1977, 1980, 1985, 1986, 1990. For the years 1995 and 2000, the data consists of countries from the OECD (Australia, Austria, Belgium, Canada, Czech Republic, Denmark, Finland, France, Germany, Greece, Hungary, Ireland, Italy, Japan, South Korea, Luxembourg, the Netherlands,

[20] Although more disaggregated data would be desirable, this was the most disaggregated, comparable cross-country input-output data that was available.

Variable	Obs	Mean	Std. Dev.	Min	p5	p50	p95	Max
Agr. Emp. Share	104	0.099	0.125	0.006	0.013	0.058	0.440	0.667
Agr. Oil Int.	104	0.045	0.026	0.007	0.013	0.041	0.094	0.134
Ind. Oil Int.	104	0.045	0.028	0.009	0.014	0.036	0.093	0.164
Ser. Oil Int.	104	0.021	0.015	0.004	0.007	0.016	0.051	0.075

Table 5: Sectoral oil intensity summary data

New Zealand, Norway, Poland, Portugal, Slovak Republic, Spain, Sweden, Turkey, UK and the USA) as well as Argentina, Brazil, China, Israel, India, Indonesia and Russia and South Africa. Summary data is presented in Table 5.

11 Calibration Appendix

11.1 Aggregate Values

Labor Force, Output, Productivity I use data from Penn World Tables (PWT) version 7.1 (see Heston et al. (2012)) to construct annual time series of PPP-adjusted GDP in constant prices, PPP-adjusted GDP per worker in constant prices, and total employment for OECD and non-OECD countries. For each country, I construct total employment L using the variables Population (POP), PPP Converted GDP Chain per worker at 2005 constant prices (RGDP-WOK), and PPP Converted GDP Per Capita (Chain Series) at 2005 constant prices (RGDPCH) as $L = (RGDPCH \times POP)/RGDPWOK$. For each country, I then construct PPP-adjusted GDP in constant 2005 prices, using the variables "PPP Converted GDP Per Capita (Laspeyres), derived from growth rates of c, g, i, at 2005 constant prices (RGDPL)" and "Population (POP)" as $GDP = RGDPL \times POP$. To obtain the values for the OECD and the non-OECD, I simply sum over the countries in each group. The series for PPP-adjusted GDP per worker in constant prices is then computed as $y = GDP/L$.

Oil Consumption Oil consumption data for the world and the OECD comes from BP (2013). To obtain non-OECD oil consumption I subtract OECD oil consumption from world oil consumption.

11.2 Sectoral Values

Definition of Sectors I define sectors according to the standard ISIC III classification. Agriculture corresponds to categories 1-5 (agriculture, forestry, hunting, and fishing). Industry corresponds to categories 10-45 (mining, manufacturing, construction, electricity, water and gas) and services refers to categories 50-99 (wholesale, retail, transport, government, financial etc).

Sectoral Labor Force To calculate sectoral employment in each country-block I combine PWT data with sectoral employment share data from ILO (2005), ILO (2013) and OECD.Stat (2013). The "ALFS Summary tables" (OECD.Stat, 2013) dataset is a subset of the Annual Labour Force Statistics database which presents annual labor force statistics and broad population series for 34 OECD member countries plus Brazil. Using this data set I can calculate the proportion of the civilian labor force in agriculture, industry and services for each of the OECD countries for 1970-2010, sL_i, $i = A, I, S$. Similarly ILO (2005) and ILO (2013) provide estimates of the proportion of global labor force in agriculture, industry and services for select years between 1970-2010, sL_i, $i = A, I, S$. The missing data in this last series is obtained by linear interpolation. Since the global data and the OECD data comes from different sources, in order to remain consistent across regions, I multiply the above implied shares by the aggregate labor force data found from PWT to obtain sectoral employment levels, $L_i = sL_i \times L$. Then, to construct sectoral employment in the OECD, I simply sum the labor force estimates for each of the countries in the OECD. To construct sectoral employment in the non-OECD, I subtract the sectoral employment levels of the OECD from the global sectoral employment levels.

