

# Market Power and Production (Mis)Allocation A Study of the World Oil Market\*

John Asker  
UCLA  
NBER

Allan Collard-Wexler  
Duke University  
NBER

Jan De Loecker  
KU Leuven and Princeton University  
NBER and CEPR

May 16, 2017

## Abstract

This paper provides estimates of the extent to which market power is a source of production misallocation. Productive misallocation, and the resulting welfare loss, occurs through more production being allocated to higher cost units of production, and less production to the most efficient production units in the economy, conditional on aggregate quantity. We examine the role of productive misallocation due to market power in the global market for crude oil, using detailed microdata on over 13,000 oil fields from 1970 to 2014, in which OPEC, a notable international cartel, is alleged to exert considerable market power. We find evidence for substantial productive inefficiency due to market power, ranging from 0.3 to 1.6 trillion USD, depending on model specification.

PRELIMINARY - PLEASE DO NOT CITE - DO NOT POST

**Key words:** *Market Power, Productive Inefficiency, Misallocation, Cartels, Oil*

**JEL Code:** *L13, L42, K21*

---

\*We would like to thank Russell Cooper and participants in seminars at ANU RSE, Harvard, KU Leuven, Northwestern, Princeton, Penn State, Texas A&M, and Yale for their comments. Excellent research assistance was provided for by Ardis Zhong, Sherry Wu and Yilin Jiang. Financial Assistance from Princeton University's Industrial Organization Group, NSF (SPS #220783) and FWO Odysseus Grant is greatly appreciated. Contact information: johnasker@econ.ucla.edu, collardwexler@gmail.com and jan.deloecker@kuleuven.be.

# 1 Introduction

This paper provides estimates of the extent to which market power is a source of production misallocation. This production misallocation, and the resulting welfare loss, occurs through more production being allocated to higher cost (less productive) units of production, and less production to the most efficient production units in the economy.

In quantifying the potential magnitude of this productive inefficiency, we examine the global market for crude oil, in which OPEC, a notable international cartel, is alleged to exert considerable market power at times. This is made possible by leveraging detailed annual on production, reserve and cost data on all (about 13,000) major global production units from 1970 through 2014, covering all relevant geography and technologies used to produce oil.

Figure 1 presents a stylized, graphical, representation of the environment we consider. Consider a producer with market power, with constant marginal cost  $MC_1$ , coupled with a competitive, price-taking, fringe with an aggregated marginal cost curve given by  $MC_f$ . To be precise, marginal cost for the fringe is equal to  $MC_f(q - q_1)$ .

The market price is equal to  $P$  and the quantity produced by the (low cost) producer with market power,  $q_1$ , is less than total production  $Q = q_1 + q_2$ . where  $q_2$  is the production of the fringe. In this setting, the production done by the fringe,  $q_2$ , is done at a higher cost than socially optimal: Indeed, all production should be done by the low cost producer. The welfare cost of this production misallocation is the shaded area.

The objective of this paper to identify and quantify the size of (the analog of) this shaded trapezoid in the global oil market. In the case of the oil market, this calculation is more complex, due to the finite reserve of oil, which creates inter-temporal linkages of supply and cost. We do, however, observe oil reserves by field, and with the help of a flexible and rich dynamic framework, we compute the productive inefficiency.

It is important to note two things. First, this production distortion is distinct from the deadweight loss generated by the lost gains from trade from having quantity distorted below to the social optimum, distance  $Q^{SP} - Q$ . Second, the production distortion here could be the outcome of a cartel agreement or the natural consequence of the unilateral exercise of market power, whether by a single monopolist facing a fringe or via Cournot (quantity) competition or some combination of both. In any case, the exercise of market power is at the root of the production misallocation. As a consequence this paper is intentionally agnostic as to the efficacy of the OPEC cartel in its methodological approach. As long as market power is being exercised in some fashion, whether unilaterally or in concert, the measurement approach employed will be valid. This is a property of leveraging cost data, rather than the more common approach of

combining an estimated demand system with a conduct assumption to recover implied costs.

Differences in the costs of production across firms have received enormous attention in the economics literature. We now know that firms that compete in narrowly defined product markets can have very different levels of *productivity* (see Syverson (2011)). This heterogeneity has proved to be key key in predicting of the impact of competition on many outcomes such as the effect of trade or of a new technology on firm and industry performance.<sup>1</sup>

The allocation of resources has also been extensively studied in the recent macroeconomic literature, which finds evidence for substantial misallocation in plant-level data, based on very different input use at different plants and large differences in plant-level productivity. This literature has highlighted the potential macroeconomic returns from improving the apparent misallocation of resources within an economy (See, for a survey, Hopenhayn (2014)). The ongoing macroeconomic literature on misallocation has considered a variety of distortions that affect the allocation of inputs across plants. However, this misallocation literature typically relies on a combination of a demand system (CES) and market structure (monopolistic competition) which eliminate the effects of market power, and in this paper we focus exclusively on the role of market power, while allowing for other sources of dispersion in cost of production in any given period.

While the presence of market power, and the underlying strategic interactions among producers, is typically assumed away as it is hard to convincingly observe and measure market power directly in the data, we leverage the specifics of the world oil market. First of all, oil is considered a homogeneous product and the presence of differences in observed cost of production is immediately informative about differences in actual production cost, and as such does not reflect differences in product quality. Second, oil is extracted across a wide range of countries, using a variety of technologies, which give rise to natural cost differences. Third, this market has been affected by the presence of a well-known cartel (OPEC) over a long period of time, whose internal organization is well understood. Fourth, we directly observe the cost of production by oil field, and such we do not have to model the extent of market power, or how exactly the cartel is organized, or disorganized, but rather we only need to simulate the counterfactual aggregate supply curve in a counterfactual competitive environment. This approach avoids to model the existence and workings of the cartel, which in the context of the world oil market is a complicated matter, if not alone for the variety of the geo-political and environmental aspects.

There is an ongoing discussion as to why some countries are poorer than others and

---

<sup>1</sup>See Goldberg et al. (2010); Melitz (2003); Edmond et al. (2015); Atkeson and Burstein (2010); Syverson (2004); Collard-Wexler and De Loecker (2015); De Loecker (2011)

how poorer countries can catch up to the richer countries. The existing misallocation literature might be thought of as attacking this question using a ‘top-down’ approach, noting the macro potential of various mechanisms to improve productive efficiency.<sup>2</sup> This project takes a more ‘bottom-up’ approach by examining a specific market and looking for misallocative channels, in just that market, that may have macro significance - the implication being that if the effects can be large in one market, then similar conduct across many other sectors in an economy can have similarly significant macro consequences. By focusing on a globally significant industry, and measuring the production inefficiencies resulting from its cartelized state, this project takes a first step to understanding the potential for endemic market power abuses to retard growth in less developed countries. This is a question that has been raised in recent times in policy circles, notably at the World Bank (Martinez Licetti and Goodwin, 2015). As far as we know, little research is able to speak to this issue. This project is a step in rectifying that gap.

## 2 Analytical framework

In a static environment the definition of a productive inefficiency is intuitive and precise. As illustrated in figure 1, it is the difference between the realized cost of production and the cost of producing the same quantity, had all firms been price takers. This productive inefficiency is the shaded area in figure 1.

In the empirical setting confronted here, the global oil market, a purely static approach is inappropriate, due to the finite nature of the resource being extracted. Thus a definition of productive inefficiency appropriate for a dynamic context needs to be adopted.

**Definition 1** *Productive inefficiency is the net present value of the difference between the realized costs of production, and the cost of production had the realized production path been produced by firms taking prices as exogenous.*

That is, the competitive benchmark is derived by holding the production in each year fixed and shifting demand for that year inward until a competitive industry would have produced, in equilibrium, the observed production. The path of costs of production thus generated is the counterfactual benchmark against which realized costs are compared to measure the extent of any production inefficiency due to market power.

Given the finiteness of the resource, it is clear that at some finite end date all resources will have been extracted. Hence, the source of inefficiency, in an industry such as this, is via sub-optimal inter-temporal substitution of production among production

---

<sup>2</sup>E.g. Asker et al. (2014).

units. Given this, the central economic object of interest is the order of extraction of assets that a competitive industry would have undertaken. The central result of this section is to provide a characterization of that order.

In addition to characterizing the extraction policy of the counterfactual competitive industry, this section builds the underlying cost function that is used in to guide measurement and modeling. It also relates the competitive solution to the social planner's solution and provides the algorithm used to compute the competitive solution. As usual, the data and empirical setting imposes some additional measurement issues that are discussed in the sections directly related to empirical analysis. This section focuses on the details of the theoretical structure common to the entire paper.

## 2.1 Modeling Preliminaries: Costs

In modeling costs the production unit is the field, denoted  $f$ , which is the unit of observation in the most disaggregate data that we have access to. For some of these fields, such as some offshore oil platforms, a field is an oil well. However, for most of the onshore oil fields, a field is composed of many different oil wells. Fields make input choices in order to minimize costs, conditional on a given level of production.

Let the production function for a field  $f$  in year  $t$ , be given by:

$$q_{ft} = \min \{ \alpha_{ft} K_{ft}, \gamma_{ft} L_{ft} \} \quad \text{s.t.} \quad q_{ft} \leq R_{ft}, \quad R_{ft} = R_{f,t-1} - q_{f,t-1} \quad R_{f0} > 0 \quad R_{ft} \geq 0 \quad (1)$$

where  $K$  and  $L$  are fixed and variable inputs, respectively and  $R$  are reserves. Finally,  $\alpha_{ft}$  and  $\gamma_{ft}$  capture the relevant technological environment.<sup>3</sup> These coefficients are field specific and as such they subsume the differences across technology (onshore, offshore, shale, etc.). The fact that the coefficients are allowed to vary across fields also implies that they capture any Hicks-neutral productivity shocks  $\omega_{ft}$ .<sup>4</sup>

Assume the price of capital inputs is given by  $r_{ft}$  and the price of labor inputs is given by  $w_{ft}$ . These input prices are assumed to be exogenous. This means that the total cost of production, assuming cost minimization at the field level, is simply given by.<sup>5</sup>

$$C(q_{ft}) = \left( \frac{w_{ft}}{\gamma_{ft}} + \frac{r_{ft}}{\alpha_{ft}} \right) q_{ft} \quad (2)$$

---

<sup>3</sup>We write down the model with capital and labor inputs ( $K$  and  $L$ ), but of course these are meant to stand in for the different inputs in the production process for oil, such as drilling equipment, production workers, and energy.

<sup>4</sup>That is the production function could have been written as  $q_{ft} = \min(\{\alpha_{ft} K_{ft}, \gamma_{ft} L_{ft}\} \omega_{ft})$

<sup>5</sup>This is derived by noting that, cost minimization implies that,  $\alpha_f K = \gamma_f L$  this means that  $K = \frac{\gamma_f L}{\alpha_f}$ . Finally, given these conditions, the production level is set by  $q_t = \gamma_f L$ , implying  $L = \frac{q_t}{\gamma_f}$ .

Additional structure is put on the process governing the evolution of ratio of input prices to the technology parameters such that  $\frac{w_{ft}}{\gamma_{ft}} = \frac{w_f}{\gamma_f} \mu_{ft}$ . This allows for variation in either a firm-specific Hicks-neutral productivity shock, common variation across inputs in the input price to technology ratios, or a combination of both.

This yields the following cost function:

$$C(q_{ft}) = \left( \frac{w_f}{\gamma_f} + \frac{r_f}{\alpha_f} \right) \mu_{ft} q_{ft} \quad (3)$$

Unit cost is then given by:

$$c_{ft} = MC(q_{ft}) = AC(q_{ft}) = \begin{cases} c_f \mu_{ft} & \text{if } q_{ft} \leq R_{ft} \\ +\infty, & \text{otherwise.} \end{cases} \quad (4)$$

where  $\mu_{ft}$  is governed by a martingale process such that  $E(\mu_{ft+k} | \mu_{ft}) = \mu_{ft}$  for  $k \geq 1$ .<sup>6</sup>

That is, costs have a hockey stick shape, constant marginal costs up to a capacity constraint given by reserves. From a measurement point of view, the constant returns to scale assumption on the components of the Leontief production function provides economic assumptions under which average cost and marginal cost are equal.<sup>7</sup>

The  $\mu_{ft}$  term captures the convolution of long-run trends in technological change, and changes in the absolute or relative cost of inputs or technology parameters ( $\gamma$  and  $\alpha$ ).<sup>8</sup> The process determining  $\mu_{ft}$  is assumed to be exogenous, which is an assumption with economic content and underscores the partial equilibrium nature of the exercise being conducted here. In an alternate, broader, context,  $\mu_{ft}$  is an equilibrium object. In particular, if lower cost fields get extracted first in the competitive counterfactual (as is the case), and these lower cost fields have lower input intensity, and the inputs are specialized, such that they are not readily deployable in some other sector, then this reallocation of production could change the equilibrium value of  $\mu_{ft}$ .

## 2.2 Production paths in competitive equilibrium

Oil is a finite resource, hence extraction path for a firm is the result of a dynamic optimization problem. Consequently, as noted above, the welfare impact of a production distortion is a function of the inter-temporal substitution of extraction times for different fields. Quantifying this welfare impact requires characterizing the path of extraction in a competitive equilibrium when demand is shifted such that the quantity produced in each year is equal to that observed in the real-world equilibrium.

<sup>6</sup>This implies that we assume the same process for both Hicks-neutral productivity shocks and input prices.

<sup>7</sup>Another implication is that costs are invariant to changes in demand conditions.

<sup>8</sup>Or, equivalently, a Hicks neutral productivity measure.

In competitive equilibrium all producers take prices as given. Let  $\delta$  be the common discount factor. Next note that  $E(\mu_{ft+k} | \mu_{ft}) = \mu_{ft}$  and that it is without loss of generality to employ the normalization  $\mu_{f1} = 1$ . Thus, for a given price path (or expectation thereof) a price taking producer solves the following problem:<sup>9</sup>

$$\max_{\{q_{ft}\}} \sum_{t=1}^T \delta^{t-1} (p_t - c_f) q_{ft}, \quad (5)$$

subject to

$$R_f \geq \sum_{t=1}^T q_{ft}, \quad \text{and } q_{ft} > 0 \quad \forall t \in \{1, \dots, T\}. \quad (6)$$

Proposition 1 and corollary 1 together establish that the lowest cost fields are extracted first in any competitive equilibrium.<sup>10</sup>

**Proposition 1** *Let marginal costs be described by equation 4. Consider two fields,  $\underline{F}$  and  $\bar{F}$ , with  $c_f$  equal to  $\underline{c}$  and  $\bar{c}$ , respectively. In any competitive equilibrium if  $\underline{c} < \bar{c}$ , then if  $\underline{R}_t > 0$ ,  $\bar{q}_t > 0$  implies  $\underline{q}_t > 0$ .*

**Proof.** Toward a contradiction assume not. Consider two periods such that, w.l.o.g.,  $t = t_2 > t_1 = 1$ . Consider a single unit of production for both  $\underline{F}$  and  $\bar{F}$ , such that  $q = \bar{q} = 1$ . Employ the normalization  $\mu_{f1} = 1$ . Hence,  $E(\mu_{ft} | \mu_{f1}) = 1$ . Thus  $E(c_f \mu_{ft} | c_f \mu_{f1}) = c_f$ .

Then, there must exist periods such that

$$\delta^{t-1} (P_t - \underline{C}) \geq (P_1 - \underline{C}) \quad (7)$$

and

$$\delta^{t-1} (P_t - \bar{C}) \leq (P_1 - \bar{C}) \quad (8)$$

where at least one inequality is strict. Assume for exposition, that the inequality in equation 8 is strict.

From equation 7,

$$\delta^{t-1} (P_t - \bar{C}) + \delta^{t-1} (\bar{C} - \underline{C}) \geq (P_1 - \bar{C}) + (\bar{C} - \underline{C}) \quad (9)$$

Since  $\delta^{t-1} (\bar{C} - \underline{C}) < (\bar{C} - \underline{C})$ , this implies  $\delta^{t-1} (P_t - \bar{C}) \geq (P_1 - \bar{C})$  which is a violation of equation 8. ■

<sup>9</sup>If prices are not known, it is assumed that all producers have the same expectations. In the maximum the price process is assumed to be known merely to keep notation simple.

<sup>10</sup>This is a result that is well established in the existing literature. It is given that the analysis rests on this result, is provided for clarity and stated for the exact cost environment employed in this paper.

