Structural Transformation and the Oil Price[1]

Radoslaw (Radek) Stefanski
University of Oxford and OxCarre

September 22, 2011

Abstract

What part of the high oil price can be explained by structural transformation in China and India? 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 China and India on world oil demand and the oil price in the OECD. Structural transformation in China and India accounts for 26% 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@economics.ox.ac.uk

1 Introduction

The average real oil price between 1970-2009 was 37 USD, approximately 3 times higher than the oil price in the forty years preceding 1970. Figure 1 suggests that this increase has coincided with a structural transformation away from agriculture towards industry and services in China and India. Employment in the agricultural sector in these countries declined from nearly 80% of the labor force in 1970 to 40% by the mid-2000's and has accounted for half of the decline in world employment share in agriculture over the period.[2] Whilst many other factors (such as the oil shocks in the 1970's) have undoubtedly contributed to a higher oil price, the above observation nonetheless prompts several questions. What part of the higher price has been driven by structural transformation in China and India? How would the oil price have evolved if China and India had not began the industrialization process? Finally, will continued structural transformation in these countries contribute to a permanently higher oil price?

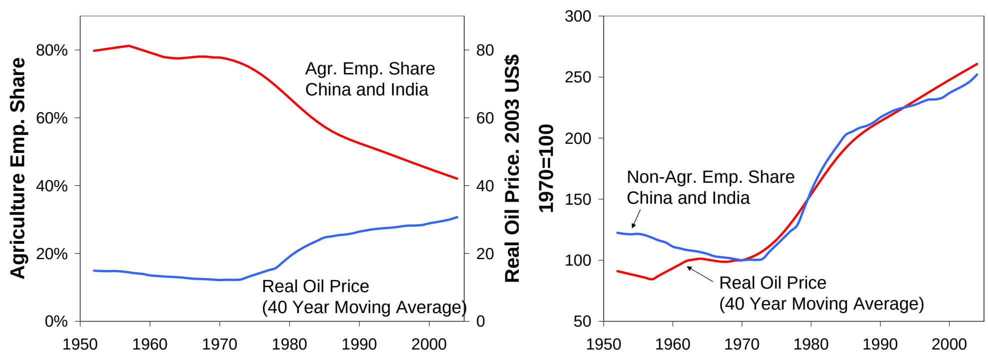

(a) Agriculture employment share in China/India and the real oil price.

(b) Indices of non-agriculture employment share in China/India and the real oil price.

Figure 1: Structural transformation in China and India and the real oil price (40 Year MA).

My answer to the above questions is based on an inverted-U shaped relationship I identify 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 spend a higher fraction of their income on oil than countries at the beginning or end of transition. If China and India 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

[2] World employment share in agriculture fell from 56% in 1970 to 36% in the mid-2000's (ILO, 2003). Assuming that China and India's share in the world's total labor force is 1/3, if Chinese/Indian agricultural employment had remained at 80%, world employment share in agriculture would only have fallen by half to 46%.

decline. The large size of China and India 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 higher oil intensity in the middle stages of industrialization. First, I show that 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 is initially large at the start of industrialization, the rising oil intensity in agriculture contributes significantly to rising aggregate oil intensity. When the economy however is dominated by non-agriculture at the end of the industrialization process, 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 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 industrializing China and India 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 one.

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.

Structural transformation in China and India is found to account for up to 26% of the increase in the oil price in the OECD over the 1970-2010 period. If China and India had not structurally transformed at the speed they did, the oil price in the OECD in 2010 would be 21% lower. What's more, the upward price pressure from structural transformation in China and India should continue in the coming decades. Importantly, this is *not* a permanent effect and should pass as China and India's industrialization comes to a close. In the long run, structural transformation in China and India can actually contribute to an oil price that rises at a slower rate than if China and India had not industralized. 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 hence miss this important non-linearity in the evolution of the oil price. As such, the qualitative take away from this paper is that 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 China and India will affect economies in the rest of the world.

Section 2 establishes the facts with respect to oil markets, China and India's industrialization and the shape of aggregate oil intensity. Section 3 and 4 describe and calibrate the model. Sections 5 - 6 present the quantitative predictions of the model, compare them with the data and performs two counterfactual experiments to gauge the impact of structural transformation on the oil price. Finally, section 7 examines the role of the supply side of the oil market and considers the impact of the past US industrialization, whilst section 8 concludes.

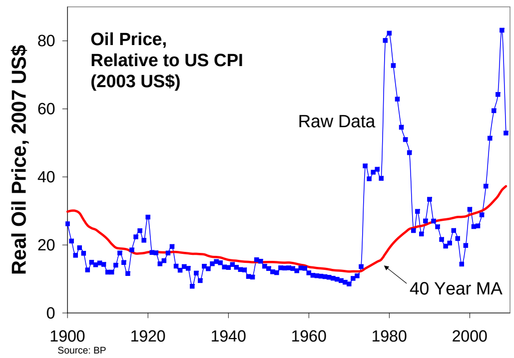

Figure 2: Real Oil Price and its Trend (Relative to the US CPI).

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 and that this has been accompanied by a large structural transformation in China and India and a 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. 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. Finally, I show some facts about oil production and reserves which justify later modeling choices.

2.1 The Oil Price, China and India

Oil Price The focus of the paper is the high level of the long run oil price after 1970. The curve labeled "Raw Data" in Figure 2, shows the 1900-2009 average annual oil price in 2003 US dollars (BP, 2008). The price pre-1970 was declining at a relatively slow and stable rate. The period after 1970 is characterized by a large increase in both the volatility and the long run level of the oil price.[3] Finding a suitable long run trend of the oil price after 1970 is non-trivial exactly for the reason that short and medium run movements in the price have often been dominant. Nonetheless, the average oil price for the 40 year period, 1970-2009, was approximately 37 USD. The average oil price in the 40 years preceding that (1930-1969), was 12 USD. This represents a 200% increase in the long run oil price. Taking a 40 year moving average of the raw data, emphasizes the higher price, post-1970. These long run trends are clearly imperfect and influenced by the oil shocks of the 1970's and the more recent oil price

[3] The standard deviation of the oil price (normalized by its mean) in the 40 years after 1970 was 3.7 times higher than the normalized standard deviation in the 40 year preceding 1970.

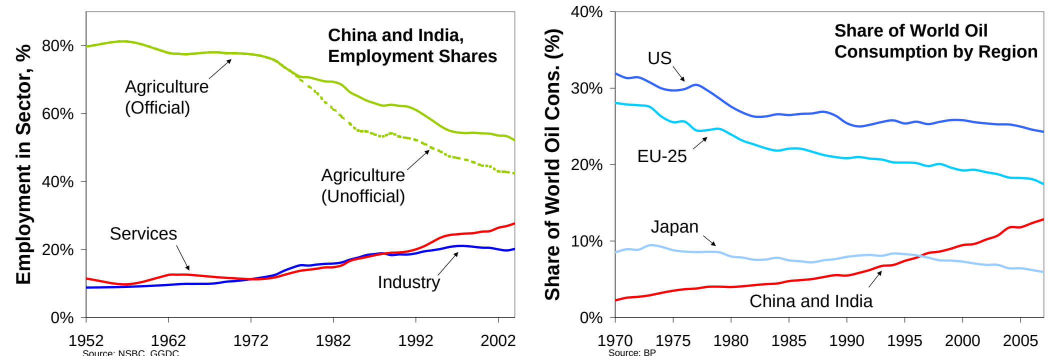

(a) Employment Shares by Sector in China and India (1952-2004)

(b) Changing World Oil Shares (1970-2007)

Figure 3: Structural transformation in China and India and oil consumption by region.

spikes. Notice however, that even in 1998 - the post-1970 oil price trough - the oil price was still more than 68% higher than in 1970. Finally, notice that a similar increase in the price level has been documented by Dvir and Rogoff (2009).

The increase in volatility post 1970 is an important feature of the data and one addressed theoretically by Dvir and Rogoff (2009), who argue that it arises from a combination of persistent demand shocks and uncertainty regarding access to supply. In their model storage of oil can be speculative and can hence increase volatility when demand is subject to persistent growth shocks. This paper differs from theirs in several key aspects. First, it takes a stand on the source of the persistent demand shocks - the growth and industrialization of China and India. Second, it takes the model to the data and quantifies the impact of Chinese/Indian industrialization on the oil price. Finally, the focus of this paper is the higher long run price level post-1970, rather than the higher price volatility. Consequently, this paper can be seen as complimentary to theirs.