Sectoral productivity To calculate sectoral productivity growth rates for the OECD, I follow Duarte and Restuccia (2010) and combine data from the PWT with data from the UN (2008). First, for each country and each sector (agriculture, industry and services), I compute sectoral shares of total value added using constant-price (2005) value-added by sector data from the UN:

$$sVA_t^i = \frac{VA_t^i}{\sum_{i=A,I,S} VA_t^i}.$$

Then, I multiply these shares by the series of aggregate GDP in 2005 (international) prices computed above from the PWT, which allows me to obtain series of value added by sector for each country. Finally, I sum the resulting sectoral value added in international prices over countries, to obtain sectoral value added for the OECD as a whole. I then use this data to calculate sectoral growth rates for the OECD. Notice, that I do not use this method to calculate productivity *levels* or growth rates in the non-OECD. Instead, I use the model to back out sector-specific productivity levels and growth rates. The reason is the same as in Duarte and Restuccia (2010): "there is substantial evidence that the PPP-conversion factors differ systematically across sectors in development."

Import Shares The sectoral import share for the non-OECD is derived from Input-Output tables constructed by the OECD (2006) for 2005. For each sector, I calculate the value of exports

	1995			1995, 2000		
	(1)	(2)	(3)	(4)	(5)	(6)
	Agr	Ind	Ser	Agr	Ind	Ser
$\log\left(\frac{w^i}{p^O}\right)$	1.22***	0.58***	0.44***	1.25***	0.59***	0.44***
	(0.11)	(0.06)	(0.06)	(0.08)	(0.05)	(0.04)
t				0.16***	0.06**	0.07***
				(0.04)	(0.03)	(0.02)
Observations	26	26	26	52	52	52
R^2	0.85	0.81	0.66	0.84	0.73	0.71

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 6: Direct estimation of sectoral oil-labor elasticity of substitution. Columns (1)-(3) refer to the baseline model and show results for 1995 cross-country data. Columns (4)-(6) refer to a model with population and energy efficiency growth and show results for 1995 and 2000 cross-country data. (Source: OECD, UN, BP).

from the OECD to the non-OECD and I take these values as the sectoral imports of the non-OECD. The sectoral import share of the non-OECD is then obtained by dividing the value of imports by the sectoral value added obtained from the same source. I find that the non-OECD import shares are 3.99% in agriculture, 45.34% in industry and 12.02% in the service sector.

Sectoral Oil Consumption In this section, I estimate oil consumption by sector in the OECD and the non-OECD in 1970. In order to do this (and for lack of sectoral oil consumption data in the non-OECD and the OECD), I use the regressions presented in equation 5 for the year 2000 and without time dummy variables. These regressions describe what fraction of a sector's value added is devoted to oil at any point in the structural transformation and they are robust over time. According to these regressions, countries that employ 65% of their work force in agriculture - approximately the share employed by the non-OECD in 1970 - spend approximately 1.5% of their agriculture value-added, 8.8% of their industry value-added and 6.8% of their service value-added on oil. Since I have data on the value-added of non-OECD agriculture, industry and services in 1970, this allows me to estimate the total value of oil used by each sector in 1970 - $p_{O,1970}O^C_{s,1970}$. This allows me to calculate what fraction of total oil consumption in the non-OECD was consumed by which sector, $\frac{p_{O,1970}O^C_{s,1970}}{\sum_s p_{O,1970}O^C_{s,1970}} = \frac{O^C_{s,1970}}{\sum_s O^C_{s,1970}}$. Given data on the non-OECD's total oil consumption, I can then estimate the quantity of oil consumed by each sector in the non-OECD and hence its share in total world oil consumption. I find the proportion of sectoral oil consumption in the OECD in the same manner.

11.3 Oil Production Derivation

The oil production function in section 3 can be shown to arise from the following two equations describing the evolution of oil reserves and oil extraction choices:

$$R_{t+1} = g_R \left(\frac{\widetilde{p}_t^O}{\widetilde{p}_{t-1}^O} \right)^{\varepsilon} R_t - O_t \text{ and } O_t = \eta \left(\frac{\widetilde{p}_t^O}{\widetilde{p}_{t-1}^O} \right)^{\varepsilon} R_t. \tag{9}$$