**Corollary 1** *In any competitive equilibrium if  $\underline{c} < \bar{c}$ , then if  $\bar{R} > 0$ ,  $\underline{q}_t > 0$  does not imply  $\bar{q}_t > 0$ .*

**Proof.** This follows the line of argument used above, noting that  $\delta^{t-1} (\underline{C} - \bar{C}) > (\underline{C} - \bar{C})$ . ■

An immediate implication of this is that, in the absence of other distortions, when low cost fields are not being exploited prior to higher cost fields coming on line, then this is an indication of the exercise of market power. As is usual, firms with market power have an incentive to delay extraction to push prices higher. Any residual demand that results will be absorbed by fringe producers, with higher unit costs.<sup>11</sup>

Unsurprisingly, given the first welfare theorem, the production plan resulting from the competitive equilibrium coincides with that of the social planner than seeks to minimize the social cost of producing that production plan.

**Lemma 1** *In the social planner's production plan, which minimizes the net present value of costs subject to satisfying a aggregate production path, coincides with that of the competitive equilibrium.*

**Proof.** The proof is straightforward, and proceeds via contradiction. ■

Following proposition 1, the production path resulting from a competitive equilibrium, that generates known aggregate production (equivalently consumption) levels in each year ( $Q_t$ ), can be computed using the following algorithm.

**Algorithm 1** *Sorting Algorithm for Optimal Extraction Decisions*

1. Start in year  $t = 1$ .
2. Set the field index  $i$  to order fields from lowest to highest marginal cost given costs  $c_f \mu_{ft}$ , such that a lower  $i$  corresponds to a lower cost.
3. Start with  $i = 1$ .
4. Drain field  $i$  until remaining reserves equal zero ( $R_{it} - q_{it} = 0$ ) or the aggregate production target is met ( $\sum_{j=1}^i q_{jt} = Q_t$ ). Update remaining reserves for this field (set  $R_{i,t+1} = R_{it} - q_{it}$ ).
5. If  $\sum_{j=1}^i q_{jt} < Q_t$ , set  $i = i + 1$ , and go to step 4.
6. Set  $t = t + 1$ .

---

<sup>11</sup>See Sweeney (1992) for an extended discussion of the comparison of competitive equilibrium and equilibrium with market power in these settings.

7. If remaining reserves are positive for any field and  $t < T$ , go to step 2
8. Else, STOP.

This algorithm is used to generate the counterfactual production path, against which the observed production path is compared to measure the extent of production misallocation.

### 3 Data and institutions

The analysis in this paper focuses on the upstream oil industry (that part of the industry concerned with extraction), as opposed to activity further downstream (such as refining). Data on the upstream oil industry was obtained from Rystad Energy (Rystad hereafter), an energy consultancy based in Norway. The specific data product is called the U-Cube, or upstream, database. The data cover the operations of each oil field around the world.<sup>12</sup> For each field the data include production and different operating and capital costs, as well as the characteristics of the field such as the geology and reserves.<sup>13</sup>

The distinction between a production unit (field) and its smaller components (wells) is important since in our data we observe cost and production information at the field level. A field, in the data, is defined as a geologically homogenous area with the same management and ownership structure. In practice, there appears to be some judgement involved in the application of this definition. Fields vary considerably in the number of wells and associated infrastructure.<sup>14</sup> For instance, in the data, the Gullfaks offshore field in Norway is decomposed into two separate oil fields, Gullfaks, which has three oil platforms, and Gullfaks South, which has a single platform.<sup>15</sup> On the other hand, in the data, the Ghawar Uthmamiyah onshore field, which is one of the largest fields in the world, is composed of many hundred wells.<sup>16</sup>

The institutional features of the upstream oil industry are important for understanding the measurement issues that arise in handling the data. As a consequence, in what follows, the data and institutions (equivalently, the data generating process) are introduced simultaneously.

An understanding of the data generating process is also crucial in establishing that production units (oil fields) can have very different costs for exogenous (geological)

---

<sup>12</sup>Very small fields in some countries may not make it into the database, being judged to be inconsequential.

<sup>13</sup>The data are described in greater detail in the data appendix.

<sup>14</sup>For many fields the number of wells can be ascertained via a google search.

<sup>15</sup>Note that a single offshore platform may be connected to many several well-heads.

<sup>16</sup>Saudi Aramco does not publish an exact number, but satellite images suggest a number in the high hundreds at least.

reasons. That is, a Norwegian deposit that exists in deep water far offshore or a Canadian tar sands deposit will have very different average (equivalently, marginal) costs of production as compared to the larger onshore deposits in Saudi Arabia.

The Rystad data covers the oil global industry. As a result it is collected from a variety of original and secondary sources, ranging from high quality government reports in countries such as Norway and the United States, through company reports for large private companies, to interviews with shipping companies and oil service companies. Collating and reconciling these sources is a difficult process, particularly in politically unstable areas such as, for instance, Syria from 2012 onward.<sup>17</sup> Thus in some countries, while aggregate production at the country level may be observable from things like tanker movements, Rystad uses engineering models to approximate micro production and costs where numbers are not directly reported. This is unavoidable when attempting to study a global industry that has strategic importance.<sup>18</sup>

### 3.1 Crude oil

Crude oil is the oil that is delivered to refineries for processing into the various hydrocarbon products used by end consumers.<sup>19</sup> These products range from gasoline and other fuel oils (the majority of refinery production) to bitumen, lubricants, propane, naphthas and some waxes (such as parafin).<sup>20</sup>

Due to variation in local geology, the nature of crude in a deposit will vary. The two most important dimensions of heterogeneity in crude is density and sulphur content. Density is commonly measured in degrees API and sulphur content as a percentage by weight.<sup>21</sup> Crudes typically have a density between 10° and 50° API.<sup>22</sup> Most refineries are geared toward processing crudes in the 30° to 40° API range, with some variation across refineries. Different refineries will have different sulphur tolerances as well. Hence, crudes that lie outside the 30° to 40° API range, or which have very low or high

---

<sup>17</sup>ISIL controlled portions of Syria in this period, including oil producing assets.

<sup>18</sup>Even in more stable areas reporting standards vary over country and time. See the data appendix for more details.

<sup>19</sup>The sources of this industry description are, where not otherwise noted, Downey (2009),

<sup>20</sup>Table A.4 in the online appendix shows the total production of each type of fuel for 2014. Specifically we count as crude Condensate, all levels of API crude from extra heavy to light, and crudes that higher sulphur content/sour, as well as NGL. We exclude biodiesel, synthetic crudes, cold to liquid production and Bitumen-based production. The latter collectively account for about two percent of oil production in 2014.

<sup>21</sup>Crudes are often referred to as heavy or light, and sweet or sour, referring to their density and sulphur content respectively.

<sup>22</sup>Common benchmark crudes are Arabian Light, Brent and WTI which have densities of 34°, 38.3°, and 39.6° API, respectively. In measuring crude output we include conventional crudes (API < 50) and condensates, which are gaseous in the deposit but liquify after extraction. Condensates have API > 50.

sulphur contents can trade at a discount (or a premium depending on market conditions) and may need to be mixed in with other crudes to meet refinery specifications.

The heterogeneity in crudes leads to a series of measurement issues. The first is how to measure the quantity associated with a deposit in terms comparable across deposits. The data measures output in energy equivalent barrels, where the benchmark is one barrel of Brent Crude. Hence, the measure of quantity accounts for the compositional heterogeneity of crudes. The second issue is that different crudes trade at different premia and discounts related to their composition. Thus the choice of a price index need to be consistent with the measure of quantity. The price of Brent crude is the price measure used here to be consistent with the production measure.

## **3.2 Crude oil production**

Once a deposit is discovered it needs to be exploited. A deposit will be located in a field, which is a deposit, or set of deposits, sitting within a common geological structure. The manner in which a field is exploited will depend on its location and underlying geology. Every deposit will be exploited by drilling production wells. Beyond that, the most basic distinction is between onshore and offshore fields.

### **3.2.1 Onshore production**

Production of an onshore deposit typically involves a range of stages or techniques.<sup>23</sup> These are referred to as primary, secondary and tertiary recovery (methods). Primary recovery can be further broken down into two stages. Flush production occurs early in the production life of a well when the subsurface pressure is sufficient to drive the crude up through the drilled bore without any assistance. While flush production persist, the production team may undertake efforts to maintain well pressure by recharging the underlying aquifer or engaging in limited gas re-injection. After flush production ends, or if the underlying pressure is insufficient to ever experience flush production, settled production occurs. During this phase, pumps of one form or another are used to lift the oil up through the bore to the collection point above the surface. Depending on the geology, there are many ways to configure the pump and pipe at this stage. The commonly observed “nodding donkey” pump is an example of settled production technology.

Secondary recovery differs from settled production in that, rather than pulling the oil out, methods are used to push. These methods center around the injection of gas or water, often extracted with the crude, into the deposit to maintain or increase well

---

<sup>23</sup>References: Downey, Morgan (2009), *Oil 101*, Wooden Table Press, La Vergne, Tennessee.

pressure. the efficacy of these approaches depends on the underlying geology.<sup>24</sup>

Tertiary recovery refers to a range of alternate extraction approaches, often used in tandem with primary and secondary recovery techniques. These include fracturing the rock to increase its permeability ('fracking'), injecting mixing agents into the deposit to increase viscosity and more aggressive approaches to gas injection to increase well pressure.

It is important to note that every deposit will have geological features that dictate a different (and at times simultaneous) combination of primary, secondary and tertiary recovery methods over the course of a well's lifespan.

Other forms of onshore deposit require other production methods. For instance, tar sands, a significant proportion of which are found in Venezuela and Canada, are heavy crudes found close to the surface mixed into loose rock or sand. These crudes are recovered through surface mining and then require cleaning (to remove sand and soil) and pre-processing (to lift the API to refinery appropriate levels). As a result tar sands can have extremely high production costs.

### **3.2.2 Offshore production**

Off-shore deposits have additional production challenges.<sup>25</sup> The extent of these challenges are determined by the water depth, the distance from land and the weather. Water depth creates both pressure at the well head (at the ocean floor) and temperature differentials between the subsurface deposit (hot) and the deep water (cold). Both lead to substantial engineering problems. For instance, to avoid frigid water changing the composition and viscosity of the crude, the bore is heated to keep the crude at a steady temperature. In deep water wells anti-freeze is also often added at the well head. Distance from land affects the way labor can be housed, transported and rotated and how the crude can be stored and transported back to land. The weather presents a series of additional challenges related platform stability, production interruptions and safety due to variation in storms, hurricanes and sea states across the globe.

### **3.2.3 Two examples**

The cost differences that arise from different operating environments are well illustrated via example. Consider the North Ward Estes field near Wickett in Ward County, Texas (an onshore conventional field) and Tract 174 in Grand Isle Block 43 located offshore of Lafourche Parish on the Louisiana coast (an offshore field in less than 100ft

---

<sup>24</sup>For instance, a porous bedrock may lead to re-injected water leaking away, with no gain in well pressure.

<sup>25</sup>The first out-of-sight-of-land offshore oil platform was established in 1947, and the first platform in > 1000ft of water was established in 1979.

of water). The unit cost (comprising all of operating and capital expenditures) of the onshore field from 1970 onward was \$7.57 per barrel.<sup>26</sup> The offshore field had a unit cost of \$19.74.<sup>27</sup>

### 3.2.4 Mapping theory to the data generating process

The theoretical structure described in section 2 is built around a fixed proportion (Leontief) production function, which implies a constant returns to scale technology in which average costs equal marginal costs. This simplifies many measurement and computational problems. Actual crude production involves many different technologies which may be used simultaneously or at different stages. As extraction moves from one technological mix to another, for instance as an aged field moves into tertiary production, cost structures may change. The economic content of this is to possibly introduce some curvature to a field or well's cost function, particularly as the last few percent of the local reserves are extracted.<sup>28</sup>

Given this, careful consideration of the nature of the leontief assumption is warranted. As in every production process some fixed costs and scale effects undoubtedly exist in this industry. For instance, a well needs to be located and run, regardless of the amount of oil recovered. This is analogous from the delivery problem of a national retailer, that must send a truck out to deliver inventory to a store, regardless of whether the truck is 50% or 100% full. At the micro level the well or truck may look like a fixed cost, but at a more aggregate level may be better treated as a component of variable costs. In judging whether the leontief assumption is helpful, it is helpful to keep in mind the level of aggregation at which the analysis is being done. The analysis is industry wide, aggregating the equivalent of an industry supply curve over 13,248 fields. The Leontief technology assumption makes this supply curve a step function. Modeling each well, and aggregating up, would at best put a small amount curvature in each step, which given the level of aggregation would be difficult to notice for the typical field.<sup>29</sup>

---

<sup>26</sup>Unit costs are computed by deflating all expenditures from 1970 (inclusive) by the US GDP deflator using 2009 as the base year. Total expenditures are divided by total production from 1970 (inclusive).

<sup>27</sup>Both fields were selected to be in the southern US states to keep currency and input market conditions as similar as possible.

<sup>28</sup>Kellogg (2014) and Anderson et al. (2014) both provide a detailed analysis of the operation of individual wells in Texas onshore fields. Both complement the current paper, by illustrating the mechanics of the production process at a much lower level of aggregation than is used in this paper, and than is possible with the Rystad data.

<sup>29</sup>Anderson et al. (2014) argue that at the micro-level the decision of how to run a well is analogous to deciding when to "tap a keg" in the sense of, once the keg is tapped, the flow is constant. Hence, the economically operative decision occurs at the extensive margin. As such, the distinction between marginal and average costs becomes moot.

### 3.3 Measurement of reserves

Central to much of the discussion in this paper is the notion of reserves. The reserve is the unextracted, but recoverable, quantity of oil remaining in the ground in a field.<sup>30</sup> The most reliable way to measure the reserve at a point in time is to see the entire production life of a field. The total extracted is the maximal reserve. Most fields are not fully exploited in the data. Hence, industry reserve estimates need to be used. The oil industry reports reserves at different levels of extraction probability. There are three levels. P90 (or P1) is the quantity able to be recovered with a 90% probability given current technical and economic conditions. The P90 reserve is the asset value able to be reported on company balance sheets under U.S. GAAP. Clearly, this definition means that reserves will fluctuate with the oil price. In the data used here, reserves are measured and reported assuming an oil price of \$70 (in 2014 dollars), which is closest to the historical average price for oil. P50 (or P1 + P2) are the reserves recoverable with a 50% probability. Finally P10 or (P1 + P2 + P3) are total reserves recoverable with a 10% chance. The level of P90, P50 and P10 can vary significantly within a field. For instance in the North Ward Estes field discussed above, P90, P50 and P10 in 1975 were estimated at 26.6, 56.4 and 66.4 million barrels. In this paper, unless stated otherwise, the reserve number used for a field is P50. The precision with which reserves are measured varies by the production stage of a field, and the country in which the field exists. Untapped deposits have less precise reserve estimates, since there is no actual production data or well pressure data to rely on. Once a field starts producing, reserves become easier to estimate, particularly as pressure starts to change, as the pressure gradient of a field as the resource is depleted is relatively well understood, conditional on geology. A further confounding factor is that oil reserves are strategic assets, with most industry sources commenting that various countries will inflate reserve figures for political reasons. This is particularly relevant for OPEC countries due to the way OPEC has computed quotas at various times in its history.<sup>31</sup>

In this paper, in descriptive discussions (prior to section 5) P50 values at an oil price of \$70 a barrel are used to report reserves. In section 5, a field's reserves in 1970 are computed as the sum of: i) the actual production history from 1970 to 2014, and ii) the P50 value at an oil price of \$70 a barrel in 2014.

#### 3.3.1 The global oil market

The global upstream market for oil is characterized by a range of actors. The buyers are refineries. The producers are oil companies, which are either state run enterprises,

---

<sup>30</sup>Depending on the geology, between 25% and 75% of the crude will remain in the deposit after production has concluded.

<sup>31</sup>This paper takes the Rystad reserve data as the best estimate available.

substantially state run, or independent. Until the early to mid 1980's most oil was traded using long term contracts at fixed price. By the mid-1980's most crude was traded in long term contracts at prices that referenced benchmark spot prices (such as NYMEX WTI or ICE Brent). By this time there was an active spot market for crude.

The state run (nationalized) oil companies, can be split into those that are run by OPEC states and those that are from non-OPEC states. Every OPEC country has its own nationalized company which controls production, albeit at times contracting with independents to run specific facilities. For instance, Saudi Arabia operates Saudi Aramco, Kuwait operates the Kuwait Petroleum Company, and Ecuador operates Petroecuador. OPEC, being a significant and persistent feature of the world market, is discussed in more depth in the next section.

