

OPEC's Impact on Oil Price Volatility: The Role of Spare Capacity

Axel Pierru*

King Abdullah Petroleum Studies and Research Center

James L. Smith

Southern Methodist University

Tamim Zamrik

King Abdullah Petroleum Studies and Research Center

Abstract

OPEC claims to hold and use spare production capacity to stabilize the crude oil market. We study the impact of that buffer on the volatility of oil prices. After estimating the stochastic process that generates shocks to demand and supply, and assessing OPEC's limited ability to accurately measure and offset those shocks, we find that OPEC's use of spare capacity has reduced price volatility, perhaps by as much as half. We also apply the principle of revealed preference to infer the implicit loss function that rationalizes OPEC's investment in spare capacity and compare it to other estimates of the cost of crude oil supply shortfalls.

Acknowledgments: The authors are grateful to Fakhri Hasanov for assistance with the Oxford Economics software package and to Lutz Kilian for comments on a preliminary version presented at ISEFI 2016 (Paris).

* Corresponding author. KAPSARC, P.O. Box 88550, Riyadh 11672, Saudi Arabia, email: axel.pierru@kapsarc.org.

Keywords: oil, price volatility, spare capacity, OPEC, revealed preference

JEL Codes: Q41, Q02, L11

1. Introduction

Spare production capacity plays a central role in the world oil market, and the spare capacity held by OPEC members in particular is significant for its potential ability to stabilize the market price. Indeed, to “ensure the stabilization of oil markets” is part of OPEC’s statutory mission¹. Often, the question has been raised whether OPEC’s spare capacity is large enough—or too large (Fattouh, 2006; *Petroleum Economist*, 2008a,b; *Saudi Gazette*, 2013). Our purpose in this paper is to shed light on the factors that have influenced OPEC’s calculation of the volume of spare capacity required to achieve its mission, and to estimate the extent to which OPEC’s utilization of spare capacity has stabilized the price of crude oil. We believe the present paper is the first attempt to fit a structural model to the behavior of OPEC’s spare capacity in pursuit of these questions.

Looking beyond OPEC as a whole, we also focus on the spare capacity held by four particular OPEC members: Saudi Arabia, Kuwait, Qatar, and the UAE. For lack of a better name, like Reza (1984) and Alhajji and Huettner (2000) we will refer to these four as the OPEC Core. They are distinguished by a perception that, unlike many other members, they have engaged most purposefully in attempts to balance the market by adjusting their production to offset demand and supply shocks. The volume of spare capacity held by the Core comprised 85% of OPEC’s total spare capacity during the period of our study.

The economic significance of efforts to stabilize the price of oil hardly requires explanation. The market is exposed to substantial shocks that disrupt both supply and demand. Whether from war, natural disasters, labor strikes, port closures, political sanctions, or terrorism, the production and delivery of oil to the market is insecure and subject to frequent and unpredictable disruptions. The demand for oil likewise is buffeted by the vagaries of the

¹ The complete mission statement is available at:
http://www.opec.org/opec_web/en/about_us/23.htm (accessed on July 28, 2015).

global economy and suffers too from other types of disruptions (e.g., the short-term substitution of diesel for nuclear energy after Fukushima, or diesel for coal prior to the 2008 Summer Olympics). The impact of each disruption is magnified by relatively low elasticities of demand and supply, which means sharp price movements may be required—especially in the short term—to restore equilibrium in the market.

Previous research (Jaffe and Soligo, 2002; Parry and Darmstadter, 2003; Kilian, 2008; Baumeister and Gertsman, 2013a; Brown and Huntington, 2015) has identified various economic costs associated with oil price volatility. Some of these costs are borne directly by the consumers and producers of crude oil and related products. They take the form of shocks to factor prices and revenue streams that make long-term business planning more difficult. Many private remedies exist to mitigate these shocks, including precautionary inventories, hedging, and long-term contracts. Less direct is the impact of oil price shocks on the business cycle and overall health of the economy. Viewing macroeconomic stability and national security as a public good, national governments and various multilateral agencies have attempted to manage these costs collectively, one example being the International Energy Agency's strategic petroleum reserve program that requires member nations to maintain 90 days of net oil import volumes in public storage.

In light of the various private and public incentives that motivate multiple entities to manage oil price risks, it is clear that OPEC's mission to stabilize the oil market is but one part of a larger picture. OPEC's role is unique, however, because it aspires to reduce price volatility directly—by acting as a swing producer that offsets physical shocks to supply and demand—rather than simply mitigating the cost of price shocks after they have occurred. This strategic capability to reduce price volatility at its source is lacking in private commercial inventories and government stockpiles. No privately-owned inventories are large enough to impact the market price (and if several private entities collaborated in the effort, they could be charged

with illegal efforts at price fixing). Moreover, the public-good aspect of price stabilization transcends the incentives of individual economic agents. Government stockpiles, although certainly large enough (if released) to impact the price, are seldom used, perhaps because they tend to be reserved for use during “emergencies” and because the rules for releasing volumes to the market (or taking volumes off the market) are vague and controversial.

It is well to consider whether the purpose of OPEC’s spare capacity is indeed to stabilize the market price. We find support not only in OPEC’s own mission statement, but also in the obvious and persistent efforts by some OPEC members to raise or lower production to offset unexpected shocks to global demand and supply. Many examples can be cited (e.g., production cuts during the global economic downturn in 2001, production increases which accompanied the unusual buildup of global demand in 2003-2004 and supply disruptions in 2011-2012). Such examples are typical of a “swing producer” and are indicative of OPEC’s commitment to stabilize the market. The CEO of Saudi Aramco acknowledged as much when reporting (*Petroleum Economist*, 2013) that “in the past two years alone, we have swung our production by more than 1.5 million barrels a day (mmb/d) in order to meet market supply imbalances.” Quite often Saudi Arabia is singled out as the ultimate swing producer, the supplier of last resort with sufficient wherewithal (physical and financial) to assume this duty². Accordingly, in addition to studying the impact of OPEC and its four Core members, we also perform a separate analysis of Saudi Arabia’s role in stabilizing the market.

One may speculate about why OPEC should be concerned with stabilizing the price of oil. Various possibilities exist, including a desire to develop the reputation of a reliable supplier, or to mitigate fluctuations in sales revenue and domestic economic growth, or to avoid

² See, for example, Fattouh and Mahadeva’s (2013) review of the literature, in which Saudi Arabia is singled out as the dominant swing producer. Nakov and Nuño (2013) show that both the size of Saudi Arabia’s spare capacity and the volatility of its monthly output greatly exceed that of other OPEC members.

triggering investment in alternative energy sources, etc. We take no particular stand and our analysis and empirical results do not hinge on any particular interpretation of OPEC's ulterior motives. We assume only that OPEC wants to reduce price volatility for whatever may be the reason.

In principle, spare capacity could be used to advance objectives besides price stabilization. Available capacity could be tapped opportunistically, for example, to skim extra revenue when prices are high. However, that extra production would also tend to lower the price and reduce volatility. Another possible use of spare capacity would be to deter entry or expansion by non-OPEC producers of crude oil. Although the value of successful entry deterrence might be large, it does not appear that OPEC has managed its spare capacity in a way that has actually deterred entry. Since 2000, total world production of crude oil has grown by 17% whereas the output of non-OPEC producers has grown even more—by 21%. OPEC production rose by only 13% over this interval.³

Relatively few papers have formally addressed OPEC's role in stabilizing the price of oil. De Santis (2003) attributes price volatility under OPEC's old production quota regime specifically to the inelasticity of Saudi Arabian supplies. Any physical disruption, he argues, would create a short-term price spike that could only be dissipated by longer term supply adjustments. De Santis assumes the absence of spare capacity which begs the question of how such a precautionary buffer would be sized and managed—or what would be its impact on price volatility.

Nakov and Nuño (2013) take the opposite approach, assuming that Saudi Arabia can and does adjust its output in response to each monthly demand shock in the manner of a Stackelberg dominant producer. By offsetting positive (negative) shocks with an increase

³ Production includes crude oil and lease condensate, as reported by the U.S. Energy Information Administration.