The world enters period t with a stock of oil reserves, R_t, accessible given period t's technology. Next period's reserves, R_{t+1}, depend on 1) how much oil is extracted this period, O_t, 2) on exogenous growth in the ability to locate, extract or process reserves captured by the growth factor, g_R and 3) on changes in the price of oil relative to a Paasche consumption index of the oil producer, $\widetilde{p}_t^O$. This last factor can be interpreted as exploration or as oil fields becoming accessible only after large increases in oil price levels. Notice also that exploration and technological progress show up in reserves only in the subsequent period. The amount of oil extracted in a period, O_t, depends on the size of reserves, R_t, on the exogenous parameter η and on changes in the oil price, $\widetilde{p}_t^O$. The responsiveness of exploration and extraction to changes in price depends on the parameter ε assumed equal across activities. This assumption implies that world oil output is given by $O_t = g_O^t (\widetilde{p}_t^O)^{\varepsilon} B_O$ where, $B_O \equiv \eta(\widetilde{p}_{-1}^O)^{-\varepsilon} R_0$ and $g_O \equiv g_R - \eta$. Notice that throughout I assume that $\eta \left(\frac{\widetilde{p}_t^O}{\widetilde{p}_{t-1}^O} \right)^{\varepsilon} < 1$ and that oil producers take oil prices as given.

11.4 Sectoral Oil-Labor Elasticities

First, I estimate an equation (7) using cross-sectional 1995 data and a panel of 1995 and 2000 data for: Austria, Brazil, Canada, China, Denmark, Finland, France, Germany, Greece, Hungary, India, Indonesia, Italy, Japan, Netherlands, Norway, Poland, Portugal, Russian Federation, Slovak Republic, Spain, Sweden, Turkey, UK and the US. The sectoral oil consumption data is obtained from the Input-Output tables constructed by the OECD (2006). In my model, wages correspond to value added per capita. This data is taken from UN (2008). Finally, 1995 and 2000 oil prices are taken from BP (2013). Taking the log of equation (7), I obtain:

$$\log \left(\frac{O_{s,t}^i}{\bar{L}_{s,t}^i} \right) = \sigma_s \log \left(\frac{\eta_s}{1-\eta_s} \right) + \sigma_s \log \left(\frac{w_t^i}{\widetilde{p}_t^O} \right) + t \log \left(g_E^{\sigma_s - 1} \right).$$

I estimate this equation using OLS and the results are shown in columns (1)-(3) of Table 6. The slope parameter on the wage to oil price ratio is the elasticity of substitution between oil and labor in a particular sector. The elasticity of substitution between oil and labor is approximately 1.22 for agriculture, 0.58 for industry and 0.44 for services. Using 2000 data and 1995-2000 panel data, I find similar results. The calibrated values in the body of the paper are thus quite close to these empirically estimated parameters.

Second, I compare my calibrated elasticity values to those found in the literature. Broadly speaking, the approach for estimating these values in the literature is similar to mine. For example, Berndt and Wood (1975) use time-series data (1947-71) to estimate the factor share functions (arising from a transcendental logarithmic production function but similar - in principle - to equation 7) in US manufacturing for four inputs - capital, labor, energy and materials - using iterative three-stage least squares. Griffin and Gregory (1976) perform a similar analysis but using cross-sectional and panel manufacturing data. For agriculture, Shankar et al. (2003) estimate the Allen partial elasticity of substitution between energy and labor to be 4.58 in Hungary. This is higher than in the calibration, but of the same order of magnitude. Furthermore, the authors use a short time period and one that included significant political upheaval in Hungary. A broader study by Salhofer (2000) performs a simple meta-analysis of 35 studies of European agriculture that examine elasticities of substitution between different inputs. He finds the average elasticity of substitution between hired labor and the "purchased inputs" category (defined as energy, fuel, pesticides and seed) is 1.3. This is close to the calibrated value. Next, in industry numerous studies find Allen partial elasticities of substitution between energy and labor to be less than one. Berndt and Wood (1975) estimate this elasticity for the US to be 0.65. Griffin and Gregory (1976) estimate the elasticity for numerous advanced European countries and for the US to be between 0.72 and 0.87. Kemfert (1998) as well as Kemfert and Welsch (2000) estimate this elasticity for Germany to be 0.871. These values are again of similar magnitude to the value 0.72 found in my calibration. Finally, Koschel (2000) finds elasticity in the German service sector to be 0.28. This again roughly matches the magnitude in our model estimates.