Outside OPEC nationalized (or substantially state run) companies exist in Mexico, Brazil, Russia, China, Malaysia, Norway and India, and in several other smaller producing nations. In other major producing countries (like the USA, UK or Canada) production is conducted by private (independent) companies. These private companies can be divided into the five (as at 2014) oil majors (ExxonMobil, Chevron, BP, Royal Dutch Shell and Total) all having revenues in excess of 100 billion US dollars, and other independent companies (e.g. Woodside Petroleum or Encana).

Table 1 shows the production shares, for the period 1970-2014, of the seven largest OPEC and non-OPEC countries. The US has the largest production, closely followed by Russia and Saudi Arabia. While these three countries have the largest production, it is important to bear in mind that production occurs in different ways within each country. The US is very decentralized, having many private firms, while Saudi has a nationalized oil company (Saudi Aramco).

### 3.3.2 The OPEC Cartel

OPEC, in 2014 (the limit of the data), comprised the countries of Algeria, Angola, Ecuador, (Indonesia), Iran, Iraq, Kuwait, Libya, Nigeria, Qatar, Saudi Arabia, UAE, and Venezuela. The membership has varied slightly over time, with the core Middle East membership being unchanged from its inception in 1960.<sup>32</sup> The first 30 years of OPEC are well documented in Yergin (1991) and Crémer and Salehi-Isfahani (1991).

Figure 2 shows the market share of OPEC and price of crude from 1970 to 2014. Before OPEC starting coordinating extensively on price reductions, it had a share global

---

<sup>32</sup>The original membership in 1960 was Iran, Iraq, Kuwait, Saudi Arabia, and Venezuela. Other members are listed together with the year they first joined OPEC, and (if appropriate) years in which membership was suspended or terminated: Qatar (1961), Indonesia (1962, suspended 1/09), Libya (1962), the United Arab Emirates (1967), Algeria (1969), Nigeria (1971), Ecuador (1973, suspended 12/92-8/07), Gabon (1975, terminated 1/95) and Angola (2007). See [www.opec.org/opec\\_web/en/about\\_us/25.htm](http://www.opec.org/opec_web/en/about_us/25.htm) accessed 29 August 2016.

production fluctuating around 48 percent. This fell to a low point of 29.2 percent in 1985 after reductions in production during the late 70s and early 80s. Following that, OPEC's share of production rose to 40.6 percent in 1993 and has stayed relatively constant since then.

For its first ten years, OPEC was ineffective in moving the oil price. This changed in 1973 when, following the outbreak of war in the Middle East, the OPEC member countries were successful in raising the oil price four-fold, albeit with (according to secondary sources) little coordination among themselves. Until 1982, the cartel did little more than try to set price guidance for its members, with little hope that that guidance would be adhered to.

In March 1982 OPEC took its first step in evolving toward its modern form. At this time the cartel introduced country-specific production quotas for the first time. Despite these initial quotas being ineffective, quotas have continued to be a defining feature of OPEC's operation. Notably in the 1985, after becoming frustrated with defections by other OPEC members, Saudi Arabia expanded production in an effort to discipline production. This may have worked had it not been for a slowdown in oil demand in the 1980's and the first Gulf War in the early 1990s.

During the 1990's expansion in oil demand outstripped the ability of non-OPEC producers to expand supply and the power of the OPEC cartel appears to have increased.<sup>33</sup> By November 1997 the cartel was in a position to exert substantial influence on the market price, but appears to have either broken down or mis-judged global demand, expanding production to coincide with the Asia crisis. This led oil prices to decline from \$35 per barrel to \$10 per barrel.<sup>34</sup> Finally, in March 1999, with the cooperation of Russia, Norway and Oman, OPEC countries were able to cut back production and move the price back into the \$30-\$40 per barrel range.

In 2000, the cartel explicitly announced a target price band of \$22-\$29 per barrel. The quota system was adjusted to include automatic adjustments should the reference price fall outside this interval. This heralded the modern era of OPEC and the most sophisticated coordinating mechanism seen to date. Unfortunately for the cartel, the second Gulf War and an expansion in demand from developing countries, in particular, made the price band unsustainable. In January 2005 it was suspended. The quota system, however, lived on.<sup>35</sup> Although the quotas are not immediately transparent, it is clear that they are highly asymmetric.<sup>36</sup>

By July 2008 the Brent spot price had increased to an intra-day high of over \$140 per

---

<sup>33</sup>See Kohl (2005) for a description of this period.

<sup>34</sup>All prices in the text of this section are nominal (not deflated).

<sup>35</sup>See Kohl (2005) and Fattouh (2007)

<sup>36</sup>We are currently working on compiling accurate quota magnitudes from OPEC publications and checking these against other industry sources. See various monthly publications of the Monthly Oil Report (Organization of Petroleum Exporting Countries (2008)).

barrel, a level in stark contrast to the price band in (ineffective) operation a mere 4 years earlier. Through to the end of 2014, oil price remained high by historical standards.<sup>37</sup>

The efficacy of the OPEC cartel over time has been, at best, variable. Especially in the early 2000s, its policies mirrored those of collusive structures seen in other industries. In other periods, the ability of OPEC to coordinate its members' production seems less clear. In this paper, the question of whether OPEC is best characterized as a political vehicle for Saudi Arabia and other gulf countries, or as a long running industrial cartel, does not require an answer. To the extent that someone in OPEC has market power and distorts production, the measurement approach adopted in section 5 will account for it.<sup>38</sup>

### 3.4 Data

Table 2 presents summary statistics for the 13,248 active fields in the data. The average field produces 3.4 million barrels per year and has reserves of 99 million barrels (the medians, are 0.2 and 3.7 respectively). There is wide variation in field size, with the 5th percentile field producing less than 1,000 barrels, and the 95th percentile field producing 11 million barrels. The largest annual production for a field observed in the data is the Samotlor field in Siberia in 1980 with almost 1.2 billion barrels produced that year. Almost 19 percent of fields are offshore.<sup>39</sup> The analysis presented in this paper is restricted to fields that were active at some point between 1970 and 2014.<sup>40</sup>

Given the Leontief production function, yielding the cost function given by equation 4, the average and marginal cost of oil production are the same. Hence, the marginal cost of production is recovered by dividing the total cost of production by the reported production (in million bbl/day). The total cost of production is obtained by summing over the ten cost categories as listed in Table 2. These are defined in Table 3.

Taxes and royalties, whether financial or in-kind, are potentially important to the extent that they are relevant costs for price taking firms. That is, for computing the counterfactual production path, they may be relevant.<sup>41</sup> However, conditional on the

---

<sup>37</sup>Despite decreasing to below \$35 a barrel in late December of 2014.

<sup>38</sup>This is a consequence of working with cost data directly. By contrast, if the more standard IO approach were adopted, in which demand estimates were combined with a pricing model to back out costs, then this issue would likely be a significant hurdle, if not fatal, to the credibility of any estimates.

<sup>39</sup>The largest field in terms of cumulative production is the Ghawar field in Saudi Arabia.

<sup>40</sup>Of the 66,920 fields in the Rystad data, 45,687 of these did not produce between 1970 and 2014. The cost and reserves data exists on these fields is solely based on engineering and geological modeling, and these fields are not used in the paper. More detail on the sample selection can be found in Table A.1 in the Online Appendix.

<sup>41</sup>Not all taxes that are recorded appear distortionary. For instance, taxes on profits are excluded from this exercise. Only those taxes or royalty related line items that have a possible distortionary impact are

production paths, when measuring productive inefficiency due to market power they are excluded, as they are not costs that are relevant to the social planner (that is, they are not true resource costs). Importantly, this means that the resources cost of production is measured absent taxes and royalties. Royalties and Government Profit Oil might at first glance contain the user cost of land, paid as share of gross production. This item is relevant for the many oil fields in our data. However the level of this rent is too high in the data to be credible entirely comprised of the user cost of land. Further, the opportunity cost of using the land (or seabed) for drilling oil is close to zero in the majority of major oil deposits due to the inhospitable nature of the local environment.

It is unclear where the tax and royalty measures in table 3 are actually distortionary. While the definitions suggest that they are distortionary, in practice they are used to by governments to extract revenues from oil deposits. These governments have a wide range of additional instruments with which to extent control on production to mitigate any distortionary effects of royalties that, in effect, operate, to some extent, like revenue sharing arrangements. Given this, and the resulting ambiguity that surrounds the distortionary impact of taxes, the baseline cost specification omits taxes and royalties from all cost measures. As a robustness check, an alternative cost specification, that includes tax and royalty measures, is also used to compute an alternate production path, and resulting distortion measure, in section 5. The baseline cost measure and the tax and royalty inclusive cost measure are compared in table 5 found in section 4.

This data is unusually rich compared to most studies of the oil market, which either use detailed micro data on a small subset of oil fields, such as Covert (2015) or Kellogg (2014)'s studies of recent activity in North Dakota Shale and Texas, or look at the global oil market with data aggregated to the country level (such as Kilian (2009)).

## 4 Descriptive evidence of production inefficiency

Central to the existence of a productive inefficiency is the existence of cost dispersion between productive units, and the capacity of low cost unit to expand production to displace the production of high cost units. This section documents these features in the data. It also provides reduced form evidence consistent with the existence of market power by OPEC, and by Saudi Arabia in particular. We begin by focusing on a small number of major oil producing nations. By focusing on a small number of countries more features of the underlying data can be illustrated. Attention is then shifted the entire global market in which the aggregate data is shown to mirror the patterns observed in the more detailed country-level analysis.

Figures 3 and 4 show moments of the distribution of production costs, for each included in the tax-inclusive costs. These are the ones that appear in Table 3.

year from 1970-2014, for each of Saudi Arabia, Kuwait, Venezuela, and Nigeria (OPEC Countries, Figure 3; and, United States, Russia, Canada, and Norway (non-OPEC Countries, Figure 4).<sup>42</sup> Panel (a) of Figure 3 examines Saudi Arabia. The solid black line is the oil price. Below that, for each year, is a black bar. This shows the range of costs lying between the 5th and 95th percentile, where the unit of observation is the barrel. That is, 90% of barrels produced by Saudi Arabia in a year have a unit cost lying in the range indicated by the black bar. The grey bars combines with the black bar indicates the range of costs between the 1st and 99th percentile. Where circles are shown, these indicate the maximum unit cost for the country.<sup>43</sup>

An examination of Figure 3 illustrates the tight range of costs for Saudi Arabian and Kuwaiti production. For both countries, costs per barrel rarely exceed \$10. Further, costs are stable relative to the oil price.

By contrast, costs in Venezuela and Nigeria are much higher and exhibit much greater dispersion. This is an important feature of the data, suggesting that, even within OPEC, scope for efficiency gains due to reallocation of production exist. If OPEC were run as an efficient cartel this feature would not exist.<sup>44</sup> Given the many internal and external political challenges faced by OPEC, which mirror those faced by any real world cartel, it is unsurprising that it fails to act as a theoretically efficient cartel might.<sup>45</sup>

Figure 4 allows the within-OPEC patterns in Figure 3 to be compared to those in non-OPEC countries. Panels (a) and (b) of Figure 4 shows the US and Russia, the two biggest oil producers between 1970 and 2014.<sup>46</sup> Both the US and Russia have more dispersion in costs than that observed in Saudi Arabia or Kuwait, although a significant proportion of production, particularly prior to 2000, has equivalent costs. Importantly, the more expensive production in both countries occurs at cost levels more than twice than that that characterizes production in Saudi or Kuwait. This is particularly pronounced in the years following 2000, and particularly in the US, where the ramp up in high cost production follows the rise in the oil price and is largely driven by unconventional onshore production (mostly shale).<sup>47</sup> Canada mirrors the US, with a similar

---

<sup>42</sup>Costs are computed using the baseline specification. That is, all taxes and royalties are excluded.

<sup>43</sup>Objects with values above \$100 are not shown

<sup>44</sup>That is, allocations within OPEC would be determined by a constrained social planner, with the production path having the same features as in Lemma 1

<sup>45</sup>See Asker (2010) for an example of a inefficient real world cartel. Marshall and Marx (2012) provide a overview of a large number of cartels and the related theoretical and empirical work that organizes our understanding of them.

<sup>46</sup>See Table 1.

<sup>47</sup>In 2014, 2039 million barrels were produced by shale (out of 4173 million barrels produced in the United States). These shale deposits had costs of 32.6 dollars per barrel, while onshore fields have costs of 7.4 dollars per barrel (where we report production weighted averages). In contrast, in 2005, shale only accounted for 24 million out of 2480 million barrels produced in the United States, and costs for onshore

ramp up in costs following 2000.

Norway is distinct from the three other countries in Figure 4, by virtue of having the vast majority of its production offshore. This accounts for the late start in production. Deepwater offshore drilling technology only became commercially viable in the late 1970s. The spikes in the ranges of unit costs reflect the starting years of oil rigs, the low production levels that the first year of production often brings, and the large scale of the infrastructure involved.<sup>48</sup> Interestingly, the rise in the oil price following 2000 brings an increase in the dispersion of costs, albeit in a much more muted way relative to the US and Canada.

The comparison of the dispersions in production costs between Saudi Arabia, Kuwait and the other six countries in Figures 3 and 4 illustrate the considerable scope for reallocation that exists.<sup>49</sup> Further, the low costs that Saudi Arabia and Kuwait enjoy make it clear that, in a competitive equilibrium, these countries would be exhausting their deposits, subject to physical limits on extraction speeds, before the more speculative fields observed in the upper portions of other countries costs distribution come online (see Proposition 1).

The extent to which the result in Proposition 1 is useful in interpreting the data rests on the plausibility of the following counter-example: if the low cost fields in Kuwait are constrained by reserves, while those in Canada are not, then it is not surprising that there is no scope for low cost countries to expand production. In Table 4 we show reserves in different regions of the world, as well as the ratio of reserve to production (that is, the number of years that a region could produce at the current rate) for 2014. Outside OPEC, the ratio of reserves to production is 10. However, in OPEC countries, the ratio of reserves to production is 20. Hence, the data are consistent with the members of the cartel are restricting production, relative to reserves, more than producers outside the cartel. As one might expect from the literature and historical commentary surveyed in section 3, the data are consistent with OPEC, and Saudi Arabia in particular having market power.<sup>50</sup>

---

fields were 7.3 dollars per barrel. Thus, much of the large increase in costs in the United States is driven by the increased production from shale.

<sup>48</sup>The lumpiness observed here is inevitable. A large offshore project will involve many wells coming on line in the same year. If production starts late in the year, little production will be recorded, despite a large expenditure on infrastructure.

<sup>49</sup>Dispersion in production costs is ubiquitous, and has been documented in a variety of settings ranging from manufacturing to services – see Syverson (2011) for an overview of the literature. Compared to the reported dispersion in productivity (measured by TFP) in the studies cited in Syverson, the dispersion in oil production is high; there is a 1:9 ratio between the 10<sup>th</sup> and 90<sup>th</sup> percentile of cost. This is markedly higher than most industries, and is especially surprising since for the oil industry, measurement is not contaminated by variation in output prices, which is a common issue in the literature (De Loecker, 2011).

<sup>50</sup>Other Middle East states, like Kuwait, also behave in ways consistent with market power. We focus on OPEC and Saudi since OPEC is the joint vehicle, and Saudi has the largest reserves and production.

The patterns observed in comparing the eight countries in Figures 3 and 4, are reflected in Table 5, which compares production, reserves and costs over time for Saudi Arabia, OPEC and all non-OPEC countries. Unit costs are reported using both the baseline specification, which omits taxes and royalties, and the alternative specification that includes taxes and royalty payments.<sup>51</sup> Considerable scope for reallocation exists in each period, with the scope increasing as time goes on.<sup>52</sup> This is the case regardless of whether the baseline or alternative cost specification is used.<sup>53</sup>

## 5 Quantifying the extent of misallocation

This section quantifies the extent to which market power can plausibly account for the misallocation observed in the production of crude oil. This is done by using the model described in section 2 to compute a counterfactual production path, which is then compared to the actual production path to quantify the cost of misallocation. This requires the model to be parameterized. The details of this parameterization are found in the subsection below. Following that, the logic by which misallocation attributable to market power is identified is discussed. The results then follow, together with a series of robustness tests.

### 5.1 Model parameterization

Algorithm 1 is used to compute the competitive allocation (production path) in the counterfactual model described in section 2. The inputs required are the aggregate production levels,  $Q_t$ , field-level total reserves,  $R_{i,t=1}$ , and field-year costs,  $c_{f\mu_{ft}}$ . The remaining element required is a social annual discount factor, needed to compute a net present value of any accumulated distortions. This is set at 0.95.