(decrease) in its own output, Saudi Arabia effectively reduces price volatility, although that result is a by-product and not the objective of its behavior. The Stackelberg framework is a very insightful approach that seems appropriate to the structure of the world oil market, but one that presumes the dominant producer can perfectly anticipate the magnitude of each shock. Substantial misjudgments in that regard, if acted upon, could in fact lead to an increase in volatility, and the possibility of mistakes may hold the producer in abeyance.

Golombek, Irarrazabal, and Ma (2016) also estimate a structural model where OPEC assumes the role of a dominant firm who prices strategically against a competitive fringe. They find strong evidence that variations in OPEC (and Core) production influence the market price and that OPEC has used this influence to exercise substantial market power. However, they do not investigate the role or significance of spare capacity or the impact of OPEC's actions on price volatility.

Fattouh (2006) provides evidence that an increase in volatility and the frequency of price spikes are in a general way due to reduced spare capacity held by OPEC and other producers, but he does not pursue the argument to the point of a formal model or empirical estimates. Difiglio (2014) recognizes OPEC's role in stabilizing prices via spare capacity and reviews reasons why similar efforts to offset disruptions using consuming nations' own strategic petroleum reserves have not been very successful. So far, the literature has not provided any formal model of a buffer capacity that is used to continuously stabilize the price of oil, which is the goal of our paper.

In Section 2 we develop a dynamic structural model of a producer using his spare capacity to stabilize the market price of its output. The model includes three components: an autoregressive stochastic process by which the residual demand for OPEC oil is shocked each month, a separate stochastic process by which OPEC attempts to estimate the size of such shocks and offset them by regulating production from its buffer stock, and finally a loss

function that describes the benefits that rationalize the observed size of OPEC's chosen buffer. We estimate the model's parameters using observed price, production and spare capacity data. In section 3, we derive an analytical formula for the marginal value of spare capacity used to stabilize the price. We adopt in Section 4 the assumption that OPEC has equated the marginal costs and benefits of its spare capacity and invoke the principle of revealed preference to identify the loss function that appears to have motivated OPEC's investment in spare capacity. In section 5, our estimate of OPEC's implicit loss function is compared to an independent estimate of the size of economic losses due to oil supply disruptions derived from a well-known macroeconomic model of the global economy. That comparison indicates the extent to which OPEC's investment in spare capacity is commensurate with the interests of the global economy. We examine the degree to which OPEC's management of spare capacity has actually damped price volatility in Section 6 based on a counterfactual reconstruction of what "unstabilized" prices would have looked like. Concluding observations are presented in Section 7.

2. A model of price stabilization using spare capacity

2.1 Model assumptions

Since there is nothing specific to OPEC in the structure of the model, we develop the framework in the context of a generic Producer who elects to develop and deploy spare capacity to stabilize the market price of his output. Implicit is the notion that Producer has sufficient production to impact the market price. We assume that demand for Producer's output in any period follows a lognormal distribution due to the arrival of shocks that follow a known autoregressive process. Given the structure of our model, lognormal shocks are consistent with normally distributed percentage changes in both quantity demanded and price (the latter is a

standard assumption in the finance literature) and implies a Gaussian likelihood function.⁴ We further assume that Producer wishes to stabilize price around a certain target level and that he creates a buffer of spare capacity to be used in this endeavor, but also that he is unable to accurately estimate the size of the shocks and therefore offsets them only imperfectly.

Let Q_t represent the demand for Producer's output in period t . We assume that current demand depends on current and past prices:

$$Q_t(P_t, P_{t-1}, \dots, P_{t-K}) = a_t \left(\prod_{k=0}^K P_{t-k}^{\omega_k} \right) e^{S_t} \quad (1)$$

where ω_0 is the short-run (monthly) elasticity of demand, $\sum_{k=0}^K \omega_k$ is the long-run elasticity, a_t is an exogenous, time-varying scaling parameter, and S_t represents a random shock that affects the demand for Producer's crude. We place no restrictions on the number of lags or the rate at which the influence of past prices decays.

The stochastic component e^{S_t} is caused by shocks to global demand and non-Producer supply. For application to monthly data, some shocks are likely to persist beyond 30 days. Accordingly we consider that the shocks S_t follow a first-order autoregressive process:

$$S_{t+1} = \kappa S_t + \sigma_S u_t \quad (2)$$

where $u_t \sim iid \ N(0,1)$, σ_S represents the standard deviation of innovations on the shock, and κ is the shock persistence (note that $\kappa = 1$ implies a random walk). The lower is κ , the faster will shocks dissipate. This implies that S_t follows a normal law and that, for given prices, Q_t follows a log-normal law.

Let P_t^* represent Producer's target price for the period t . It is assumed that the target price vector is determined exogenously according to many criteria that lie outside the scope of

⁴ The normal distribution also plays a key role regarding the tractability of our analysis. In particular, the expected size of production shortfalls that occur when spare capacity is exhausted takes a simple closed form in the case of lognormal shocks (see Equation 20).

our analysis.⁵ Given the price target, Producer adjusts output each period to prevent deviations of the market price from P_t^* . In the vernacular of the oil market, P_t^* is the price that Producer chooses to “defend.” And, let Q_t^* be the volume that Producer would have to produce in period t to defend the target price in the absence of shocks (i.e. $S_t = 0$). From (1) we have:

$$Q_t^* = a_t P_t^{*\omega_0} \prod_{k=1}^K P_{t-k}^{\omega_k}, \quad (3)$$

where the P_{t-k} are the prices observed in previous periods.

We assume that, in order to offset positive shocks to demand, Producer adopts a policy of maintaining a buffer sized as a fixed proportion of Q_t^* . Letting C_t represent production capacity at period t , we have:

$$C_t = B Q_t^*, \quad \text{with } B > 1. \quad (4)$$

Our goal is to identify the value of constructing a buffer and to identify its optimal size.

When estimating the size of the shock, Producer makes the error $\sigma_z z_t$, where z_t is uncorrelated with S_t and $z_t \sim iid N(0,1)$. The shock perceived by Producer is therefore $S_t + \sigma_z z_t$. Given the target price, Producer thus perceives the call on its crude to be:

$$\tilde{Q}_t = a_t (P_t^{*\omega_0} \prod_{k=1}^K P_{t-k}^{\omega_k}) e^{S_t + \sigma_z z_t}. \quad (5)$$

Although we will refer to z_t as estimation error, it is in fact a composite of various random factors (e.g., political, operational, logistical, etc.) that, in addition to estimation error, might impact or constrain Producer’s production in any given month. From (3) and (5) we have

$$\tilde{Q}_t = Q_t^* e^{S_t + \sigma_z z_t}. \quad (6)$$

Spare capacity X_t is the difference between total installed capacity and the perceived call on crude:

$$X_t = \max\{0, C_t - \tilde{Q}_t\} \quad (7)$$

⁵ The target price might be as determined by a strategic dominant firm, as in Nakov and Nuno (2013) and Golombek, Irarrazabal, and Ma (2016), but might also incorporate additional geopolitical, social, financial, and intertemporal considerations.

Figure 1a illustrates price formation when the buffer is large enough to fully absorb the shock. \bar{P}_t represents the (undamped) price that would have been obtained if Producer had not used spare capacity to offset shocks. Figure 1b illustrates price formation when the buffer is not sufficient to fully absorb the shock, in which case P_t exceeds the target price due to the lack of sufficient production capacity.