China and India's Structural Transformation Figure 3(a) shows how China and India's (government reported) employment share in agriculture has fallen from nearly 80% in 1970 to 50% by 2004. Meanwhile, the share of employment in industry and services has risen from approximately 10%, to 20% in industry and 30% in services according to Timmer and de Vries (2007) and NBSC (2006). Brandt and Zhu (2010) and Rawski and Mead (1998) argue that official Chinese data significantly underestimates the outflow of labor from agriculture. Using household survey data from the Research Centre for the Rural Economy, Brandt and Zhu (2010) construct unofficial estimates of employment in agriculture which show an even larger decline in agricultural employment illustrated by the dashed line in Figure 3(a). Regardless of which

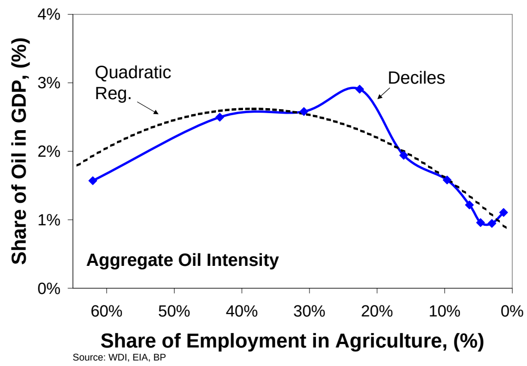

Figure 4: Share of oil in GDP vs. share of employment in agriculture for a panel of the 100 largest countries by population for the years, 1980-2005. Line indicates the decile averages of the data.

data set is used, in absolute terms, this is one of the largest inter-sectoral movement of labor in history. The industrialization process has also had a large impact on oil demand. Figure 3(b) shows how China and India's share in world oil consumption rose by approximately 13 percentage points between 1970-2007. Meanwhile, the share of the EU-25, the US and Japan in world oil consumption fell by approximately 21 percentage points (BP, 2008).

2.2 The Demand Channel of Structural Transformation

Aggregate Oil Intensity Next, I describe the channel through which structural transformation influences demand. 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 GDP are relatively structurally undeveloped whereas countries that have lower agriculture shares are more structurally developed. The index itself is fairly unimportant. I could alternatively consider a country's share of GDP arising from agriculture or its income per capita - any yardstick that is positively correlated with a structural transformation is appropriate and was checked to give similar results.

The share of GDP spent on oil (or 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 4. The smooth, continuous line plots decile averages of aggregate oil intensity versus the share of employment in agriculture for a panel of the 100 largest countries (for the years 1980-2005). The pooled data is first sorted according to employment share in agriculture, then divided into ten groups, and finally the average employment share in agriculture and the average oil intensity of each group is

COEFFICIENT	(1) Agg. Oil Int.	(2) Agg. Oil Int.	(3) Agg. Oil Int.	(4) Agg. Oil Int.
agrShare	0.0951***	0.0971***	0.0515***	0.05492***
	(0.0046)	(0.0044)	(0.0051)	(0.0049)
agrShareSq	-0.1235***	-0.1266***	-0.0921***	-0.0852***
	(0.0068)	(0.0066)	(0.0078)	(0.0073)
Time FE	no	yes	no	yes
Country FE	no	no	yes	yes
Observations	1273	1273	1273	1273
R^2	0.261	0.331	0.857	0.878
Argmax	0.385	0.383	0.280	0.322

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 1980-2005. Column (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.

calculated and shown in the above graph. For details of data construction, see Appendix 9.1. The inverted-U shape of aggregate oil intensity is clearly visible. The peak oil intensity occurs between 21% and 31% employment share in agriculture.

To test the robustness of this result, I run four additional quadratic regressions. The first, relates aggregate oil intensity to the share of employment in agriculture and it's square. The results are shown in column (1) of Table 1 and plotted as the dashed line in Figure 4. The remaining regressions control for time and country fixed effects and the results are shown in columns (2)-(4) of Table 1. In each of the four regressions, the coefficients are highly significant and an inverted-U shape exists as can be seen from the negative sign of the squared term.

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 changes over a structural transformation.

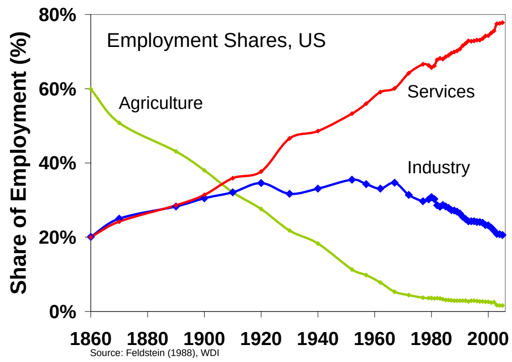

Figure 5: Employment share by sector in the US, 1860-2004.

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.[4] Figure 5 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 a systematic change in the oil intensity of individual sectors (agriculture, industry and services) is another feature of structural transformation. 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 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

[4] Here, and in the rest of the paper unless noted otherwise, I divide sectors according to the standard ISIC III classification. Agriculture is defined to correspond 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).

	(1)	(2)	(3)
COEF.	Oil Int. Agr.	Oil Int. Ind.	Oil Int. Ser.
agrSh.	-0.0492***	0.0599***	0.0750***
	(0.0176)	(0.0205)	(0.0092)
Constant	0.0519***	0.0411***	0.0154***
	(0.0044)	(0.0051)	(0.0022)
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.

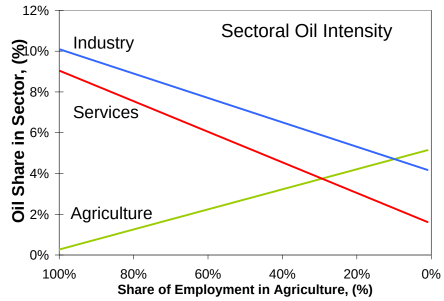

(b) Regression lines. Extended from 0%-100% employment for illustrative purposes.

Figure 6: Changes in sectoral oil intensity over structural transformation.

for 1970-2000. See Appendix 9.2 for construction details and summary statistics. Figure 6 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 6(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 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.

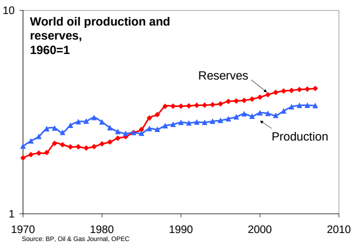

(a) Production and Reserves, World.

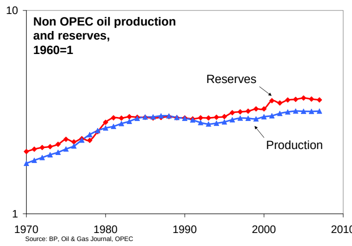

(b) Production and Reserves, Non-OPEC.

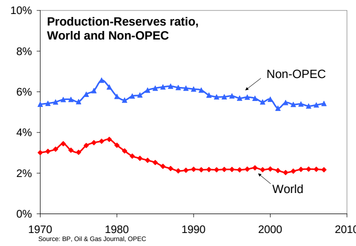

(c) Production-Reserve Ratio, World & Non-OPEC.

Figure 7: Oil Production and Reserves Indices in World and non-OPEC Countries.

2.4 Oil Production

The primary focus of this paper is the impact of changing demand on the oil price. As such, I largely abstract from many supply-side issues such as uncertainty, imperfect competition and storage.[5] Furthermore, similarly to Dvir and Rogoff (2009), 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 that the exhaustibility channel seems to play only a very limited role in driving the oil price. Figure 7 highlights this by showing that despite enormous oil extraction over the last 40 years, oil reserves in both OPEC and non-OPEC countries have steadily increased, so that the ratio of world production to reserves has remained roughly constant.[6] On top of this, measured oil reserves only include discovered, technologically and economically feasible, conventional oil. The USGS estimated in 1994 that there are an additional 3.5 trillion barrels of currently unfeasible conventional oil and a further 1.7 trillion barrels of undiscovered oil. There are also vast quantities of discovered *unconventional* oil such as the 3.5-4 trillion barrels from tar-sand in Venezuela and Canada. All together this “non-reserve” oil is nearly seven times the size of measured world oil reserves and 220 times the world oil consumption in 2007 (BP, 2008). The barrier thus, seems not to be the scarcity of oil, but rather technological and geological constraints which make it impossible or prohibitively expensive to extract much of the oil that is in place. Improvements in technology can thus allow new reserves to be found or existing discoveries to be feasibly extracted.