Aggregate production is observed in each year from 1970-2014.<sup>54</sup> For years following 2014, global production (equivalently, demand) is assumed to grow at a rate of 1.3% per year, which is the (geometric) average growth rate observed for 1970-2014.

Reserves are described in section 3.3. As is also described in that section, in the baseline specification reserves are measured using P50 reserve figures, assessed at a

---

<sup>51</sup>See the discussion in section 3.4 .

<sup>52</sup>This is not surprising, as if distortions persist and low cost OPEC deposits remain significant, in later periods higher cost deposits should come online as lower cost, non-OPEC, deposits get exhausted. Hence, the potential for gains from reallocation should get larger over time.

<sup>53</sup>Saudi Arabia does not extract all its oil revenue via income taxation, which is not included as it is characterized as a profit tax and hence does not distort production.

<sup>54</sup>It is assumed that markets clear within the year, so that annual demand and production are equivalent. There is relatively little storage of extracted crude oil relative to total demand. Indeed, the best way to store oil is not to pump it out in the first place.

price per barrel of \$70 in 2014 dollars. Since reserves fluctuate somewhat over time for a given field, the actual production up to 2014 is added to the P50 reserve level in 2014 to give the reserve level for a given field available in 1970.<sup>55</sup>

Field-level costs are the central input required by algorithm 1. Recovering  $c_f \mu_{ft}$  from the cost data is the central aspect of generating this input. However, some auxiliary modeling elements, that bear on costs, are also relevant. The auxiliary elements are dealt with first. Then the recovery of  $c_f \mu_{ft}$  is discussed.

The first auxiliary element is that the path of field discovery is assumed to be exogenous. Hence, for a field discovered in 1980, the cost of production is infinite prior to that date. Similarly, fields that are never observed to produce between 1970 and 2014 are excluded. This is equivalent to assuming that the cost of these fields are infinite.

The second auxiliary element is the imposition of a limit on the proportion of  $R_{i,t=1}$  that can be extracted in each year. The model in section 2 assumes that any amount of oil may be extracted, up to the limit of available reserves, in any year. This is clearly a simplification. A range of engineering and geological factors can limit the proportion of reserves that can be extracted from a field in any given year, not least of which is the need to maintain a minimum level of pressure in the well so as to make extraction feasible – extraction that is too fast can lead to sharp drops in well pressure. The median producing field extracts 1.9 percent of its initial reserves per year, and the 95th percentile field extracts 13.3 percent.<sup>56</sup> Thus, in the main specification we will cap the extraction rate at ten percent of  $R_{i,t=1}$ . Algorithm 1 is easily adjusted to accommodate these auxiliary model elements, and we will present robustness checks where the extraction rate is alternatively chosen to be two percent, or unrestricted.<sup>57</sup>

We now turn to recovery of  $c_f \mu_{ft}$  from the cost data. Unit costs for a field-year are measured as described in section 3.4. These unit costs, denoted  $c_{ft}$ , need to be decomposed into three elements. The first element is the time-invariant marginal cost,  $c_f$ . The second element is a technology-year specific cost shifter,  $\mu_{st}$ , where  $s$  indexes the technology (onshore and offshore). The third element is measurement error,  $\exp(\epsilon_{ft})$ . That is,  $c_{ft} = c_f \mu_{ft} = c_f \mu_{st} \exp(\epsilon_{ft})$ . In the counterfactual, production undertaken by field  $f$  in year  $t$  is taken to have occurred at cost  $c_f \mu_{st}$  per barrel.

The technology-year specific cost shifter,  $\mu_{st}$ , is estimated as

$$\ln \hat{\mu}_{st} = \sum_{f \in s} \kappa_{ft} \ln c_{ft}, \quad (10)$$

<sup>55</sup>For some fields, we see reserves increasing over time, most likely because of new discoveries inside the field, and improvements in technology which makes more oil recoverable. If we had used reserves reported in 1970, this would lead us to the uncomfortable position of having more oil extracted in the period 1970-2015 than reported reserves in 1970, at least for certain regions of the world.

<sup>56</sup>Figure A.4 in the Online Appendix presents a histogram for the extraction rate of oil.

<sup>57</sup>All that is required is that the algorithm keep track of activity in a year and set prices to be infinite once the relevant field-level limits are reached. Proposition 1 and corollary 1 are similarly unaffected.

where  $\kappa_{ft}$  is the quantity weight of a field in a given year's total output,  $\kappa_{ft} = \frac{q_{ft}}{\sum_{f \in s} q_{ft}}$ . Observations are weighted by production, as opposed to giving all fields equal weighting, since a field is an already aggregated unit of production, with the extent of aggregation varying across fields

The time-invariant marginal cost,  $c_f$ , is then estimated, allowing for measurement error, using the following (within-field) regression:<sup>58</sup>

$$(\ln c_{ft} - \hat{\mu}_{st}) = \ln \hat{c}_f + \epsilon_{ft}. \quad (11)$$

Estimation is conducted using weighted least squares, with the weights being the proportion of total field output being done in that year.

Where confidence intervals are reported, they are computed via a bootstrap. Specifically, a two-step bootstrap routine is employed. In the first step, for each resample  $k$  we take the true dataset and resample field-year observations  $ft$ , and compute  $\mu_{st}^k$ . Next, in the second step, for each field in the true dataset, the field-years are resampled. This allows  $c_f^k$  to be estimated using the  $\mu_{st}^k$  from the first step.<sup>59</sup> This in turn allows  $c_f^k \mu_{st}^k$  to be computed and using in algorithm 1 to compute counterfactual predictions  $\theta^k$ . The goal of this procedure is to capture the estimation error in both the field-technology coefficient  $\mu_{s,t}$  but also in the field specific coefficient  $\mu_f$ . Fifty bootstrap iterations are used.

## 5.2 Identification of misallocation costs attributable to market power

There are two challenges to identifying the economic impact of misallocation plausibly attributable to market power in the oil market. First, it is unlikely that every instance of misallocation can be attributed to market power. Second, the data do not extend past 2014, which means that we do not see extraction paths in the data beyond this point.

In the absence of any other source of distortion, measuring distortions due to market power would be straightforward. The net present value as at 1970 of the cost of the observed production path would be compared to the net present value of the competitive equilibrium production path. The difference between the two would be the misallocative effect of market power measured as a stock in 1970 (we will present numbers deflated to 2014 dollars to make dollar numbers comparable across the paper). In practice, this approach would likely subsume a range of distortions that may be attributable to sources other than market power.

<sup>58</sup>Fields with only one year of production are assumed to have no measurement error.

<sup>59</sup>The field-year observations used to compute  $\mu_{st}^k$  are resampled independently from those used to compute each  $c_f^k$ . The practical reason to do this is that there are a few large fields composed of tens of thousands of individual oil wells, such as Saudi Arabia's Ghawar fields, that have large, and central, effects in the counterfactual exercises.

To focus the measurement on market power, it is necessary to articulate where market power is held. In the context of the global oil, given the evidence presented in sections 3 and 4, market power could be exercised by Saudi Arabia, some intermediate subset of OPEC or by OPEC as a whole. We present results for Saudi and OPEC, these being the extremes of the extent of effective coordination among OPEC countries.

When, for illustrative purposes, OPEC is considered the repository of market power, this still leaves distortions outside and within OPEC to consider. Given this, we proceed by solving a series of constrained social planner problems. First, we solve for the competitive allocation holding each country's production level in each year fixed. This removes internal distortions likely not attributable to OPEC.

Next, OPEC production is held fixed, and the competitive allocation is solved for again. This means the undistorted market is free to reallocate production both within an country and across countries, subject to keeping OPEC production in each year the same as is observed in data. Lastly, the unconstrained competitive allocation is computed, which we call the (world) optimal solution. This allocation is only required to deliver the global production observed in each year in the data. We take the difference between the competitive allocation, holding OPEC production constant, and the optimal solution to be the distortion attributable to OPEC.

Almost surely, this measure of misallocation is conservative. In particular, it removes the distortions that emerge within OPEC itself that may be due to the political constraints that need to be met for OPEC to exercise any market power. That is, Saudi Arabia and Kuwait likely need to assign a positive quota to Venezuela in order to give them some rents from complying with the overall OPEC production plan. In most years, an efficient cartel would not have Venezuela producing. In the computation described above, distortions of this sort are not counted. Given that many real cartels are observed to use inefficient mechanisms to coordinate, at least some of the misallocation within OPEC should be attributable to the coordinated exercise of market power (see Marshall and Marx (2012) for an extensive overview, and Asker (2010) for a specific example). In addition, some of the within or across country distortions seen in countries outside OPEC may be due to strategic responses to OPEC production plans. To understand the extent to which this can further increase the misallocation attributable to market power, a competitive allocation in which only the the country allocations within OPEC are held fixed. This is then compared to the world optimal solution.

Lastly, the censoring in the data, such that production paths past 2014 are not observed, needs to be addressed. Given that oil is a finite resource, the central source of misallocative cost will be due to fields that are cheap to exploit being delayed, such that the resulting gains from trade occurs in the future and hence is discounted. This means that the future actual path of production matters for a measure of misallocation, as the more the exploitation of cheap resources are delayed, the greater the misalloca-

tion.<sup>60</sup> In the face of the inevitable censoring, we take a conservative approach. To project the path of “actual” production out past 2014 we compute the competitive solution, taking the stocks in each country at the end of 2014 as initial state variables. This means that there is no distortion in the path of actual production after the end of the data. As a result, the misallocation numbers we report are an underestimate of the true magnitude.<sup>61</sup>

### 5.3 Results

We rely on the sorting algorithm introduced above to quantify the role of market power in distorting the efficient allocation of resources. We do this under three distinct scenario’s. First, we take a simple static view, and calculate the productive inefficiency for each year in the data, by applying the algorithm to each year, given the relevant marginal cost estimates, paired with information on reserves. Second, we turn to the dynamic productive inefficiency, as defined in section 5.5, and calculate the net present value of the production distortion, from 1970 onward. We do this for two different sample periods, 1970-2014 and 1970-2050. Third, we perform similar counterfactual simulations but we add various (quantity) constraints to isolate the inefficiency stemming from within the cartel, both across and within its members, and outside the cartel. This approach is used to obtain conservative estimates of the productive inefficiency due to market power, by eliminating deviations from cost minimization that are plausible unrelated to either coordinated, or unilateral, market power effects. Finally, we turn to a few robustness checks and extensions. Every counterfactual cost that is computed is an estimate, given that we rely on estimated marginal cost, and therefore we report standard errors, and these are calculated as discussed in section 5.1.

### 5.4 Static productive inefficiency

Table 6 presents the market share of the twenty largest oil producing countries in 2014, and those that would be predicted by our competitive model starting in 2014 (hence the term static). Saudi Arabia produced 13 percent of the world’s oil, but the competitive

---

<sup>60</sup>Our measure of distortions is also distorted even further because we only consider the contribution of fields that have produced in the data from 1970 to 2014. There are many fields that have yet to come online, and the costs of these fields, based on Rystad’s estimates for on unexploited oil reserves, are reported as being considerably cheaper in Saudi Arabia and Kuwait than in the rest of the world. However, incorporating these fields would require us to take a different approach to measurement, relying strongly on the accuracy of Rystad’s cost forecasting model.

<sup>61</sup>Another option would be to use a fully specified structural model of OPEC and other producers to simulate this forward. Among other thing, this would require capturing, in a parsimonious model, the geo-political aspects of OPEC, and world oil production generally, which seems beyond the scope of the current exercise. Instead, we opt for an approach which introduces a clear conservative bias.

model would have its production rise to 41 percent. This pattern of substantial underproduction relative to the competitive outcome can be found across all OPEC members around the Persian Gulf, with Kuwait having a competitive share of 16 percent, but an actual share of 3 percent, while these numbers are 9 and 6 percent respectively for Iran, and the OPEC countries around the Persian Gulf would produce 74 percent of the world's oil in 2014 under a competitive model, versus the 26 percent they produce in actuality. However, we predict that OPEC members outside the Persian Gulf would only produce 4 percent of global output in a competitive outcome versus the 14 percent that they currently produce, while the market share of non-OPEC countries is 21 percent in a competitive outcome versus an actual market share of 61 percent. It is not surprising that the market share of OPEC members would be much higher under competition: the cartel exercises market power by restricting output. Moreover, it is not clear if OPEC members outside the Persian Gulf can be effectively punished for violating the cartel arrangement, so we might expect these countries to restrict their output less.

In Figure 5 we plot the yearly productive inefficiency over our sample period. In particular, in each year we calculate the productive inefficiency by comparing the actual cost of production to the optimal one. We decompose the portion of this inefficiency that is due to within country misallocation of production, between country misallocation of production inside and outside of the OPEC cartel, and finally misallocation of production between OPEC and non-OPEC members.

In 2013, the total cost of production was 240 billion dollars, while the optimal cost of production would have been 121 billion dollars. Of the 119 billion dollars of excess costs, 37 billion occurs because of misallocation within each country, which we have little to say about. However, 33 billion dollars of this misallocation occurs because too much output is produced outside the OPEC cartel. Furthermore, 7 billion dollars of misallocation is due to misallocation of production between countries inside OPEC, such as too much output produced by Venezuela and too little by Saudi Arabia, and is likely due to the organization of the OPEC cartel. Thus, in 2013, about one third of productive inefficiency is due to cartelization of the oil industry.

Moreover, the pattern of misallocation reported in Figure 5 is instructive. Total inefficiency peaks in 1982 and in 2013, echoing the periods when prices for oil rise to 100 dollars a barrel. This is an essential ingredient into our approach: whenever demand for oil increases, and there is little response from OPEC members, which is indicative of productive distortions in the market.

This calculation does not, of course, take into account the essential dynamics inherent to the oil market, due to the finite reserves of any given oil field. In particular, we would argue that the static counterfactual is surely overstating the welfare effects of OPEC, since the low cost reserves in OPEC cannot be extracted indefinitely. How-

ever, the cross-sectional and time-series pattern of the static analysis is informative as to when the presence of market power affect welfare the most, through productive inefficiency. We turn to the calculation of the proper dynamic productive inefficiency next.

## 5.5 Dynamic productive inefficiency

We now turn to the calculation of the dynamic productive inefficiency per our definition. We calculate the net present value of the cost of production of the entire observed quantity path in our, 1970-2014, starting in 1970. We also consider a longer time period, 1970-2050, where we forecast our demand for oil with 1.3 p.a. growth rate.<sup>62 63</sup>

### 5.5.1 Unconstrained Counterfactual

Table 7 presents the counterfactual results, both for the period 1970 to 2050 in column 1, and 1970 to 2015 in column 2, with bootstrapped standard errors in parenthesis. Before we discuss these welfare effects, it is useful to simply plot the competitive counterfactual market share of OPEC, and compare it to the actual market share, which we do in figure 6. The entire analysis is precisely to let market shares reallocate (both intra- and inter-temporally) to minimize total cost of production. Given that we rely, in the baseline results, on a constant marginal cost estimate, the driving force behind the cost decrease comes from the market share reallocation. Notice that in competitive equilibrium, OPEC would produce over ninety percent of output in the earliest years, from 1970 to 1980. Diving more deeply into which fields would produce in the competitive equilibrium in these early years, over ninety percent of world output in the 1970s would come from three fields: Ghawar Shedgum and Ghawar Uthmaniyah in Saudi Arabia, and Greater Burgan in Kuwait (see Table A.9 in the appendix). Indeed, it is Saudi Arabia and Kuwait that are thought to be the leading members of OPEC.

In column 2 of table 7, the first row (actual) presents the cost of oil extracted from 1970 to 2014, using our estimates of marginal costs, as 2.29 trillion 2014 dollars, from the perspective of an agent in 1970.<sup>64</sup> For the period 1970 to 2050, this cost is 2.73 trillion dollars. However, our prediction from the competitive equilibrium for this industry would be a far lower cost of 1.27 trillion and 2.00 trillion, respectively. This implies a cost savings between 1.02 to 0.73 trillion dollars, depending on the time frame. This is a sizable magnitude, in particular if we compare this to the overall value of production across the time period.

---

<sup>62</sup>In practice, given our estimates of reserves for fields that have produced between 1970 and 2014, that is not counting untapped fields, oil production ceases in about 2035.

<sup>63</sup>1.3 p.a. is the geometric average of growth of oil demand from 1970 to 2014.