2.2 Estimating the estimation error based on observed price volatility

To stabilize the price, Producer supplies \tilde{Q}_t , i.e. the perceived call on its output. The resulting price P_t is therefore such that: $a_t(\prod_{k=0}^K P_{t-k}^{\omega_k})e^{S_t} = \tilde{Q}_t$, which after using (5) gives:

$P_t = P_t^* e^{\frac{\sigma_z z_t}{\omega_0}}$, or equivalently:

$$\ln(P_t) = \ln(P_t^*) + \frac{\sigma_z z_t}{\omega_0} \quad (8)$$

In other words, the deviation of the price from its target is due to the estimation and/or execution error (z_t) that prevents Producer from precisely offsetting the shock. We can therefore use observed price volatility to estimate the size of this error. The conventional measure of volatility, vol , is based on the variance of returns (percentage change in price). From (8) we have:

$$vol^2 = var \left[\ln \left(\frac{P_t}{P_{t-1}} \right) \right] = \sigma_{TP}^2 + 2 \left(\frac{\sigma_z}{\omega_0} \right)^2. \quad (9)$$

The first term in this expression is the variance of the periodic percentage changes in Producer's target price: $\sigma_{TP}^2 = var \left(\ln \left(\frac{P_t^*}{P_{t-1}^*} \right) \right)$. Solving (9) for the standard deviation of Producer's estimation error gives:

$$\sigma_z = \frac{|\omega_0|}{\sqrt{2}} \sqrt{vol^2 - \sigma_{TP}^2}. \quad (10)$$

Assuming that $\sigma_{TP}^2 = 0$ therefore provides an upper bound on σ_z . Of course, the term σ_{TP}^2 would vanish only if the target price were increasing by a constant percentage each month. Upon reviewing the development of the crude oil market during our sample period, it may not

be unreasonable to assume $\sigma_{TP}^2 \cong 0$.⁶ This assumption, along with an estimate of the residual demand elasticity, allows us to approximate (and bound) the standard deviation of Producer's estimation error:⁷

$$\hat{\sigma}_Z \leq vol \times \frac{|\omega_0|}{\sqrt{2}} \quad (11)$$

For our purposes, we use the average monthly Brent crude oil spot price series published by the U.S. Energy Information Administration to estimate vol as the standard deviation of log-returns of the average monthly price over our sample period (September 2001 to October 2014). This gives $vol = 8.58\%$. The value ω_0 in (11) represents the short-run (monthly) elasticity of residual demand for Producer's oil, and is by construction equal to $[\varepsilon_D - (1 - \rho)\varepsilon_S]/\rho$, where ε_D and ε_S represent the short-run elasticity of global demand and non-Producer supplies, and ρ is the Producer's market share of global output.

Our estimation procedure is therefore sensitive to ε_D and ε_S , the short-run elasticities of global demand and non-Producer supply. Given the range of estimates found in the literature, our analysis will be subjected to sensitivity analysis. The literature sees both global demand and non-OPEC supply to be highly inelastic in the very short-run. Hamilton (2009) proposed a short-run global demand elasticity of -6%, but noted that it might be higher or lower. Based on observed price movements following specific disruptions of the market, Smith (2009) suggested short-run demand and supply elasticities of -5% and +5%. Baumeister and Peersman (2013b) provide corroborating evidence based on a time-varying parameter vector autocorrelation analysis of global crude oil demand and supply. Their estimates of the quarterly demand elasticity fall between -5% and -15% throughout our sample period, and their

⁶ Despite the tremendous disruption caused by the 2008/2009 financial crisis, a simple monthly regression of $\ln(\text{price})$ against time, from September 2001 to October 2014, produces an R^2 of 89%, which is indicative of exponential growth at a fairly constant rate.

⁷ Let λ measure the portion of observed volatility due to changes in the target price. Thus, $\sigma_{TP}^2 = \lambda \times vol^2$, in which case (11) takes the general form: $\hat{\sigma}_Z = vol \times \frac{|\omega_0|}{\sqrt{2}} \times \sqrt{1 - \lambda}$. As we show later, however, our main results and conclusions are robust with respect to the presumed value of λ .

estimates of the quarterly supply elasticity are of the same magnitude. Because our data are monthly, we consider a global demand elasticity ranging from -1% to -5% to be consistent with this literature, and for values within this range we take $\varepsilon_S = |\varepsilon_D|$. Kilian and Murphy (2014) suggest a much higher estimate of the short-run elasticity of demand (-26%) based on a structural vector autoregression that attempts to account for monthly changes in the global volume of speculative crude oil inventories. Therefore, we also include a sensitivity case where the monthly demand elasticity is -26% and the monthly supply elasticity is 0%, to match Kilian and Murphy's assumptions.

To calculate market share (ρ) for each group of producers, we compute the average crude oil supply per month over the sample period. Our supply data are from the IEA Monthly Oil Data Service. Production from the Neutral Zone is not included in Saudi production (but included in OPEC Core production). For OPEC, we use IEA's "OPEC Historical Composition" series. Table 1 provides the implied elasticities of residual demand. Table 2 shows the corresponding estimates of the standard deviation of the estimation error, in both relative and absolute terms, calculated from (11). The absolute estimation errors (barrels per day) attributed to Saudi Arabia, the Core, and OPEC are roughly equal in size. The values range between 0.07 and 1.16 mmb/d, with all values below or equal to 0.4 mmb/d if global demand elasticity does not exceed -5%. Since Saudi production is smaller than that of the Core, which in turn is smaller than that of OPEC, the relative size of the estimation error (as a percentage of producer's output) respectively decreases, which seems sensible to us.

The precision of these estimates can be calculated using the Chi-Square distribution. A 95% confidence interval for σ_Z^2 is given by: $\left[\frac{(n-1)\hat{\sigma}_Z^2}{K_{.975}}, \frac{(n-1)\hat{\sigma}_Z^2}{K_{.025}} \right]$, where $K_{.975}$ and $K_{.025}$ are cutpoints from the Chi-Square distribution with $n - 1$ degrees of freedom. Based on the 158 monthly observations in our sample, the 95% confidence interval for σ_Z is: $[0.901\hat{\sigma}_Z, 1.124\hat{\sigma}_Z]$.

2.3 Estimation of other model parameters

Because $C_t = \tilde{Q}_t + X_t$, and after using (4) and (6) to substitute for C_t , we have:

$$-\ln\left(1 + \frac{X_t}{\tilde{Q}_t}\right) = S_t + \sigma_z z_t - \ln(B) \quad (12)$$

The left-hand side of (12) is observable. The right-hand side represents the perceived autoregressive shocks to Producer's demand (cf. (2)) with unknown parameters B (buffer size), σ_S (volatility of demand shocks), and κ (shock persistence). Given monthly data on actual production (\tilde{Q}_t) and spare capacity (X_t), along with our previous estimate of σ_z , maximum likelihood estimates of B , σ_S , κ , and the covariance matrix are obtained by the procedure described in Appendix 1. We ignore the data censoring represented by (7) which occurs if the shock exceeds the size of the buffer. However this should not matter since in our sample only Saudi spare capacity ever reached zero, and during three months only. The frequency of censoring is therefore very low, which reflects the fact that the spare capacity has almost always been sufficient to meet the perceived call on production.

Figure 2 shows the monthly variation in spare capacity of OPEC, Saudi Arabia, and the OPEC Core. Our data come from the International Energy Agency (IEA) and represent what they call “effective” spare capacity.⁸ The monthly spare capacity data for Saudi Arabia and OPEC were provided directly by IEA in an Excel file⁹. To build the series for the OPEC Core, we collected¹⁰ the data for Kuwait, UAE and Qatar from monthly issues of IEA's *Oil Market Report*. Because spare capacities are not reported on a regular basis prior to September 2001,

⁸ According to the IEA, spare capacity is defined as “capacity levels that can be reached within 30 days and sustained for 90 days.” Effective spare capacity captures the difference between nominal capacity and the fraction of that capacity actually available to markets (Munro, 2014).

⁹ Email from Steve Gervais (IEA) on January 7th, 2015.

¹⁰ We had five missing data for the non-Saudi members of OPEC Core. We consider that, because of a typo, the values for November 2002 and 2010 are those reported for October in the December's Oil Market Report (as these values differ from those reported for October in the November's report). The three other missing data are for June 2002, April 2003 and March 2007. We interpolate the missing values with the formula: $X_{C,t} = X_{S,t} + \frac{X_{O,t} - X_{S,t}}{2} \left[\frac{X_{C,t-1} - X_{S,t-1}}{X_{O,t-1} - X_{S,t-1}} + \frac{X_{C,t+1} - X_{S,t+1}}{X_{O,t+1} - X_{S,t+1}} \right]$, where $X_{S,t}$, $X_{C,t}$ and $X_{O,t}$ represent the spare capacity of Saudi Arabia, OPEC Core and OPEC, respectively, in month t .

our sample extends from September 2001 to October 2014 (158 observations for each series). These are the primary data with which we estimate the stochastic process governing shocks to the call on OPEC production. The estimates and their standard errors are reported in Table 3 for the case where the global demand elasticity is assumed to be -1%. The estimates and standard errors corresponding to the other elasticity cases are nearly identical to these and are therefore relegated to the appendix.