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

[6] The decline of this ratio in the 1980s was driven entirely by largely fictitious reserve revaluations in Kuwait, Saudi Arabia and the UAE in response to OPEC quota revisions (OECD, 2005).

3 The Model

The model is constructed to capture a shift of labor across sectors and changing sector specific oil intensities which 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 China/India on the oil price.

3.1 The Economic Environment

There are three countries - China/India (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; and (3) the size of their labor force, L^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:

$$\max_{\{A_j^i, I_j^i, S_j^i\}_{j=C,D}} \left(\alpha_A (c_A[A_C^i, A_D^i] - \bar{A})^{\frac{\rho-1}{\rho}} + \alpha_I c_I[I_C^i, I_D^i]^{\frac{\rho-1}{\rho}} + \alpha_S c_S[S_C^i, S_D^i]^{\frac{\rho-1}{\rho}} \right)^{\frac{\rho}{\rho-1}} \tag{3}$$

$$\text{s.t.} \sum_{j=C,D} \left(p_A^j A_j^i + p_I^j I_j^i + p_S^j S_j^i \right) = \begin{cases} w^i & \text{if } i = C, D \\ \frac{p^O O^i}{\bar{L}^i} & \text{if } i = O \end{cases}.$$

In the above, α_s is the utility weight on sector $s = A, I, S$ and $\sum_s \alpha_s = 1$, $\bar{A}$ is a subsistence level in agriculture, 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 population. 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_j^i, at price p_s^j. Goods from the same sector but different countries are then bundled together to produce final sectoral goods using the Armington aggregator, $c_s[C, D] = \left(\nu_s^i C^{\frac{\gamma-1}{\gamma}} + (1 - \nu_s^i) D^{\frac{\gamma-1}{\gamma}} \right)^{\frac{\gamma}{\gamma-1}}$, where ν_s^i 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, $\nu_s^i = 1 - \nu_s^{i'}$ and consumers in $i = O$ value consumption goods from C and D equally, $\nu_s^O = 0.5$.

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:

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

where, $F_s[L, O, t] = (\eta_s(g^t_E O)^{\frac{\sigma_s - 1}{\sigma_s}} + (1 - \eta_s)L^{\frac{\sigma_s-1}{\sigma_s}})^{\frac{\sigma_s}{\sigma_s-1}}$. Countries differ only 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, vary across sectors but remain constant across countries. Finally, I assume there exists energy specific technical progress, g_E, which is the same across countries and sectors.

Oil production I assume a simple reduced from 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, whilst $\widetilde{p}^O_t$ is the price of oil relative to a Paasche consumption index of the oil producer, $\widetilde{p}^O_t$.[7] Over time oil output can increase due to growth in productivity but it is also responsive to changes in the oil price where ε is the price elasticity of oil supply. In Appendix 9.5 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 $s = A, I, S$:

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

Competitive Equilibrium For every t, a competitive equilibrium is: (1) A set of consumption good prices $\{p^i_{s,t}\}_{s=A,I,S}$ and wages $\{w^i_t\}$ for $i = C, D$ as well as oil prices $\{p_{O,t}\}$; (2) household allocations $\{s^i_t\}_{s=A,I,S}$, for $i = C, D, O$; and (3) firm allocations $\{L^i_{s,t}\}_{s=A,I,S}$ and $\{O^i_{s,t}\}_{s=A,I,S}$ for $i = C, D$, such that: (a) Given prices, (1), households' allocations, (2), solve the households problem (3); (b) Given prices, (1), firms' allocations, (3), solve the firms problem in Equation (4) ; and (c) good, labor and oil markets clear. Standard arguments ensure that an equilibrium exists and is unique.

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

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

(Kongsamut et al., 2001). I assume there exists a subsistence level of agricultural consumption, $\bar{A}$. At low levels of TFP, a higher proportion of the labor force must be devoted to agriculture to produce the required output. As TFP in agricultural grows, the subsistence output can be produced with fewer workers, releasing them to other sectors. 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. When calibrating the model, the elasticity is chosen to match the observed flow of labor towards services, given observed TFP growth rates.

Sectoral Intensities Changing sector specific oil intensity is captured by a non-unitary oil-labor elasticity of substitution in production. To see this, 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 respectively. 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 (total) oil labor ratio of China/India and the OECD over the 1970-2003 period decreased by less than 7%, or 0.2% 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.

4 Calibration of the Model

The model is calibrated to match: 1) the levels of sectoral employment, total factor productivity and oil consumption in China/India and the OECD in 1970 as well as 2) the growth rates of sectoral productivity, labor force, world oil output and the oil price between 1970 and 2003.

Parameter	Values China	OECD	Oil Prod.	Target
$B^i_{A,1970}$	0.12	5.11	–	Initial Prod. in A
$B^i_{I,1970}$	0.58	17.67	–	Initial Prod. in I
$B^i_{S,1970}$	0.50	20.70	–	Initial Prod. in S
g^i_A	1.027	1.029	–	Prod. growth in A
g^i_I	1.045	1.020	–	Prod. growth in I
g^i_S	1.034	1.010	–	Prod. growth in S
g_N	1.018	1.018	1.018	C+D Labor force growth, 1970-2003
$\bar{L}^i$	2.35	1	0.05	Size of Labor force, 1970
g_E	1.013	1.013	-	Change in Real Oil Price, 1970-2003
g_O	–	–	1.002	World oil output growth, 1970-2003
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.

Productivity and Labor Force Parameters Ideally, I would obtain total factor productivity by finding the residual, $B^i_{s,t} = Y^i_{s,t} / \left(\eta_s (g^E O^i_{s,t})^{\frac{\sigma_s - 1}{\sigma_s}} + (1 - \eta_s) L^i_{s,t}{}^{\frac{\sigma_s - 1}{\sigma_s}} \right)^{\frac{\sigma_s}{\sigma_s - 1}}$, from the data - 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 and only for two years for China and India - 1995 and 2000. Consequently, I choose $B^i_{s,t}$ in the model to match *labor* productivity data (see Appendix 9.6). Calculating TFP in this way I can find the annualized growth rate of the productivity for all countries and for all sectors. I normalize the labor force in D to 1, and set the size of the labor force in C to match the size of the labor force in China and India relative to the OECD in 1970. The labor force in O is set to be small, at 5% of the D labor force. I assume that labor force growth, g_N, is equal across all countries and set it to the growth rate of combined C and D labor force between 1970-2003. The calibration results are given in the first part of Table 2.

Oil Production and Oil Efficiency In the literature, 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. In the baseline I set the price elasticity of oil supply to be 0.33 - the mid-point between the highest and lowest estimate in the literature. Appendix 9.10 tests the robustness of this choice on the results. Next, I normalize B_O to 1 and then choose g_O so that the model matches world oil output growth over the 1970-2003 period of 1.3%. Finally, I choose $g_E = 1.013$ to match the increase in the real oil price (relative to the OECD 2003 consumer price index) observed in the data over the 1970-2003 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 Chinese/Indian industrialization towards the increase. The

Parameter	Values	Target
$\bar{A}$	0.09	Empl. in Agr. in C
α_A	0.0001	Empl. in Agr. in D
α_I	0.05	Empl. in Ind. in C
ρ	0.28	Empl. in Ind. in D
σ_A	2.52	Oil cons. in Agr. in D, 1970
σ_I	0.72	Griffin and Gregory (1976)
σ_S	0.69	Oil cons. in Agr. in D, 1970
η_A	0.013	Oil cons. in Agr. in C, 1970
η_I	0.026	Oil cons. in Ind. in C, 1970
η_S	0.013	Oil cons. in Ser. in C, 1970
ν_A^i	0.90	Trade share in Agr., $i = C, D$
ν_I^i	0.69	Trade share in Ind., $i = C, D$
ν_S^i	0.94	Trade share in Ser., $i = C, D$
γ	1.59	Change in share in world oil consumption

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

estimate of g_E found in this way is entirely plausible and is identical to that found by Hassler et al. (2011). The calibration results are given in the second part of Table 2.