<sup>64</sup>The sum of costs of oil extraction from 1970 to 2014 are 3.07 trillion.

### 5.5.2 Constrained Counterfactuals: decomposing distortions

The above comparison is useful to quantify the wedge between observed and optimal (given our model setup) cost of production. However, measured costs could reflect a host of other constraints or distortions such as actual technological and economic constraints not appropriately measured in our cost estimates, or other distortions related to policy.<sup>65</sup> MIGHT BE BETTER IN THE STATIC SECTION SAY For example, violations of cost minimization inside a region, e.g. OPEC, do affect the total cost of production observed in the market, but are not inducing productive inefficiencies, in the sense of the definition employed in our approach. This particular failure of cost minimization is of course related to the market power of OPEC, but it is best seen as a failure of coordination within the cartel. We refer to this inefficiency as *cartel inefficiency*. It, of course, negatively affects welfare and it is relevant in the overall calculation of damages due to market power. The same observation can be made for failures of cost minimization across regions of the rest of world. Take for example the unmeasured cost of delivering oil to the refineries, which are not fully taken into account in our measures of cost.<sup>66</sup> We therefore engage in a variety of model-based decompositions where we introduce constraints in the total production of a given region or country.

The first constraint we consider is to hold fixed the total quantity produced by the members of OPEC, as observed in our data in each year. We therefore let the total cost of production be optimized within, and outside OPEC, and we compare this benchmark, Optimal OPEC, to the Optimal as defined above. This comparison allows to decompose the difference between the actual and the optimal into three distinct components: 1) cartel inefficiency, 2) productive inefficiency due to market power, and 3) distortions outside OPEC. The second constraint performs a similar counterfactual but now holds every country's production as observed in the data, and such we can decompose the failures of cost minimization within and across countries. Finally, we hold fixed the total production of Saudi Arabia. As discussed in section XX, all oil fields in Saudi Arabia are under full ownership of the nationalized oil company Saudi Aramco, and with a 13 percent share of world production, it is expected to impact market-wide aggregates.

When we consider constrained counterfactuals we no longer compare the simulated costs to the actual cost (the first row in Table 7), but compare the constrained optimal to the unconstrained optimal. As expected, the implied productive inefficiency is smaller, as we no longer contribute all the departures from cost minimization to market power. In the case of the OPEC quantity constrained we now obtain a productive

---

<sup>65</sup>These distortions are in fact featured more prominently than market power, the focus of our analysis, in the literature on misallocation.

<sup>66</sup>While we do observe transportation cost, it does not capture the full cost of delivery to the final demand – i.e. the refinery.

inefficiency of 120 billion USD (2.12 – 2.00). When considering country-level quantity distortions the inefficiency increases to 320 billion. The size of the productive inefficiency under the Saudi constraint is, as expected, the lowest, and amounts to 70 billion dollars. Moreover, the inefficiency attributable to the cartel including both the constraint on production attributed to OPEC, and the constraint on production between members of OPEC in the Persian Gulf and outside the Gulf, adds up to 160 billion dollars, about an 8 percent increase in costs over optimal costs.

## 6 Alternative specifications and robustness

We discuss alternative specifications of the cost function and verify whether our main results are sensitive to those. In particular, we consider the presence of start-up cost and a non-linear marginal cost schedule, and we revisit our dynamic counterfactual simulations as above.

### 6.1 Start-up Costs

We distinguish start-up cost from (per period) fixed cost of operation. The latter are appropriately accounted for in our measurement of (unit) cost.<sup>67</sup> Start-up cost, however, become a potential complicating factor for the baseline analysis when fields require a large initial capital expenditure to come online. Expenditures of this form, linked to “switching the field on”, are one time expenditures that are sunk in the first year of operation (and possibly in the year prior to operation). In this section, robustness to the existence of these start-up costs is evaluated.

Start-up costs can be an issue in two parts of the analysis. First, they may alter the counterfactual production path. That is, a high start-up cost, low marginal cost, field may delay its initial production date in a competitive equilibrium relative to that predicted in the baseline model. Second, these fixed costs are welfare relevant and, to the extent to which they are not being counted in the measurement, they may offset the otherwise conservative nature of the calculations executed using the baseline model.

As a preliminary measure, the extent to which start-up costs are potentially an issue is investigated by examining several dimensions of the data. First, we examine the proportion of production between 1970 and 2014 that is provided by fields that were producing in 1970. For these fields all start-up costs will already be sunk, and these fields collectively are responsible for 628,866 million bbl out of the 1,081,171 million bbl produced over the entire sample period (1970-2014), or about 58% of total production.

---

<sup>67</sup>See section X.

In terms of total cost of production, they account for over 45%. This implies that for a large part of the fields, the presence and magnitude of any startup cost is irrelevant.

Second, for the fields that became active during our sample period, we observe recorded facility and well capital expenditures, and other cost components, before they actually produce oil. This suggests a measure of startup cost: the sum of expenditures prior to production – i.e.  $C_f^S = \sum_{s=1}^{p_f} C_{f,t-s}$  and  $C_{f,t-s} = \sum_k C_{f,k,t-s}$  and  $k$  indexes the various cost expenditure categories as discussed in section XX, and  $p_f$  is the first period we observe expenditures of a given field  $f$ . While this measure of startup cost is clearly subject to a sample selection bias – i.e., we only observe positive startup cost for 43,363 field-year observations, out of a total of 405,112 field-year observations. However, conditional on this sample selection, the variation in startup costs is as expected: e.g. offshore deep water startup costs are the highest among the various technologies. While the magnitude of the startup cost matter, for our counterfactual simulations it is in particular relevant whether fields of lower (unit) cost are subject to lower startup cost – i.e., whether the correlation between cost of production and startup cost is positive or negative. The correlation is computed by correlating the startup cost,  $C_f^S$ , to the total cost of production over a field’s production lifetime, over 13,387 oil fields, and it is 0.47. The Spearman rank correlation, on the same sample, is 0.86, and suggests a very strong monotone relationship between a field’s initial startup cost and subsequent total cost of production.<sup>68</sup> The strong positive correlation, both in magnitudes and in ranking, implies that the existence of startup costs will not affect the sorting algorithm employed to simulate the counterfactuals in our analysis. The sorting algorithm guarantees that the low-cost fields are called to produce first, and these fields are also characterized with lower startup costs, which preserves the ordering, and the optimal dynamic path of production.

We further assess sensitivity to start-up costs by re-computing the baseline welfare figures in a number of ways. First, we recompute the baseline, forcing each field to come online in the first year of observed production. In effect, this means that all fields in the counterfactual are assumed to sunk any start-up costs in the same years as in the actual data. Thus, all the returns to reallocation occur on the intensive margin. This shuts down any welfare gains from reallocation on the extensive margin (i.e. moving production across time). As such it is a lower bound on the actual returns from reallocation, taking into account any start-up costs that may exist. It is useful to note that this bound does not require knowledge of the exact level of these costs since, by construction, in comparing observed paths to any counterfactual the fixed costs will exactly cancel.

Figure 8 investigates the size of these fixed costs by computing, within a field, cu-

---

<sup>68</sup>This relationship is robust to including country, technology and total production of the field, in a regression framework.

mulative costs and cumulative production. The magnitude of these startups costs can be identified from the proportion of expenditures incurred prior to the start of production. The figure then aggregates this relationship weighting by field level production, and breaking out fields by onshore, offshore shelf, offshore deepwater, and shale/oil sands. For a conventional onshore field, about 20 percent of costs are incurred before the first barrel is produced, while for a offshore deepwater field, this number is closer to 40 percent, which is as expected, since offshore oil fields have much larger startup costs from building oil platforms. In contrast, shale has much smaller startup costs (even if this relationship is far noisier given the more limited number of shale fields), which is logical given that many individual wells are drilled in the same field, and extraction time is relatively short, on the order of two years versus the ten to fifteen years of production for a typical offshore oil platform. Our assumption of constant marginal costs would be born out exactly if the cumulative costs were on the 45 degree line. Figure 8 shows that at the end of a field's production, also indicates that there is some increasing marginal costs at the end of a field's life, as the right tail of the cost distribution is slightly convex.

## 6.2 Marginal cost curvature

The quantification of the productive inefficiency relies on a measure of marginal cost that only varies over *time*, and not cumulative quantity, to the extent that there are common technology-year shocks to cost, the  $\mu_{s,t}$  in equation (A.6). An obvious concern is that any given field cost schedule is up (or down) ward sloping over a certain range of quantities. Before we consider an alternative measurement model for cost, and discuss the associated counterfactuals, it is important to keep in mind two features of our approach. First, we consider the entire global market for crude oil, and aggregate production across thousands of oil fields, which implies that within-field curvature is likely to be less important. Second, we already impose an extreme form of curvature by limiting the speed of extraction, albeit common across all oil fields. There is, however, still room for within-field cost curvature to affect the optimal allocation of production, and therefore impact the extent of the productive misallocation.

IDEAS: - DEMONSTRATE MOST OF VARIATION IS GONE AFTER  $\mu_{s,c,t}$  - ADD COUNTRY AND TRY COHORT - ADD COHORT/COUNTRY DUMMIES IN COST EXPRESSION - ESTIMATE CURVATURE BY TECHNOLOGY

## 6.3 Uncertainty

this has to be about differences in expectations of future prices, or uncertainty simply in the cost process, which again we ruled out except for common  $\mu_{s,t}$  and that is assumed to be a martingale.

-TRUNCATING THE DISTRIBUTION OF COST -WITHIN COUNTRY STUFF: OTHER STORIES

- should also have a better discussion of fixed v marginal in the initial discussion of costs/production/data

## 6.4 Discount Factor

Proposition 1 goes through if discount factors are negatively correlated with cost of production. I.e. if fields with lower cost have a lower discount factor (value the future less), the sorting algorithm is preserved. Delta goes down as cost goes down  $-\frac{\partial \delta}{\partial c} > 0$ .

## 6.5 Robustness

Table 8 shows our dynamic counterfactuals for various alternative modeling assumptions, such as alternative measures of costs, alternative assumptions on extraction rates and reserves, and different assumptions on measuring reserves. The results discussed above are both qualitatively and quantitatively robust to changing the maximum extraction speed (up and down), lowering and increasing the proven reserves, and using time-invariant estimates of cost of production. The only results which change substantially concern maximal extraction rates: our results become very sensitive once we look at an extraction rate of 2 percent of reserves (low extraction rate), since the sorting algorithm needs to resort to very high cost oil fields. Moreover, we can obtain actual costs that are lower than optimal costs if low cost fields have extraction rates above 2 percent in the data.

## 7 Conclusion

This paper provides estimates of the extent to which market power is a source of production misallocation. This production misallocation, and the resulting welfare loss, occurs through more production being allocated to higher cost (less productive) units of production, and less production to the most efficient production units in the economy. In quantifying the potential magnitude of this productive inefficiency, we examine the global market for crude oil, in which OPEC, a notable international cartel, is alleged to exert considerable market power at times.

We find evidence for substantial productive inefficiency due to market power, of which we attribute 160 billion dollars to the activities of the cartel, depending on model specification. The results from this study indicate that market power can affect aggregate outcomes, here the total cost of production in the world oil market and subsequent

the price of oil, which in turn affects a host of economic decisions and macroeconomic aggregates.

## References

- Anderson, Soren T, Ryan Kellogg, and Stephen W Salant**, "Hotelling under pressure," *Journal of Political Economy*, 2014, *forthcoming*.
- Asker, John**, "A study of the internal organization of a bidding cartel," *The American Economic Review*, 2010, pp. 724–762.
- , **Allan Collard-Wexler, and Jan De Loecker**, "Dynamic Inputs and Resource (Mis)Allocation," *Journal of Political Economy*, 2014, 122 (5), 1013–1063.
- Atkeson, Andrew and Ariel Tomás Burstein**, "Innovation, Firm Dynamics, and International Trade," *Journal of Political Economy*, 2010, 118 (3), 433–484.
- Collard-Wexler, Allan and Jan De Loecker**, "Reallocation and Technology: Evidence from the US Steel Industry," *American Economic Review*, 2015, 105 (1), 131–171.
- Crémer, Jacques and Djavad Salehi-Isfahani**, *Models of the oil market*, Vol. 2, Taylor & Francis, 1991.
- De Loecker, Jan**, "Product differentiation, multiproduct firms, and estimating the impact of trade liberalization on productivity," *Econometrica*, 2011, 79 (5), 1407–1451.
- Edmond, Chris, Virgiliu Midrigan, and Daniel Yi Xu**, "Competition, Markups, and the Gains from International Trade," *American Economic Review*, 2015, 105 (10), 3183–3221.
- Fattouh, Bassam**, "The drivers of oil prices: the usefulness and limitations of non-structural model, the demand-supply framework and informal approaches," *Oxford Institute for Energy Studies Working Paper*, 2007.
- Goldberg, Pinelopi Koujianou, Amit Kumar Khandelwal, Nina Pavcnik, and Petia Topalova**, "Imported Intermediate Inputs and Domestic Product Growth: Evidence from India," *The Quarterly Journal of Economics*, 2010, 125 (4), 1727–1767.
- Hopenhayn, Hugo**, "Firms, Misallocation, and Aggregate Productivity: A Review," *Annual Review of Economics*, 2014, 6, 735–770.
- Kellogg, Ryan**, "The effect of uncertainty on investment: evidence from Texas oil drilling," *The American Economic Review*, 2014, 104 (6), 1698–1734.
- Kilian, Lutz**, "Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market," *American Economic Review*, 2009, 99 (3), 1053–1069.
- Kohl, Wilfrid L**, "The Perfect Storm," *Harvard International Review*, 2005, 26 (4), 68.
- Marshall, Robert and Leslie Marx**, *The Economics of Collusion*, Cambridge, MA: MIT Press, 2012.

**Martinez Licetti, Martha and Tanja Goodwin**, “Bad news for cartels, good news for the poor in Latin America,” ["http://blogs.worldbank.org/psd/bad-news-cartels-good-news-poor-latin-america-0"](http://blogs.worldbank.org/psd/bad-news-cartels-good-news-poor-latin-america-0) 2015.

**Melitz, Marc J**, “The impact of trade on intra-industry reallocations and aggregate industry productivity,” *Econometrica*, 2003, 71 (6), 1695–1725.

**Organization of Petroleum Exporting Countries**, “Monthly Oil Market Report,” June 2008.

**Sweeney, James**, “Economic theory of depletable resources: An introduction,” in Allen Kneese and James Sweeney, eds., *Handbook of Natural Resource and Energy Economics*, Vol. 3, CHECK, 1992, chapter 17.

**Syverson, Chad**, “Product substitutability and productivity dispersion,” *Review of Economics and Statistics*, 2004, 86 (2), 534–550.

– , “What Determines Productivity?,” *Journal of Economic Literature*, 2011, 49 (2), 326–365.

**Yergin, Daniel**, *The prize: The epic quest for oil, money & power*, Simon and Schuster, 1991.

# 8 Tables and Figures

Figure 1: Production distortions due to market power





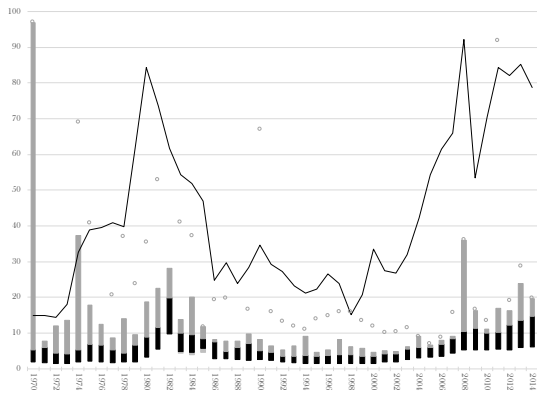
Figure 2: OPEC market share and oil price (1970-2014)




Notes: The vertical axis on the left is in dollars and corresponds to the annual average oil price, which is indicated by the black line. This price series is deflated with the US GDP deflator (base year 2009). The OPEC market share in each year is indicated by the this, dashed, black line. The vertical axis on the right indicates the level of the market share. Countries are included OPEC in all years if they had ever had active membership between 1970 and 2014

Figure 3: Production costs and price (1970-2014):  
OPEC countries


(a) Saudi Arabia




(b) Kuwait



(c) Venezuela




(d) Nigeria




Notes: Each panel plots the dispersion of the costs of production (by barrel) in a country, and the price of oil. The vertical axis is \$/barrel, from 0 to 100 in increments of 10. The horizontal axis is in years, from 1970 to 2014. Costs are indicated by the bars and circles. The (grey and black) bar indicates the range of costs within the 1st and 99th percentile of production. That is, the cheapest, and most expensive, 1% of barrels produced in the year are excluded. The black portion of the grey indicates the 5th to 95th percentile range. Circles indicate the maximum cost per barrel incurred in a year. Where a cost exceeds \$100 per barrel, it is not shown (the vertical axis is truncated at 100) – this accounts for many of the maxima not being visible, for instance. The oil price is indicated by the black line. All series have been deflated with the US GDP deflator (base year 2009). All costs are measured according to the baseline specification

Figure 4: Production costs and price (1970-2014):  
Other Countries


(a) US




(b) Russia



(c) Canada




(d) Norway




Notes: The notes for Figure 3 also apply to this figure. The Norwegian panel reflects little meaningful production prior to 1978.