The estimates of B , the size of the buffer, and σ_S , the magnitude of innovations on the shock, exhibit a common pattern: the greatest values are obtained for Saudi Arabia, the lowest for OPEC. This is consistent with the view that Saudi Arabia is the key swing producer who absorbs more shocks than any other OPEC producer. In all cases, the estimated size of the Saudi buffer is about 21% of the expected call on Saudi Arabia's output, whereas for the Core (15%) and OPEC as a whole (9%) it is smaller.

To appreciate the size of the estimated buffers, we first determine Q^* , the average call on Producer's crude over our sample period. Q^* is the average of $Q_t^* = \frac{\bar{Q}_t + X_t}{B}$. For an elasticity of global demand of -1%, this gives $Q^* = 8.82$ mmb/d for Saudi Arabia, 14.91 mmb/d for the Core, and 30.02 mmb/d for OPEC as a whole. The average size of the buffer in physical terms is then calculated by multiplying these values by $B - 1$. As one would expect and aggregation requires, the larger is the group of countries, the bigger is the buffer: 1.94 mmb/d for Saudi Arabia, 2.27 mmb/d for OPEC Core, and 2.64 mmb/d for OPEC. The Saudi figure is consistent with the many pronouncements that have emanated from the Kingdom that put their intended buffer between 1.5 and 2 mmb/d (see for instance *Petroleum Economist* (2005, 2012) and H.E. Ali Al-Naimi's address at CERAWeek (2009) and remarks at the 12th International Energy Forum (2010)).

Regarding the estimated speed at which shocks dissipate (Table 3), the estimated half-life is roughly 25 months ($\kappa = 0.973$). Although differences in the estimates of κ appear small

and are not statistically significant across the elasticity cases, the implied half-life is considerably shorter (15 months) for the case of -26% demand elasticity (see appendix).

3. Incremental value of spare capacity

We assume that Producer incurs costs in any period when the perceived call exceeds available production capacity. The resulting shortfall in production equals the portion of the perceived call that Producer is not able to meet. However, the cost associated with a production shortfall is not measured by foregone revenue, but is much broader. Specifically, we assume that Producer's "optimum" production level depends on many things, economic as well as perhaps some non-economic factors, all of which go into determining the target price. The cost of a shortfall is therefore determined by Producer's displeasure or disutility due to the inability to pursue its optimum course.¹¹

We denote production shortfalls by $O_t \stackrel{\text{def}}{=} \max\{0, \tilde{Q}_t - C_t\}$. From (4), (6) and (7), the shortfall can be written equivalently as $O_t = \max\{0, (e^{S_t + \sigma_z Z_t} - B)Q_t^*\}$. The probability of a shortfall depends on the size of the buffer and is given by:

$$\varphi_t(B) \stackrel{\text{def}}{=} \text{pr}(O_t > 0|B) = \int_{\ln(B)}^{\infty} g_t(\xi) d\xi \quad (13)$$

where $g_t(\cdot)$ represents the marginal density of $S_t + \sigma_z Z_t$ based on the information set at time $t = 0$.

The expected size of the shortfall is:

$$E[O_t|B] = \int_{\ln(B)}^{\infty} (e^{\xi} - B)Q_t^* g_t(\xi) d\xi \quad (14)$$

whereas the conditional expectation, given that a shortfall occurs, is:

$$E[O_t|B \cap O_t > 0] = \frac{E[O_t|B]}{\varphi_t(B)}. \quad (15)$$

¹¹ Indeed, there may be no foregone revenues associated with a production shortfall. For example, if residual demand for OPEC's oil is inelastic, then the shortfall would actually produce greater revenue than intended. Nevertheless, by definition of the target price, OPEC would presumably regret not being able to prevent the market price from exceeding its target.

We postulate a quadratic loss function that reflects the present value of Producer's disutility due to future shortfalls:

$$L = \alpha \sum_{t=1}^T \frac{(O_t)^2}{(1+r)^t} \quad (16)$$

where r is the real risk-adjusted periodic discount rate and α is a latent preference parameter that reflects the weight that Producer attaches to shortfalls. The loss function is increasing in the square of the size of individual shortfalls and additive regarding their occurrence.

The planning horizon is defined by T . We treat T as the service life of a designated production facility kept for spare. The value of the buffer to Producer is determined by its ability to reduce the expected loss resulting from shortfalls. As shown in Appendix 3, the incremental value, v , of spare capacity is given by:

$$v = -\frac{\partial E[L|B]}{\partial B} = 2\alpha \sum_{t=1}^T \frac{E[O_t|B]Q_t^*}{(1+r)^t} \quad (17)$$

Note that the value of expanding the buffer does not depend on the functional form of $g_t(\cdot)$, only on $E[O_t|B]$, which is itself the product of $\varphi_t(B)$ (the probability of a shortfall) and $E[O_t|B \cap O_t > 0]$, as well as the length of the planning horizon, the expected call, and α . In the next section, we show how all of these parameters can be estimated from existing data. Of particular interest is the estimated value of α because that will allow us to calibrate Producer's loss function and compare the cost of shortfalls as perceived by Producer (whether OPEC, OPEC Core, or Saudi Arabia) to independent estimates of the global economic cost of supply disruptions. That comparison, in turn, will provide an indication of the extent to which OPEC's stabilization policy is commensurate with the interests of the global economy.

As shown in Appendix 3, an immediate implication of (17) is that the expected loss and the value of incremental spare capacity are both decreasing in the size of the buffer. To evaluate (17), we need to calculate:

$$E[O_t|B] = Q_t^*(E[e^{S_t + \sigma_z Z_t} | S_t + \sigma_z Z_t > \ln(B)] - B) \times \varphi_t(B). \quad (18)$$

Since we are considering a long-term policy of maintaining a buffer of optimal size, we will use the covariance-stationary process that satisfies (2) (see Hamilton (1994) p. 53). $S_t + \sigma_z Z_t$ therefore follows a normal law with mean zero and variance $\sigma^2 = \sigma_z^2 + \frac{\sigma_S^2}{1-\kappa^2}$. We make the additional simplifying assumption that the call on Producer's crude in the absence of shocks remains stable and equal to Q^* .

We use the following fact about the mean of a truncated lognormal distribution (Johnson et al., 1994, p.241):

$$E[e^{S_t + \sigma_z Z_t} | S_t + \sigma_z Z_t > \ln(B)] = e^{\sigma^2/2} \frac{\Phi(\sigma - \ln(B)/\sigma)}{\varphi(B)} \quad (19)$$

where: $\varphi(B) = 1 - \Phi(\ln(B)/\sigma)$ and where $\Phi(\cdot)$ represents the cumulative distribution of the standard normal law. Therefore, from (14) we have:

$$E[O_t | B] = Q^* \left(e^{\sigma^2/2} \Phi\left(\sigma - \frac{\ln(B)}{\sigma}\right) - B \left(1 - \Phi\left(\frac{\ln(B)}{\sigma}\right)\right) \right). \quad (20)$$

Upon substituting (20) into (17), we obtain the parametric form of the incremental value of spare capacity:

$$v = \left(e^{\sigma^2/2} \Phi\left(\sigma - \frac{\ln(B)}{\sigma}\right) - B \left(1 - \Phi\left(\frac{\ln(B)}{\sigma}\right)\right) \right) \sum_{t=1}^T \frac{2\alpha(Q^*)^2}{(1+r)^t}. \quad (21)$$

4. Revealed preference for spare capacity

We now show how (21) and the principle of revealed preference can be used to infer the value of α , the behavioral parameter that reveals how much importance Producer attaches to the avoidance of production shortfalls.

Denote by h the marginal cost to provide one barrel per day of additional spare capacity. This is the capital expenditure to construct the capacity, plus maintenance cost, less any net revenue generated when that incremental barrel of spare capacity is used to increase production. If the cost is K to construct a production facility with peak production rate R , then the capital cost per daily barrel of spare capacity is given by $k = K/R$. In addition, we have

to account for the maintenance costs (which are incurred even when spare capacity is not used) and net financial gains (which are generated only when barrels are released from spare capacity). Any release generates marginal revenue that may be either positive or negative depending on the elasticity of residual demand. The net financial gain from each release is the marginal revenue minus the operating cost of producing the barrel. Over the life of the facility, the present value of the incremental maintenance cost is represented by m , while the expected net present value of financial gains is represented by f . Therefore, from the Producer's perspective the present value cost of an incremental barrel of spare capacity is $h = k + m - f$.