Structural Transformation Parameters Parameters $\bar{A}$, α_A, α_I and ρ are chosen to match China/India's and the OECD's (official) employment distribution across agriculture and industry in 1970. The first two parameters influence employment levels in agriculture. A high subsistence level in agriculture, $\bar{A}$, means that China/India - with their relatively low TFP in agriculture - must devote a large share of their labor force to agriculture. For the OECD, where the TFP in agriculture is significantly higher, this parameter plays a smaller role and instead, employment in agriculture is primarily determined by the utility parameter, α_A - the more consumers enjoy agriculture, the higher the employment in agriculture in the OECD. The utility weight on industry, α_I, plays a similar role in influencing industrial employment in China/India. Finally, the employment in industry in the OECD is determined by the parameter ρ, the elasticity of substitution between agriculture, industry and services. Given the OECD's higher TFP in industry and services in 1970, consumers in the OECD will consume more industry and more service goods than consumers in China/India. Exactly *how* much more of each type of good is consumed in the OECD, is influenced by the elasticity of substitution between goods in the utility. This, in turn, influences the quantity of workers employed in industry. The calibration results are given in the first part of Table 3.

Oil Intensity Parameters Oil parameters, σ_s and η_s are chosen to match the oil consumption of every sector in China/India and the OECD in 1970. From a firm's first order conditions, the sectoral oil-labor ratio in each country and at each point in time is given by:

$$\frac{O^i_{s,t}}{L^i_{s,t}} = \left(\frac{\eta_s}{1-\eta_s}\right)^{\sigma_s} \left(\frac{w^i_t}{p^O_t}\right)^{\sigma_s} \left(g_E^{\sigma_s - 1}\right)^t. \tag{7}$$

Since the oil price in both countries is the same, the oil labor ratios at each point in time, depend only on the ratio of wages between the two countries. At every point in time, a country with a higher value added per-capita will use more oil per worker in every sector (recall that in the model, value added is equal to wage income). The sector specific elasticity of substitution between oil and labor, σ_s, determines *how* much more oil per worker richer countries use. The higher the elasticity, the higher the oil-labor ratio used for a given wage-oil price ratio, in a given sector. Since sectoral employment in each country is pinned down by the structural transformation parameters, by choosing oil-labor elasticity and share parameters, I can set the oil consumption of each sector, in each country in 1970.

The calibration proceeds as follows. Since total world oil supply in the first period is exogenous,[8] only five of the six oil parameters are needed to match oil consumption in each sector at a point 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 9.7 for details). The results from the third step of the calibration are given in the second part of Table 3. The calibration implies that industry has the highest oil use ($\eta_I = 0.026$), whilst services ($\eta_S = 0.013$) 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 $\sigma_A = 2.51$ and (gross) complements in industry and services with elasticities $\sigma_I = 0.72$ and $\sigma_S = 0.69$.

In the above calibration, I effectively use variation in cross-country wage to oil price ratios to calibrate the model, rather than variation across time in a single country. This approach is chosen for two reasons. First, there is only limited variation in wage to oil price ratios in time series data. In 1970, the wage-oil price ratio was approximately 47 times higher in the US than in China/India (WDI). On the other hand in US data between 1970-2003, the highest wage-oil price ratio was only 9 time higher than the lowest.[9] Furthermore, much of this variation came from the relatively short lasting but very sharp oil shocks of the 1970's, which may not be relevant for the long term substitutions examined in this paper. In particular, using smoothed oil prices (the 30 year moving average) in the US over the same period, the highest relative wage to oil price ratio was only 2 time higher than the lowest.

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

[9] More formally: $\frac{w^{US}_{1970}/p^O_{1970}}{w^C_{1970}/p^O_{1970}} = 47$ whilst $\frac{\max_t w^{US}_t/p^O_t}{\min_t w^{US}_t/p^O_t} = 9$.

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

Finally, I perform two exercises to check the robustness of these results. First, in Appendix 9.8, I re-estimate these elasticities using equation (7) directly. I use 1995 cross-sectional data and estimate the elasticity of substitution between oil and labor to be 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 are quite close to these empirically estimated parameters.

Second, I compare my results to 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) estimates this elasticity for the US to be 0.65. Griffin and Gregory (1976) estimates 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.

Trade Parameters The home bias parameters in $i = C, D$ are chosen to match trade flows between these countries. The home bias parameters in O are set to 0.5 - the oil producer shows no preference for C's or D's good. Finally, I choose the elasticity of substitution between C's and D's goods to match the change in the share in world oil consumption of China/India and the OECD over the 1970-2003 period. The results from the third step of the calibration are given in the third part of Table 3.

5 Baseline Simulation

This section examines the qualitative results from the baseline simulation. A direct comparison between the data and the model is made in Appendix 9.9. Since the model is run from 1970 to 2100, I assume that future sectoral productivity growth rates in both regions, and the future effective productivity growth rates in the oil industry stay at their 1970-2003 levels.

The top two panels of Figure 8 show the changing shares of employment over time. In China/India, employment share falls in agriculture, forms an inverted-U in industry and rises in 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 and it's share of employment in industry falls. The middle two panels of Figure 8 show sector specific oil intensities. Oil intensity is rising in agriculture, but falling in industry and services. This development is analogous to the development of oil intensities in the data shown in Figure 6(b). The bottom two panels of Figure 8 show how changing sectoral intensities and sectoral size translate into changes in aggregate oil intensity. As the economy shifts away from agriculture towards industry and services, aggregate oil intensity rises, since sectoral oil intensity in industry and services is significantly higher than in agriculture and oil intensity in the largest sector - agriculture - is rising. Later, as the economy shifts from industry towards services, aggregate oil intensity falls since oil intensity in services is lower than in industry and oil intensity is falling in the largest two sectors - industry and services. Since China/India is in the early stages of transition, aggregate oil intensity first rises and then falls forming an inverted-U. Since the OECD is in the final stages of its structural transformation, we only see the second part of the inverted-U - a falling aggregate oil intensity. Below, I show that the key to the inverted-U result are the different sectoral oil-labor elasticities across sectors which cause oil intensity to rise in agriculture, and fall in non-agriculture.

Figure 9(a) shows the shares in world oil consumption of regions. Faster growth and rising aggregate oil intensity result in China/India's share in world oil consumption increasing and the