Figure 5: Decomposing Static Distortions




Note: Static Distortions for each year presented in 2014 dollars, with the total height of each bar representing the difference between the actual cost of production and the optimal cost of production. Within country distortion is the difference between the actual cost of production and the optimal cost of production holding country level production fixed. Between country distortion is the difference between the optimal cost of production holding country level production fixed, versus the optimal cost holding either the production inside or outside of OPEC fixed. OPEC versus Non-OPEC distortion is the difference between the cost of production holding production inside and outside OPEC fixed, versus optimal production.



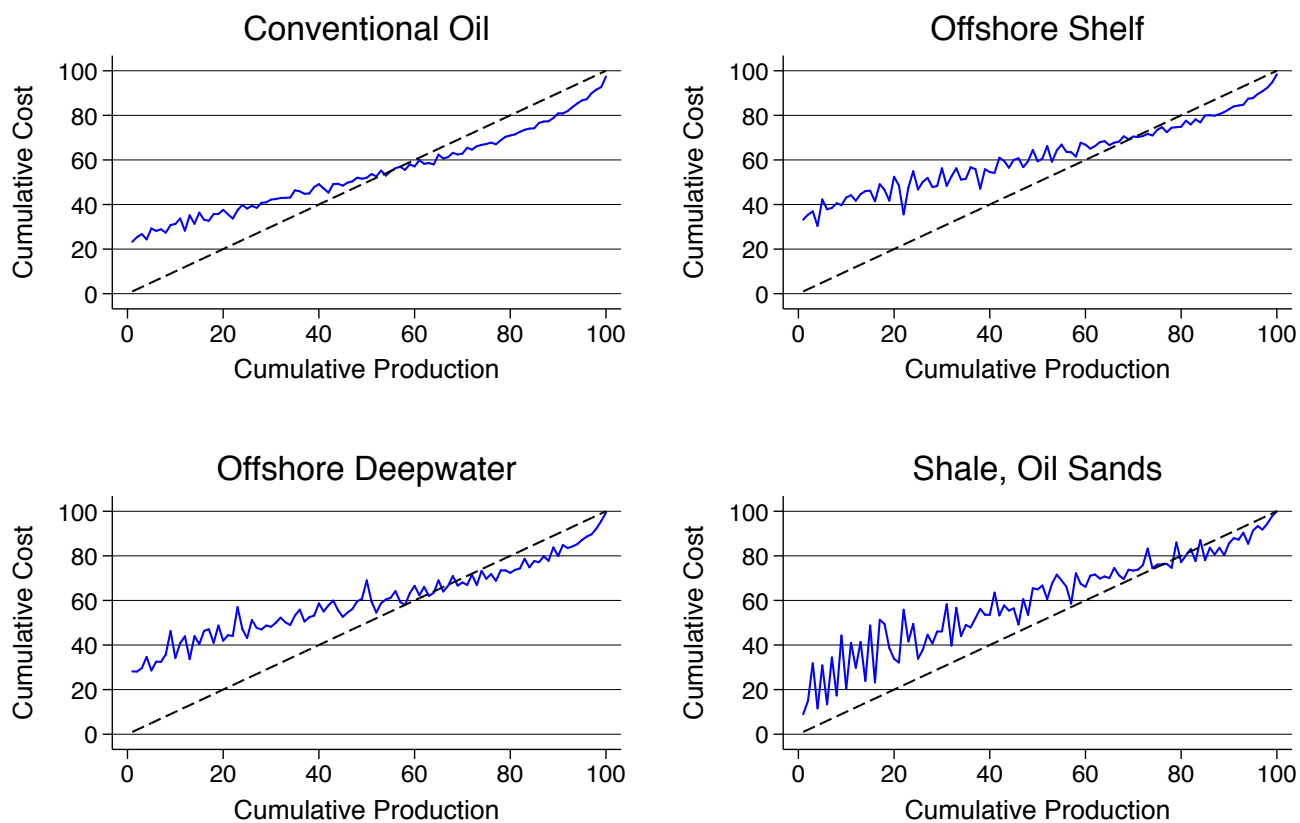
OPEC Share - Counterfactual presents the share of production accounted for by OPEC.

Figure 6: OPEC Market Share 1970-2014



Note: Share of total production accounted for by fields that start production after 1970.

Figure 7: Production of Entrant Fields



Note: Cumulative Production is measured as cumulative production for a field in year  $t$  divided by the total production observed over a field's lifespan. Cumulative Costs are defined likewise for costs. Conventional Oil is 72% of production, Offshore Shelf is 21%, Offshore Deepwater is 6%, Shale is 1%. Only fields that start producing after 1970 are used for this figure.

Figure 8: Costs over the Field Lifecycle

Table 1: Largest crude producers, % of global production 1970-2014

OPEC		Non-OPEC	
Saudi Arabia	11.8%	United States	14.4%
Iran	5.4%	Russia	13.0%
Venezuela	3.8%	China	4.1%
UAE	3.1%	Mexico	3.7%
Nigeria	2.8%	Canada	3.3%
Iraq	2.7%	UK	2.4%
Kuwait	2.6%	Norway	2.4%

Notes: Global production from 1970-2014 was 1,156 billion barrels. Collectively these 14 countries account for 85.4% of global production.

Table 2: Summary statistics, by field-year

Variable	mean	median	5%	95%
Field-year characteristics:				
Production (mB/year)	3.43	0.22	0.00	10.92
Reserves (mB)	99.49	3.71	0.03	239.78
Discovery Year	1965	1967	1911	1999
Startup Year	1971	1974	1916	2005
Off-shore	0.19			
Costs: (\$m)				
Exploration Capital Expenditures	0.61	0.00	0.00	0.41
Well Capital Expenditures	9.10	0.49	0.00	35.32
Facility Capital Expenditures	5.14	0.21	0.00	16.85
Production Operating Expenditures	10.41	0.46	0.00	38.47
Transportation Operating Expenditures	2.27	0.13	0.00	7.01
SGA Operating Expenditures	2.65	0.22	0.00	8.85
Taxes Operating Expenditures	1.41	0.00	0.00	1.09
Royalties	18.19	0.40	0.00	45.36
Government Profit Oil	15.59	0.00	0.00	21.00

Notes: Only fields with active production during 1970-2014 are included. There are 66,920 fields in the Rystad data. 13,248 of these fields have active production. Reserves data exists for 13,298 fields. As a result in section 5 11,457 fields are used. All numbers are in \$US deflated by the US GDP deflator for 2009. mB indicates million barrels. The unit of observation is the field-year.

Table 3: Definitions of cost components

Exploration Capital Expenditures:	Costs incurred to find and prove hydrocarbons: seismic, wildcat and appraisal wells, and general engineering costs.
Well Capital Expenditures	Capitalized costs related to well construction, including drilling costs, rig lease, well completion, well stimulation, steel costs and materials.
Facility Capital Expenditures	Costs to develop, install, maintain and modify surface installations and infrastructure.
Abandonment Cost	Costs for decommissioning a field.
Production Operating Expenditures	Operational expenses directly related to the production activity. The category includes materials, tools, maintenance, equipment lease costs and operation related salaries. Depreciation and other non-cash items are not included.
Transportation Operating Expenditures	Represents the costs of bring the oil and gas from the production site/processing plant to the pricing point (only upstream transportation). The category includes transport fees and blending costs.
SGA Operating Expenditures	Operating expenses not directly associated with field operations. The category includes administrative staff costs, office leases, related benefits (stocks and stock option plans) and professional expenses (legal, consulting, insurance). Only exploration and production related SG&A are included.
Taxes Operating Expenditure	Local US taxes that are directly related to production. The category includes ad valorem taxes (county based) and severance taxes (state based).
Royalties	The sum of all gross taxes, including royalty and oil and export duties.
Government Profit Oil	The production sharing agreement equivalent to petroleum taxes, but paid in kind (that is, the government contracts with a company to develop and operate the field, but retains rights a proportion of the production). Government Profit Oil reduces the company's entitlement production and is thus treated as a royalty effect in company reports.

Source: Rystad U-Cube External Use Documentation.

Table 4: Reserves and production, 2014

	Reserves (mB)	Share of world reserves (%)	Reserves/(Annual production) (%)
Non-OPEC	218,054	50	10
OPEC	220,561	50	19
Saudi Arabia	74,194	17	18
Kuwait	15,723	4	16

Data are for 2014. Total reserves for the world in 2014 were 438 billion barrels. The ratio of reserves-to-production was 14. OPEC countries are listed in section 3.3.2. Countries are included OPEC in all years if they had ever had active membership between 1970 and 2014. Reserves are reported using P50 measures at a world price of \$70 per barrel.

Table 5: Unit costs across the global oil industry (1970-2014)

	1970-1979	1980-1989	1990-1999	2000-2014
Number of active fields	4,766	7,088	9,760	12,085
Mean oil price	20	40	21	59
Mean global production (mB/year)	20,861	21,489	23,984	26,298
OPEC	9,979	7,289	9,606	11,249
Mean global reserves (mB)	737,928	728,532	661,815	517,559
OPEC	392,912	365,891	328,914	254,730
Unit costs (Baseline specification):				
95th percentile Saudi Arabia	5.8	13.6	4.4	10.4
Median Saudi Arabia	2.3	5.6	2.3	5.4
95th percentile OPEC	6.7	18.6	7.6	20.1
Median OPEC	2.4	5.9	2.8	6.1
95th percentile non-OPEC	6.7	15.6	9.2	28.2
Median non-OPEC	3.6	7.0	4.1	9.7
Unit costs (including taxes and royalty payments):				
95th percentile Saudi Arabia	5.8	13.6	4.4	10.4
Median Saudi Arabia	2.3	5.6	2.3	5.4
95th percentile OPEC	30.2	53.6	21.1	79.1
Median OPEC	2.8	13.6	6.5	12.0
95th percentile non-OPEC	26.3	40.1	20.3	75.3
Median non-OPEC	9.1	14.8	9.1	24.0

Notes: The unit cost is computed as per section 3.4 and top-coded at \$1,500. The unit of observation for unit cost is the barrel. Percentiles and medians are calculated at the barrel level. All prices and costs are deflated with the US GDP deflator (base year 2009). Reserves are reported using P50 measures at a world price of \$70 per barrel. OPEC countries are listed in section 3.3.2. Countries are included OPEC in all years if they had ever had active membership between 1970 and 2014. Only fields active between 1970 and 2014 are included.

Table 6: Static Counterfactual for 2014

<u>Country</u>	<u>Actual Output Share</u>	<u>Competitive Output Share</u>
<u>Persian Gulf OPEC</u>	0.258	0.744
Iran	0.057	0.091
Iraq	0.029	0.069
Kuwait	0.030	0.155
Qatar	0.009	0.015
Saudi Arabia	0.133	0.414
United Arab Emirates	0.031	0.075
<u>Other OPEC</u>	0.135	0.044
Algeria	0.021	0.015
Indonesia	0.020	0.002
Libya	0.025	0.012
Nigeria	0.028	0.006
Venezuela	0.041	0.009
<u>Non-OPEC</u>	0.607	0.212
Brazil	0.014	0.001
Canada	0.023	0.006
China	0.045	0.002
Kazakhstan	0.010	0.000
Mexico	0.023	0.013
Norway	0.027	0.009
Russia	0.144	0.047
United Kingdom	0.022	0.001
United States	0.132	0.013
Rest of the World	0.136	0.044

Note: Rest of the world refers to countries that are outside the top 20 largest oil producers between 1970 and 2014.

Table 7: Dynamic Counterfactual Results

<u>Total Cost in Trillions of Dollars</u>	<u>1970-2100</u>	<u>1970-2014</u>
Actual	2.73 (0.19)	2.29 (0.12)
Optimal	2.00 (0.16)	1.27 (0.08)
Optimal Cost Holding OPEC Allocation Fixed	2.12 (0.18)	1.42 (0.08)
Optimal Cost Holding OPEC Persian Gulf and OPEC not Persian Gulf Fixed	2.16 (0.17)	1.48 (0.08)
Optimal Cost Holding Country Allocation Fixed	2.32 (0.18)	1.67 (0.08)
Optimal Cost Holding Saudi Allocation Fixed	2.07 (0.17)	1.35 (0.08)
Ratios with resp Ratio Actual to Optimal	37% (7)	72% (14)
Ratio Optimal to Optimal OPEC Allocation Fixed	6% (2)	12% (4)
Ratio Optimal to Optimal OPEC Persian Gulf and OPEC not Persian Gulf Fixed	8% (2)	17% (4)
Ratio Optimal to Optimal Country Allocation Fixed	16% (3)	31% (6)
Ratio Optimal to Optimal Saudi Allocation Fixed	4% (2)	6% (3)

Notes: Bootstrapped standard errors in parenthesis. Total Costs from 1970 to 2015 in Net Present Value (assuming a 5 percent discount rate) in 2014 dollars. Country allocation fixed holds the total production per country by year fixed, OPEC allocation holds the production of OPEC and non-OPEC countries fixed by year, and likewise for Saudi and Gulf (Gulf is Saudi Arabia, Kuwait, United Arab Emirates, Iraq, Iran, Qatar). We use an extraction rate of 10 percent of reserves, the p50 measures of reserves (proven and probable), a demand growth rate of 1.3 percent per year.

Table 8: Counterfactual Results: Robustness

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Base	$c_f$	taxes	highextr	lowextr	noextrac	lowres	highres	start date
<u>1970-2050 Counterfactual</u>									
<u>Cost in Billions of dollars</u>									
Actual Cost	2734	3047	4876	2758	2656	2752	2733	2733	2725
Optimal Cost Country	2317	2540	4134	2276	2330	2288	2315	2315	2363
Optimal Cost OPEC Fixed	2118	2287	3486	2056	3128	2074	2116	2115	2180
Optimal Cost	2001	2143	3241	1961	2065	1986	2000	1999	2040
Optimal Cost SA	2072	2232	3583	2007	2183	2034	2071	2070	2117
Optimal Cost AP	2106	2276	3717	2042	2251	2063	2104	2104	2161
Optimal Cost ME	2162	2345	3661	2092	2710	2108	2160	2160	2230
Optimal Cost InOut	2164	2347	3718	2101	3212	2114	2162	2161	2232
<u>Difference as Percent of Optimal Cost</u>									
Difference OPEC-Optimal	6	7	8	5	52	4	6	6	7
Difference Country-Optimal	16	19	28	16	13	15	16	16	16
Difference SA-Optimal	4	4	11	2	6	2	4	4	4
Difference AP-Optimal	5	6	15	4	9	4	5	5	6
Difference ME-Optimal	8	9	13	7	31	6	8	8	9
Difference Actual-Optimal	37	42	50	41	29	39	37	37	34
Difference InOut	8	9	15	7	56	6	8	8	9
<u>1970-2014 Counterfactual</u>									
<u>Cost in Billions of dollars</u>									
Actual Cost	2188	2530	3938	2188	2188	2188	2188	2188	2187
Optimal Cost Country Fixed	1666	1924	3006	1632	1853	1635	1665	1665	1723
Optimal Cost OPEC Fixed	1417	1630	2096	1361	2684	1370	1417	1417	1477
Optimal Cost	1269	1457	1853	1229	1587	1241	1268	1268	1313
Optimal Cost SA Fixed	1348	1552	2284	1297	1714	1313	1347	1347	1393
Optimal Cost AP Fixed	1396	1609	2466	1346	1787	1361	1395	1395	1452
Optimal Cost ME Fixed	1480	1705	2352	1415	2267	1428	1480	1480	1545
Optimal Cost InOut	1482	1707	2415	1423	2773	1430	1481	1481	1547
<u>Difference as Percent of Optimal Cost</u>									
Difference OPEC-Optimal	12	12	13	11	69	10	12	12	12
Difference Country-Optimal	31	32	62	33	17	32	31	31	31
Difference SA-Optimal	6	7	23	6	8	6	6	6	6
Difference AP-Optimal	10	10	33	10	13	10	10	10	11
Difference ME-Optimal	17	17	27	15	43	15	17	17	18
Difference Actual-Optimal	72	74	113	78	38	76	73	73	67
Difference InOut	17	17	30	16	75	15	17	17	18

Notes:  $c_f$  uses a time invariant cost estimate, taxes includes distortionary taxes in our measure of economic costs, lowextr, highextr, and no extraction rate use a 0.02, 0.5, and 1 extraction rate, lowres uses a forecasted price of 70 dollars a barrel and p90 reserve measure, highres uses a forecasted price of 130 dollars a barrel and pmean reserve measure, startd constraints fields to only be useable after they started production in the data rather than their time of discovery.