If, consistent with the principle of revealed preference, we assume the historical size of the buffer has been optimized, then the marginal cost of the buffer must equal the marginal benefit. Since v is derived for a buffer defined in relative terms, at the optimized buffer size we must have: $hQ^* = v$,¹² which implies:

$$k + m - f = \left(e^{\sigma^2/2} \Phi\left(\sigma - \frac{\ln(B)}{\sigma}\right) - B \left(1 - \Phi\left(\frac{\ln(B)}{\sigma}\right)\right) \right) \sum_{t=1}^T \frac{2\alpha Q^*}{(1+r)^t}. \quad (22)$$

A rational agent therefore sizes the buffer based on four factors: the size and persistence of shocks to demand, the precision with which agent can estimate those shocks, the importance attached to resulting shortfalls (as represented by the parameter α of the loss function), and the net cost of developing, maintaining, and operating spare capacity. Given the estimated size of the buffer, (22) allows us to calibrate the loss function that would rationalize OPEC's investment in spare capacity:

$$\alpha = \frac{k+m-f}{\left(e^{\sigma^2/2} \Phi\left(\sigma - \frac{\ln(B)}{\sigma}\right) - B \left(1 - \Phi\left(\frac{\ln(B)}{\sigma}\right)\right) \right) \sum_{t=1}^T \frac{2Q^*}{(1+r)^t}}. \quad (23)$$

¹² $h = k + m - f$ represents the cost of expanding the buffer by one barrel per day, which on average corresponds to $1/Q^*$ in percentage terms, whereas v represents the benefit from expanding the buffer by 1 percent.

We have previously discussed all parameter estimates that appear in the denominator of (23) and now turn to the cost parameters in the numerator. The spare capacity that exists within OPEC can be drawn from many sources, including increased liftings from producing fields as well as additional production from idle facilities (if any). We assess the costs associated with incremental production via the simplifying assumption that it all comes from a dedicated buffer facility that is reserved for that specific purpose. Although this may depart somewhat from reality, it provides a useful proxy for the more complicated and diffuse costs that may actually be incurred.

The capital cost of spare capacity (k) can be estimated using data from the most recent oil field development in Saudi Arabia, the Manifa field that is located in a shallow offshore setting. According to Henni (2013), the total capital cost to develop Manifa's production capacity of 900,000 barrels per day (which corresponds to our parameter R) is \$15.8 billion (which corresponds to our parameter K). Therefore, the capital cost per daily barrel of production capacity is given by $k = K/R = \$17,500$.

We assume that maintenance cost remains constant throughout time and take it from QUESTOR, IHS's cost estimating software package that is a petroleum-industry standard. For an idle facility located in Saudi Arabia with 40 wells and combined production capacity of 200,000 barrels per day, QUESTOR estimates annual maintenance cost to be \$410 per daily barrel, or \$34.17 on a monthly basis. Thus, for an annual real discount rate of 4%, in line with Pierru and Matar's (2014) findings for Saudi Arabia, the present value of maintenance costs over the 240-month life of the facility (which corresponds to our parameter m) would be:

$$\sum_{t=1}^{240} \frac{34.17}{(1.04)^{\frac{t}{12}}}, \text{ which gives } m = \$5,670 \text{ per daily barrel of capacity.}$$

The incremental barrel of buffer capacity would only be used if there otherwise would be a shortfall. To calculate the expected incremental revenue, we consider the average price (\$45.24 per barrel) observed during the three months when Saudi Arabia ran out of spare

capacity (from August to October 2004) and use that price (along with the usual formula for marginal revenue) to determine the parameter f as the sum of expected monthly net financial gains over the 240-month life of the facility discounted at 4%, assuming \$2 production cost per barrel¹³. We thus have:

$$f = \sum_{t=1}^{240} \frac{\varphi(B)30 \left(\left(1 + \frac{1}{\omega_0}\right)^{45.24-2} \right)}{(1.04)^{\frac{t}{12}}}.$$

The resulting estimates of the financial gains due to an incremental barrel of buffer are shown in Table 4. Note that the “gains” are negative if residual demand is inelastic because any increase in production would lower the price and also revenue.

5. Assessment and discussion of the implicit loss function

After substituting the parameter estimates discussed above into (23), we obtain the estimated values of α shown in Table 5. These values represent the implicit weight that Producer attaches to the avoidance of production shortfalls.

Using the estimated values of α , it is possible to evaluate the loss function that rationalizes Producer’s choice of the buffer. Consider, for example, OPEC Core, for whom the expected size of a shortfall (when one occurs) is roughly half a million barrels per day.¹⁴ By substituting the estimated values of α from Table 5 into the loss function (16), we can calculate the cost the Core attaches to a half-million barrels per day shortfall lasting for 6 months. We get a cost of \$24.87 billion if the elasticity of global demand is assumed to be -1% , \$12.80 billion if elasticity is -3% , \$8.38 billion if elasticity is -5% , and \$1.34 billion if elasticity is -26% .

¹³ *Petroleum Intelligence Weekly* (2011) reports Saudi Aramco’s group-wide average production cost as falling between \$2 and \$3 per barrel. Also note that we assume that the marginal revenue is received for each of thirty days within any month affected by a shortfall.

¹⁴ From (15) and (20) we derive the expression of the (conditional) expected shortfall size: $E[O_t|B \cap O_t > 0] = \frac{Q^*}{\varphi(B)} (e^{\sigma^2/2} \Phi(\sigma - \ln(B)/\sigma) - B\varphi(B))$. For OPEC Core this gives a size ranging between 0.50 and 0.53 mmb/d when global demand elasticity is lower or equal to -5% (0.82 mmb/d when elasticity is -26%).

These results mean little when standing alone, but are of considerable interest when compared to independent estimates of the economic cost that such a supply disruption would impose on the global economy. For this purpose, we apply Oxford Economics *Global Economic Model* to simulate the negative impact on global GDP of disruptions of varied size and duration. Although one may question whether global GDP losses have any per se relevance to members of OPEC, the size of those losses does provide a direct measure of the value of establishing reliable supplies of crude oil. To the degree that OPEC wants to be viewed as a reliable supplier, the size of GDP losses caused by supply disruptions may therefore at least indirectly influence the size and management of OPEC's spare capacity. Appendix 4 provides the information on the Oxford Economics model and the procedure we followed, and the results obtained are given in Table A4.

For differently sized six-month disruptions, Figure 3 compares the global cost inferred from the Oxford Economics model with the inferred cost that rationalizes OPEC Core's choice of buffer (assuming different values for elasticity of global oil demand). According to the model, a six-month production shortfall of 0.5 mmb/d that is assumed to occur at the beginning of 2015 would have reduced the present value of global GDP over the next five years by some \$22.36 billion, relative to the reference scenario. This loss lies near the top of the range of perceived costs that we believe the OPEC Core may have attributed to such a shortfall. Similar results hold for Saudi Arabia and for OPEC as a whole. Even if global demand is assumed to be relatively elastic in the short term (-5%), the costs that rationalize OPEC Core's buffer comprise some 40% of the "global" cost of disruptions. This does not imply that the buffer is too small, only that the OPEC Core may for whatever reason be motivated to address only a portion of the potential damage caused by oil shocks. We repeat an earlier point: OPEC and its members are but one piece of a much larger picture when it comes to neutralizing the impact of oil shocks. Whether it is reasonable to believe that 60% of the burden should be left for

individual consumers, producers, government agencies, and multilateral organizations, (not to mention the other members of OPEC) to deal with, we are unable to say. However, if global demand is assumed to be highly inelastic in the short term (-1%), the costs that rationalize the size of the Core's buffer actually exceed the level of global costs projected by the Oxford Economics model. We also note that whatever motivated OPEC Core's decision to build its observed buffer, that decision has sheltered the global economy from potential disruptions. Let us for instance consider $\varepsilon_D = -5\%$. According to (19) and (20), if there had been no buffer ($B = 1$), the probability of a production shortfall and its expected size would have increased from 3.3% to 50% and from 0.53 mmb/d to 1.09 mmb/d, respectively.