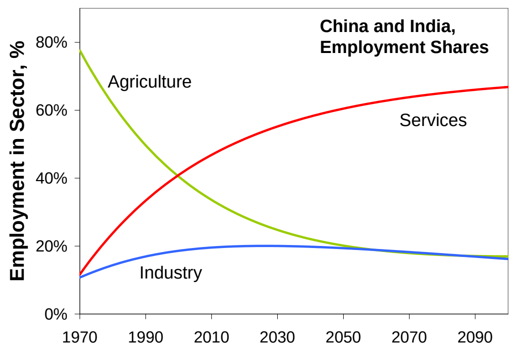

(a) China/India, Employment Shares

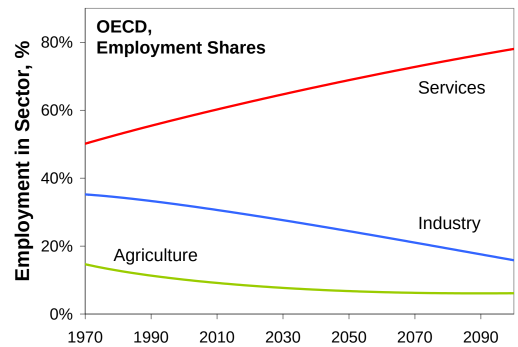

(b) OECD, Employment Shares

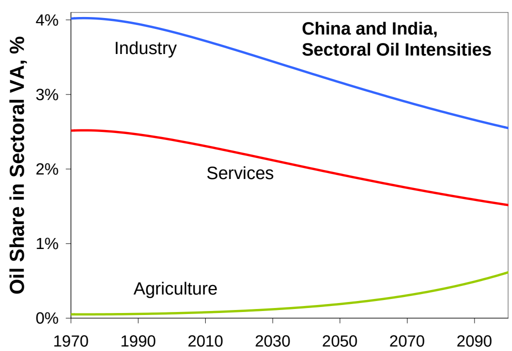

(c) China/India, Sectoral Oil Intensities

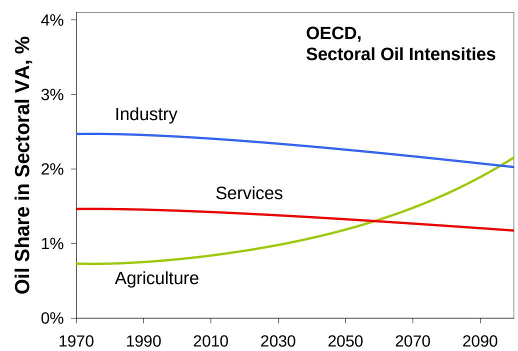

(d) OECD, Sectoral Oil Intensities

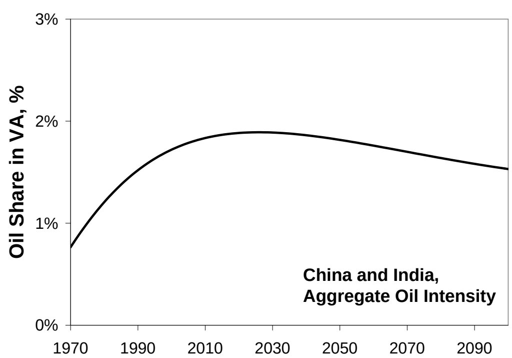

(e) China/India, Aggregate Oil Intensities

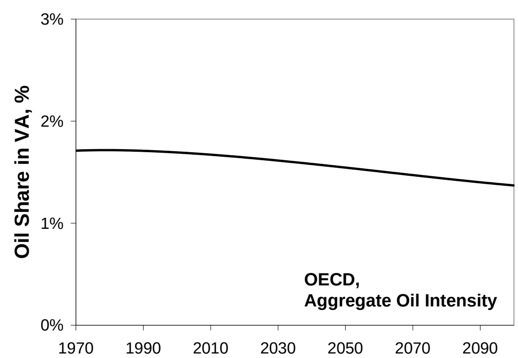

(f) OECD, Aggregate Oil Intensities

Figure 8: Simulation result for employment shares, sectoral oil intensities and aggregate oil intensities.

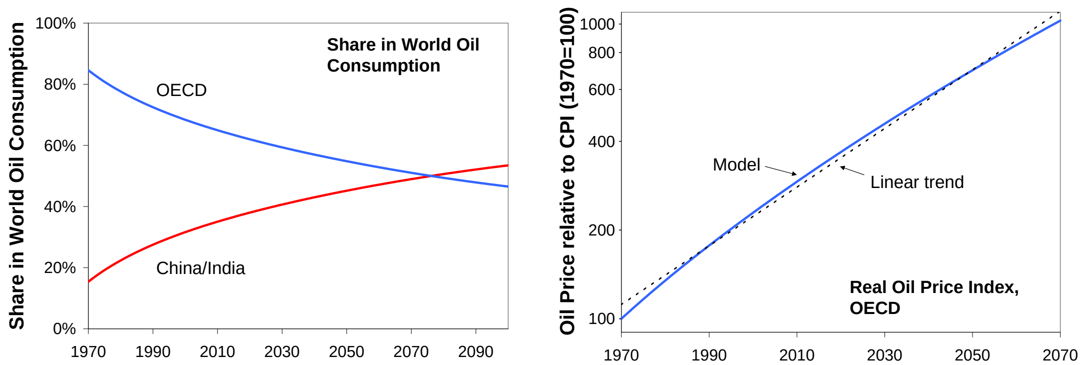

(a) Oil Consumption Shares (b) OECD Oil Price

Figure 9: Share of oil consumption by region and OECD oil price.

OECD's share falling. Finally, Figure 9(b) shows the rising oil price (relative to the OECD's 2003, fixed basket consumer price index) that arises from the model. The price is shown on a log scale, so that slopes can be interpreted as growth rates. Also included is a linear regression line which captures the average growth rate of the oil price over the period. Notice that the rate of increase of the oil price is not constant. The oil price first increases faster than average and then slower than average resulting in an oil price that is hump shaped relative to trend. I argue in the next section that this non-linearity is caused by the hump shape in aggregate oil intensity in China/India. As China/India spends first a higher then a lower fraction of its income on oil, the oil price rises and then falls (relative to trend) in the OECD.

Sources of the hump shape As was argued in section 2.3, the inverted-U aggregate oil intensity 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 that of non-agriculture is falling. When the economy is dominated by agriculture, this contributes to rising aggregate oil intensity. When the economy becomes dominated by non-agriculture, this contributes to 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). Which of these two forces are key to generating the hump-shaped aggregate intensity?

To answer this, I perform three counterfactuals. First, I replace oil use parameters across

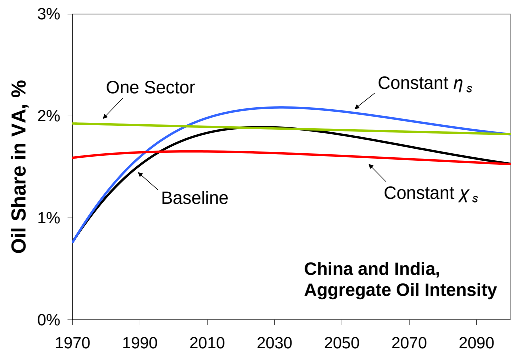

Figure 10: Sources of the hump shaped oil intensity. Constant η_s, χ_s and a one sector version of the model.

sectors by their weighted averages, $\eta_s = 0.017$. Second, I replace sectoral oil-labor elasticity parameters with their weighted averages, $\sigma_s = 0.97$.[10] Finally, I take a one sector version of the model, where *all* sector specific parameters are replaced by their weighted averages, so that structural transformation is switched off in the model (the details of this procedure are discussed in the next section). Figure 10, shows the implication of these counterfactuals for aggregate oil intensity. Without either of the first two channels, we still observe a hump shaped aggregate oil intensity, although it is evident that the major driver of the inverted-U are the different elasticities of substitution across sectors rather than sectoral variation in oil use. In a one sector version of the model with no structural transformation however, oil intensity declines monotonically and (log) linearly.

6 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 China and India's growth and structural transformation 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 China and India in all sectors and compare the resulting price to that obtained in the baseline model. The result is shown in Figure 11(a), and labeled 'No growth'. This experiment allows me to measure the total effect that growth and structural transformation in China and India have had on the oil

[10] The weights are taken to be OECD period zero employment shares.

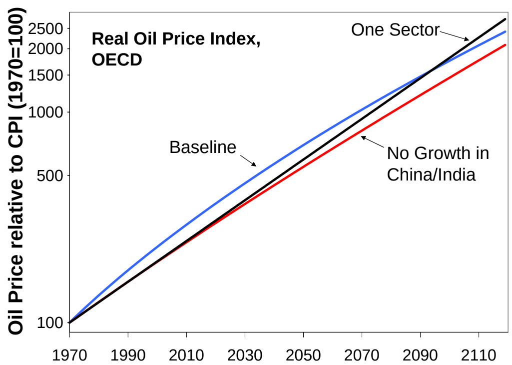

(a) Real Oil Price Index, OECD

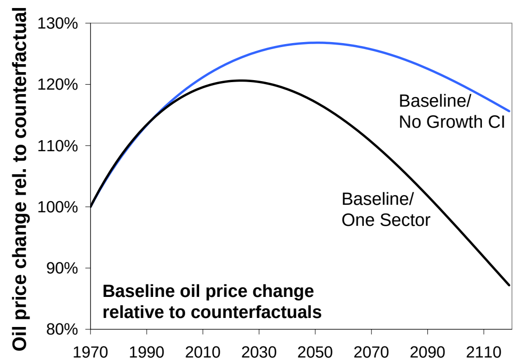

(b) Baseline oil price change relative to counterfactuals

Figure 11: Counterfactual experiments, prices.

price. If China and India 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 11(b), the China/India effect results in oil prices in the model that are nearly 27% higher at their peak in 2051 than they would be without the China/India effect. Furthermore, according to the model, the long run oil price in the OECD is 21% (or roughly 6.57 USD) higher in 2010, than if China and India had not structurally transformed at the speed that they did. Since the long run oil price increased from from approximately 12 to 37 USD between 1970 and 2010, growth and structural transformation in China and India can account for 26% of the observed increase in the oil price.