12 Online Appendix [NOT TO BE PUBLISHED]

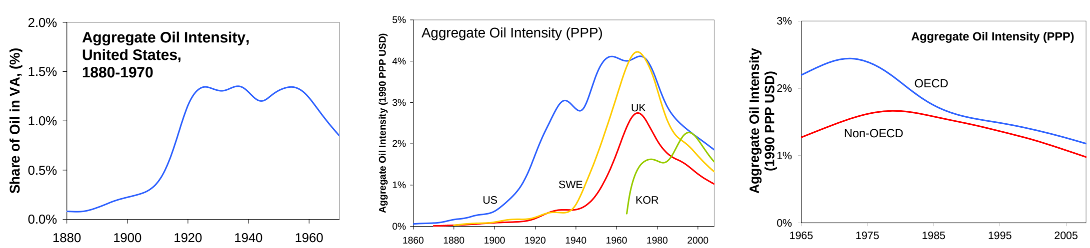

(a) USA (current)

(b) USA, Sweden, UK, S. Korea (PPP)

(c) Non-OECD and OECD

Figure 13: Aggregate oil intensity measured using constant oil price and 1990 PPP GDP.

Time Series Evidence I provide some evidence of a hump-shaped oil intensity within countries over time. One has to be cautious when interpreting this data due to the existence of confounding factors that may drive aggregate oil intensity over time, but are largely exogenous to the country such as the oil shocks in the 1970s and 1980s or the transition to different energy technologies. The US between 1880 and 1970 largely avoids the above pitfalls: its oil market was almost entirely closed[21] and relatively free of government intervention whilst oil had already been adopted as a viable energy source. Figure 13(a) shows that a hump-shaped aggregate oil intensity existed for US during this period. Another way to try to get around some confounding factors - like the volatile oil price post-1970 - is to use the constant price, PPP measure of oil intensity introduced above. Figure 13 shows oil intensities for the US (1860-2008), the UK (1870-2008), Sweden (1880-2008) and South Korea (1965-2008). A strong hump-shape pattern emerges in all cases. One can also re-do the same calculation for OECD and non-OECD countries for 1965-2008 - again we see a clear hump-shape pattern in non-OECD intensity and a roughly declining intensity over the period in the OECD. Nonetheless, we should be cautious about concluding too much from these time-series graphs. Whilst the aggregate intensity of different countries peaks at different times, there is a notable decline in intensity post-1970 which could be a consequence of the massive oil shocks of that period. Thus, whilst looking at time-series evidence is enlightening, the panel-data approach that control for time- and country-fixed effects is still preferable.

In the above, oil consumption (Table Db166) and nominal GDP data (Table Ca10) for the US comes from Carter et al., eds (2006). Historical oil consumption is denominated in British Thermal Units (BTU). To convert the price of oil per barrel, into price per BTU I have used the estimated conversion factor of 5.8 million BTU's per crude barrel of oil as given by the EIA. Notice that this conversion is only an estimate and can vary over time and with type of oil. Historical oil consumption is also an estimate derived from trade data. As such, the levels of oil intensity found here should be taken with a grain of salt. Data for oil consumption before 1965 for Sweden comes from Kander and Lindmark (2004) and from Fouquet and Pearson (1998) for the UK. Oil consumption data for all countries and regions after 1965 comes from (BP, 2013). PPP measures of GDP come from Maddison (2007) with 1990 being the reference year. Oil price data in 1990 comes from BP (2013).

Alternative Measures of Structural Transformation Next, I consider alternative measures of structural transformation. Agricultural employment shares represent the preferred yardstick due to its widespread availability and simplicity, but none of the results hinge on this

[21] Pre 1970, 80-100% of US oil consumption was produced domestically. By 1980 this share was down to 50%.

	agrvaSh	agrvaShSq	mfgvaSh.	servaSh	servaShSq	logGDPpc	sqlogGDPpc
Agg. Int	0.061***	-0.126***	0.022***	0.092***	-0.114***	0.073***	-0.004***
	(0.006)	(0.011)	(0.004)	(0.015)	(0.015)	(0.003)	(0.000)
Time FE	yes		yes	yes		yes	
Obs.	2,053		1,884	2,049		2,186	
R^2	0.130		0.097	0.143		0.311	

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 7: Regressions of aggregate oil intensity versus shares of agricultural, manufacturing and service value added in total value added and log of GDP per capita (2005, PPP USD) and their squares.

particular choice. Table 7 re-runs the regressions of Table 1 using value added shares of agriculture, manufacturing and services as well as log of PPP GDP per capita (in 2005 USD) as yardsticks of industrialization. All the results go through as before.[22]

Finally, notice that the model lends itself to a particular measure of structural transformation. Combining equations 6 and 7 and taking the ratio across countries gives the following relationship: $\frac{N^c_{s,t}}{N^w_{s,t}} = \left(\frac{w^c_t}{w^w_t}\right)^{\sigma_s - 1}$, which describes how changes in relative wages influence relative sectoral oil intensity. In a one sector model, this relation also describe the log-linear evolution of *aggregate* intensity with the relative wage. In Table 8 I regress aggregate oil intensity relative to OECD aggregate oil intensity on value added per capita relative to OECD value added per capita (and its square) for our baseline sample of countries. The results suggest the relationship is highly *non-linear* and follows a strong hump shape.