# **Online Data Appendix to *Market Power and Production (Mis)Allocation: A Study of the World Oil Market* – Not For Publication**

## **A.1 Data collection and sources**

Compiling data on the oil industry involves confronting issues common to evaluating the performance of any global industry. Oil extraction is completed by many different government and non-government entities (companies and otherwise), across many different countries, in a wide variety of geological and environmental settings (e.g. on-shore and off-shore extraction). In addition, sales of crude oil are made between many different buyers and sellers in a largely decentralized market. As a consequence, there is no centralized data collection protocol that leads to a unified dataset of the sort commonly used in industry studies that focus on a specific product class in a specific geographic location, or studies based on census datasets collected in some countries by statistical agencies.

### **A.1.1 Collection**

In this study, we employ the data sold on the commercial market by one such firm, a Norwegian consultancy called Rystad. Various parts of Rystad’s data product has been used in other economic studies, including.... [CITE IMF, GREENSTONE, OTHERS? WHO ARE THEIR CLIENTS]. In this data appendix we describe the Rystad data, and in particular how Rystad collects the data, measures the variables of interest.

We obtain data under a license from Rystad for their U-Cube (Upstream Cube) Database. In this section we provide a short description of how this dataset is assembled by Rystad. The U-Cube is a bottom-up database: one that starts from individual oil fields and aggregates them up to obtain country and global production.

The data from Rystad concerns a large number of fields from 1970 to 2014. Since this dataset is used to forecast future oil prices, it includes currently producing fields, but also fields that could possibly come online in the future. Table A.1 shows the data cleaning steps we perform, along with the total number of fields after each step. The data has 66,000 unique fields, but only 21,000 produce at any point in time. These non-producing fields are used by Rystad for forecasting purposes — they are estimates of the production of a field that has not started production yet. Of the 21,000 producing fields, about 2,000 are gas only fields which we drop from our analysis. Linking these fields to measures of reserves leaves us with 13,000 fields, but 92 percent of production.

### **A.1.2 Data Sources**

There are two types of data that are used to construct field-level data. Geological and Lifecycle data, and economic data. Rystad keeps a database of the type and geology of the field. Some of this information is about the physical aspects of the assets such as whether the asset is an oil sands, shale, or an offshore oil platform of a certain depth, or whether the oil field produces gas,

or heavy or sour crude oil. This information is complemented with more detail on the exact geology of the oil field. As well, Rystad keeps track of the discovery date and depletion of these oil fields.

Much of the field data is obtained from government and company sources. For instance, in the United States, there is detailed data on field level production from information on the royalties that oil firms pay for their leases. As well, many companies publish information on production levels and reserves in different fields. The second type of data are economic. These data are mainly sourced from three different places including company and government Reports, and information from oil service firms through documents and interviews. Table A.2 shows information on the data sources for different regions in the world. While data in the United States comes from direct measurement, data in Saudi Arabia is mainly extrapolated from the geological attributes of fields in this country.

While there is a considerable amount of data on oil fields, most of the economic and production data is extrapolated from similar oil fields. For instance, there is no data on fields in Saudi Arabia since Saudi ARAMCO does not publish information on its operations. Therefore, Rystad uses information on costs of oil fields that are comparable, principally those in Iraq, to infer the costs of production in Saudi Arabia. *this should be a deeper Saudi case study*

## A.2 Measurement of main variables

### A.2.1 Production

Table A.4 presents total production of different fuels in 2014 in millions of barrels. All units of fuel are converted to crude barrel equivalent, since different fuels like Condensate and Gas have different energy production by volume. Thus, in the paper we refer to production as a homogeneous product measured in crude oil equivalents. We drop production of gas products, namely Gas, LNG and NGL, and flared or injected Gas. While gas and oil are often recovered jointly in one oil field, we observe expenditures of the various categories broken down by oil and oil, based on the production share of both products. [ALLAN CHECK]

Some of these fields are fairly close to production units, such as individual offshore oil rigs, while other assets are quite large, such as the Ghawar field, with hundreds of rigs. Indeed, in 2014, the tenth percentile oil asset produces 8,000 barrels per year, while the 90th percentile produces 1.9 million barrels per year.

Figure A.1 shows the time series of total production from 1970 to 2014, as well as the total number of producing fields.

### A.2.2 Costs and Revenues

Revenues are computed by applying prices to each oil product, and average prices per product are shown in Table A.7 for 2014. In this year, the average price of oil is about 80 dollars a barrel, but the standard deviation of the price is 15 dollars, mainly due to geographic variation in prices. Notice as well that there is limited variation in prices across different types of oil.

Table A.5 presented total costs, revenues, and rents, for the oil industry in 2014. Revenues were 2.5 trillion dollars in 2015, of which 1.8 trillion was rents of various sorts, and 0.7 trillion

was costs. Costs are decomposed into several categories.

Exploration expenditures are the costs associated with acquiring acreage, doing seismic, and drilling wildcats or appraisal wells to discover and delineate oil and gas fields.

The capital expenditures ( Capital Expenditures) include all development costs related to facilities and drilling of wells. The initial capital expenses are related to establishing the facility and necessary infrastructure, as well as pre-drilling costs, often termed development Capital Expenditures. Throughout the field life capital expenses include drilling of more development wells (well Capital Expenditures) and modifications done to the facility, processing system or e.g. subsea infrastructure.

Facility Capital Expenditures is a sum of several costs including: Development cost - associated with the construction and implementation of the facility required for the processing and production of the field; Modification cost- expenditure related to larger maintenance and improvements required to keep the facility operational.; and, Subsea brownfield are the brown-field costs related to subsea installations. Well Capital Expenditures is the cost related to the drilling of wells required to maintain and increase the production levels.

Transportation Operating Expenditures represents the costs associated with bringing the oil and gas from the production site/processing plant to the pricing point (only upstream transportation). The category includes transport fees and blending costs.

Selling, General and Administrative Expenses represent operation expenses not directly associated with field operations. The category includes administrative staff costs, office leases, related benefits (stocks and stock option plans), and professional expenses (legal, consulting, insurance).

Production Operating Expenditures represents operational expenses directly related to the production activity. Production Operating Expenditures includes 3 elements. These elements are the lease, fixed and variable costs. The lease cost is used if the asset operator has chosen to e.g. lease an FPSO instead of developing their own. In these cases there will be no fixed cost as the Lease cost includes all the fees required for the structural operation cost. The Lease cost applies only to assets with leased FPSO (or lease penalties). The fixed cost represents the base or invariable cost associated with running the facility. These costs can be very different for offshore and onshore assets, from high offshore to in some cases negligible or nothing onshore. The variable cost is the cost associated with the actual production per barrel, such as the cost for electricity, machinery, salaries, tariffs, and other per barrel expenses. Furthermore we observe *rents*, decomposed into the following components: Firm and Government Profits, Royalties, Operating Taxes and Income.

### **A.2.3 Reserves**

Central to much of the discussion in this paper is the notion of reserves. The reserve is the unextracted, but recoverable, quantity of oil remaining in the ground in a field. The most reliable way to measure the reserve at a point in time is to see the entire production life of a field. The total extracted is the maximal reserve. In reality, most fields are not fully exploited in the data. Hence, industry reserve estimates need to be used. The oil industry reports reserves at different levels of extraction probability. There are three levels. P90 (or P1) are is the quantity able to recovered with a 90% probability given current technical and economic conditions. The P90 reserve is the asset value able to be reported on company balance sheets under U.S. GAAP.

Clearly, this definition means that reserves will fluctuate with the oil price. In the data used here, reserves are measured and reported assuming a constant oil price of 70 dollars a barrel into the future. However, Rystad also constructs reserves given a price of oil of 100 and 130 dollars a barrel, and we will present robustness checks to these different measures. P50 (or P1 + P2) are the reserves recoverable with a 50% probability. Finally P10 or (P1 + P2 + P3) are total reserves recoverable with a 10% chance. The level of P90, P50 and P10 can vary significantly within a field.

In this paper, the reserve number used for a field is P50. The precision with which reserves are measured varies by the production stage at which a field is at, and the country in which the field exists. Untapped deposits have less precise reserve estimates, since there no actual production data or well pressure data to rely on. Once a field starts producing reserves become easier to estimate particularly as pressure starts to change, as the pressure gradient of a field, conditional on geology, as the resource is depleted is relatively well understood. A further confounding factor is the oil reserves are strategic assets, with most industry sources commenting that various countries will inflate reserve figures for political reasons. This is particularly relevant for OPEC countries due to the way OPEC has computed quotas at various times in its history. This paper takes the reserve data as the best estimate available, understanding that the data quality is ,at least to some extent, subject to a market test. That said, we comment on the robustness of our estimates to errors in reported reserves.

Table A.8 shows total reserves in the world in 1970, 1990, 2000, and 2010. They stand at 439 trillion barrels in 2014, if one considers P50 and a forecasted price of oil of \$70 a barrel.

### A.3 Tables and Figures

Table A.1: Data Cleaning and Sample Frame

	Fields	Total Production in Trillions of Barrels
All Fields	66,920	1.81
Drop non-production fields	21,233	1.81
Drop non-oil fields	19,803	1.16
Drop missing reserves	13,248	1.07

Table A.2: Data Source by Region of the World

Region	Government webpage	Annual Report	Investor Presentation	Company Press Releases	Articles on oil	Other document	Analyst estimate	Modelled
Australia	57	1	1	11			18	11
Caribbean		9					3	88
Central America	57							43
Central Asia		8	1	6	1	6	1	77
East Africa		6		3	7		1	20
East Asia		26	3	2		4	2	62
Eastern Europe	23	3		1			1	45
Middle East	2	11	1	1	4	2	14	64
North Africa	3	13	3	3	4	3	9	63
North America	26	8	1	1		1	4	60
Oceania		9						10
Russia		31	1	2	1	6	12	47
South Africa								100
South America	65	2		1	2			30
South Asia	21	1	8	2	11		2	54
South East Asia	17	12	2	1		1	1	66
Southern Europe	57	1	1	9				32
West Africa	23	23	2	3	5	3	1	40
Western Europe	87				1		1	11
Total	29	12	1	2	1	1	4	49





Figure A.1: Total Production and Number of Oil Producing Assets

Table A.3: Data Source by Year


Region	Government webpage	Annual Report	Investor Presentation	Company Press Releases	Articles on oil	Other document	Analyst estimate	Modelled
1970	0.44	0.00	0.00	0.00	0.02	0.00	0.13	0.41
1971	0.45	0.00	0.00	0.00	0.02	0.00	0.12	0.40
1972	0.44	0.00	0.00	0.00	0.02	0.03	0.11	0.39
1973	0.43	0.00	0.01	0.00	0.03	0.03	0.10	0.41
1974	0.44	0.00	0.00	0.00	0.02	0.03	0.12	0.38
1975	0.38	0.00	0.00	0.00	0.03	0.04	0.13	0.42
1976	0.38	0.00	0.00	0.00	0.03	0.03	0.14	0.43
1977	0.36	0.00	0.00	0.00	0.03	0.04	0.15	0.42
1978	0.34	0.00	0.00	0.00	0.03	0.01	0.20	0.43
1979	0.33	0.00	0.00	0.00	0.02	0.03	0.20	0.41
1980	0.33	0.00	0.00	0.00	0.03	0.04	0.23	0.37
1981	0.34	0.00	0.00	0.00	0.03	0.07	0.20	0.35
1982	0.32	0.00	0.00	0.00	0.04	0.04	0.23	0.37
1983	0.32	0.00	0.00	0.00	0.04	0.02	0.25	0.38
1984	0.32	0.00	0.00	0.00	0.04	0.02	0.25	0.38
1985	0.31	0.00	0.00	0.00	0.03	0.02	0.24	0.39
1986	0.34	0.01	0.00	0.00	0.03	0.02	0.26	0.34
1987	0.32	0.01	0.00	0.00	0.03	0.02	0.23	0.40
1988	0.32	0.01	0.00	0.00	0.03	0.02	0.23	0.39
1989	0.31	0.01	0.00	0.00	0.03	0.02	0.23	0.41
1990	0.28	0.01	0.00	0.00	0.03	0.02	0.30	0.36
1991	0.29	0.01	0.00	0.00	0.05	0.02	0.29	0.35
1992	0.29	0.01	0.00	0.00	0.05	0.02	0.29	0.35
1993	0.31	0.01	0.00	0.00	0.05	0.01	0.27	0.34
1994	0.34	0.01	0.00	0.00	0.04	0.01	0.27	0.32
1995	0.39	0.01	0.00	0.00	0.04	0.01	0.25	0.29
1996	0.37	0.01	0.00	0.00	0.04	0.02	0.28	0.28
1997	0.37	0.01	0.00	0.00	0.03	0.02	0.30	0.27
1998	0.40	0.01	0.00	0.00	0.04	0.03	0.25	0.27
1999	0.37	0.01	0.00	0.00	0.03	0.03	0.28	0.28
2000	0.34	0.02	0.00	0.00	0.04	0.01	0.30	0.29
2001	0.34	0.04	0.00	0.00	0.03	0.03	0.27	0.29
2002	0.33	0.04	0.00	0.00	0.03	0.03	0.30	0.26
2003	0.36	0.04	0.00	0.00	0.03	0.03	0.25	0.27
2004	0.36	0.05	0.00	0.00	0.03	0.02	0.26	0.27
2005	0.37	0.08	0.00	0.00	0.03	0.02	0.22	0.27
2006	0.36	0.09	0.00	0.01	0.02	0.01	0.23	0.27
2007	0.35	0.10	0.00	0.01	0.02	0.01	0.24	0.27
2008	0.41	0.11	0.00	0.00	0.03	0.01	0.17	0.26
2009	0.40	0.15	0.01	0.01	0.03	0.01	0.16	0.25
2010	0.37	0.16	0.01	0.01	0.03	0.01	0.16	0.25
2011	0.27	0.22	0.02	0.02	0.04	0.03	0.15	0.26
2012	0.27	0.20	0.02	0.05	0.03	0.02	0.13	0.28
2013	0.25	0.20	0.01	0.04	0.03	0.04	0.11	0.32
2014	0.21	0.12	0.01	0.08	0.06	0.03	0.09	0.40
Total	0.34	0.05	0.00	0.01	0.03	0.02	0.21	0.33

Table A.4: Total Production in Millions of Barrels by Fuel in 2014

Fossil Fuel	Total	Fields with Positive Production
Biodiesel	189	51
Bioethanol	593	41
Bitumen	303	46
CTL	66	3
Condensate	2213	3767
Extra Heavy Oil	1688	403
Gas	18226	10088
Heavy Oil 15-19	568	190
Heavy Oil 20-23	864	416
LNG	2020	145
Light	3145	1928
NGL	3155	6581
Refinery Gains	872	68
Regular	19045	7836
Sour (> 3%)	661	243
Synthetic crude	373	17
Unsold Gas (flared)	356	1564
Unsold Gas (injected)	1809	1357
Total	56147	34744



(a) Erha Field



(b) Heimdal Field

Note:

P90, P50, and P10 estimates of reserves presented.

Figure A.2: Resources for the Erha Offshore Field

Table A.5: **Summary Statistics for Costs in Millions of dollars in 2014**

Costs	733,863
Abandonment	2,901
Exploration	67,326
Facility Capital Expenditures	121,288
Well Capital Expenditures	254,893
Production Operating Expenditures	204,678
SGA Operating Expenditures	45,145
Transportation Operating Expenditures	37,633
Revenues	2,535,762
Rents	1,801,899
Bonuses	10
Firm Profit	392,858
Government Profit	297,029
Income Tax	739,072
Royalties	343,063
Taxes Operating	29,867
Private Rent	392,868
Public Rent	1,409,032

Notes: SGA Costs refer to Selling, General and Administrative Expenses.