6. OPEC's measured impact on price volatility

We evaluate the set of counterfactual prices, \bar{P}_t , that would have been obtained if Producer had produced Q_t^* instead of using spare capacity to offset shocks:

$$a_t \left(\prod_{k=0}^K \bar{P}_{t-k}^{\omega_k} \right) e^{S_t} = a_t P_t^{*\omega_0} \prod_{k=1}^K \bar{P}_{t-k}^{\omega_k}.$$

The left-side of this equation represents demand under the set of counterfactual prices whereas the right-side represents the expected call given past counterfactual prices, as defined in (3). It follows that:

$$\ln \left(\frac{\bar{P}_t}{\bar{P}_{t-1}} \right) = \ln \left(\frac{P_t^*}{P_{t-1}^*} \right) + \frac{S_{t-1} - S_t}{\omega_0},$$

which implies:

$$\text{var} \left(\ln \left(\frac{\bar{P}_t}{\bar{P}_{t-1}} \right) \right) = \sigma_{TP}^2 + \frac{1}{\omega_0^2} \text{var}(S_t - S_{t-1}) \quad (24)$$

The covariance stationary process for S_t that satisfies (2) is such that:

$\text{var}(S_t) = \text{var}(S_{t-1}) = \frac{\sigma_S^2}{1-\kappa^2}$, with $\text{cov}(S_t, S_{t-1}) = \frac{\kappa\sigma_S^2}{1-\kappa^2}$. Eq. (24) therefore gives:

$$\text{var} \left(\ln \left(\frac{\bar{P}_t}{\bar{P}_{t-1}} \right) \right) = \sigma_{TP}^2 + \frac{2\sigma_S^2}{(1+\kappa)\omega_0^2}. \quad (25)$$

Producer's action stabilizes the price if and only if $var\left(\ln\left(\frac{\bar{P}_t}{\bar{P}_{t-1}}\right)\right) > var\left(\ln\left(\frac{P_t}{P_{t-1}}\right)\right)$.

According to (9) and (25), this will occur if and only if: $\frac{2\sigma_S^2}{(1+\kappa)\omega_0^2} > \frac{2\sigma_Z^2}{\omega_0^2}$, which requires:

$$\sigma_Z^2 < \sigma_S^2 / (1 + \kappa). \quad (26)$$

This condition highlights the fact that the Producer's ability to stabilize the price depends only on the precision of its estimate relative to the volatility and persistence of shocks; it does not depend on the elasticity of demand. Our estimate of the precision of Producer's estimate, however, is conditioned on the presumed elasticity, as shown in Table 2. Condition (26) for price stabilization therefore appears to be satisfied under certain of our elasticity scenarios (e.g., Saudi Arabia and OPEC Core when the price elasticity of global demand is presumed to be -1%, and for OPEC when the elasticity is presumed to be -1% or -3%), but not in others (see Table 3 and Appendix Tables A1-A3).

We now conduct an independent counterfactual experiment to examine OPEC's actual impact on price volatility over our sample period. This counterfactual experiment does not use any of our previous estimates or results.

The observed price P_t is such that: $a_t(\prod_{k=0}^K P_{t-k}^{\omega_k})e^{S_t} = \tilde{Q}_t$. But, if Producer had been content only to meet its expected call, Q_t^* , then the quantity $\frac{\tilde{Q}_t + X_t}{B}$, instead of \tilde{Q}_t , would have been put on the market each month. The counterfactual prices \bar{P}_t therefore satisfy:

$$a_t(\prod_{k=0}^K \bar{P}_{t-k}^{\omega_k})e^{S_t} = \frac{\tilde{Q}_t + X_t}{B}.$$

By combining these two equations we obtain: $B \prod_{k=0}^K \bar{P}_{t-k}^{\omega_k} = \left(1 + \frac{X_t}{\tilde{Q}_t}\right) \prod_{k=0}^K P_{t-k}^{\omega_k}$,

which, after taking logs, lagging once, and then differencing, implies:

$$\sum_{k=0}^K \omega_k \ln\left(\frac{\bar{P}_{t-k}}{\bar{P}_{t-k-1}}\right) = \sum_{k=0}^K \omega_k \ln\left(\frac{P_{t-k}}{P_{t-k-1}}\right) + \ln\left(1 + \frac{X_t}{\tilde{Q}_t}\right) - \ln\left(1 + \frac{X_{t-1}}{\tilde{Q}_{t-1}}\right). \quad (27)$$

Conditional on the presumed elasticities (ω_k), which are the only unobservable variables on the right-hand side, the counterfactual price series can be calculated recursively from starting values and actual prices using the following transformation of (27):

$$\ln\left(\frac{\bar{P}_t}{\bar{P}_{t-1}}\right) = \ln\left(\frac{P_t}{P_{t-1}}\right) + \frac{1}{\omega_0} \ln\left[\left(1 + \frac{X_t}{\bar{Q}_t}\right) / \left(1 + \frac{X_{t-1}}{\bar{Q}_{t-1}}\right)\right] + \sum_{k=1}^K \frac{\omega_k}{\omega_0} \left[\ln\left(\frac{P_{t-k}}{P_{t-k-1}}\right) - \ln\left(\frac{\bar{P}_{t-k}}{\bar{P}_{t-k-1}}\right)\right].$$

For our calculations, we consider the same range of values for short-run elasticities as before, assuming in all scenarios that both global oil demand and non-Producer supplies have a long-run elasticity of -0.3. We consider $K = 48$, $\omega_1, \dots, \omega_{48}$ being assumed to be identical and equal to the difference between the long-run and monthly elasticities of residual demand divided by 48. When performing the recursion, counterfactual and observed prices are considered to be identical before the start of the sample. The first 48 counterfactual log-price returns are discarded when the counterfactual volatility is computed. We report the resulting counterfactual monthly volatility in Table 6. In all scenarios, the counterfactual volatility exceeds the (factual) volatility of 8.41% historically observed between September 2005 and October 2014, which is to say that OPEC's utilization of spare capacity appears to have damped price movements. Using the same methodology, we have performed an extended sensitivity analysis showing that this result still holds when different values for K and long-run elasticities are considered. This indicates that condition (26) must be satisfied in practice. It also casts doubt on the presumption that the short-run elasticity of global demand deviates much from zero, since values that exceed $|\varepsilon_d| = 3\%$ imply estimates of σ_Z that fail condition (26), which would be a contradiction of Table 6.¹⁵

Our counterfactual experiment does not reveal the specific identity of the “Producer” whose actions have succeeded in stabilizing the market, be it Saudi Arabia, the Core, or OPEC

¹⁵ Even if we treat the reported values of $\hat{\sigma}_Z$ as upper bounds on the estimation error and factor out the contribution due to volatility of the target price, per footnote 8, the adjusted estimates of σ_Z would still violate condition (26) and contradict the results shown in Table 6 unless the elasticity of global demand is close to zero. We note that Nakov and Nuño (2013) experience similar difficulty simulating historical price and output volatilities when a high demand elasticity (−0.26) is imposed on their model.

acting all together. Based on other evidence, however, one may doubt that OPEC as a whole has played this role. Many OPEC members are reported to produce continuously at full capacity. Whether it has been Saudi Arabia or the OPEC Core acting as swing producer makes little difference, at least according to the estimates shown in Table 6. The counterfactual volatilities are similar in both scenarios. Our estimates indicate that Saudi/Core intervention has damped oil price volatility by some 21% to 28% if the monthly elasticity of global demand is thought to be -3%, or by 60% to 69% if the elasticity is thought to be -1%.¹⁶ Differences in assumptions regarding the elasticity of global demand and non-OPEC supply translate into big differences regarding the extent to which OPEC appears to have stabilized the market price. One's view of the impact of OPEC's efforts to stabilize the price clearly runs inversely to one's opinion about short-run elasticities of demand and supply, which makes the elasticity a worthy subject of further research.

7. Concluding remarks

The present paper is the first attempt to fit a structural model to the behavior of OPEC's spare capacity. Although discussions of oil price dynamics frequently mention this factor, as yet there has been no quantitative investigation of the determinants of the size or impact of OPEC's spare capacity. To that end, we have constructed a model having three main components: an autoregressive stochastic process by which the residual demand for OPEC oil is shocked, a separate stochastic process by which OPEC attempts to estimate the size of such shocks and offset them by regulating production from its buffer stock, and finally, an implicit loss function that rationalizes the observed size of OPEC's chosen buffer—and which can be compared to independent assessments of the global economic cost of oil supply disruptions.

