CENTRE FOR APPLIED MACROECONOMICS AND COMMODITY PRICES (CAMP)

CAMP Working Paper Series
No 10/2019

OPEC's crude game

The supply curve in a dynamic, strategic environment

Even Comfort Hvinden

OPEC's crude game

The supply curve in a dynamic, strategic environment *

Even Comfort Hvinden †

November 25, 2019

The market behavior nationalized oil companies in the Organization of Petroleum Exporting Countries (OPEC) is starkly time-varying. I rationalize OPEC's behavior in an infinitely repeated game of Cournot competition with imperfect monitoring, capacity constraints to output, and demand evolving as a Markov chain. I adapt the methodology of Abreu, Pearce, and Stacchetti (1990) to derive optimal symmetric equilibria. High-powered incentives are created by the threat of output wars, the severity of which is endogenously determined by current and future expected market conditions. Implied price elasticities of supply increase in magnitude and may change sign under constrained incentive creation. The key empirical implication is that unanticipated changes to OPEC's strategic environment will persistently alter their behavior and create breaks in the joint stochastic distribution of equilibrium prices and quantities.

*This work is part of the research activities at the Centre for Applied Macro- and Commodity Prices (CAMP) at the BI Norwegian Business School. I thank Jørgen Juel Andersen, Arthur van Benthem, Andrew Comfort, Thomas Størdal Gundersen, Bård Harstad, Martin Blomhoff Holm, Ragnar Enger Juelsrud, Plamen Nenov, Alessia Russo, Juuso Toikka, Ragnar Torvik, Dario Sidhu, Daniel Spiro, and various seminar participants for helpful comments. The MATLAB code is available upon request.

†E-mail: even.c.hvinden@gmail.com.

1 Introduction

"OPEC is strong when prices are weak, and weak when prices are strong."

Sadek Boussena, OPEC conference president, 1989-1990, cited in Bret-Rouzaut and Favennec (2011)

The market for crude oil is large, volatile and of considerable importance to the global economy. Understanding the oil price-quantity relationships has therefore been a subject of long-standing interest in the literature. Among others, a stand must be taken on the Organization of Petroleum Exporting Countries' (OPEC) ability to affect market outcomes through their output decisions. It is widely agreed that OPEC members are endowed with considerable market power, strive actively to increase their profits by coordinating output restraint, and face no legal constraints on collusion. Yet empirical evidence suggests output constraint has been only partially successful, with OPEC conduct varying starkly over time. This has led to calls for richer models of oligopolistic competition that may shed light on evolving OPEC behavior.[1] This paper takes a first step in this direction, asking: How does non-cooperative oligopolists' supply respond optimally to changes in the market environment?

To shine light on this question I consider a model of quantity competition where production capacity is finite, output is imperfectly observed, and the inverse demand function is dynamic, evolving according to a Markov chain. I then study the properties of *optimal* symmetric public Markov perfect equilibria. In equilibrium, the oligopolists' ability to restrain output when cooperating, and the frequency, intensity, and duration of output wars when not, are jointly and endogenously determined with respect to current- and future expected market conditions. Output wars optimally feature short periods of elevated supply and significantly depressed prices. The output wars' frequency and duration is increasing, and intensity decreasing, in the residual demand for the oligopolists' product. Under cooperative periods an increase in demand will affect both the one-shot deviation pay-off and the (probabilistic) cost of initiating an output war. The net change in incentive power depends on the prevailing monitoring quality and the expected future profitability. The incentive power of punishments wilts under contemporaneous demand pressure when information is poor or if there are expectations of diminishing future profits. A second-order effect of falling cartel discipline is a further reduction in the power of incentives. This "unraveling" of incentive power generates concave and even non-monotonic price-quantity relationships in cooperative periods, with price elasticities of supply diverging locally to positive or negative

[1] For example, concluding their review of OPEC's output policies and past modeling efforts, Fattouh and Mahadeva (2013) write that "[the] evolution of OPEC behavior indicates that OPEC's conduct is not constant. [...] This also explains the failure of empirical studies to reach more concrete conclusions: Although some [models] may fit the data quite well in specific time periods, they fail miserably in [others]. Hence, this review emphasizes the importance of relying on dynamic models that allow for changes in OPEC behavior." See also the concluding remarks in Griffin (1985).

infinity.

The main insight is that optimal oligopolistic output is generally not observationally equivalent with competitive behavior when incentives for output restraint are dynamically constrained. The magnitude and sign of the supply elasticity then depends sensitively and non-linearly on monitoring technology and the expected evolution of future demand. Unanticipated changes to OPEC's strategic environment may yield persistent and substantial shifts in their behavior.

The implication for applied work is that historical price-quantity relationships under strategic competition, time-averaged over long samples, are an uninformative summary statistic of current and future behavior. Indeed, existing empirical estimates of aggregate supply elasticities are contentious and sensitive to choices of sample period, model specification, and identifying assumptions.[2] My analysis suggests that exploiting the state-contingent properties of OPEC's output choices will improve inference and our understanding of the crude oil market.

The paper relates to three strands of literature. First and foremost I join in a long-standing effort to apply models of imperfect competition to shed light on OPEC behavior. A closely related work is Rauscher (1992), who analyzes OPEC's supply when cartel discipline is exogenously assumed proportional to underlying demand. Other notable contributions are Salant (1976), Hnyilicza and Pindyck (1976), Greene (1991), Huppmann (2013), Nakov and Nuño (2013), and Behar and Ritz (2017). The common theme in this work is that a representative OPEC producer competes inter-temporally with a non-OPEC fringe. The non-cooperative aspects of OPEC members' interaction is not modeled and variation in cartel discipline is absent or exogenously imposed. In contrast, my focus here is on how output discipline is endogenously determined by underlying conditions.

Second, the analysis informs a long-running debate on the identification of supply- and demand shocks in the oil market, see Kilian (2009), Kilian and Murphy (2014), Aastveit, Bjørnland, and Thorsrud (2015), with more recent contributions by Caldara, Cavallo, and Iacoviello (2019), Fueki et al. (2018), Baumeister and Hamilton (2019a). My analysis implies that under oligopolistic competition in a rich strategic environment, structural relationships will not be identified under the assumption of fixed supply elasticities.

Third, this paper relates to an expansive literature on the determinants of collusion, see Green and Porter (1984) on monitoring quality, Brock and Scheinkman (1985), Fabra (2006) on capacity constraints, Rotemberg and Saloner (1986), Haltiwanger and Harrington Jr (1991), Wilson and Reynolds (2005) on the effect of transitory demand shocks on firm behavior, and Abreu, Pearce, and Stacchetti (1990) on optimal equilibria under imperfect monitoring. I contribute to this literature not by originally recognizing the significance of these individual components, but by combining them in a common framework with a rich action set. The model I consider is quite generalizable, so the insights from this paper may therefore be of broader interest.

[2] See Kilian and Murphy (2014), Baumeister and Hamilton (2019a), Kilian and Zhou (2019), Baumeister and Hamilton (2019b), and Kilian (2019).

I proceed as follows. Section 2 considers the stylized facts of OPEC's time-varying behavior. I critically discuss the literatures' existing interpretations to motivate my analysis. Sections 3 and 4 present model primitives and solution concept, respectively. The endogenous cartel discipline is analyzed in Section 6. I conclude with suggestions for further research in Section 7.

2 OPEC's market power and time-varying behavior

A handful of oil companies, mainly but not exclusively the nationalized oil producers of OPEC member nations, are widely viewed as enjoying considerable market power.[3][4] There is strong evidence that OPEC's member states have systematically restrained production, but that the extent of collusion is less than perfect and that their conduct is temporally unstable.[5] For a stylized illustration, consider Figure 1, plotting twelve-month changes in monthly OPEC crude oil output and log real crude oil prices between January 1985 and October 2019.[6] I have highlighted by shaded bars four significant episodes of oil price drops: 1986, 1997, 2008, and 2014. Measured across the entire 1985-2019 sample, OPEC output and price developments are linearly uncorrelated. However, this masks significant and sign-varying correlation in sub-samples. Between 2002 and 2014, OPEC's out-

[3]OPEC members' produce at lower cost, higher capacity, and greater flexibility relative to their competitors, and thus may unilaterally affect equilibrium prices. Al-Qahtani, Balistreri, and Dahl (2008) comprehensively review the evidence of cartel behavior accumulated up to 2008. Among others, empirical studies that reject both the price-taking and price-setting hypotheses of OPEC behavior in favor of a dominant firm, competitive fringe set-up are Alhajji and Huettner (2000), Spilimbergo (2001), Hansen and Lindholt (2008), and Golombek, Irarrazabal, and Ma (2018). See also Huppmann and Holz (2015).

[4]For the purposes of this paper, the identity of oligopolistic firms is held fixed. The question of which companies join oligopolistic agreements may be an avenue of future research. For example Rosneft, a nationalized Russian oil producer, is a plausible non-OPEC candidate for a dominant player. It has been reported that Russia coordinated output cuts with OPEC following the 2014 price fall, leading to the so-called "OPEC+" format, e.g. Astakhova, Olga and El Gamal, Rania: Russia, Saudi Arabia agree OPEC+ format should be extended, accessed August 21 2018 from www.reuters.com.

[5]Various econometric techniques have been applied to explicitly estimate the time-variation in supply behavior. These include regression switching models, unit-root econometrics, structural estimation of dynamic Stackelberg competition with non-OPEC firms, and sample splitting, see Almoguera, Douglas, and Herrera (2011), Barros, Gil-Alana, and Payne (2011), Kolodzeij and Kaufmann (2014), Huppmann and Holz, 2012, Ratti and Vespignani (2015) respectively. See in particular Baumeister and Peersman (2013) who estimates a time-varying parameter vector- autoregressive model aggregating OPEC and non-OPEC output, but explicitly attributes variation in estimated supply behavior to changes in OPEC's ability to cooperate. See also the discussion Dees et al. (2007). All the aforementioned contributions find strong evidence of persistent time-variation in OPEC behavior. Finally Dvir and Rogoff (2009) and Dvir and Rogoff (2014) consider long-run variation in market power, studying samples that predate the formation of OPEC.

[6]Monthly data on crude oil production is from the International Energy Agency's Monthly Oil Data Service. To construct a real oil price series I have deflated the U.S. crude coil composite acquisition cost by refiners from the Energy Information Agency with the average all-item CPI from the Federal Reserve Bank of St. Louis data service.

put and price changes are strongly and positively correlated. Particularly, during the price collapse accompanying the 2008 global financial crisis, OPEC rapidly restricted, and only gradually increased, output. Contrast this to their actions during the 1986, 1997 and 2014 episodes, where OPEC production sharply increases in the face of collapsing prices. The correlation here is sharply negative. Why does OPEC's behavior varies so readily over time? To what end is OPEC apparently flooding the market, as suggested by the sometimes negative correlation during steep price drops?

Two interpretations of OPEC's capricious behavior have predominated in the literature. One conceives OPEC, or a subset of OPEC members, as a representative actor engaged in intertemporal competition *à la* Stackelberg vs. a price-taking, non-OPEC fringe. The relatively higher volatility in OPEC's behavior is exogenously imposed.[7]. This paper complements those analyses by providing a microfoundation for endogenously determining OPEC's excess output volatility. The second views OPEC as a failed cartel, unable to cooperate since 1986 and behaving as effective price-takers.[8] I briefly review the evidence upholding this view and argue that there are no theoretical grounds to support rejecting the (long-standing) hypothesis that OPEC is a relevant force crude oil market outcomes.

The failed-cartel hypothesis views the exertion of market power as a temporary departure from a competitive equilibrium. The existence of OPEC is largely irrelevant, and there is no need to distinguish between an OPEC-led output cut and, for example, production disruptions due to a hurricane off the U.S. Gulf coast.[9] This view is based chiefly on an interpretation of the evidence Almoguera, Douglas, and Herrera (2011), who study the extent of time-varying cartel discipline in OPEC output. Their analysis presents evidence that following a largely cooperative period between 1974 and 1986, aggregate OPEC output became non-cooperative until 2004, the end of their sample. Baumeister and Kilian (2016, p. 145) write that "[the 1974-1986 period] is the first time in its history (and the only time) that OPEC took a proactive role in trying to influence the price of oil [...]", claiming that the post-1986 non-cooperative period has been absorbing. The notion that cooperation among nationalized oil companies permanently collapsed after 1986 is not consistent with record-high ($> 100\%$) compliance in recent output cuts.[10] Moreover, Baumeister and Kilian claim that the inability to cooperate is *predicted* by theory, writing that "[...] OPEC agreements to jointly restrict oil production in an effort to prop up the price of oil proved ineffective, with many OPEC members cheating on OPEC agreements, as predicted by the economic theory of cartels (for

[7]See the analyses in Greene (1991), Nakov and Nuño (2013), and Bornstein, Krusell, and Rebelo (2019).