The Role of Structural Transformation In the second counterfactual, I replace all sector specific parameters with weighted sectoral averages. Since sector specific parameters are the same across sectors, the model collapses to a standard, open economy, one sector growth model. Country specific parameters are weighed by each country's period zero employment shares so that $g_s^C = 1.030$, $B_{s,1970}^C = 0.22$, $g_s^D = 1.016$ and $B_{s,1970}^D = 17.35$. Parameters that are common across countries but differ across sectors are weighed by OECD period zero employment shares so that $\sigma_s = 0.97$, $\eta_s = 0.017$ and $\nu_s = 0.85$. Finally, I recalibrate $g_O = 1.005$ to maintain observed world oil output growth of 1.3% a year over the 1970-2003 period. In this manner, I turn off the impact of structural transformation in China/India but keep growth effects. The resulting price index, is shown in Figure 11(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 11(b), which shows the change in oil prices in the baseline relative to the one sector world. Notice, that prices increased 21% more by 2024 in the baseline than in the

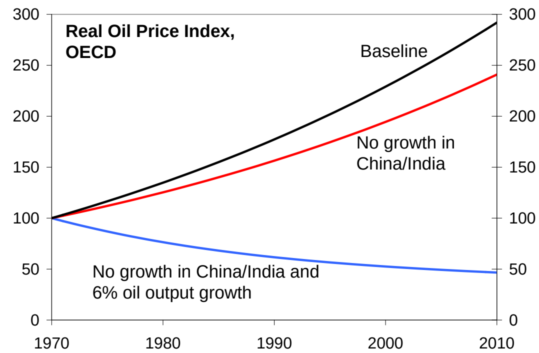

Figure 12: The Role Oil Output Growth

one sector model.[11] 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. Finally, notice that until approximately 2010, the increase in prices in the baseline relative to the one sector world and the baseline relative to the zero growth world, are essentially the same. This indicates, that practically all the China/India driven increase in the price of oil observed thus far, has come from structural effects rather than from growth effects. 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.

7 Extensions and Generalization

The Role Oil Output Growth If Chinese/Indian growth and industrialization explain only a quarter of the 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 oil output fell from approximately 6% between 1900 and 1970 to 1.3% 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 China/India did not grow over the 1970-2010 period. In addition however, I choose $g_O = 1.064$ so that world oil output in the model grew at 6%. The result are shown in Figure 12. Without Chinese/Indian industrialization and with the higher oil output growth, the price of oil over the period would

[11] 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.

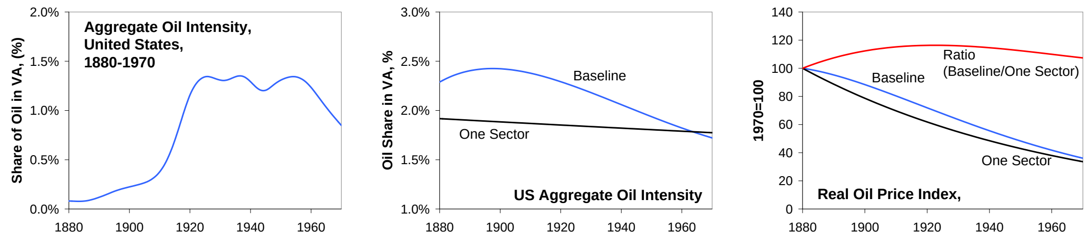

(a) Agg. Oil Int., USA (b) Agg. Oil Int., Model (c) Oil Price in OECD

Figure 13: US Oil intensity in model and data and it's impact on the oil price, 1880-1970.

have fallen by approximately 53%. The decline in oil output growth was key in generating a higher oil price post-1970. Nonetheless, industrialization in China and India still played a significant role: relative to the counterfactual, 20% of the higher price in 2010 can be explained by Chinese/Indian industrialization.

Not all humps created equal Figure 13(a) shows that a hump shaped aggregate oil intensity also existed for US time series data between 1880 and 1970.[12] Did structural transformation play a role in generating this hump and what impact did it have on the oil price? To answer this question, I extend the model back to 1880. I assume that China/India did not grow over the period and set $g_O = 1.049$ to match US oil consumption growth during the 1900-1970 period of approximately 4.5% a year (EIA). I also set $g_E = 1$ to match the change in the long run US oil price over this period. All other parameters stay the same. Figure 13(b) shows the results for aggregate oil intensity in the US. Qualitatively, structural transformation captures the hump shape in the US and implies a peak in the early 1900's - somewhat earlier than what was observed. Quantitatively, the model does not capture the enormous initial increase in intensity. The above calibration is, of course, very rudimentary as the original model is fitted to OECD data and not US data. More importantly however, the model does not take into account the enormous change in fossil fuel mix that took place between 1880 and 1960. Oil's share in fossil fuel energy in the US increased from 2% to 47% with a shift from coal to oil energy. This may explain a large portion of the initial intensity increase. Importantly, changing fuel mix is not a problem for more contemporary data - at least for the period under consideration. After 1960 the share of petroleum in fossil fuel energy has remained roughly constant in both the US and in China/India (WDI, BP). Whilst future shifts in fuel mix could contribute to oil intensity falling at a faster rate, the primary purpose of the model is to analyze the period between 1970 and the present day when oil's importance remained relatively fixed.

[12] For data construction details see Appendix 9.4.

Finally, the impact of US industrialization on the oil price can be seen in Figure 13(c). Given the higher oil output growth rate, the oil price fell over the period. The contribution of the hump shape on oil intensity however, is still the same. A one sector version of the model, calibrated as before, predicts a larger drop in the price of oil. The ratio of oil price changes between the baseline and the one sector model is again hump shaped. This highlights the robustness of the previous price result. Although the observed path of the oil price may be driven by many factors, this does not diminish or change the role of structural transformation.

8 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 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 these 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 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 - China and India's impact on the oil price. In particular, the model developed here predicts that as long as the structural transformation in China/India and other developing nations follows past patterns, the upward pressure on oil prices from China and India's industrialization will continue for many decades but 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 Chinese/Indian industrialization as economies become dominated by services. As such, China and India's impact on the oil price is not necessarily permanent.

9 Appendix

9.1 Aggregate Oil Intensity

Variable	Obs	Mean	Std. Dev.	Min	p5	p50	p95	Max
Agr. Emp. Share	1273	0.201	0.196	0.001	0.013	0.125	0.606	0.915
Agg. Oil Int.	1273	0.035	0.027	0.009	0.028	0.041	0.083	0.261
Agg. Oil Int. (smoothed)	1273	0.022	0.016	0.004	0.007	0.017	0.055	0.115

(a) Summary Data

	(1)	(2)	(3)	(4)
COEFFICIENT	Agg. Oil Int.	Agg. Oil Int.	Agg. Oil Int.	Agg. Oil Int.
agrShare	0.1901***	0.1819***	0.1445***	0.1008***
	(0.0104)	(0.0091)	(0.0189)	(0.0128)
agrShareSq	-0.2472***	-0.2349***	-0.1499***	-0.1483***
	(0.0154)	(0.0135)	(0.0286)	(0.0193)
Time FE	no	yes	no	yes
Country FE	no	no	yes	yes
Observations	1273	1273	1273	1273
R^2	0.217	0.419	0.596	0.826
Argmax	0.384	0.387	0.482	0.340

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

(b) Regressions

Table 4: Part (a) shows summary statistics for aggregate oil intensity and agricultural employment share data. Part (b) is an analog of Table 4(b) in the text, with the oil price data left un-smoothed.