By estimating the parameters of this model with monthly data, we obtain plausible results regarding the size and persistence of demand and supply shocks that impact the global

¹⁶ The 21% reduction in volatility is calculated from Table 6 as $(10.6\% - 8.4\%)/10.6\%$, etc.

oil market, plausible estimates of the precision (or lack thereof) of OPEC's ability to estimate and offset shocks, and plausible estimates of the scope of OPEC's implicit concern for the economic costs that price shocks impose on the global economy. We also perform a counterfactual experiment to calculate the historical impact of OPEC's use of spare capacity. Depending on one's particular beliefs regarding the short-run elasticity of global demand and supply for oil, OPEC's impact may be viewed as large or small—but in all cases OPEC appears to have at least partially offset shocks and stabilized the price. Under plausible assumptions regarding the elasticity of demand, OPEC's stabilizing influence appears to have been very substantial, with indications that Saudi Arabia may have acted as a supplier of last resort and, relative to the size of the residual demand for its oil, absorbed more shocks than the other OPEC members.

In this paper, we have abstracted from the potential impact of price volatility on the formation of production capacity outside OPEC and on fuel substitution in demand. In addition, our study has focused on the past. We do not overlook the strategic change within OPEC in late 2014 to rebalance the market, but that episode followed the end of our sample period. In any event, OPEC appears to have resumed its role in helping to stabilize the market, albeit at a lower target price.

List of References

- Alhajji AF, Huettner D. 2000. OPEC and world crude oil markets from 1973 to 1994: cartel, oligopoly, or competitive? *Energy Journal* **21(3)**: 31-60. DOI: 10.5547/ISSN0195-6574-EJ-Vol21-No3-2
- Baumeister C, Peersman G. 2013a. Time varying effects of oil supply shocks on the US economy. *American Economic Journal: Macroeconomics* **5(4)**: 1-28. DOI: 10.1257/mac.5.4.1
- Baumeister C, Peersman G. 2013b. The role of time-varying price elasticities in accounting for volatility changes in the crude oil market. *Journal of Applied Econometrics* **28**: 1087-1109. DOI: 10.1002/jae.2283
- Brown S, Huntington HG. 2015. Evaluating U.S. oil security and import reliance. *Energy Policy* **79(C)**: 9-22.
- De Santis RA. 2003. Crude oil price fluctuations and Saudi Arabia's behaviour. *Energy Economics* **25**: 155-173.
- Difiglio C. 2014. Oil, economic growth and strategic petroleum stocks. *Energy Strategy Reviews* **5**: 48-58. DOI: 10.1016/j.esr.2014.10.004
- Fattouh B. 2006. Spare capacity and oil price dynamics. *Middle East Economic Survey* **49(5)**.
- Fattouh B, Mahadeva L. 2013. OPEC: what difference has it made? Oxford Institute for Energy Studies, MEP 3.
- Golombek R, Irarrazabal AA, Ma L. 2016. OPEC's market power: An empirical dominant firm model for the oil market, working paper, September 12.
- Hamilton JD. 1994. *Time Series Analysis*. Princeton University Press.
- Hamilton JD. 2009. Causes and consequences of the oil shock of 2007-08. *Brookings Papers on Economic Activity*.
- Henni A. 2013. Huge Saudi field begins production ahead of schedule. *Journal of Petroleum Technology* **65(6)**: 56-59.
- International Energy Forum (Cancun). 2010. Global energy markets: reducing volatility and uncertainty. Remarks by H.E. Minister of petroleum Ali Al-Naimi, available at <https://www.ief.org/resources/files/events/12th-ief-ministerial-cancun-mexico/presentation-he-ali-l-al-naimi.pdf>.
- Jaffe AM, Soligo R. 2002. The role of inventories in oil market stability. *Quarterly Review of Economics and Finance* **42**: 410-415.
- Johnson N, Kotz S, Balakrishnan N. 1995. *Continuous Univariate Distributions*, Second edition, Wiley.

Karp L, Newberry D. 1991. OPEC and the U.S. oil import tariff. *Economic Journal* **101**: 303-313. DOI: 10.2307/2233820

Kilian L. 2008. Exogenous oil supply shocks: how big are they and how much do they matter for the US economy? *Review of Economics and Statistics* **90**(2): 216-240. doi:10.1162/rest.90.2.216

Kilian L, Murphy DP. 2014. The role of inventories and speculative trading in the global market for crude oil. *Journal of Applied Econometrics* **29**(3): 454-478. DOI: 10.1002/jae.2322

Minister of Petroleum Ali Al-Naimi address at CERAWeek. 2009. Available at <http://www.saudiembassy.net/announcement/announcement02100901.aspx>.

Munro D. 2014. 'Effective' OPEC spare capacity. *IEA Journal*, 7-13.

Nakov A, Nuño G. 2013. Saudi Arabia and the oil market. *The Economic Journal* **123**: 1333-1362. DOI: 10.1111/ecoj.12031

Parry IW, Darmstadter J. 2003. The costs of U.S. oil dependency. Resources for the Future, Discussion Paper 03-59.

Petroleum Economist. 2005. Saudi's Al-Naimi: deliverability, not availability, is oil's problem. October 1.

Petroleum Economist. 2008a. OPEC on the attack. August 1.

Petroleum Economist. 2008b. Saudi Arabia rethinks expansion plans. November 25.

Petroleum Economist. 2012. Saudi Arabia remains the supply-side optimist. October 30.

Petroleum Economist. 2013. We need new thinking for the new realities we face. October 15.

Petroleum Intelligence Weekly. 2011. Understanding Saudi Arabia's spare capacity. March 7.

Pierru A, Matar W. 2014. The impact of oil price volatility on welfare in the Kingdom of Saudi Arabia: implications for public investment decision-making. *Energy Journal* **35**: 97-116. DOI: 10.5547/01956574.35.2.5

Reza AM. 1984. The price of oil and conflict in OPEC. *Energy Journal* **5**: 29-33. DOI: 10.5547/ISSN0195-6574-EJ-Vol5-No2-2

Saudi Gazette. 2013. Maintaining spare capacity relevant? November 10.

Smith JL. 2009. World oil: market or mayhem? *Journal of Economic Perspectives* **23**(3): 145-164. DOI: 10.1257/jep.23.3.145

Figure 1a: Price formation when buffer is sufficient to fully absorb shocks




Figure 1b: Price formation when buffer is not sufficient to fully absorb shocks




Figure 2: Spare capacity (million barrels/day)





Figure 3: Estimated cost of oil supply shortfalls

Dashed lines: inferred from OPEC Core's behavior; Solid line: inferred from Oxford Economics model.



	Saudi Arabia	OPEC Core	OPEC	Elasticity of global demand
Avg. production (mmb/d)	8.66	14.56	29.57	
Avg. market share	11.7%	19.7%	40.1%	
Implied monthly elasticity of residual demand (ω_0)	-0.16	-0.09	-0.04	-1%
	-0.48	-0.27	-0.12	-3%
	-0.80	-0.46	-0.20	-5%
	-2.22	-1.32	-0.65	-26%

Table 1: Implied short-run elasticity of residual demand (mmb/d=million barrels per day)

Elasticity of global demand	Saudi Arabia		OPEC Core		OPEC	
	$\hat{\sigma}_z$	mmb/d*	$\hat{\sigma}_z$	mmb/d*	$\hat{\sigma}_z$	mmb/d*
-1%	0.97%	0.084	0.55%	0.081	0.24%	0.072
-3%	2.92%	0.253	1.66%	0.242	0.73%	0.215
-5%	4.86%	0.421	2.77%	0.403	1.21%	0.358
-26%	13.47%	1.162	8.01%	1.162	3.94%	1.162

Table 2: Estimation error based on observed monthly price volatility

* $\hat{\sigma}_z$ times average own crude oil production.