[8]See Kilian (2009), Baumeister and Peersman (2013), Kilian and Murphy (2014), and Baumeister and Kilian (2016).

[9]The structural vector autoregression models in this tradition aggregate OPEC and non-OPEC output into a representative, global producer. For instance Kilian and Murphy (2014) describe oil supply shocks in their model as incorporating "supply disruptions associated with exogenous political events in oil-producing countries as well as unexpected politically motivated supply decisions by OPEC members and other flow supply shocks."

[10]See for instance Wingfield, Brian *et al*: OPEC's allies unite on oil cuts, accessed September 12 2019 from www.bloomberg.com.

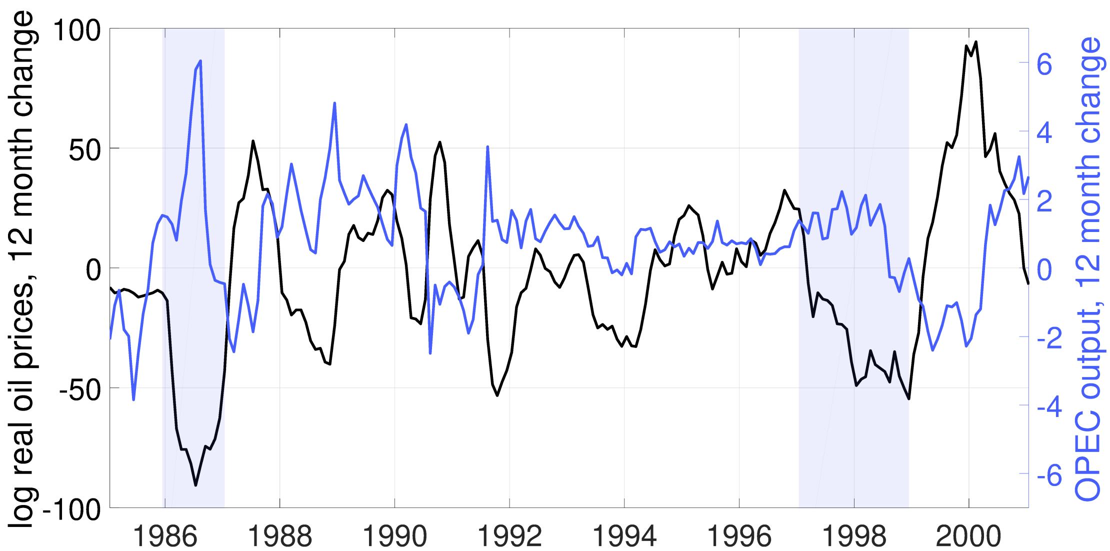

(a) 1985 - 2001

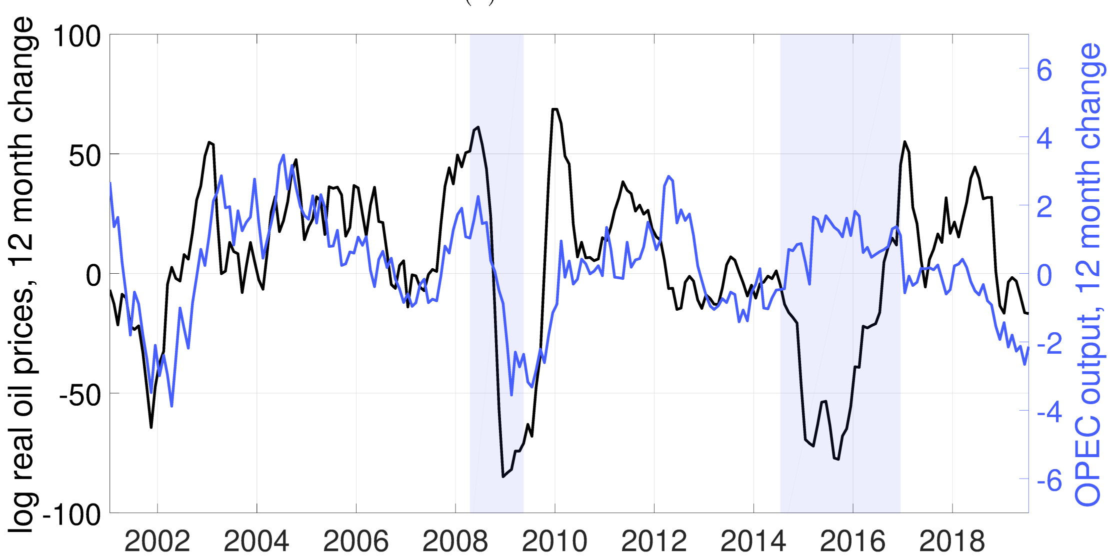

(b) 2001 - 2019

Figure 1: Log real oil prices, September 2019 US dollars per barrel. OPEC crude oil output, millions of barrels per day. Twelve-month change. January 1985 - September 2019. Highlighted historical episodes: 1986 and 1997 output wars, global financial crisis of 2008, and 2014- price fall. Source: International Energy Agency Monthly Oil Data Statistics, Energy Information Agency, Federal Reserve Bank of St. Louis

example, Green and Porter, 1984) [...].” This view is at odds with the standard interpretation of public equilibria in games of imperfect monitoring. By definition, there is no cheating in equilibrium. Interpreting equilibrium-path output wars, Green and Porter (1984, 88–89, my italics) write that “[we show] collusive conduct may [...] result in a pattern of industry performance marked by *recurrent* episodes in which price and profit levels sharply decrease. *Thus we reject the received view that performance of this type necessarily indicates an industry where firms are engaging in a sequence of abortive attempts to form a cartel.*” On the contrary, they continue, in the presence of imperfect monitoring “[...] we point out [... the] necessary appearance of [equilibrium path output wars] if collusion is to take place.” Acknowledging that imperfect monitoring does not predict a necessary collapse of cooperation raises the question of what behavior is, in fact, implied. To make progress a stand must be taken on the relevant properties of OPEC’s environment.

Industry experts and economic historians have argued that the steep price declines in 1986 and 1997 were explicitly due to intentional market flooding by leading OPEC members, punishing peers for quota violations.[11] The incidence of such punishments, or output wars, has directed attention to imperfect monitoring models as a salient framework to study OPEC behavior. The intuition is that OPEC agreements which successfully restrict total production create an incentive for individual producers to cheat. But since OPEC members do not perfectly observe each others’ actions, they cannot know with certainty whether an unexpected, adverse price development resulted from out-of-equilibrium play or not. Incentive compatibility is maintained by equilibrium path punishments.[12] It has been argued that this framework captures salient properties of OPEC’s strategic environment.[13]

I add to imperfect monitoring three additional determinants of strategic competition. First, as the existence and relevance of considerable demand variation is acknowledged to the point of self-evidence, I specify a dynamic residual demand environment.[14] Non-OPEC output is not explicitly modeled, but may be interpreted as subsumed in the dynamic residual demand function. Second, I do not select an equilibrium arbitrarily but derive optimal symmetric equilibria, adapting

[11] See the accounts in Noreng (2006), Downey (2008), Yergin (2011), and also Coll (2012).

[12] This is a general result in the theory of repeated games under imperfect monitoring. See for instance Mailath and Samuelson (2006, p. 233).

[13] See the discussions in Barsky and Kilian (2004) Almoguera, Douglas, and Herrera (2011), and Fattouh and Mahadeva (2013). The idea is that data on crude output is of varying quality and available after a long lag. The imperfect monitoring of OPEC’s output is publicly and transparently endorsed by the International Energy Agency (IEA), see “OPEC Crude Production” in the IEA glossary, accessed October 12 2018 from www.iea.org. The following quotes by Neil Atkinson, chief analyst at IEA, is illustrating: “OPEC, [accounting] for about one-third of global oil output, is a “big black hole [in terms of data],” Mr. Atkinson said. Wary of disclosure that could lead to embarrassments like owning up to cheating on agreed production ceilings, the OPEC member states have not “produced or published reliably transparent data for [many] years.” See Reed, Stanley: Satellites Aid the Chase for Better Information on Oil Supplies, accessed October 12 2018 from www.nytimes.com.

[14] See for instance the Energy Information Agency: What drives crude oil prices? or the many econometric analyses of the crude oil market cited above.

results from Hotelling (1931).[15] The restriction to a simple, symmetric, and short-run framework is motivated by a desire for parsimony. I propose that studies of producer heterogeneity, cartel entry- and exit, alternative sources of informational frictions, capacity investment, and asymmetric equilibria are a promising avenue for future research.

3 Model primitives

Time is discrete, indexed by t over an infinite horizon. The common discount factor is $\delta \in (0, 1)$. Two symmetric, dominant producers compete in homogeneous output. Player $i \in \{1, 2\}$ produces $\tilde{x}_i = x_i + \tilde{h}_i \leq x^{\max}$ from the set $\mathcal{X}$, with $\mathcal{X} = \{0, \epsilon, 2\epsilon, 3\epsilon, ..., x^{\max}\} \subset \mathbb{R}_+$ an evenly ϵ-spaced grid. Quantities x_i are observable, but players may freely sell additional $\tilde{h}_i$ units unobserved by their competitor at no further cost. Let $\tilde{\boldsymbol{x}} \in \mathcal{X}^2$ be the action profile. Production costs are quadratic $c(x) = \kappa x^2$ with $\kappa > 0$. Let $\boldsymbol{\iota}'$ a two-vector of ones so that $\boldsymbol{\iota}'\tilde{\boldsymbol{x}}_t$ is total time-t output. The inverse demand function is defined

$$p(\tilde{\theta}, y, \tilde{\boldsymbol{x}}) := \frac{\tilde{\theta} y}{\xi + \beta \boldsymbol{\iota}' \tilde{\boldsymbol{x}}} \tag{1}$$

parameterized with $y > 0$, $\beta > 0$ and $\xi > 2\beta x^{\max}$.[16] In context it is natural to interpret (1) as representing, in reduced form, the oligopolistic producers' residual demand, subsuming exogenously given competitive output, market growth, and other developments. The unobserved random variable $\tilde{\theta}$ is log-normally distributed, $\ln \tilde{\theta} \sim N\left(-\sigma_\theta^2/2, \sigma_\theta^2\right)$ with independent realizations over t. The distribution and density F_θ and f_θ are commonly known. The parameterization implies expectation 1 and variance $e^{\sigma_\theta^2} - 1$. Prices are then conditionally log-normally distributed

$$\ln p(y, \boldsymbol{x}) \sim N(\ln y - \ln(\kappa + \beta \boldsymbol{\iota}' \boldsymbol{x}) - \sigma_\theta^2/2, \sigma_\theta^2) \tag{2}$$

on $\mathbb{R}_+$ with distribution $F_p(\cdot|\boldsymbol{x})$ and density $f_p(\cdot|\boldsymbol{x})$ and support $\mathbb{R}_+$ independent of actions. The parameter σ_θ governs a mean-preserving spread of the distribution and has a natural interpretation as monitoring quality. Taking expectations over $\tilde{\theta}$, the ex-ante inverse-demand function is then

$$\mathbb{E}_\theta p(\tilde{\theta}_t, y, \tilde{\boldsymbol{x}}) := p(y, \tilde{\boldsymbol{x}}) := \frac{y}{\xi + \beta \boldsymbol{\iota}' \tilde{\boldsymbol{x}}} \tag{3}$$

and is everywhere inelastic, with inverse elasticity

$$-\frac{\partial p(\tilde{\theta}, y, x)}{\partial x} \frac{x}{p} = \frac{\beta x}{\xi + \beta x} < 1 \tag{4}$$

[15] See the discussions in Anderson, Kellogg, and Salant (2018) and Bornstein, Krusell, and Rebelo (2019).

[16] The latter restriction ensures that ξ has a monotonic impact on optimal quantities, which is shown following Equation (7) below.

tending to a unit-elastic demand only in the limit as $x \to \infty$. The property (4) may be taken to tractably encode the assumption that no oligopolistic producer is able to single-handedly push the market into an elastic demand region, thereby simplifying the exposition. Taking expectations over $\tilde{\theta}$ the ex-ante profit function is

$$\pi_i(y, \boldsymbol{x}_t) = \frac{yx_i}{\xi + \beta \boldsymbol{\iota}'\tilde{\boldsymbol{x}}} - \kappa \tilde{x}_i^2 \tag{5}$$

strictly concave with second-order derivative $\partial^2/\partial^2 x^2 \pi_i(y,x) = {-\xi y \beta}/{(\xi+\beta x)^2} - 2\kappa < 0$ everywhere. Fixing demand y, let

$$\boldsymbol{x}^n(y) := \left\{ \boldsymbol{x} : x_i = \arg\max_{x \in \mathcal{X}} \ \pi_i\left(y, x, x_{-i}\right) = x_{-i} \ \forall \, x \in \boldsymbol{x} \right\} \tag{6}$$

be the action profile constituting a symmetric pure strategy, stage-game Nash equilibrium. Because output is discrete, the existence of (6) is not ensured for all $y \in \mathbb{R}$. I restrict attention to demand levels y where every element in $\mathcal{X}$ is a stage-game Nash equilibrium

Assumption 1. ***Existence of stage-game Nash equilibrium.*** *Let* $\mathcal{Y} := \{y : \boldsymbol{x}^n \in \mathcal{X}^2\}$.

which is straightforwardly implemented. Fixing a set of welfare weights $\boldsymbol{\alpha} = (\alpha, 1-\alpha)'$, the highest feasible pay-off is given by

$$\boldsymbol{x}(y, \boldsymbol{\alpha}) := \arg\max_{\boldsymbol{x} \in \mathcal{X}^2} \ \alpha \pi_1(y, \boldsymbol{x}) + (1-\alpha)\pi_2(y, \boldsymbol{x}) \tag{7}$$

where, unless otherwise stated, I take $\alpha = 0.5$, where $\boldsymbol{x}^m := \boldsymbol{x}(y, 0.5)$ is the jointly profit-maximizing or monopoly output.