12.1 Total Factor Productivity

Ideally, I would calculate TFP as a residual directly from the data using the expression: $B^i_{s,t} = \frac{Y^i_{s,t}}{\left(\eta_s (g^E O^i_{s,t})^{\frac{\sigma_s - 1}{\sigma_s}} + (1-\eta_s) L^i_{s,t}{}^{\frac{\sigma_s-1}{\sigma_s}}\right)^{\frac{\sigma_s}{\sigma_s-1}}}$, where $Y^i_{s,t}$ is a country i's sector s gross output, $O^i_{s,t}$ is its oil use and $L^i_{s,t}$ is its labor force. Data on sectoral oil use and gross output however, is available only for limited countries. Consequently, I estimate sectoral productivity growth data using labor productivity data which is given by $\bar{D}^i_{s,t} = \frac{V^i_{s,t}}{L^i_{s,t}}$, where $V^i_{s,t}$ is the value added of sector s in country i at time t in constant prices, calculated above. In the model, labor productivity is defined as:

$$D^i_{s,t}(B^i_{s,t}) \equiv \frac{p^i_{s,2005} B^i_{s,t} \left(\eta_s (g_E O^i_{s,t})^{\frac{\sigma_s-1}{\sigma_s}} + (1-\eta_s) L^i_{s,t}{}^{\frac{\sigma_s-1}{\sigma_s}}\right)^{\frac{\sigma_s}{\sigma_s-1}} - p^O_{2005} O^i_{s,t}}{L^i_{s,t}}, \tag{10}$$

[22] Notice that I regress oil intensity versus manufacturing share and not by manufacturing share squared. This is because manufacturing tends to follow a hump-shape with structural transformation.

	Rel. Agg. Int	Rel. Agg. Int	Rel. Agg. Int	Rel. Agg. Int
logRelw	-0.683***	-0.673***	-0.914***	-0.894***
	(0.029)	(0.028)	(0.034)	(0.033)
logRelwSq	-0.106***	-0.107***	-0.026***	-0.034***
	(0.007)	(0.007)	(0.008)	(0.008)
Time FE	no	yes	no	yes
Country FE	no	no	yes	yes
Obs.	1,177	1,177	1,177	1,177
R^2	0.472	0.517	0.924	0.938

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 8: Regressions of aggregate oil intensity (normalized by OECD intensity) on current price GDP per capita (normalized by OECD GDP per capita).

where the term in the numerator is the value added of a sector s in country i at time t in 2003 prices (notice that this is just the gross output of sector s less intermediate inputs - the oil used in the sector). TFP levels, $B^i_{s,t}$, are chosen so that the model's implied labor productivity, matches observed labor productivity, i.e. $D^i_{s,t}(B^i_{s,t}) = \bar{D}^i_{s,t}$.

12.2 Trade Shares

Figure 14 plots the *aggregate*, constant price import consumption shares in the model and the data (WDI (2013)). In the non-OECD, the model captures the trend of the data although it overestimates trade initially and underestimates it in the second part of the sample. In the OECD the model does well in the second part of the sample matching both the trend and the level of trade, but overestimates trade shares initially. There are several reasons for the imperfect match. First, the model is calibrated to sectoral rather than aggregate trade flows (and only to that of the non-OECD at that). Second, I keep a symmetry in preferences across home and foreign goods in the model as there is no intrinsic reason to suspect that cross-country preferences over home and foreign goods vary systematically (ceteris paribus). As such, when faced with two very different trade patterns, there is only so much asymmetry that a fundamentally symmetric model can generate. Third, there is a large increase in global trade over this period that models far more sophisticated than mine have failed to match. Finally, the baseline model does not take into account the massive changes in non-OECD trade policy over this period (particularly in the non-OECD). With respect to the last factor, I can choose trade policies in the model that take the form of time-varying, country-specific wedges, τ_i for $i = C, D$, on all imported consumption goods to match each country's non-fuel import share, with all revenues being lump sum rebated

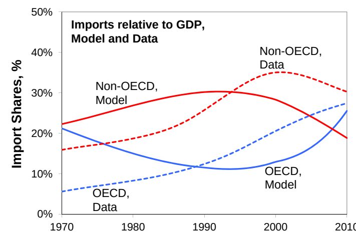

Figure 14: Import shares: Imports relative to GDP in data and model.