I construct a panel data set that consists of the worlds largest 100 countries (by population in the year 2000) over the 1980-2005 period. 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 the WDI. 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 the WDI. Total country-specific oil consumption (in Quadrillions of BTU) and current oil price (dollars per million BTU) data come from the EIA. Summary statistics for the data are shown in Table 4(a). The oil price used is the 30-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 1980's. The shape of the curve however, remains unchanged if observed oil price data is used. Table 4(b), replicates the regressions found in Table 1 using un-smoothed data. The data still exhibit a strong, inverted-U shape.

9.2 Sectoral Oil Intensity Data

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

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[13] 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, 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.

9.3 Productivity Data

Productivity measures are constructed using data from the UN. I consider two regions: 1) China and India and 2) the OECD (which, due to data availability, I take to be Australia, Austria, Belgium, Canada, Denmark, Finland, France, Greece, Italy, Japan, Netherlands, New Zealand, Norway, Portugal, Spain, Sweden, Turkey, UK, US). First, for each region I construct aggregate value added measured in constant price (2003) US dollars. Since UN data is given in constant 1990 USD prices and current prices, I use the current price data to re-base the 1990 data to 2003 prices for each country. Summing over countries within each region, I obtain aggregate value added measured in constant price (2003) US dollars for each region, $V = \sum_j V^j$. Next, I construct constant price sectoral value added data for each region. To do this, I construct current price sectoral shares of total value added, v_s^j, for each country and sector, s, using the series of value-added by sector from the UN. I multiply these shares by constant price aggregate value added in each country calculated above, to obtain $V_s^j = v_s^j V^j$. I then sum over all countries in a region to obtain a region's sectoral value added, $V_s = \sum_j V_s^j$. Finally, to calculate

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

sectoral productivity, I need to construct sectoral employment data. To maintain consistency across countries, I extract a country's aggregate labor force from the Penn World Tables, L^j. I then calculate sectoral employment shares, e^j_s, for individual countries from SourceOECD for the OECD, NBSC for China and Timmer and de Vries (2007) for India. I multiply a country's sectoral employment share by the country's labor force to obtain sectoral employment, $L^j_s = L^j e^j_s$. Summing over countries within a region, I obtain a region's sectoral labor force, $L_s = \sum_j L^j_s$. Labor productivity in a sector of a particular region is then, $\bar{D}_s = V_s/L_s$.

9.4 US Long Run Data

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.

9.5 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:

$$R_{t+1} = g_R \left(\frac{\widetilde{p}^O_t}{\widetilde{p}^O_{t-1}} \right)^{\varepsilon} R_t - O_t \text{ and } O_t = \eta \left(\frac{\widetilde{p}^O_t}{\widetilde{p}^O_{t-1}} \right)^{\varepsilon} R_t. \tag{8}$$

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 how much oil is extracted this period, O_t, on exogenous growth in the ability to locate, extract or process reserves captured the growth factor, g_R, and on changes in the price of oil relative to a Paasche consumption index of the oil producer, $\widetilde{p}^O_t$. 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, $\widetilde{p}^O_t$. 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 $O_t = g^t_O (\widetilde{p}^O_t)^{\varepsilon} B_O$ where, $B_O \equiv \eta (\widetilde{p}^O_{-1})^{-\varepsilon} R_0$ and $g_O \equiv g_R - \eta$. Notice that throughout I assume that $\eta \left(\frac{\widetilde{p}^O_t}{\widetilde{p}^O_{t-1}} \right)^{\varepsilon} < 1$ and that oil producers take oil prices as given.

9.6 Total Factor Productivity Calibration

Labor productivity in the data 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 in Appendix 9.3. In the model, labor productivity is defined as:

$$D^i_{s,t}(B^i_{s,t}) \equiv \frac{p^i_{s,2003} B^i_{s,t} \left(\eta_s (g_E O^i_{s,t})^{\frac{\sigma_s - 1}{\sigma_s}} + (1-\eta_s) {L^i_{s,t}}^{\frac{\sigma_s - 1}{\sigma_s}} \right)^{\frac{\sigma_s}{\sigma_s - 1}} - p^O_{2003} O^i_{s,t}}{L^i_{s,t}}, \tag{9}$$

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. $D^i_{s,t}(B^i_{s,t}) = \bar{D}^i_{s,t}$.

9.7 Oil Consumption Calibration

In this section, I estimate oil consumption by sector in China/India and the OECD in 1970. In order to do this (and for lack of sectoral oil consumption data in China/India and the OECD), I use the regressions presented in equation 6 (but 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 80% of their work force in agriculture - approximately the share employed by China/India in 1970 - spend less than 1% of their agriculture value-added, 9% of their industry value-added and 7% of their service value added on oil. Since I have data on the value-added of Chinese/Indian 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 China/India 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 a country's total oil consumption, I can then estimate the quantity of oil consumed by each sector in each country and hence its share in total world oil consumption

Since I model the entire world as only the OECD and China/India, I need to make an assumption on how the oil that is not consumed by either the OECD (as defined in the text) or China/India in the data, is allocated across countries C and D in the model. In the data, China/India and the OECD (as defined in the text) consumed only 68% of the world's total oil consumption in 1970. Furthermore, this number did not stay constant over time, but fell to 60% by 2003 BP. This implies that I cannot assume that the growth rate of oil consumption in China/India and the OECD is the same as the growth rate of world oil output without adjusting the data. Consequently, I divide world oil consumption in the data into two groups: 1) The

	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 estimated 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).

OECD and the Former Soviet Union and 2) China/India and other emerging economies.[14] I match the oil consumed by the first group in the data to $D's$ oil consumption in the model and the oil consumption by the second group in the data to C's oil consumption in the model. This adjustment ensures that total oil consumption of C and D in the model grows at the same rate as world oil consumption in the data (by construction).

9.8 Estimating Sectoral Oil-Labor Elasticities

In this section, I estimate 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 (2008). Taking the log of equation (7), I obtain:

$$\log\left(\frac{O^i_{s,t}}{\bar{L}^i_{s,t}}\right) = \sigma_s \log\left(\frac{\eta_s}{1-\eta_s}\right) + \sigma_s \log\left(\frac{w^i_t}{\bar{p}^O_t}\right) + t\log\left(g_N^{\sigma_s - 1}\right),$$

where σ_s. 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.

[14] These are all the other countries that are not in any of the other groups.

9.9 Data and Model

This appendix compares the predictions of the model directly to the data. Figure 14(a) and (b) show the employment shares in the (official) data and those predicted by the model, panels (c) and (d) shows aggregate oil intensity, whilst panel (e) shows the oil shares for both regions[15] and panel (f) shows the model and oil price data (both the 40 year moving average and the raw data). In general, the model matches the trend of employment shares in both regions, although it over-predicts the decline in agricultural employment in China/India and under-predicts it in the OECD. Similarly, the model over-predicts the increase in service employment in China/India and under-predicts the increase in the OECD. The model does relatively well in matching industry employment shares in both regions. The model captures the increase of aggregate oil intensity in China/India and the decline in the OECD - but not the volatility. Since the model was calibrated to match regional oil shares and the oil price increase, I obtain a good fit.

Wedges What are the sources of the mismatch between employment in the data and model? One possibility may be the alleged over-reporting of agricultural employment by China mentioned in the main text. Other data (such as sectoral value added) may also be misreported. Subsidies (in industry) and labor re-allocation restrictions across sectors (e.g. the Hukou system in Chinese agriculture, or the Indian caste system) could play a large role. The model may also be too simple to match the structural transformations across two very different regions.[16] To see the quantitative importance of the mismatch between the data and the model, I introduce time varying wedges in sector specific wages chosen so that the model exactly matches the respective employment data in both China/India and the OECD. These wedges represent a variety of factors, such as the above mentioned subsidies, taxes and restrictions but also any other conditions that promote or prevent labor from reallocating to a particular sector. More specifically, I assume there exists a country and sector specific wedge on wages in agriculture and industry, $\tau^i_{s,t}$, so that the effective wage earned by a worker in a given sector is given by $(1+\tau^i_{s,t})w^i_{s,t}$.[17]
The resulting wedges are shown in Figure 15(a). The model implies that the agriculture and industry sector must have been more attractive to workers in China/India and less attractive in the OECD than a model without wedges would predict. This implies subsidy-like wedges in China/India and tax-like wedges in the OECD on agriculture and industry wages. Importantly, notice from Figure 15(b) and (c) that the models with wedges predict a similar oil intensity and oil price as the baseline model without wedges. The quantitative importance of the mismatch is

[15] These are constructed in Appendix 9.7.
[16] Duarte and Restuccia (2010), for instance, find that in order to match labor reallocation across sectors in the US, they need to assume a non-homotheticity in the service sector. In my case, this would help match employment shares in the OECD, but exacerbate the problem in China/India.
[17] The wedges are assumed to be financed by a lump sum tax/transfer from households.