Notice that an increase in ξ has both a level and slope effect on the inverse price elasticity (4). Increasing ξ makes demand less elastic and reduces the change in elasticity induced by an incremental increase in output. The former (level) effect decreases optimal production (7) while the latter (slope) increases it. The restriction $\xi > 2\beta x^{\max}$ ensures that the level effect dominates.[17] The restricted demand environment may be interpreted as limiting the oligopolists' power such

[17] To verify this, suppose for the moment actions are continuous and profits differentiable in output. The first-order-condition defining (symmetric) monopoly quantities is

$$\frac{x^m(y)(\xi + 2\beta x^m(y))^2}{\xi} = \frac{y}{2\kappa}$$

where ${\partial x^m(y)}/{\partial \xi} < 0$ if the left-hand-side is increasing in ξ. Differentiating, this demands

$$\frac{\partial}{\partial \xi} \frac{x^m(y)(\xi + 2\beta x^m(y))^2}{\xi} > 0 \Rightarrow \xi > 2\beta x^m$$

so restricting $\xi > 2\beta x^{\max} \geq 2\beta x^{\mathrm{m}}$ ensures the monotonic relationship everywhere.

that output restrictions do not push the market into an elastic region and where there is always an individual incentive to increase output. The demand level $y \in \boldsymbol{y} = \{y^1, y^2\} \subset \mathcal{Y}$ may be low or high with $0 < y^1 < y^2$, and evolves as a two-state Markov chain with transition matrix $\boldsymbol{M}$

$$\boldsymbol{M} = \begin{pmatrix} m^1 & 1 - m^1 \\ 1 - m^2 & m^2 \end{pmatrix} \tag{8}$$

stationary and irreducible. In the following, state-j values of endogenous variables are denoted by a j superscript, for example x^j. Finally, the stage game proceeds in the following steps:

1. Demand $y \in \boldsymbol{y}$ is given
2. Players choose actions $\boldsymbol{x} \in \mathcal{X}$
3. Noise $\tilde{\theta}_t$, price $p(\tilde{\theta}_t, y, \tilde{\boldsymbol{x}})$, and profits $\pi(\tilde{\theta}_t, y, \boldsymbol{x})$ are realized

Players condition their actions on the demand-state y and take expectations over the idiosyncratic shock θ.

4 Solution concept

I consider symmetric, Markov public perfect equilibria of the repeated game, that is an equilibrium in strongly symmetric public strategies that condition on the observable, current-valued demand-state y.[18] I refer throughout to the somewhat more general notation and concepts of Abreu, Pearce, and Stacchetti (1990), hereafter APS, with which I assume the reader is familiar.[19] This section aims to succinctly demonstrate that the solution methods for (one-state) public perfect equilibria in APS generalize directly to a time-homogeneous Markov demand environment. That is, by demonstrating that the necessary and sufficient primitive assumptions for APS are satisfied, it is not necessary to recreate their entire line of proof. It is well-known that augmenting a repeated games of imperfect monitoring with a public correlation device leaves the solution concept essentially unchanged, see remarks 2.3.3 and 7.1.4 in Mailath and Samuelson (2006). Intuitively, the generalization to multiple, observable states is equally straightforward because public equilibria already feature a recursive Markov structure in the signal history. Thus the discounted, average pay-offs may be decomposed into a stage-game pay-off and a convex combination of continuation values for reward- and punishment phases, with weights given by transition probabilities and discount rates. The convex combination of continuation values implies that equilibrium pay-offs are interdependently vector-valued and jointly determined. The inclusion of multiple demand states simply requires continuation values to be defined through

[18]See Abreu, Pearce, and Stacchetti (1986) for optimal, symmetric public- and Maskin and Tirole (2001) on Markov perfect equilibria, respectively.

[19]For an excellent and comprehensive introductory treatment of repeated games with imperfect monitoring I refer the reader to Mailath and Samuelson (2006).

yet another recursive convex combination. This operation preserves all the necessary properties (measure, convexity, boundedness, monotonicity) demanded of the functions used in APS to construct optimal public equilibria.

I denote the entire signal space, prices and observable actions, by $\Omega := \mathcal{X}^2 \times \mathbb{R}_+$. Let a history $h_t = \{p_1, \cdots, p_{t-1}, \boldsymbol{x}_1, \cdots, \boldsymbol{x}_{t-1}\}$ be the set of commonly observed signals- and actions available at the beginning of stage t. Let $h_1 = \emptyset$ and $\mathcal{H}_t$ the set of feasible period-t histories. Because the strategies are stationary, relying only on the current signal, I omit the t subscript. Let $\sigma : \mathcal{H} \times \mathcal{Y} \to \mathcal{X}$ a stationary Markov public strategy, prescribing a set of actions for every $t = 1, 2, \cdots$ and state. I denote the by $\sigma^j(h)$ the actions prescribed after history h in state j.

Definition 1. ***Equilibrium.*** *A profile $\boldsymbol{\sigma}$ of Markov public strategies that constitute a sequential equilibrium of the repeated game for all t and $h \in \mathcal{H}$ is a Markov perfect public equilibrium.*

As in the case of a single demand state, every history will yield a well-defined continuation game, so the formulation of sequential rationality is exactly as in APS. Let $\mathcal{S}$ be the set of all Markov PPE. Assumption 1 ensures that a static Nash equilibrium exists in each demand state, so $\mathcal{S}$ is non-empty. Let $v(\boldsymbol{\sigma})$ the pay-off induced by $\boldsymbol{\sigma} \in \mathcal{S}$ and $\mathcal{V} := \{v(\boldsymbol{\sigma}) : \boldsymbol{\sigma} :\in \mathcal{S}\}$, stated in discounted, average terms. This set is bounded, above by repeated play of $\boldsymbol{x}^m(y^j)$ and below, through individual rationality, by a pay-off of 0. Consider the following property:

Definition 2. ***Bang-bang property.*** *A Markov PPE $\boldsymbol{\sigma}$ such that after any history $h \in \mathcal{H}$ the continuation values are extremal, $\boldsymbol{\phi} : \mathcal{X}^2 \times \mathbb{R}_+ \to \mathrm{ext}\mathcal{V}$, is said to be bang-bang.*

The following proposition states that the salient results from APS apply to the Markov generalization.

Proposition 1. ***Optimal equilibria.*** *The unique, efficient symmetric public Markov perfect equilibrium is in bang-bang strategies. It may be computed as the fixed-point of a set-valued contraction mapping $\boldsymbol{B}(\mathcal{V}) = \mathcal{V}$.*

The proof is in Appendix A.1. Knowing that a unique, optimal symmetric Markov public perfect equilibrium exists, I now set out to characterize it in terms of primitive variables, making possible a numerical implementation of the operator $\boldsymbol{B}$.

5 Optimal equilibrium

The optimal, bang-bang strategy may be stated as follows: Begin in the regular phase, playing $\overline{\boldsymbol{x}}_j$, where $j = \{1, 2\}$ denotes the demand state. If a player publicly defects, $x_i \in \overline{X}_j \subset \mathcal{X}_i$, or there is an adverse signal realization, $p_t \in \overline{P}_j \subset \mathbb{R}$, switch to the punishment phase. Play $\underline{\boldsymbol{x}}_j$, remaining there if $x_i \in \underline{X}_j \subset \mathcal{X}_i$ or $p_t \in \underline{P}_j \subset \mathbb{R}$. If not, switch to the reward phase. Let $\tau : \mathcal{H} \times \mathcal{Y} \to [0, 1]$ the transition probability implied by the trigger regions $\overline{X}_j, \underline{X}_j, \overline{P}_j, \underline{P}_j$. I compactly denote a bang-bang strategy by a collection $\boldsymbol{\sigma}_{bb} = \{\overline{\boldsymbol{x}}^j, \underline{\boldsymbol{x}}^j, \overline{\tau}^j, \underline{\tau}^j\}_{j=1}^2$. Average pay-offs under this strategy satisfy the following stationary system

$$\begin{aligned}\overline{\boldsymbol{v}} &= (1-\delta)\cdot\overline{\boldsymbol{\pi}} + \delta\cdot\boldsymbol{M}(\overline{\boldsymbol{\tau}}\cdot\overline{\boldsymbol{v}} + (\boldsymbol{\iota}-\overline{\boldsymbol{\tau}})\cdot\underline{\boldsymbol{v}})\\ \underline{\boldsymbol{v}} &= (1-\delta)\cdot\underline{\boldsymbol{\pi}} + \delta\cdot\boldsymbol{M}(\underline{\boldsymbol{\tau}}\cdot\overline{\boldsymbol{v}} + (\boldsymbol{\iota}-\underline{\boldsymbol{\tau}})\cdot\underline{\boldsymbol{v}})\end{aligned} \tag{9}$$

where $\boldsymbol{v}$, $\boldsymbol{\pi}$, $\boldsymbol{\tau}$ stack present-valued- and stage-game pay-offs, transition probabilities by demand states 1, 2 and $\cdot$ denotes element-wise multiplication. Stacking $\boldsymbol{V} = (\overline{\boldsymbol{v}}, \underline{\boldsymbol{v}})$ and $\boldsymbol{\Pi} = (\overline{\boldsymbol{\pi}}, \underline{\boldsymbol{\pi}})$ the system may be represented as

$$\boldsymbol{V} = (1-\delta)\boldsymbol{\Pi} + \delta\boldsymbol{P}\boldsymbol{V} \tag{10}$$

with equilibrium transition matrix

$$\boldsymbol{P} := \begin{pmatrix}\boldsymbol{M}\cdot\overline{\boldsymbol{T}} & \boldsymbol{M}\cdot(\boldsymbol{I}-\overline{\boldsymbol{T}})\\ \boldsymbol{M}\cdot\underline{\boldsymbol{T}} & \boldsymbol{M}\cdot(\boldsymbol{I}-\underline{\boldsymbol{T}})\end{pmatrix} \tag{11}$$

defined by 2×2 matrices $\overline{\boldsymbol{T}} := (\overline{\boldsymbol{\tau}}, \overline{\boldsymbol{\tau}})$, $\underline{\boldsymbol{T}} := (\underline{\boldsymbol{\tau}}, \underline{\boldsymbol{\tau}})$ and $\boldsymbol{I} = (\boldsymbol{\iota}, \boldsymbol{\iota})$. Element $p_{ij} \in \boldsymbol{P}$ denotes the probability of transitioning from state i to j, see Figure 2.

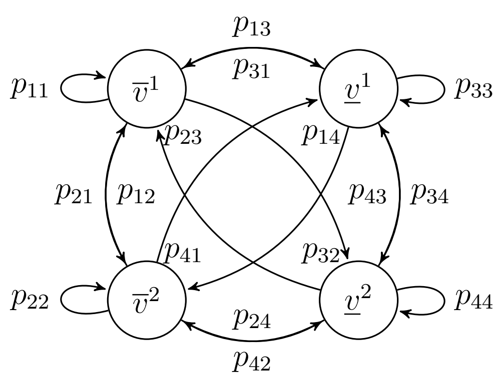

Figure 2: Equilibrium states $\overline{\boldsymbol{v}}$, $\underline{\boldsymbol{v}}$ and transition matrix $\boldsymbol{P}$.