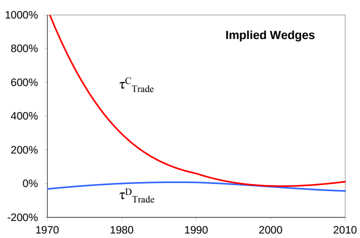

(a) Implied Wedges.

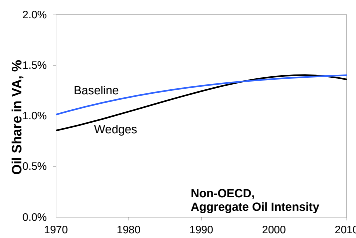

(b) Agg. Oil intensity the non-OECD

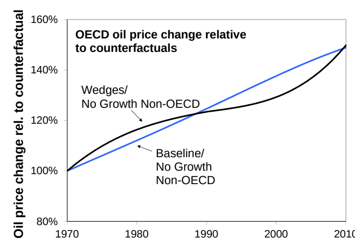

(c) Oil price rel. to model with zero growth in OECD.

Figure 15: Impact of different population growth rates.

to consumers. The results are shown in Figure 15. As expected, a massive opening up within the non-OECD is visible in the form of a decline of the wedge. The impact on intensity and prices is shown in panels (b) and (c) of the Figure and is minimal. The model predicts a slightly lower intensity in the non-OECD, but one that is qualitatively similar. The impact of the wedges on the price of oil relative to the zero-growth model is also negligible.

References

Alhajji, A. and D. Huettner, "The target revenue model and the world oil market: Empirical evidence from 1971 to 1994.," *The Energy Journal*, 2000, *21(2)*, 121144.

Berndt, E. and D. Wood, "Technology, Prices and the Derived Demand for Energy," *The Review of Economics and Statistics*, 1975, *57*, 259–268.

Blanchard, Olivier J. and Jordi Gali, *International Dimensions of Monetary Policy*, University Of Chicago Press,

BP, *BP Statistical Review of World Energy June 2008* 2013.

Brandt, Loren and Xiaodong Zhu, "Accounting for China's Growth," *IZA Discussion Papers 4764*, 2010.

Buera, Francisco J. and Joseph P. Kaboski, "Scale and the Origins of Structural Change," *No WP-08-06, Working Paper Series, Federal Reserve Bank of Chicago*, 2008.

Carter, Susan B., Scott Sigmund Gartner, Michael R. Haines, Alan L. Olmstead, Richard Sutch, and Gavin Wright, eds, *Historical Statistics of the United States, Earliest Times to the Present: Millennial Edition*, Cambridge University Press, 2006.

Caselli, Francesco, *Handbook of Economic Growth, Volume 1A.*, Elsevier B.V.,

Dahl, C. and T. Duggan, "U.S. energy product supply elasticities: A survey and application to the U.S. oil market," *Resource and Energy Economics*, 1996, *18(3)*, 24363.

Dekle, Robert and Guillaume Vandenbroucke, "A Quantitative Analysis of China's Structural Transformation," *Federal Reserve Bank of San Francisco Working Paper*, 2011.

Duarte, M. and D. Restuccia, "The Role of the Structural Transformation in Aggregate Productivity," *The Quarterly Journal of Economics*, 2010, *125(1)*, 129–173.

Dvir, Eyal and Kenneth S. Rogoff, "Three Epochs of Oil," *NBER Working Paper No. 14927*, 2009, *April.*

Echevarria, C., "Changes in Sectoral Composition Associated with Economic Growth," *International Economic Review*, 1997, *38.*

EIA, *International Energy Statistics* U.S. Energy Information Administration 2013.

Fouquet, Roger and Peter J.G. Pearson, "A thousand years of Energy use in the United Kingdom," *The Energy Journal*, 1998, *19(4)*, 1–42.

Gately, D., "OPECs incentives for faster output growth.," *The Energy Journal*, 2004, *25(2)*, 7596.

Gollin, Douglas, Stephen Parente, and Richard Rogerson, "The Role of Agriculture in Development," *American Economic Review*, 2002, *92(2)*, 160–164.