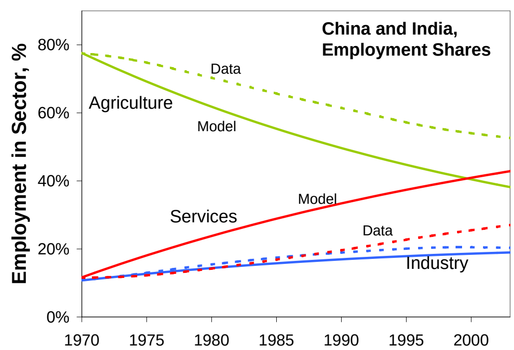

(a) China/India, Employment Shares

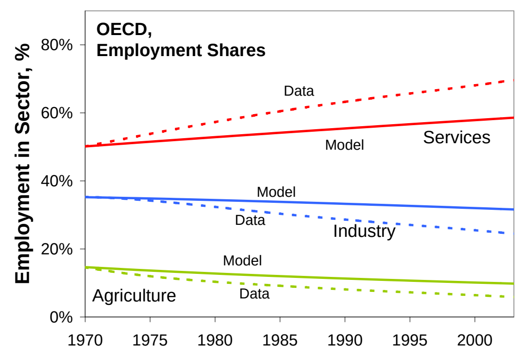

(b) OECD, Employment Shares

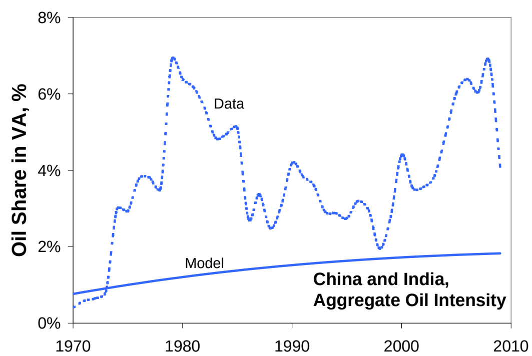

(c) China/India, Aggregate Oil Intensities

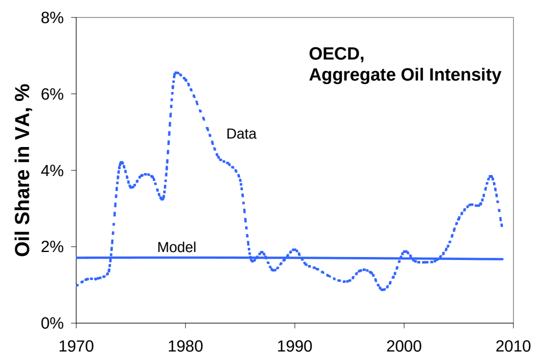

(d) OECD, Aggregate Oil Intensities

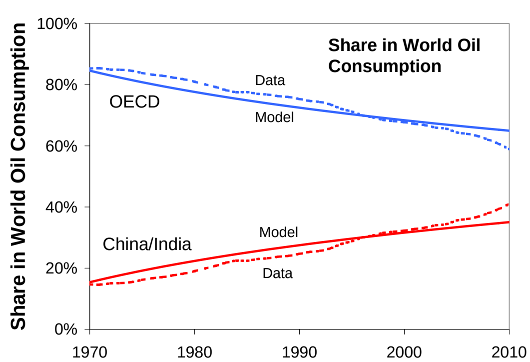

(e) Oil Consumption Shares

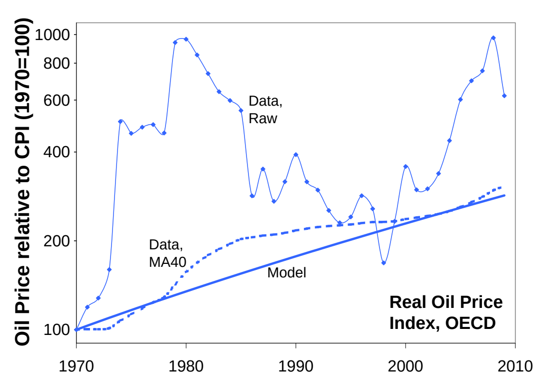

(f) OECD Oil Price

Figure 14: Simulation and Data.

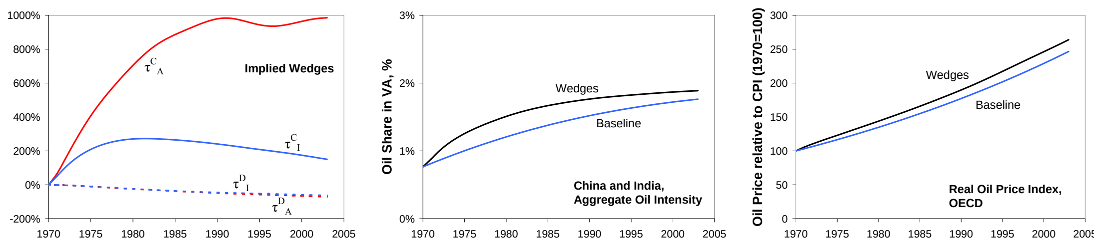

(a) Agg. Oil Int., USA (b) Agg. Oil Int., Model (c) Oil Price in OECD

Figure 15: US Oil intensity in model and data and it's impact on the oil price, 1880-1970.

relatively small. As such, without loss of generality, in the main body of the paper I focus only on the baseline model without wedges.

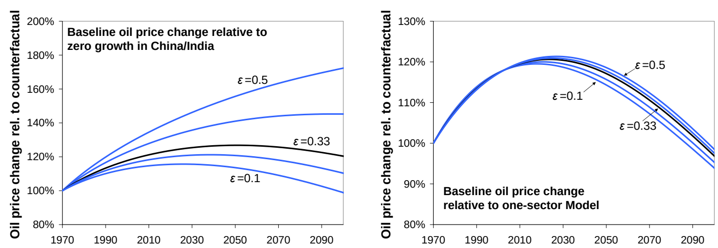

(a) Baseline oil price change relative to the no growth in CI. (b) Baseline oil price change relative to one sector model.

Figure 16: Impact of varying price elasticity of oil supply on OECD price of oil.

9.10 Sensitivity Analysis

This appendix investigates the role played by the price elasticity of substitution of oil supply, ε. I vary the elasticity so that $\varepsilon = 0.1; 0.2; 0.4; 0.5$ and re-calibrate the model to match the oil growth rate of 1.3% per year between 1970 and 2003. First, aggregate oil intensity in China/India is nearly indistinguishable between the scenarios. Next, I switch of productivity growth in China/India and use the model to measure the additional impact of Chinese/Indian growth and industrialization on the oil price in the OECD. The results are shown in Figure 16(a). Interestingly, lower price elasticity of substitution of oil supply implies a higher measured contribution of Chinese/Indian growth on the oil price. Whilst the contribution of growth and industrialization is relatively sensitive to the chosen elasticity, the impact of only industrialization is not. Figure 16(b), shows changes in the oil price in the baseline model relative to the corresponding one sector model. The differences across scenarios are small.

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* 2008.

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.

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, *EIAData.*

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.*

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.

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.*

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

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.

NBSC, "National Bureau of Statistics of China, State Council of the People's Republic of China," 2006.

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

OECD, "World Energy Outlook," *OECD Publishing*, 2005.

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

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

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.

Timmer, Marcel P. and Gaaitzen J. de Vries, "A Cross-Country Database For Sectoral Employment And Productivity In Asia And Latin America, 1950-2005," *Groningen Growth and Development Centre Research Memorandum GD-98*, 2007.

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," 2007.