Notice that equations (9) decompose the pay-off into current- and future pay-offs with weights $1-\delta$, δ. In sequential equilibria it is common knowledge that, following any (zero-probability) defection, continuation play follows the initial equilibrium. Taking any scalar element from (9), the generic IC constraint is therefore

$$\begin{gathered}(1-\delta)\pi(y^j, x_i, \overline{x}_{-i}^j) + \delta\tau(x_i, x_{-i}^j)[m^j(\overline{v}^j - \underline{v}^j) + (1-m^j)(\overline{v}^{-j} - \underline{v}^{-j})]\\ \leq\\ (1-\delta)\pi(y^j, \boldsymbol{x}^j) + \delta\tau(\boldsymbol{x}^j)[m^j(\overline{v}^j - \underline{v}^j) + (1-m^j)(\overline{v}^{-j} - \underline{v}^{-j})]\end{gathered} \tag{12}$$

for all $x_i \in \mathcal{X}$ in regular- and punishment phases, for players i and states j. Concatenate action profiles $\boldsymbol{x} \in \mathcal{X}$ $\overline{\boldsymbol{X}}_i := (\boldsymbol{x}, (\overline{x}^1, \overline{x}^2)')$, $\underline{\boldsymbol{X}}_i := (\boldsymbol{x}, (\underline{x}^1, \underline{x}^2)')$ action profiles across states (rows) where column (producers) elements are free in $\mathcal{X}$ and $-i$ plays the corresponding equilibrium action and $\overline{\boldsymbol{X}} := (\overline{\boldsymbol{x}}^1, \overline{\boldsymbol{x}}^2)$, $\underline{\boldsymbol{X}} := (\overline{\boldsymbol{x}}^1, \overline{\boldsymbol{x}}^2)$. Rearranging and stacking terms yields an attractively transparent formulation of the incentive compatibility constraint, that is for all $\boldsymbol{x} \in \mathcal{X}^2$ and $i \in \{1, 2\}$

$$\begin{aligned} \Delta\overline{\boldsymbol{\pi}}(\overline{\boldsymbol{X}}_i) &\leq \frac{\delta}{1-\delta} \cdot \Delta\overline{\boldsymbol{\tau}}(\overline{\boldsymbol{X}}_i) \cdot \boldsymbol{M}\Delta\boldsymbol{v} \\ \Delta\underline{\boldsymbol{\pi}}(\underline{\boldsymbol{X}}_i) &\leq \frac{\delta}{1-\delta} \cdot \Delta\underline{\boldsymbol{\tau}}(\underline{\boldsymbol{X}}_i) \cdot \boldsymbol{M}\Delta\boldsymbol{v} \end{aligned} \tag{13}$$

where $\Delta\overline{\pi}(\overline{\boldsymbol{X}}_i) := \boldsymbol{\pi}(\overline{\boldsymbol{X}}_i) - \boldsymbol{\pi}(\overline{\boldsymbol{X}})$, $\Delta\overline{\boldsymbol{\tau}}(\overline{\boldsymbol{X}}_i) := \overline{\boldsymbol{\tau}}(\overline{\boldsymbol{X}}_i) - \overline{\boldsymbol{\tau}}(\overline{\boldsymbol{X}})$ and $\Delta\boldsymbol{v} := \overline{\boldsymbol{v}} - \underline{\boldsymbol{v}}$. Equation (13) states that in equilibrium, the relative gain to a deviation may not exceed the expected- and discounted cost of incrementally increasing the probability of switching to, or remaining in, the punishing state. Incentives are said to have higher power the greater is $\Delta\boldsymbol{v}$, the value function differential, dynamically linked across states through the transition matrix $\boldsymbol{M}$. Thus, higher-powered incentives in any state enforce greater one-shot deviation pay-offs in all states. It is useful to solve for $\Delta\boldsymbol{v}$ in terms of primitives,

$$\Delta\boldsymbol{v} = (\boldsymbol{I} - \delta(\overline{\boldsymbol{\tau}} - \underline{\boldsymbol{\tau}})\boldsymbol{M})^{-1}(1-\delta)(\overline{\boldsymbol{\pi}} - \underline{\boldsymbol{\pi}}) \tag{14}$$

which makes transparent that incentives $\Delta\boldsymbol{v}$ are increasing in the per-period loss incurred under punishment and the difference $\Delta\boldsymbol{\tau}$ in probabilities of that loss being sustained.

$$\begin{aligned} \frac{\partial\boldsymbol{\Delta v}}{\partial\boldsymbol{\Delta\pi}} &= (\boldsymbol{I} - \delta\Delta\boldsymbol{\tau}\boldsymbol{M})^{-1}(1-\delta) > 0 \\ \frac{\partial\boldsymbol{\Delta v}}{\partial\boldsymbol{\Delta\tau}} &= (\boldsymbol{I} - \delta\Delta\boldsymbol{\tau}\boldsymbol{M})^{-1}\delta\boldsymbol{M}(\boldsymbol{I} - \delta\Delta\boldsymbol{\tau}\boldsymbol{M})^{-1} > 0 \end{aligned}$$

I now characterize the functions τ. I begin by claiming that the extremal action profiles satisfy

$$0 \leq \boldsymbol{\iota}'\boldsymbol{x}^m(y^j, \boldsymbol{\alpha}) \leq \boldsymbol{\iota}'\overline{\boldsymbol{x}}^j \leq \boldsymbol{\iota}'\boldsymbol{x}^n(y^j) \leq \boldsymbol{\iota}'\underline{\boldsymbol{x}}^j \tag{15}$$

by the following argument. First, any quantity less than $\boldsymbol{x}^m(y^j, \boldsymbol{\alpha})$ violates individual rationality. Second, there are no equilibria with output in both phases exceeding stage-game equilibrium quantities $\boldsymbol{x}^n(y^j)$, which would violate incentive compatibility (IC) by construction. Then $\overline{v} \geq \underline{v}$ and the incentive constraints (13) hold trivially for downwards (upwards) deviations in the regular (punishing) phase, as the left-hand-side is negative and the right-hand-side non-negative. In deriving transition probabilities, attention may be restricted to profitable deviations. Beginning in the punishment phase and fixing an action profile, I argue that the most severe punishment is achieved by

$$\underline{\tau}_i^j(\boldsymbol{x}) = (1 - \mathbb{1}(\boldsymbol{\iota}'\boldsymbol{x} \neq \boldsymbol{\iota}'\underline{\boldsymbol{x}}^j))\underline{\tau}^j \tag{16}$$

where $\mathbb{1}(\cdot)$ the indicator function and $\underline{\tau}^j \in [0,1]$ governs the stochastic length of the punishment. Note that (16) demands that the action profiles be in observable quantities only, so any downward deviation is immediately detected and there is no information asymmetry. Holding constant the continuation pay-offs, $\underline{\tau}^j$ is set so

$$\pi(y^j, x_i, \underline{x}^j_{-i}) - \pi(y^j, \underline{\boldsymbol{x}}^j) = \frac{\delta}{1-\delta}(\underline{\tau}^j - 0)[m_j(\overline{v}^j - \underline{v}^j) + (1-m_j)(\overline{v}^{-j} - \underline{v}^{-j})] \quad (17)$$

holds with equality for all j. If not, strictly lower pay-offs exist and the pay-off is not extremal. Notice that, by implication, $\underline{\tau} = 0$ if the punishment is in stage-game actions $\boldsymbol{x}^n$. Turn to the regular phase, and fix some $\overline{\boldsymbol{x}}^j$ to be enforced. The optimal transition function $\overline{\tau}$ minimizes wasteful equilibrium transitions (size, $\overline{\tau}$) while maintaining incentive compatibility (power, $\overline{\tau}^* - \overline{\tau}$). The trade-off between size and power of the players' statistical test is optimized with the following structure

Proposition 2. ***Trigger price, information bound.*** *The regular-phase transition probability is*

$$\overline{\tau}^j(\boldsymbol{x}) = 1 - F_\theta(\overline{p}/p(y^j,\boldsymbol{x})) \quad (18)$$

for both players i and where $0 < \overline{p}^j < \exp(-\sigma_\theta^2 3/2)p(y^j, \boldsymbol{x})$ and such that at least one of the equations in (13) holds with equality for some feasible deviation. In the case of a public defection $\overline{\tau}^j_i(\boldsymbol{x}) = 0$.

which restricts the transition probability to the convex region of F_θ, see Appendix A.2 for the proof. In deriving the transition probabilities for both regular- and punishing phases I argued that at least one of the incentive compatibility constraints must bind with equality,

Corollary 1. ***Binding constraints.*** *The incentive compatibility constraints in (13) bind with equality in each state and phase.*

a property used repeatedly below. By Proposition 1, the cartel maximizes its profits by maximizing incentive power. In terms of primitives, the equilibrium solves

$$\max_{\{\overline{\boldsymbol{x}}^j, \underline{\boldsymbol{x}}^j, \underline{\tau}^j, \overline{p}^j\}_{j=1}^2} \Delta \boldsymbol{v} \quad (19)$$

subject to the constraints (13) evaluated for every $\boldsymbol{x}^* \in \mathcal{X}^2$. The algorithm for computing the fixed-point operator $\boldsymbol{B}$ is detailed in online Appendix C.

6 Incentive power, unraveling of cartel discipline

I characterize the endogenous variation in cartel discipline and its implication for observable behavior. I show first that the return to an optimal deviation from the monopoly action increases convexly in demand y under a weak condition on ξ and that $\Delta\boldsymbol{v}$ is concave in y and reaches a maximum. Thus the left- and right-hand-sides of the incentive compatibility constraint (13) are respectively convexly increasing and concave in y. It follows that cartel discipline is decreasing in demand y as constraints to $\Delta\boldsymbol{v}$ are activated, increasing output war frequencies or local implied price elasticities of supply. I characterize how the concavity of $\boldsymbol{v}$ is modulated by the monitoring and dynamic demand environment.

Two key observable implications are illustrated through numerically solved equilibrium values. Firstly, the frequency and duration of output wars increases in demand. Second, the implied supply elasticities in the cooperative phase vary non-linearly and may go to positive- or negative infinity for regions of the parameter space.

I fix y^1, m_1, m_2 and consider a sufficient condition for monotonic changes to one-shot deviation values in y^2. To begin, suppose δ is such that the incentive compatibility constraint (13) holds when evaluated for $\boldsymbol{x}^m$ the monopoly action profile defined in (7). How does an incremental increase in y^2 affect the relative value of one-shot deviations and incentive power? Under punishment, the one-shot-deviation $\Delta\underline{\pi}^2$ falls in y^2 for any quantity. In regular play, the change in relative value of a one-shot deviation $\Delta\overline{\pi}^2$ is ambiguous if x^* is constrained by output capacity.

Lemma 1. ***Unconstrained one-shot deviation.*** *As the output increment $\epsilon \to 0$, there exists a finite ξ^* such that for all $\xi \geq \xi^*$*

$$x^* := \max_{x \in \mathcal{X}} \Delta\pi(x, x^m, y)$$

is interior to $\mathcal{X}$ and the relative value of a one-shot deviation $\Delta\pi(x, x^m, y)$ is convexly increasing in demand y.

The proof is in Appendix A.3. Recall from Section 3 that a high value of ξ limited the oligopolists' power such that the market is constrained to an inelastic region of the demand curve. If sufficiently constrained, optimal one-shot deviations will be small enough not to be capacity constrained. This will ensure one-shot deviation values increasing in demand, which is assumed in the following.[20]

It is straightforward to verify that the feasible, per-period, relative loss from incurring a punishment is decreasing when the punishing quantity is capacity constrained.

[20]Notice that when quantities are indivisible, there will generally be local, incremental reductions $\Delta\overline{\pi}$ due to asynchronous changes in $x^m(y)$ and $x^*(y)$.

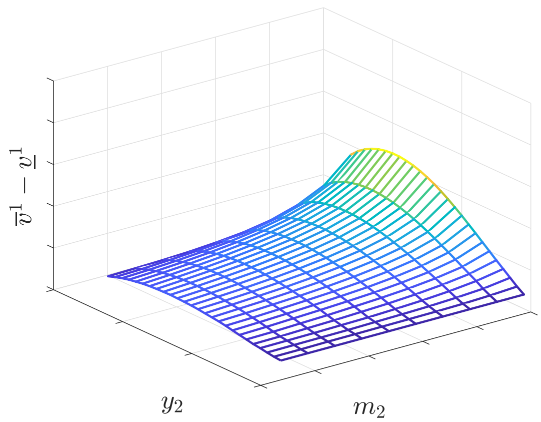

(a) Incentive power vs. persistence, low demand state.

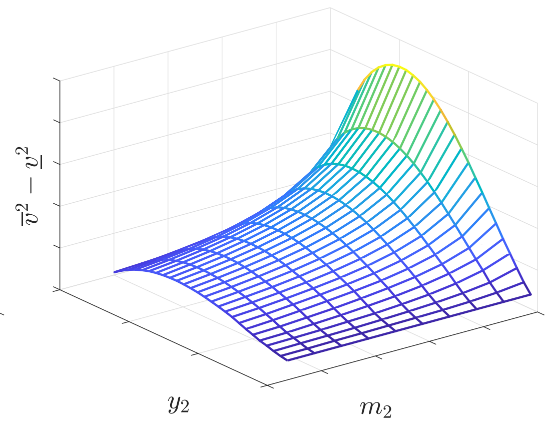

(b) Incentive power vs. persistence, high demand state.