Griffin, J. and P. Gregory, "An Intercountry Translog Model of Energy Substitution Responses," *The American Economic Review*, 1976, *66*, 845–857.

Hamilton, James D., "Oil Shocks and the Macroeconomy: The Role of Price Variability," *The Journal of Political Economy*, 1983, *91(2)*, 228–248.

Hassler, John, Per Krusell, and Conny Olovsson, "Energy-Saving Technical Change," *Institute for International Economic Studies Working Paper*, 2011, *Stockholm University.*

Heston, Alan, Robert Summers, and Bettina Aten, *Penn World Table Version 7.1* Center for International Comparisons of Production, Income and Prices at the University of Pennsylvania 2012.

Hotelling, H., "The Economics of Exhaustible Resources," *Journal of Political Economy*, 1931, *39*, 137–175.

ILO, "Economically active population 1950-2010: ILO database on estimates and projections of the economically active population (5th edition)," 2003.

_ , *World Employment Report 2004-05: Employment, Productivity, And Poverty Reduction* 2005.

_ , *Global employment trends 2013: Recovering from a second jobs dip* International Labour Office (ILO) 2013.

Kander, Astrid and Magnus Lindmark, "Energy consumption, pollutant emissions and growth in the long run: Sweden through 200 years," *European Review of Economic History*, 2004, *8*, 297–335.

Kemfert, C., "Estimated substitution elasticities of a nested CES production function approach for Germany," *Energy Economics*, 1998, *20*, 249–264.

_ **and H. Welsch**, "Energy-Capital-Labor Substitution and the Economic Effects of CO2 Abatement: Evidence for Germany," *Journal of Policy Modeling*, 2000, *22*, 641–660.

Kongsamut, P., S. Rebelo, and D. Xie, "Beyond balanced growth," *The Review of Economic Studies*, 2001, *68.*

Koschel, Henrike, "Substitution Elasticities between Capital, Labour, Material, Electricity and Fossil Fuels in German Producing and Service Sectors," *ZEW Discussion Paper No. 00-31*, 2000.

Krautkraemer, Jeffrey A., "Nonrenewable Resource Scarcity," *Journal of Economic Literature*, 1998, *36(4)*, 2065–2107.

Krichene, N., "World crude oil markets: Monetary policy and the recent oil shock.," *IMF Working Paper*, 2006, *06/62.*

Kuralbayeva, K. and Radoslaw Stefanski, "Windfalls, Structural Transformation and Specialization," *Journal of International Economics*, 2013, *90*, 273–301.

Maddison, A, *Phases of Capitalist Development*, Oxford: Oxford University Press, 1982.

Maddison, Angus, "Historical Statistics for the World Economy: 1-2008 AD," 2007.

Mork, Knut Anton, "Oil and the Macroeconomy When Prices Go Up and Down: An Extension of Hamilton's Results," *The Journal of Political Economy*, 1989, *97(3)*, 740–744.

Ngai, L. and C. Pissarides, "Structural change in a multi-sector model of growth," *The American Economic Review*, 2007, *97.*

OECD, *OECD I-O Database - Industry Classification and Concordance with ISIC Rev 3, 2006 Edition* OECD 2006.

_ , "OECD.Stat," 2012.

OECD.Stat, *ALFS Summary Tables* OECD 2013.

Rawski, T.G. and R.W. Mead, "On the Trail of China's Phantom Farmers," *World Development*, 1998, *26*, 767–781(15).

Restuccia, Diego, Dennis Tao Yang, and Xiaodong Zhu, "Agriculture and Aggregate Productivity: A Quantitative Cross-Country Analysisy," *Journal of Monetary Economics*, 2008, *55(2)*, 234–250.

Salhofer, Klaus, "Elasticities of Substitution and Factor Supply Elasticities in European Agriculture: A Review of Past Studies," *Institut fr Wirtschaft, Politik und Recht Universitt fr Bodenkultur Wien, mimeo*, 2000, *83-W-2000.*

Shankar, B., J. Piesse, and C. Thirtle, "Energy substitutability in transition agriculture: estimates and implications for Hungary," *Agricultural Economics*, 2003, *29*, 181–193.

UN, *National Accounts Statistics database* 2008.

van der Ploeg, Frederick, "Natural Resources: Curse or Blessing?," *Cesifo Working Paper*, 2010, *No. 3125.*

WDI, "World Development Indicators," 2013.