**Table A.6: Share of Revenues Going to Costs, Private, Public Rent, and Distortionary Taxes in 2014**

Country	Cost	Public Rent	Private Rent	Distortionary Taxes
Algeria	0.13	0.73	0.13	0.18
Brazil	0.37	0.40	0.22	0.20
Canada	0.48	0.30	0.22	0.20
China	0.34	0.47	0.19	0.04
Indonesia	0.23	0.62	0.16	0.02
Iran	0.15	0.58	0.27	0.00
Iraq	0.15	0.80	0.05	0.01
Kazakhstan	0.25	0.55	0.20	0.27
Kuwait	0.09	0.86	0.05	0.00
Libya	0.19	0.53	0.27	0.10
Mexico	0.24	0.68	0.09	0.00
Nigeria	0.26	0.51	0.23	0.13
Norway	0.27	0.54	0.19	0.00
Qatar	0.08	0.65	0.27	0.08
Rest of the World	0.27	0.47	0.26	0.08
Russia	0.23	0.57	0.20	0.51
Saudi Arabia	0.10	0.76	0.14	0.00
United Arab Emir	0.12	0.83	0.05	0.48
United Kingdom	0.64	0.20	0.16	0.00
United States	0.62	0.36	0.03	0.24
Venezuela	0.24	0.52	0.25	0.30


*Notes:* Cost plus public and private rent sum to one. Distortionary Taxes are Royalties and Taxes on Operating Expenditures, and are part of Public Rent.

**Table A.7: Price Per Barrel by Fuel in 2014**

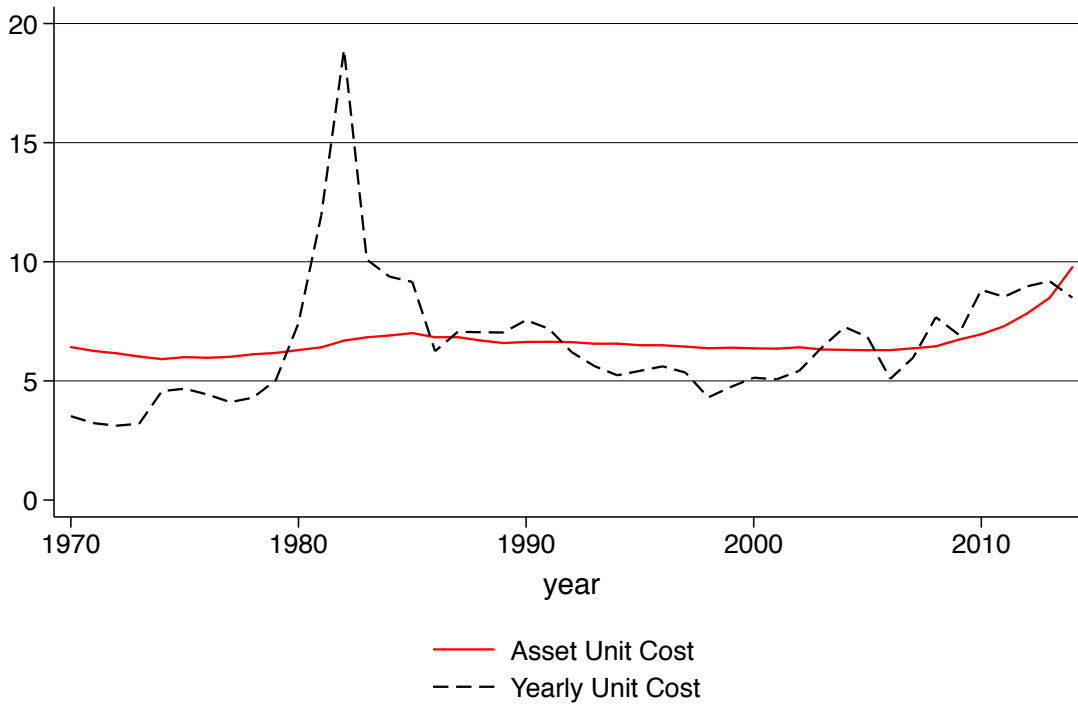
Fossil Fuel	Mean	Standard Deviation
Bitumen	61	6
Condensate	85	14
Extra Heavy Oil	78	7
Heavy Oil 15-19	78	9
Heavy Oil 20-23	82	10
Light	88	15
NGL	45	15
Regular	83	15
Sour (> 3%)	83	6
Synthetic crude	85	4

Table A.8: Reserves, Probability of Recovery and Forecasted Price

Year	Reserve in Trillions of Barrels				
	P10 \$70 barrel	P50 \$70 barrel	P90 \$70 barrel	P10 \$100 barrel	P10 \$130 barrel
1970	932	724	354	1138	1172
1990	914	699	385	1206	1264
2000	797	609	342	1121	1194
2014	572	439	248	1020	1126



(a) Producer Price Index: Oil and Gas Drilling and Services



Notes: Weighted average by year presented. Unit cost is derived as total cost per unit by year, asset cost is the production weighted cost over an asset's lifespan.

(b) Production Weighted Unit and Asset Costs

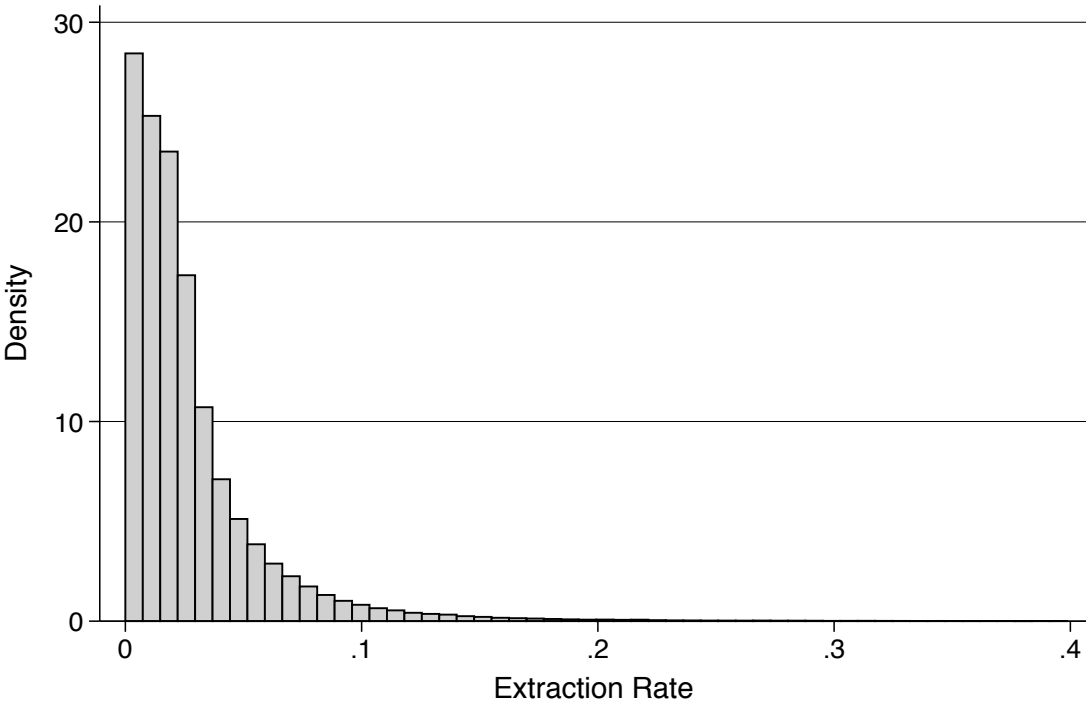


Figure A.4: Extraction Rate of Oil (Production/Reserves)  
Extraction Rate is computed as the production of oil divided by initial reserves of a field.




Figure A.5: Average Costs in Dynamic Counterfactuals

UPDATE WITH LATEST NUMBERS

Table A.9: Top Fields (over 20 percent market share) under optimal production

Asset	Year	Share	Cost Per Unit
Ghawar Uthmaniyah, SA	1970	0.81	2.37
Ghawar Uthmaniyah, SA	1971	0.82	2.37
Ghawar Uthmaniyah, SA	1972	0.83	2.37
Ghawar Uthmaniyah, SA	1973	0.94	2.37
Greater Burgan (Burgan, Magwa, Ahmadi), KW	1974	0.70	2.42
Kirkuk, IQ	1974	0.20	2.42
Greater Burgan (Burgan, Magwa, Ahmadi), KW	1975	0.73	2.42
Kirkuk, IQ	1975	0.21	2.42
Greater Burgan (Burgan, Magwa, Ahmadi), KW	1976	0.67	2.42
Greater Burgan (Burgan, Magwa, Ahmadi), KW	1977	0.22	2.42
Ghawar Shedgum, SA	1977	0.37	2.54
Ghawar Shedgum, SA	1978	0.37	2.54
Ghawar Shedgum, SA	1979	0.36	2.54
Samotlor, RU	1981	0.39	2.79
Gachsaran, IR	1982	0.21	2.77
Samotlor, RU	1982	0.42	2.79
Romashkinskoye, RU	1982	0.23	2.80
Samotlor, RU	1983	0.43	2.79
Romashkinskoye, RU	1983	0.24	2.80
Abqaiq, SA	1983	0.21	2.83
Romashkinskoye, RU	1984	0.23	2.80
Abqaiq, SA	1984	0.24	2.83
Ahwaz Asmari, IR	1984	0.22	2.98
Abqaiq, SA	1985	0.24	2.83
Ahwaz Asmari, IR	1985	0.24	2.98
Ahwaz Asmari, IR	1986	0.23	2.98
Marun, IR	1987	0.24	3.41
Marun, IR	1988	0.26	3.41
Marun, IR	1989	0.25	3.41
Cantarell Akal, MX	1989	0.22	3.48
Cantarell Akal, MX	1990	0.21	3.48
Safaniya (excl-Khafji PNZ), SA	1991	0.43	3.64
Safaniya (excl-Khafji PNZ), SA	1992	0.46	3.64
Safaniya (excl-Khafji PNZ), SA	1993	0.45	3.64

Table A.10: Alternative Specifications for Dynamic Counterfactuals

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)
	Base	$c_f$	taxes	breakeven	ppi cost	largeres	small res	lowextr	highextr	noextr	lowres	highres	startd	only 1970
Actual Cost	2734	3047	4876	16123	8114	2823	1968	2656	2752	2758	2733	2733	2725	1624
Optimal Cost Country	2317	2540	4134	9947	5257	2341	1767	2330	2288	2276	2315	2315	2363	1431
Optimal Cost OPEC Fixed	2118	2287	3486	7384	4821	2145	1981	3128	2074	2056	2116	2115	2180	1323
Optimal Cost	2001	2143	3241	4705	4567	2020	1603	2065	1986	1961	2000	1999	2040	1196
Optimal Cost SA	2072	2232	3583	5794	4742	2091	1786	2183	2034	2007	2071	2070	2117	1265
Optimal Cost AP	2106	2276	3717	6479	4823	2141	1841	2251	2063	2042	2104	2104	2161	1311
Optimal Cost ME	2162	2345	3661	8180	4922	2192	2027	2710	2108	2092	2160	2160	2230	1367
Optimal Cost InOut	2164	2347	3718	8188	4927	2194	2044	3212	2114	2101	2162	2161	2232	1371
Difference OPEC-Optimal	6	7	8	57	6	6	24	52	4	5	6	6	7	11
Difference Country-Optimal	16	19	28	111	15	16	10	13	15	16	16	16	16	20
Difference SA-Optimal	4	4	11	23	4	4	11	6	2	2	4	4	4	6
Difference AP-Optimal	5	6	15	38	6	6	15	9	4	4	5	5	6	10
Difference ME-Optimal	8	9	13	74	8	8	26	31	6	7	8	8	9	14
Difference Actual-Optimal	37	42	50	243	78	40	23	29	39	41	37	37	34	36
Difference InOut	8	9	15	74	8	9	28	56	6	7	8	8	9	15
Actual Cost 2015	2188	2530	3938	16123	6507	2281	1665	2188	2188	2188	2188	2188	2187	1365
Optimal Cost Country Fixed 2015	1666	1924	3006	9947	3053	1700	1408	1853	1635	1632	1665	1665	1723	1118
Optimal Cost OPEC Fixed 2015	1417	1630	2096	7384	2515	1417	1895	2684	1370	1361	1417	1417	1477	990
Optimal Cost 2015	1269	1457	1853	4705	2251	1287	1139	1587	1241	1229	1268	1268	1313	843
Optimal Cost SA Fixed 2015	1348	1552	2284	5794	2392	1362	1516	1714	1313	1297	1347	1347	1393	913
Optimal Cost AP Fixed 2015	1396	1609	2466	6479	2499	1407	1567	1787	1361	1346	1395	1395	1452	971
Optimal Cost ME Fixed 2015	1480	1705	2352	8180	2645	1483	1967	2267	1428	1415	1480	1480	1545	1047
Optimal Cost InOut 2015	1482	1707	2415	8188	2651	1485	1978	2773	1430	1423	1481	1481	1547	1048
Difference OPEC-Optimal 2015	12	12	13	57	12	10	66	69	10	11	12	12	12	17
Difference Country-Optimal 2015	31	32	62	111	36	32	24	17	32	33	31	31	31	33
Difference SA-Optimal 2015	6	7	23	23	6	6	33	8	6	6	6	6	6	8
Difference AP-Optimal 2015	10	10	33	38	11	9	38	13	13	10	10	10	11	15
Difference ME-Optimal 2015	17	17	27	74	17	15	73	43	15	15	17	17	18	24
Difference Actual-Optimal 2015	72	74	113	243	189	77	46	38	76	78	73	73	67	62
Difference InOut 2015	17	17	30	74	18	15	74	75	15	16	17	17	18	24
OPEC Share - Actual	38	38	38	38	38	37	44	38	38	38	38	38	38	48
OPEC Share - Counterfactual	56	58	35	77	58	55	52	59	63	70	56	57	53	68

Table A.11: Alternative Specifications for Static Counterfactual

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
	Base	$c_f$	taxes	breakeven	ppi cost	largeres	small res	lowextr	highextr	noextr	startd
Actual Cost	205	241	477	945	604	241	241	241	241	241	241
Optimal Cost Country	172	204	406	681	405	204	204	255	130	117	198
Optimal Cost OPEC Fixed	129	154	261	541	233	154	154	258	104	94	154
Optimal Cost	102	121	173	254	169	121	121	258	90	84	121
Optimal Cost SA	109	129	249	342	184	129	129	258	96	89	129
Optimal Cost AP	115	137	291	400	202	137	137	258	99	88	137
Optimal Cost ME	131	157	291	576	235	157	157	258	106	95	157
Optimal Cost InOut	132	158	291	579	240	158	158	258	106	95	158
Difference OPEC-Optimal	27	27	51	113	37	27	27	0	15	11	27
Difference Country-Optimal	69	68	134	168	139	69	68	-1	44	39	63
Difference SA-Optimal	7	7	44	34	9	7	7	0	6	5	7
Difference AP-Optimal	13	13	68	58	19	13	13	0	9	5	13
Difference ME-Optimal	29	29	68	127	39	29	29	0	17	12	29
Difference Actual-Optimal	102	99	175	272	256	99	99	-6	167	186	99
Difference InOut	30	31	68	128	41	31	31	0	17	13	31
OPEC Share Actual	39	39	39	39	39	39	39	39	39	39	39
OPEC Share Counterfactual	80	80	74	94	75	80	80	25	88	80	80
ME Share Actual	30	30	30	30	30	30	30	30	30	30	30
ME Counterfactual	75	75	72	91	68	75	75	21	87	78	75
AP Share Actual	18	18	18	18	18	18	18	18	18	18	18
AP Counterfactual	48	48	58	57	44	48	48	12	70	59	48

Table A.12: Welfare with and without Taxes

1970-2014		
	<u>No Tax</u>	<u>Tax</u>
Optimal	1.25	1.41
Optimal OPEC Persian Gulf	1.46	1.57
Optimal Country	1.64	1.64
Actual	2.16	2.16
1970-2050		
	<u>No Tax</u>	<u>Tax</u>
Optimal	2.00	1.96
Optimal OPEC Persian Gulf	2.16	2.12
Optimal Country	2.32	2.25
Actual	2.73	2.73

Note: Tax refers to competitive equilibrium solved using costs inclusive of taxes, while no tax refers to the competitive equilibrium with taxes stripped out. Production costs are evaluated using only production costs.