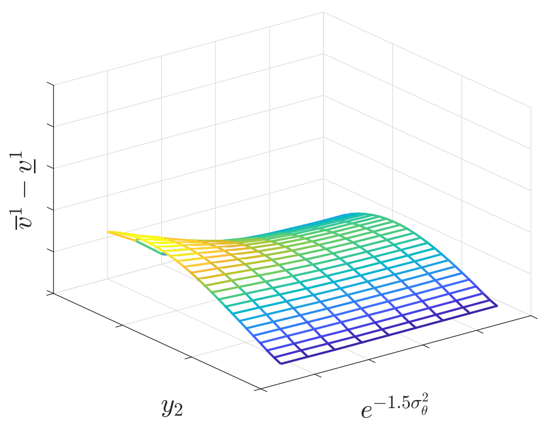

(c) Incentive power vs. monitoring quality, low demand.

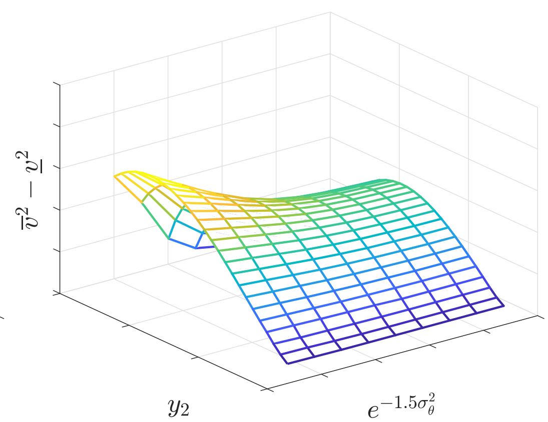

(d) Incentive power vs. monitoring quality, high demand.

Figure 3: Incentive power $\Delta \boldsymbol{v}$ against the level y_2 and persistence m_2 of the high demand state (a,b) and the noise wedge $\exp(-1.5\sigma_\theta^2)$ (c,d), see Proposition 2. See Table 1 for a complete list of parameter values.

Lemma 2. *If the capacity constraint binds in the punishing phase, the per-period loss* $\pi(y, \boldsymbol{x}^m) - \pi(y, \boldsymbol{x}^{max})$ *is decreasing in* y.

The proof is in Appendix A.6. It follows immediately that once demand pressure is sufficient that the capacity constraint is reached under punishment, severity is maintained through longer duration, depressing $\underline{\boldsymbol{\tau}}$ and raising $\Delta\tau = \overline{\boldsymbol{\tau}} - \underline{\boldsymbol{\tau}}$. Since $\Delta\boldsymbol{\tau}$ is bounded Δv reaches a maximum in y. But the one-shot-deviation is increasing everywhere in y from the feasible optimum. Thus it follows from Proposition 2 that $\overline{\boldsymbol{\tau}}$ must eventually fall to maintain incentive compatibility, diluting $\Delta\boldsymbol{v}$. The resulting total change in $\Delta\boldsymbol{\tau}$ depends on the trade-off between signal power and false positives as governed by σ_θ. The following propositions characterize the impact of demand persistence and monitoring quality on incentive creation.

Proposition 3. ***Incentive power and demand-state persistence.*** *Suppose that* $\Delta v^2 > \Delta v^1$. *Then an increase in the persistence* m_2 *of state* 2 *increases the incentive power and cartel discipline in both states,* $\partial\Delta\boldsymbol{v}/\partial m_2 > 0$ *and vice versa,* $\partial\Delta\boldsymbol{v}/\partial m_1 < 0$.

Proposition 4. ***Incentive power and monitoring quality.*** *Incentives* $\Delta\boldsymbol{v}$ *are decreasing in* σ_θ.

The proofs are in Appendices A.4 and A.5. The resulting concavity of $\boldsymbol{v}$ is illustrated in Figure 3, which plots equilibrium values of $\Delta\boldsymbol{v}$ against the level y_2 and persistence m_2 of the high demand state and information wedge $\exp(-1.5\sigma_\theta^2)$, see Proposition 2. The key observable implication of falling incentive power and increasing incentives to deviate is the unraveling of cartel discipline in constrained states. Formally:

Corollary 2. ***Unraveling of cartel discipline.*** *Suppose that* $\Delta v_2 > \Delta v_1$ *and consider a marginal increase signal noise* σ_θ, *or a decrease in demand persistence* m_2. *Then the action profile become weakly less extremal with* $\overline{\boldsymbol{x}}^2$ *increasing and* $\underline{\boldsymbol{x}}^2$ *decreasing. If quantities remain unchanged, the transition probabilities* $\overline{\boldsymbol{\tau}}$, $\underline{\boldsymbol{\tau}}$ *must increase and decrease, respectively.*

I conclude by discussing the unraveling dynamic, illustrating the supply behavior with three numerically computed equilibrium values. First, Figure 4 plots implicit "supply curves", that is, equilibrium prices $p(y, \boldsymbol{x})$ vs. quantities $\iota'\boldsymbol{x}$. Optimal equilibrium price-quantity combinations are shown for range of persistence and monitoring quality parameters with lighter colors indicating more constrained incentive creation. The limiting competitive (stage-game equilibrium) and profit-maximizing (monopoly) actions are plotted in black. Second, the approximate elasticities

$$\frac{\partial \iota'\overline{\boldsymbol{x}}}{\partial p(y, \iota'\overline{\boldsymbol{x}})} \frac{p(y, \iota'\overline{\boldsymbol{x}})}{\iota'\overline{\boldsymbol{x}}} = \frac{\partial \iota'\overline{\boldsymbol{x}}}{\iota'\overline{\boldsymbol{x}}} \frac{p(y, \iota'\overline{\boldsymbol{x}})}{\partial p(y, \iota'\overline{\boldsymbol{x}})} \approx \frac{\Delta \iota'\overline{\boldsymbol{x}}}{\iota'\overline{\boldsymbol{x}}} \frac{p(y, \iota'\overline{\boldsymbol{x}})}{\Delta p(y, \iota'\overline{\boldsymbol{x}})} \tag{20}$$

shown in Figure 5 provide a unit-free measure of local supply behavior. The expected share of time spent in each state is given by the stationary distribution $\boldsymbol{\mu}$ of the transition matrix $\boldsymbol{P}$, satisfying $\boldsymbol{\mu} := \boldsymbol{z} : \boldsymbol{P}'\boldsymbol{z} = \boldsymbol{z}$ is plotted in Figure 6.

In low-demand states, ample spare production capacity yields unconstrained incentive creation. A vanishing share of time is spent in output wars and cooperative supply behavior is similar across parameterizations. As demand y^2 increases the strategic constraints induce concavity of $\Delta\boldsymbol{v}$ in y^2 while the return to one-shot deviation values $\Delta\overline{\boldsymbol{\pi}}(\boldsymbol{x}^*)$ from $\boldsymbol{x}^m$ increases convexly. Incentive compatibility must be maintained by increasing $\Delta\boldsymbol{\tau}$ or $\overline{\boldsymbol{x}}$. But output in excess of monopoly production $\boldsymbol{x}^m$ or raising transition probabilities $\overline{\tau}$ have the second-order effect of reducing $\overline{\boldsymbol{v}}$, and by the bang-bang property, increase $\underline{\boldsymbol{v}}$. In turn the weakened incentives will necessitate further reductions in cartel discipline and increases in punishment frequency and duration. When incentive creation is sufficiently constrained the second-order effect dominates and local supply elasticities tend to positive- or negative infinity.

7 Conclusion

This paper has answered a long-standing call for richer models of imperfect competition that may rationalize OPEC's time-varying behavior. As a first step in this direction I have studied how cartel discipline is endogenously determined in optimal symmetric equilibria of an imperfect monitoring model with a dynamic environment and capacity constraints. When demand is low and monitoring quality is high, strategic competition is less salient and OPEC's behavior may be aptly summarized by a constant price elasticity. When incentive creation is constrained the magnitude and sign of local the supply elasticities depend sensitively and non-linearly on the current- and future expected strategic environment.

Returning to the motivating Figure 1, my theoretical analysis yields derives novel interpretations and testable implications that may be pursued in future research. Optimal equilibria imply short, intense, and rare output wars, all continuous empirical properties that may be tested. The variation in cartel discipline may be plausibly explained by either coordination on a new, more efficient equilibrium or changes in the fundamental, strategic environment. Further research on this topic is required. Indeed, this paper has only scratched the surface of potential strategic mechanisms through which OPEC members interact. Producers are assumed homogeneous, their preferences stable, cartel membership fixed, and monitoring public. A promising avenue of research may combine historical and empirical evidence with modern game theory in pursuit of a more refined model of crude oil supply.

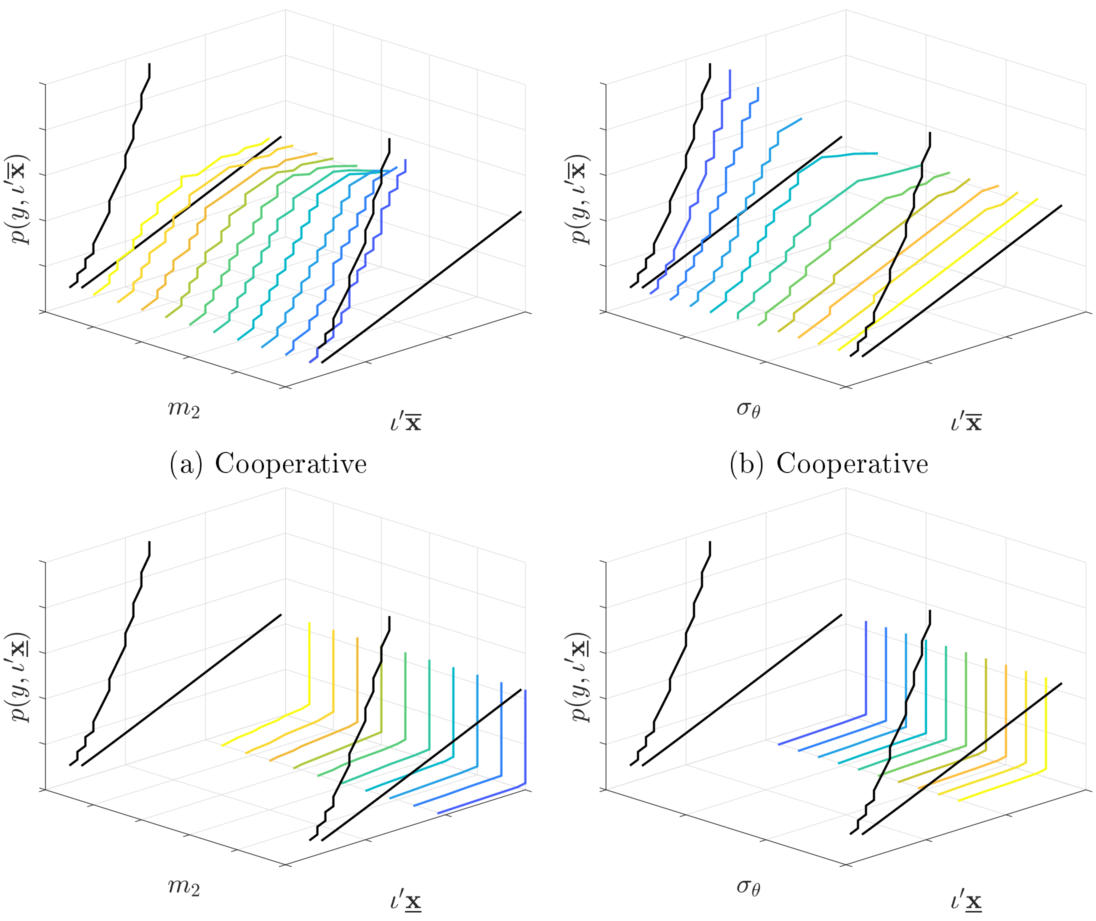

Figure 4: Equilibrium prices $p(y, \boldsymbol{\iota}'\boldsymbol{x})$ vs. quantities $\boldsymbol{\iota}'\boldsymbol{x}$ for a range of demand, persistence- and signal noise parameters y_2, m_2, σ_θ. The cooperative- and punishment phases are plotted in (a), (b) and (c), (d), respectively. Limiting monopoly- and stage game equilibria are in black. Lighter colors indicate lower persistence and monitoring quality. See Table 1, Appendix B for a complete list of parameter values.

A Proofs

A.1 Proposition 1. Optimal equilibria

Proof. Proof is by construction. I proceed by first verifying that the necessary restrictions to the stage game apply. Next, I show that generalization to Markov demand maintains the required recursive structure of the equilibrium.