	<i>Saudi Arabia</i>	<i>OPEC Core</i>	<i>OPEC</i>
σ_z	0.97%	0.55%	0.24%
κ	0.973 (0.017)	0.972 (0.016)	0.972 (0.018)
B	1.220 (0.082)	1.152 (0.061)	1.088 (0.037)
σ_s	2.3% (0.2%)	1.8% (0.1%)	1.2% (0.1%)

Note: standard errors in parentheses

Table 3: Maximum likelihood estimates if $\varepsilon_D = -1\%$

Elasticity of global demand	Saudi Arabia	OPEC Core	OPEC
-1%	-\$27,991	-\$73,682	-\$268,930
-3%	-\$6,675	-\$17,357	-\$75,641
-5%	-\$1,925	-\$9,116	-\$39,537
-26%	\$11,784	\$3,554	-\$9,315

Table 4: Present value of financial gains generated by availability of an incremental barrel of buffer capacity (f)

Elasticity of global demand	Saudi Arabia	OPEC Core	OPEC
-1%	15.66	16.78	25.01
-3%	7.98	8.63	9.95
-5%	5.84	5.65	6.77
-26%	0.39	0.90	1.64

Table 5: Parameter α

Short-run elasticity of global demand	Saudi Arabia	OPEC Core	OPEC
-1%	21.1%	26.7%	36.8%
-3%	10.6%	11.6%	13.3%
-5%	9.5%	10.0%	10.7%
-26%	8.7%	8.8%	8.9%

Table 6: Counterfactual monthly volatility without OPEC buffer. Observed (factual) volatility = 8.4%.

Appendix 1. Log-likelihood maximization

From (12), by setting $d_t = -\ln\left(1 + \frac{x_t}{\hat{Q}_t}\right)$ we have:

$$d_t = S_t + \sigma_z z_t - \ln(B)$$

$$d_{t-1} = S_{t-1} + \sigma_z z_{t-1} - \ln(B)$$

Equivalently:

$$S_t = d_t - \sigma_z z_t + \ln(B)$$

$$S_{t-1} = d_{t-1} - \sigma_z z_{t-1} + \ln(B)$$

By replacing S_t and S_{t-1} in (2) we get:

$$d_t - \sigma_z z_t + \ln(B) = \kappa(d_{t-1} - \sigma_z z_{t-1} + \ln(B)) + \sigma_S u_t$$

which gives:

$$d_t = (\kappa - 1)\ln(B) + \kappa d_{t-1} + \sigma_S u_t - \kappa \sigma_z z_{t-1} + \sigma_z z_t. \quad (\text{A1})$$

We set:

$$W_t = \sigma_S u_t - \kappa \sigma_z z_{t-1} + \sigma_z z_t = w_t \sqrt{\sigma_S^2 + \sigma_z^2(\kappa^2 + 1)} \quad (\text{A2})$$

where w_t is a standard normal variate.

(A1) can be rewritten:

$$d_t = (\kappa - 1)\ln(B) + \kappa d_{t-1} + W_t \quad (\text{A3})$$

with:

$$\text{cov}(W_{t-1}, W_t) = -\kappa \sigma_z^2 \quad (\text{A4})$$

(A2), (A3) and (A4) allow for defining the log-likelihood function as the natural logarithm of the density of a multivariate normal law with a “tridiagonal” covariance matrix (variances on the main diagonal and covariance terms on the two adjacent diagonals). The estimates are the parameter values that maximize the log-likelihood function and their standard errors are derived from the Hessian matrix of the log-likelihood function. The MATLAB code is available upon request.

Appendix 2. Robustness of results to the value assumed for ε_D

	<i>Saudi Arabia</i>	<i>OPEC Core</i>	<i>OPEC</i>
σ_z	2.92%	1.66%	0.73%
κ	0.971 (0.018)	0.971 (0.016)	0.973 (0.016)
B	1.215 (0.077)	1.149 (0.058)	1.085 (0.036)
σ_S	2.3% (0.2%)	1.7% (0.1%)	1.1% (0.1%)

Note: standard errors in parentheses

Table A1: Maximum likelihood estimates if $\varepsilon_D = -3\%$

	<i>Saudi Arabia</i>	<i>OPEC Core</i>	<i>OPEC</i>
σ_z	4.86%	2.77%	1.21%
κ	0.967 (0.019)	0.967 (0.018)	0.971 (0.017)
B	1.214 (0.070)	1.150 (0.053)	1.085 (0.033)
σ_S	2.3% (0.3%)	1.8% (0.2%)	1.1% (0.1%)

Note: standard errors in parentheses

Table A2: Maximum likelihood estimates if $\varepsilon_D = -5\%$

	<i>Saudi Arabia</i>	<i>OPEC Core</i>	<i>OPEC</i>
σ_z	13.47%	8.01%	3.94%
κ	0.953 (0.027)	0.953 (0.023)	0.956 (0.021)
B	1.218 (0.055)	1.154 (0.041)	1.087 (0.026)
σ_S	2.4% (0.5%)	1.9% (0.3%)	1.2% (0.2%)

Note: standard errors in parentheses

Table A3: Maximum likelihood estimates if $\varepsilon_D = -26\%$

Appendix 3. Incremental value of spare capacity

We calculate the expected loss as a function of the size of the buffer:

$$\begin{aligned} E[L|B] &= \alpha E \left[\sum_1^T \frac{(O_t)^2}{(1+r)^t} \right] = \alpha \left[\sum_1^T \frac{E[\max\{0, (e^{S_t + \sigma_z Z_t} - B) Q_t^*\}^2]}{(1+r)^t} \right] \\ &= \alpha \sum_{t=1}^T \frac{(Q_t^*)^2 \int_{\ln(B)}^{\infty} (e^{\xi} - B)^2 g_t(\xi) d\xi}{(1+r)^t} \end{aligned}$$

Let us now determine the value of increasing the size of the buffer. The incremental value, v , of spare capacity is given by the first derivative of the expected loss:

$$v = -\frac{\partial E[L|B]}{\partial B} = -\alpha \sum_{t=1}^T \frac{(Q_t^*)^2 \frac{\partial}{\partial B} \left(\int_{\ln(B)}^{\infty} (e^{\xi} - B)^2 g_t(\xi) d\xi \right)}{(1+r)^t}$$

Since by application of Leibniz Rule:

$$\frac{\partial}{\partial B} \left(\int_{\ln(B)}^{\infty} (e^{\xi} - B)^2 g_t(\xi) d\xi \right) = -2 \int_{\ln(B)}^{\infty} (e^{\xi} - B) g_t(\xi) d\xi$$

We obtain (17): $v = 2\alpha \sum_{t=1}^T \frac{E[O_t|B] Q_t^*}{(1+r)^t}$

This implies:

$$\frac{\partial v}{\partial B} = 2\alpha \sum_{t=1}^T \frac{(Q_t^*)^2 \frac{\partial}{\partial B} \left(\int_{\ln(B)}^{\infty} (e^{\xi} - B) g_t(\xi) d\xi \right)}{(1+r)^t} = -2\alpha \sum_{t=1}^T \frac{\varphi_t(B) (Q_t^*)^2}{(1+r)^t} < 0$$

Appendix 4. Use of Oxford Economics to simulate the impact of an oil supply shortfall

Oxford Economics' Global Economic Model (<http://www.oxfordeconomics.com/>) formed our estimates of GDP losses due to oil supply disruptions. The reference scenario is the version released in November 2014. To simulate the impact of a supply shock of given size and duration, we reduce non-OPEC oil supply accordingly, starting in the first quarter of 2015, while OPEC Oil Supply is kept the same as in the reference scenario until the fourth quarter of 2016; i.e. OPEC is not allowed to make up the shortfall in non-OPEC production. We observe global GDP for each quarter from 2015 to 2020 (virtually all impacts are realized during that interval) and the present value of cumulative GDP losses—relative to the reference scenario and expressed in real terms (2010 prices)—is computed using a 4% annual discount rate. The resulting cumulative GDP change is multiplied by Oxford Economics' world GDP deflator from year 2010 to year 2015 in order to be expressed in 2015 prices. The results are shown in Table A4.

Duration (Months)	Size (mmb/d)	Cumulative world GDP loss (Billion US dollars)
3	0.5	12.68
3	0.961	30.80
3	2	93.25
6	0.2	5.42
6	0.5	22.36
6	0.961	58.56
6	1.5	165.85
6	2	322.62
6	3	715.49
6	4	1239.94
6	5	1877.54
9	0.5	37.56
9	0.961	107.44
9	2	570.44
12	0.5	57.16
12	0.961	154.53
12	2	721.92

Table A4: Cumulative world GDP loss for various oil supply shocks (source: Oxford Economics)