The five key assumptions in Abreu, Pearce, and Stacchetti (1990, p.1045) are satisfied by construction. Action spaces are finite (1). The signal is continuously distributed with support independent of actions (2, 3). Stage-game pay-offs are continuous in the signal (4). Finally, a pure-strategy Nash equilibrium exists in the stage game (5). I now show that the introduction of Markov demand leaves unaltered the relevant computational primitives. The key object used in APS to construct a symmetric equilibrium in one-dimensional pay-offs is $\mathcal{L}(\Omega;\mathbb{R})$ the set of all bounded, Lebesgue-measurable mappings $l:\Omega\to\mathbb{R}$ from signals into the reals. Under Markovian demand the continuation values are naturally in $\mathbb{R}^2$, one for each state. But the assumption of a constant transition matrix allows a particularly simple computation of their (one-dimensional) expected value that preserves a recursive structure. Let W^1, W^2 bounded subsets of $\mathbb{R}$, $\boldsymbol{W}=W^1\times W^2$, and let $\mathcal{A}^j(\mathbb{R}^2;\boldsymbol{M})$ denote the family of convex combinations $a^j:\mathbb{R}^2\times[0,1]\to\mathbb{R}$ using the state (row) j transition probabilities in $\boldsymbol{M}$ as weights. For some $\boldsymbol{w}\in\boldsymbol{W}$, then, $a^j(\boldsymbol{w},\boldsymbol{M})=m_jw^j+(1-m_j)w^{-j}$. Consider the family $\mathcal{L}_c(\Omega;\boldsymbol{M},\mathbb{R}^2)$ of function compositions $l\circ a:\Omega\to\mathbb{R}$. Clearly, the convex combinations a are defined for any non-empty subset of $\mathbb{R}^2$. Second, the operation preserves boundedness and measure, the properties demanded of $\mathcal{L}(\Omega;\mathbb{R})$ in APS. Thus, Theorems 1 and 2 in APS apply to $\mathcal{L}_c$. Third, Theorem 3 requires that the convex combination a preserves convexity (and thus compactness), which it does by definition. Finally, Lemma 1 and Theorems 4-5 require that a is rank-preserving, that is $a^j(\boldsymbol{w},\boldsymbol{M})>a^j(\boldsymbol{w}',\boldsymbol{M})$ if $\boldsymbol{w}>\boldsymbol{w}'$, which holds by the monotonicity of a^j. Theorem 6 depends on the compactness of $\boldsymbol{w}$. The necessity of bang-bang strategies for an optimal symmetric equilibrium is shown by Theorem 7, see in particular the comment on pp. 1058. □

A.2 Proposition 2. Trigger price, information bound.

$$\overline{\tau}_i^j(\boldsymbol{x})=1-F_\theta(\overline{p}/p(y^j,\boldsymbol{x}))$$

for both players i and where $0<\overline{p}_i^j<\exp(-\sigma_\theta^2 3/2)p(y^j,\boldsymbol{x})$ and such that at least one of the equations in (13) holds with equality for some feasible deviation in $\{\overline{x}_i^j+\epsilon,\overline{x}_i^j+\tilde{h}\epsilon\}$ for at least one player. In the case of a public defection $\overline{\tau}_i^j(\boldsymbol{x})=0$.

Proof. In the reward phase a profitable deviation is upward, which may be up to $\max(\tilde{h}\epsilon,x^{\max}-\overline{x}_i^j)$ units in hidden quantities. Thus, defections must be inferred from realizations of conditionally log-normally distributed prices, $\ln p(\tilde{\theta},y,\tilde{\boldsymbol{x}})\sim N(\ln(\max(y-\beta\boldsymbol{\iota}'\tilde{\boldsymbol{x}},\xi))-\sigma_\theta^2/2,\sigma_\theta^2)$, where distribution $F_p(\cdot|\boldsymbol{x})$ is parameterized by the action profile $\boldsymbol{x}$. Notice that the players' inference problem is effectively a goodness of fit test across of "models" $\ln p(\tilde{\theta}, y, \tilde{\boldsymbol{x}})$ with unknown parameters $\tilde{\boldsymbol{x}}$. The likelihood-ratio test of the hypothesis $\tilde{\boldsymbol{x}} > \overline{\boldsymbol{x}}$ is then uniformly most powerful by the Neyman-Pearson lemma, minimizing size, given power. The log-normal distribution satisfies the monotone likelihood ratio property in total output, that is

$$\partial \left(\frac{f_p(\ln p | \hat{\boldsymbol{x}})}{f_p(\ln p | \boldsymbol{x})} \right) \Big/ \partial p < 0$$

for $\boldsymbol{\iota}'\hat{\boldsymbol{x}} > \boldsymbol{\iota}'\boldsymbol{x}$, so the likelihood of a deviation is monotonically decreasing in the realized price level. A tail test of observed prices, $\overline{p}^j \leq p(\theta, y^j, \boldsymbol{x})$ is then a sufficient statistic for the likelihood ratio. Thus the functional form of the transition function is $\Pr(p(\theta, y^j, \boldsymbol{x}) \leq \overline{p}^j) = \Pr(\theta \leq \overline{p}^j/p(y^j,\boldsymbol{x})) = F_\theta(\overline{p}^j/p(y^j,\boldsymbol{x}))$. To determine the upper bound $\overline{p}$, fix first $\overline{\boldsymbol{x}}$. Notice that incentives are provided by the conditional difference (power) $\overline{\tau}_i^j(x_i, \overline{x}_{-i}^j) - \overline{\tau}_i^j(\overline{\boldsymbol{x}}^j)$, not the level $\overline{\tau}_i^j(\overline{\boldsymbol{x}}^j)$ (false positive rate). It is never optimal for the trigger price $\overline{p}^j$ to locate $F_\theta(\overline{p}^j/p(y^j,\boldsymbol{x}))$ in the concave region, as the same power can be achieved for a strictly lower false positive ratio, increasing pay-offs while maintaining incentives. The second derivative F_θ'' changes sign at the mode,

$$\begin{aligned} \frac{\partial^2 F_\theta(z)}{\partial z \partial z} &= 0 \\ -\frac{f_\theta}{z} - f_\theta \cdot 2 \left(\frac{\ln(z)}{\sqrt{2}} \sigma_\theta + \frac{\sigma_\theta}{\sqrt{2} \cdot 2} \right) \frac{1}{z\sqrt{2}\sigma_\theta} &= 0 \\ z &= \exp(-\frac{3}{2}\sigma_\theta^2) \end{aligned}$$

where replacing z with $\overline{p}^j/p(y^j,\boldsymbol{x})$ in the final expression above yields the proposed bound, restricting F_θ to the efficient convex region. Finally, and again for fixed actions $\overline{\boldsymbol{x}}$, the trigger $\overline{p}^j$ is optimally set to the lowest level such that all incentive compatibility constraints hold, minimizing false positives. □

A.3 Proposition 1. Unconstrained one-shot deviation.

As the output increment $\epsilon \to 0$, there exists a ξ^*, finite, such that $x^* < x^{max}$ and the relative value of a one-shot deviation $\Delta\pi^*$ is convexly increasing in demand y.

Proof. Let $x^*(y)$ the stage-game best response to $x^m(y)$, the jointly symmetric profit-maximizing quantity. Consider the first-order conditions for $x^m(y)$:

$$x^m(y) = \frac{p(y, \boldsymbol{x}^m(y)) - 2\kappa x^m(y)}{-2\partial/\partial x\, p(y, \boldsymbol{x}^m(y))}$$

$$\Leftrightarrow$$

$$-\partial/\partial x\, p(y, \boldsymbol{x}^m(y)) \frac{x^m(y)}{p(y, \boldsymbol{x}^m(y))} = \frac{\beta x^m(y)}{\xi + \beta x^m(y)} = 1 - \frac{2\kappa x^m(y)}{p(y, \boldsymbol{x}^m(y))}$$

The left-hand-side increases concavely $x^m(y)$, to unity. But the marginal cost increases linearly in output, so the monopoly price must therefore be increasing convexly relative to marginal cost. I now show that the difference $x^*(y) - x^m(y)$, convexly increasing in y, is modulated by ξ. Consider the first-order conditions defining $x^m(y)$,

$$\frac{x^m(y)(\xi + 2\beta x^m(y))^2}{\xi} = \frac{y}{2\kappa}$$

and $x^*(y)$:

$$\frac{x^*(y)(\xi + \beta x^*(y) + \beta x^m(y))^2}{\xi + \beta x^m(y)} = \frac{y}{2\kappa}$$

They may be combined to form

$$\begin{gathered}\frac{x^m(y)(\xi + 2\beta x^m(y))^2}{\xi} = \frac{x^*(y)(\xi + \beta x^*(y) + \beta x^m(y))^2}{\xi + \beta x^m(y)} \\ \Leftrightarrow \\ \frac{x^m(y)(\xi + 2\beta x^m(y))^2}{x^*(y)(\xi + \beta x^*(y) + \beta x^m(y))^2} = \frac{\xi}{\xi + \beta x^m(y)}\end{gathered} \tag{21}$$

which states that, for any ξ the right-hand-side is less than unity, so $x^*(y)$ increases more than one-for-one with $x^m(y)$. Hence, $x^*(y) - x^m(y)$ increases convexly in y. Notice finally that the right-hand-side goes to 1 as $\xi \to \infty$, asymptotically restraining x^* towards x^m. Since $x^m < x^{\max}$, there exists some ξ^* such that $x^* < x^{\max}$ and the deviation profits are not restrained by the capacity constraint. □

A.4 Proposition 3. Incentive power and demand-state persistence.

Suppose that $\Delta v^2 > \Delta v^1$. Then an increase in the persistence m_2 (m_1) of state 2 (1) increases (decreases) the incentive power and cartel discipline in both states, $\partial\Delta\boldsymbol{v}/\partial m_2 > 0$ and $\partial\Delta\boldsymbol{v}/\partial m_1 < 0$.

Proof. I evaluate the derivative for m_2. A symmetric argument applies to m_1 simply reverses the sign. Consider $\frac{\delta}{1-\delta}\Delta\boldsymbol{\tau} \cdot \boldsymbol{M}\Delta\boldsymbol{v}$, the incentive compatibility constraints' (13) right-hand-side. Ignoring multiplicative constants $\frac{\delta}{1-\delta}\Delta\boldsymbol{\tau}$ and differentiating with respect to m_2 yields a system of equations

$$\frac{\partial\Delta v^1}{\partial m_2} = \frac{\delta\Delta\tau^1(1-m_1)}{1-\delta\Delta\tau^1 m_1}\frac{\partial\Delta v^2}{\partial m_2} \tag{22}$$

$$\frac{\partial\Delta v^2}{\partial m_2} = \frac{\delta\Delta\tau^2}{1-\delta\Delta\tau^2 m_2}\left(\Delta v^2 - \Delta v^1 + (1-m_2)\frac{\partial\Delta v^1}{\partial m_2}\right) \tag{23}$$

where the sign of (22) is determined wholly by the sign of (23). Inserting (23) in (22) and evaluating yields

$$\frac{\partial \Delta v^1}{\partial m_2} = \frac{\delta \Delta \tau^1 (1-m_1)}{1-\delta \Delta \tau^1 m_1} \frac{\delta \Delta \tau^2}{1-\delta \Delta \tau^2 m_2} \left(\Delta v^2 - \Delta v^1 + (1-m_2) \frac{\partial \Delta v^1}{\partial m_2} \right)$$
$$\Leftrightarrow$$
$$\frac{\partial \Delta v^1}{\partial m_2} = \frac{\Delta v^2 - \Delta v^1}{(1-\delta \Delta \tau^1 m_1)(1-\delta \Delta \tau^2 m_2) - \delta^2 \Delta \tau^1 \Delta \tau^2 (1-m_1)(1-m_2)} > 0$$
$$\Rightarrow$$
$$\frac{\partial \Delta v^2}{\partial m_2} > 0$$

where the final inequality holds under the assumption that $\Delta v^2 > \Delta v^1$. By the argument in Proposition 2, at least one incentive compatibility constraint binds in equilibrium, for each phase and state. The increase in m_2 introduces slack into these constraints, allowing re-optimization, and thereby pushing regular- and punishing pay-offs $\overline{v}^j$, $\underline{v}^j$ towards (weakly) more extremal values. □

A.5 Proposition 4. Incentive power and monitoring quality.

Incentives $\Delta \boldsymbol{v}$ are decreasing in σ_θ.

Proof. Increasing σ_θ reduces the slope of $F_\theta(\cdot)$ so the conditional transition probability

$$\frac{\partial \Delta \overline{\tau}}{\partial \sigma_\theta} = \frac{\partial}{\partial \sigma_\theta} \left(F_\theta \left(\frac{\overline{p}}{p(y^j, \overline{\boldsymbol{x}})} \right) - F_\theta \left(\frac{\overline{p}}{p(y^j, \boldsymbol{x}^*)} \right) \right) < 0$$

is less sensitive to a given deviation $\boldsymbol{\iota}'\boldsymbol{x}^* > \boldsymbol{\iota}'\overline{\boldsymbol{x}}$. This claim is easily verified by evaluating the derivative and rearranging terms

$$\frac{\exp\left[\frac{\ln\left(\frac{\overline{p}}{p(y^j,\mathbf{x}^*)}\right)+\frac{1}{2}\sigma_\theta^2}{\sqrt{2}\sigma_\theta}\right]^2}{\exp\left[\frac{\ln\left(\frac{\overline{p}}{p(y^j,\overline{\mathbf{x}})}\right)+\frac{1}{2}\sigma_\theta^2}{\sqrt{2}\sigma_\theta}\right]^2} < 1 < \frac{\left(\frac{\ln\left(\frac{\overline{p}}{p(y^j,\overline{\mathbf{x}})}\right)}{\left(\sqrt{2}\sigma_\theta\right)^2} + \frac{1}{2^{\frac{3}{2}}}\right)}{\left(\frac{\ln\left(\frac{\overline{p}}{p(y^j,\mathbf{x}^*)}\right)}{\left(\sqrt{2}\sigma_\theta\right)^2} + \frac{1}{2^{\frac{3}{2}}}\right)} \tag{24}$$

where I use that $F_\theta(z) = 0.5 + \pi^{-1} \int_0^u \exp(-x^2)dx$ for $u = 0.5(\ln z + 0.5\sigma_\theta)\sigma_\theta^{-2}$. By Corollary 1, at least one incentive compatibility constraint binds under regular play, so the increase in σ_θ renders the initial equilibrium incentive incompatible upon impact. In response, quantities $\overline{\boldsymbol{x}}$ or trigger price $\overline{p}$ must increase, decreasing $\overline{v}$, in turn increasing $\underline{v}$, and thus also Δv. □

A.6 Lemma 2. Punishment severity.

If the capacity constraint binds in the punishing phase, the per-period loss $\pi(y, \boldsymbol{x}^m) - \pi(y, \boldsymbol{x}^{max})$ *is decreasing in* y.

Proof. Notice first that the punishment pay-off $\pi(y, \boldsymbol{x}^{\max})$ increases linearly at a rate ${}^{x^{\max}}/_{\xi+\beta\boldsymbol{\iota}'\boldsymbol{x}^{\max}}$ in y. The marginal growth of monopoly profit $\pi(y, \boldsymbol{x}^m)$ in y is non-decreasing by individual rationality, since by maintaining current output it increases minimally at the rate ${}^{x^m(y)}/_{\xi+\beta\boldsymbol{\iota}'\boldsymbol{x}^m(y)}$. Thus, $\pi(y, \boldsymbol{x}^m)$ grows convexly. But there exists some y' such that $x^m(y') = x^{\max}$, with $\pi(y, \boldsymbol{x}^{\max})$ approaching $\pi(y, \boldsymbol{x}^m)$ from below. For the profits to meet, the growth of $\pi(y, \boldsymbol{x}^m)$ must then be lower than $\pi(y, \boldsymbol{x}^{\max})$ for all $y < y'$, and hence $\pi(y, \boldsymbol{x}^m) - \pi(y, \boldsymbol{x}^{\max})$ is decreasing in y. □

B Tables, figures

Parameter	**Description**	**Value**
$\mathcal{X}$	Elements in action set	31
$x^{\max}$	Output capacity	3
κ	Marginal cost shifter	1
δ	Discount factor	0.9
β	Demand parameter	3
ξ	Demand parameter	10
σ_θ	Standard deviation, signal noise	0.15
$\boldsymbol{\sigma}_\theta$	Range of signal noise	$\{0.05, 0.1, \cdots, 0.5\}$
$\exp(-{}^3/_2\sigma_\theta^2)$	Noise wedge	0.97
$\exp(-{}^3/_2\boldsymbol{\sigma}_\theta^2)$	Range of noise wedge	$\{0.99, \cdots, 0.69\}$
y_1, y_2	Demand level, reference states	$(5.2, 50.9)$
$\boldsymbol{y}_2$	Range of demand levels	$\{5.2, \cdots, 209.3\}$
m_1, m_2	Persistence, reference states	$(0.9, 0.9)$
$\boldsymbol{m}_2$	Range of persistence parameters	$\{0.1, 0.2, \cdots, 1\}$

Table 1: Parameter values for numerical solutions.

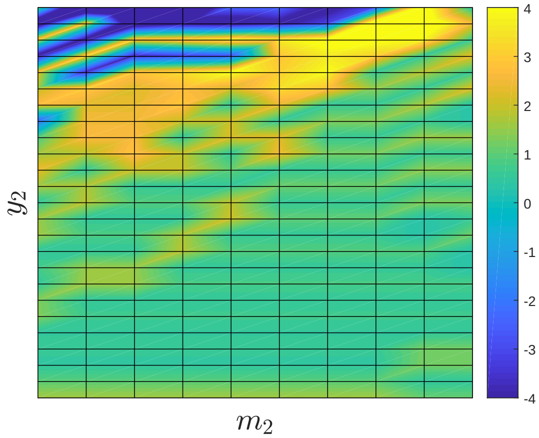

(a) Persistence

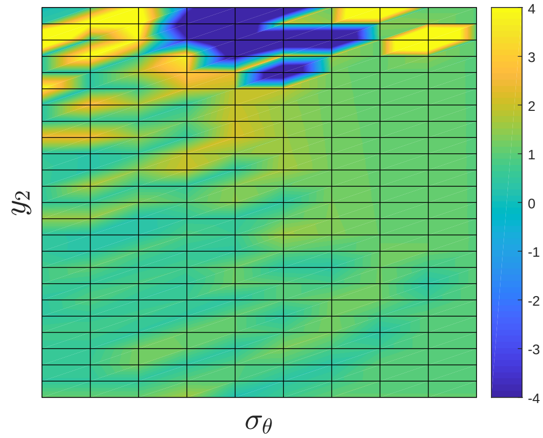

(b) Signal noise

Figure 5: Heat map plots of approximate price elasticities of supply $\Delta\iota'\overline{\boldsymbol{x}}(\iota'\overline{\boldsymbol{x}}^2)^{-1}p(y^2,\iota'\overline{\boldsymbol{x}}^2)(\Delta p(y^2,\iota'\overline{\boldsymbol{x}}^2))^{-1}$ vs demand level y_2, persistence m_2, and signal noise σ_θ.

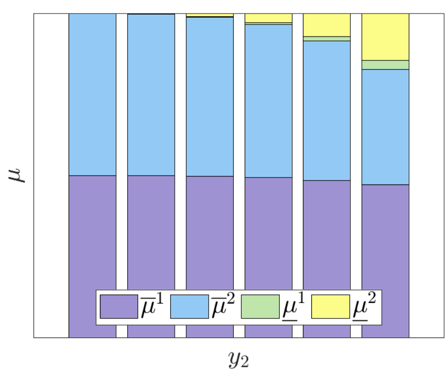

(a) Low signal noise

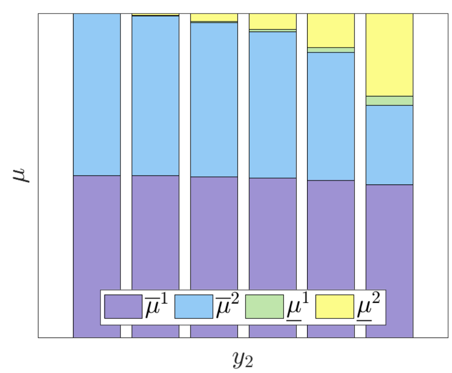

(b) High signal noise

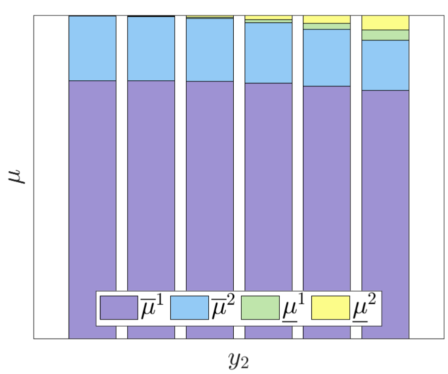

(c) Low persistence

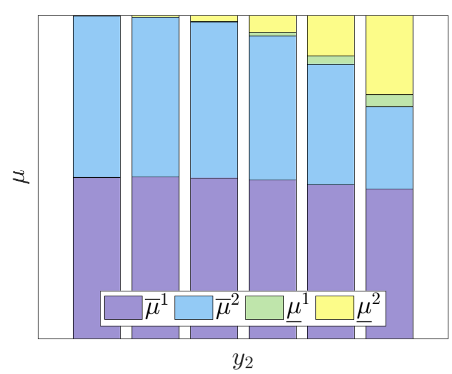

(d) High persistence

Figure 6: Stationary distribution $\boldsymbol{\mu} := \boldsymbol{z} : \boldsymbol{P}'\boldsymbol{z} = \boldsymbol{z}$ of the transition matrix $\boldsymbol{P}$, Equation 11.

C Algorithm: Not for publication

I implement the algorithm proposed by Abreu, Pearce, and Stacchetti, 1990, Section 5. The set of equilibrium pay-offs $\mathcal{V}$ is computed by repeatedly iterating $\boldsymbol{B}$ on a set of initial values $\mathcal{V}^*$ satisfying $\mathcal{V} \subset \mathcal{V}^*$. I describe the initialization procedure, then give an overview of the main steps in the computation, and finally detail the exact zero-finding algorithm.

C.1 Preliminaries, initialization

Select a tolerance ζ. Fix state variables σ_θ, $(y^1, y^2) \in \mathcal{Y}$. Let

$$\overline{w}_0^j := \max_{\boldsymbol{x} \in \mathcal{X}^2} \pi(y^j, \boldsymbol{x})$$
$$\underline{w}_0^j := \min_{\boldsymbol{x} \in \mathcal{X}} \pi(y^j, \boldsymbol{x})$$

be initial values, $j = \{1, 2\}$. Let the corresponding action profiles by $\overline{\boldsymbol{x}}_0^j$ and $\underline{\boldsymbol{x}}_0^j$. Adapting the notation used to define the identities Equations (9) and (9), compactly denote producer i's continuation value under action profile $\boldsymbol{x}_i$

$$\begin{aligned} \overline{\boldsymbol{v}}(\boldsymbol{x}_i, \overline{\boldsymbol{x}}_{-i}, \overline{\boldsymbol{\tau}}, \boldsymbol{W}_t) &= (1-\delta) \cdot \overline{\boldsymbol{\pi}}(\boldsymbol{x}_i, \overline{\boldsymbol{x}}_{-i}) \\ &+ \\ \delta \cdot (\overline{\boldsymbol{\tau}}(\boldsymbol{x}_i, \overline{\boldsymbol{x}}_{-i}) &\cdot \overline{\boldsymbol{w}}_t + (\boldsymbol{\iota} - \overline{\boldsymbol{\tau}})(\boldsymbol{x}_i, \overline{\boldsymbol{x}}_{-i}) \cdot \underline{\boldsymbol{w}}_t) \end{aligned} \tag{25}$$

$$\begin{aligned} \underline{\boldsymbol{v}}(\boldsymbol{x}_i, \underline{\boldsymbol{x}}_{-i}, \underline{\boldsymbol{\tau}}, \boldsymbol{W}_t) &= (1-\delta) \cdot \underline{\boldsymbol{\pi}}(\boldsymbol{x}_i, \underline{\boldsymbol{x}}_{-i}) \\ &+ \\ \delta \cdot (\underline{\boldsymbol{\tau}}(\boldsymbol{x}_i, \underline{\boldsymbol{x}}_{-i}) &\cdot \overline{\boldsymbol{w}}_t + (\boldsymbol{\iota} - \underline{\boldsymbol{\tau}}(\boldsymbol{x}_i, \underline{\boldsymbol{x}}_{-i})) \cdot \underline{\boldsymbol{w}}_t) \end{aligned} \tag{26}$$

where exogenous continuation values $\overline{w}^j$ and $\underline{w}^j$, transition probabilities $\overline{\tau}^j$, $\underline{\tau}^j$ and action profiles $\overline{\boldsymbol{x}}^j$ and $\underline{\boldsymbol{x}}^j$ are stacked in two-vectors. The set of feasible and in individually rational deviations are $\{\overline{x}^j + \epsilon, ..., x^{max}\}$ and $\{0, \epsilon, ..., \underline{x}^j - \epsilon\}$ in the reward- and punishment state respectively.

C.2 Iteration

Index the iterations by $t = \{0, 1, 2, ...\}$, with $t = 0$ denoting the initial values. Value functions are said to have converged when $|\overline{w}_t^j - \overline{w}_{t-1}^j| \leq \zeta$ and $|\overline{w}_t^j - \overline{w}_{t-1}^j| \leq \zeta$. Starting from the initialization value $j = 0$, iterate the following steps until convergence:

1. **Compute candidate transition probabilities.** Search, state for every individually rational $\boldsymbol{x}$ in $\mathcal{X}^2$, for a trigger $\overline{p}^j(\boldsymbol{x})$ and transition probability $\underline{\tau}^j$ satisfying incentive compatibility for all deviations with equality for at least one deviation

$$\overline{\boldsymbol{v}}(\boldsymbol{x}_i, \overline{\boldsymbol{x}}_{-i}, \overline{\boldsymbol{\tau}}, \boldsymbol{W}_t) = \overline{\boldsymbol{v}}(\overline{\boldsymbol{x}}_i, \overline{\boldsymbol{x}}_{-i}, \overline{\boldsymbol{\tau}}, \boldsymbol{W}_t) \tag{27}$$

$$\underline{\boldsymbol{v}}(\boldsymbol{x}_i, \underline{\boldsymbol{x}}_{-i}, \underline{\boldsymbol{\tau}}, \boldsymbol{W}_t) = \underline{\boldsymbol{v}}(\underline{\boldsymbol{x}}_i, \underline{\boldsymbol{x}}_{-i}, \underline{\boldsymbol{\tau}}, \boldsymbol{W}_t) \tag{28}$$

and

$$\overline{\boldsymbol{v}}(\boldsymbol{x}'_i, \overline{\boldsymbol{x}}_{-i}, \overline{\boldsymbol{\tau}}, \boldsymbol{W}_t) \leq \overline{\boldsymbol{v}}(\overline{\boldsymbol{x}}_i, \overline{\boldsymbol{x}}_{-i}, \overline{\boldsymbol{\tau}}, \boldsymbol{W}_t)$$

$$\underline{\boldsymbol{v}}(\boldsymbol{x}'_i, \underline{\boldsymbol{x}}_{-i}, \underline{\boldsymbol{\tau}}, \boldsymbol{W}_t) \leq \underline{\boldsymbol{v}}(\underline{\boldsymbol{x}}_i, \underline{\boldsymbol{x}}_{-i}, \underline{\boldsymbol{\tau}}, \boldsymbol{W}_t)$$

for all other $\boldsymbol{x}_i \neq \boldsymbol{x}_i$. Gather transition probabilities and quantities solving the equalities in vectors and matrices $\overline{\boldsymbol{P}}^j_t$, $\overline{\boldsymbol{X}}^j_t$ and $\underline{\boldsymbol{T}}_j$, $\underline{\boldsymbol{X}}_j$.

3. **Update continuation values.** Evaluate continuation values for every combination $(\overline{\boldsymbol{x}}^j_t, \overline{\boldsymbol{p}}^j_t) \in \{\overline{\boldsymbol{X}}^j_t, \overline{\boldsymbol{P}}^j_t\}$ and $(\underline{\boldsymbol{x}}^j_t, \underline{\tau}^j_t) \in \{\underline{\boldsymbol{X}}^j_t, \underline{\boldsymbol{T}}^j_t\}$ and select extremal continuation values:

$$\overline{\boldsymbol{w}}_{t+1} = \max_{(\overline{\boldsymbol{x}}^j_t, \overline{\boldsymbol{p}}^j_t) \in \{\overline{\boldsymbol{P}}^j_t, \overline{\boldsymbol{X}}^j_t\}} \overline{\boldsymbol{v}}(\boldsymbol{x}_i, \overline{\boldsymbol{x}}_{-i}, \overline{\boldsymbol{\tau}}, \boldsymbol{W}_t) \tag{29}$$

$$\underline{\boldsymbol{w}}^j_{t+1} = \min_{(\underline{\boldsymbol{x}}^j_t, \underline{\tau}^j_t) \in \{\underline{\boldsymbol{T}}_j, \underline{\boldsymbol{X}}_j\}} \underline{\boldsymbol{v}}(\boldsymbol{x}_i, \underline{\boldsymbol{x}}_{-i}, \underline{\boldsymbol{\tau}}, \boldsymbol{W}_t) \tag{30}$$

Due to discounting, we have $\overline{\boldsymbol{w}}^j_{t+1} \leq \overline{\boldsymbol{w}}^j_t$ and $\underline{\boldsymbol{w}}^{\div}_{j+1} \geq \underline{v}^{\div}_j$. Using $\overline{v}^{+}_{j+1}$ and $\underline{v}^{\div}_{j+1}$ as new values, return to Step 1.

References

Aastveit, Knut Are, Hilde C Bjørnland, and Leif Anders Thorsrud (2015). "What drives oil prices? Emerging versus developed economies". In: *Journal of Applied Econometrics* 30.7, pp. 1013–1028.

Abreu, Dilip, David Pearce, and Ennio Stacchetti (1986). "Optimal cartel equilibria with imperfect monitoring". In: *Journal of Economic Theory* 39.1, pp. 251–269.

— (1990). "Toward a theory of discounted repeated games with imperfect monitoring". In: *Econometrica: Journal of the Econometric Society*, pp. 1041–1063.

Al-Qahtani, Ayed, Edward Balistreri, and Carol Dahl (2008). "Literature review on oil market modeling and OPEC's behavior". In: *Manuscript, Division of Economics and Business, Colorado School of Mines*.

Alhajji, Anas F and David Huettner (2000). "OPEC and world crude oil markets from 1973 to 1994: cartel, oligopoly, or competitive?" In: *Energy Journal* 21.3, pp. 31–60.

Almoguera, Pedro A, Christopher C Douglas, and Ana María Herrera (2011). "Testing for the cartel in OPEC: non-cooperative collusion or just non-cooperative?" In: *Oxford Review of Economic Policy* 27.1, pp. 144–168.

Anderson, Soren T, Ryan Kellogg, and Stephen W Salant (2018). "Hotelling under pressure". In: *Journal of Political Economy* 126.3, pp. 984–1026.

Barros, Carlos Pestana, Luis A Gil-Alana, and James E Payne (2011). "An analysis of oil production by OPEC countries: Persistence, breaks, and outliers". In: *Energy Policy* 39.1, pp. 442–453.

Barsky, Robert B and Lutz Kilian (2004). "Oil and the Macroeconomy since the 1970s". In: *Journal of Economic Perspectives* 18.4, pp. 115–134.

Baumeister, Christiane and James D Hamilton (2019a). "Structural interpretation of vector autoregressions with incomplete identification: Revisiting the role of oil supply and demand shocks". In: *American Economic Review* 109.5, pp. 1873–1910.

— (2019b). "Structural Interpretation of Vector Autoregressions with Incomplete Identification: Setting the Record Straight". In: *Manuscript, University of California San Diego*.

Baumeister, Christiane and Lutz Kilian (2016). "Forty years of oil price fluctuations: Why the price of oil may still surprise us". In: *Journal of Economic Perspectives* 30.1, pp. 139–60.

Baumeister, Christiane and Gert Peersman (2013). "The role of time-varying price elasticities in accounting for volatility changes in the crude oil market". In: *Journal of Applied Econometrics* 28.7, pp. 1087–1109.

Behar, Alberto and Robert Ritz (2017). "OPEC vs US shale: Analyzing the shift to a market-share strategy". In: *Energy Economics* 63.C, pp. 185–198.

Bornstein, Gideon, Per Krusell, and Sergio Rebelo (2019). *Lags, Costs, and Shocks: An Equilibrium Model of the Oil Industry*. Tech. rep.

Bret-Rouzaut, Nadine and Jean-Pierre Favennec (2011). *Oil and Gas Exploration and Production: Reserves, costs, contracts*. 3rd ed. Paris: Editions Technip.

Brock, William A and Jose A Scheinkman (1985). "Price setting supergames with capacity constraints". In: *The Review of Economic Studies* 52.3, pp. 371–382.

Caldara, Dario, Michele Cavallo, and Matteo Iacoviello (2019). "Oil price elasticities and oil price fluctuations". In: *Journal of Monetary Economics* 103, pp. 1–20.
Coll, Steve (2012). *Private empire: ExxonMobil and American power*. Penguin.
Dees, Stephane et al. (2007). "Modelling the world oil market: Assessment of a quarterly econometric model". In: *Energy Policy* 35.1, pp. 178–191.
Downey, Morgan (2008). *Oil 101*. Wooden Table Press.
Dvir, Eyal and Kenneth Rogoff (2014). "Demand effects and speculation in oil markets: Theory and evidence". In: *Journal of International Money and Finance* 42, pp. 113–128.
Dvir, Eyal and Kenneth S Rogoff (2009). *Three Epochs of Oil*. Working Paper 14927. National Bureau of Economic Research.
Fabra, Natalia (2006). "Collusion with capacity constraints over the business cycle". In: *International Journal of Industrial Organization* 24.1, pp. 69–81.
Fattouh, Bassam and Lavan Mahadeva (2013). "OPEC: What Difference Has It Made?" In: *Annual Review of Resource Economics* 5.1, pp. 427–443.
Fueki, Takuji et al. (2018). "Identifying oil price shocks and their consequences: the role of expectations in the crude oil market". In: BIS Working Papers No. 725.
Golombek, Rolf, Alfonso A Irarrazabal, and Lin Ma (2018). "OPEC's market power: An empirical dominant firm model for the oil market". In: *Energy Economics* 70, pp. 98–115.
Green, Edward J and Robert H Porter (1984). "Noncooperative collusion under imperfect price information". In: *Econometrica: Journal of the Econometric Society*, pp. 87–100.
Greene, David L (1991). "A note on OPEC market power and oil prices". In: *Energy Economics* 13.2, pp. 123–129.
Griffin, James M (1985). "OPEC behavior: a test of alternative hypotheses". In: *The American Economic Review* 75.5, pp. 954–963.
Haltiwanger, John and Joseph E Harrington Jr (1991). "The impact of cyclical demand movements on collusive behavior". In: *The RAND Journal of Economics*, pp. 89–106.
Hansen, Petter Vegard and Lars Lindholt (2008). "The market power of OPEC 1973–2001". In: *Applied Economics* 40.22, pp. 2939–2959.
Hnyilicza, Esteban and Robert S Pindyck (1976). "Pricing policies for a two-part exhaustible resource cartel: The case of OPEC". In: *European Economic Review* 8.2, pp. 139–154.
Hotelling, Harold (1931). "The economics of exhaustible resources". In: *Journal of political Economy* 39.2, pp. 137–175.
Huppmann, Daniel (2013). *Endogenous Shifts in OPEC Market Power: A Stackelberg Oligopoly with Fringe*. Discussion Papers of DIW Berlin 1313. DIW Berlin, German Institute for Economic Research.
Huppmann, Daniel and Franziska Holz (2012). "Crude oil market power—a shift in recent years?" In: *Energy Journal*, pp. 1–22.
— (2015). *What about the OPEC Cartel?* Tech. rep. DIW Roundup: Politik im Fokus.

Kilian, Lutz (2019). "Facts and Fiction in Oil Market Modeling". In: *Federal Reserve Bank of Dallas Working Paper* 1907.
— (2009). "Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market". In: *American Economic Review* 99.3, pp. 1053–69.
Kilian, Lutz and Daniel P Murphy (2014). "The role of inventories and speculative trading in the global market for crude oil". In: *Journal of Applied Econometrics* 29.3, pp. 454–478.
Kilian, Lutz and Xiaoqing Zhou (2019). "Oil Supply Shock Redux?" In: *manuscript, University of Michigan.*
Kolodzeij, Marek and Robert K Kaufmann (2014). "Oil demand shocks reconsidered: a cointegrated vector autoregression". In: *Energy Economics* 41, pp. 33–40.
Mailath, George J and Larry Samuelson (2006). *Repeated games and reputations: long-run relationships.* Oxford University Press.
Maskin, Eric and Jean Tirole (2001). "Markov perfect equilibrium: I. Observable actions". In: *Journal of Economic Theory* 100.2, pp. 191–219.
Nakov, Anton and Galo Nuño (2013). "Saudi Arabia and the oil market". In: *The Economic Journal* 123.573, pp. 1333–1362.
Noreng, Øystein (2006). *Crude power: politics and the oil market.* Vol. 21. IB Tauris.
Ratti, Ronald A and Joaquin L Vespignani (2015). "OPEC and non-OPEC oil production and the global economy". In: *Energy Economics* 50, pp. 364–378.
Rauscher, Michael (1992). "Cartel instability and periodic price shocks". In: *Journal of Economics* 55.2, pp. 209–219.
Rotemberg, Julio J and Garth Saloner (1986). "A supergame-theoretic model of price wars during booms". In: *American Economic Review* 76.3, pp. 390–407.
Salant, Stephen W (1976). "Exhaustible resources and industrial structure: A Nash-Cournot approach to the world oil market". In: *Journal of Political Economy* 84.5, pp. 1079–1093.
Spilimbergo, Antonio (2001). "Testing the hypothesis of collusive behavior among OPEC members". In: *Energy Economics* 23.3, pp. 339–353.
Wilson, Bart J and Stanley S Reynolds (2005). "Market power and price movements over the business cycle". In: *The Journal of Industrial Economics* 53.2, pp. 145–174.
Yergin, Daniel (2011). *The Prize: The epic quest for oil, money & power.* New York: Simon and Schuster.

Centre for Applied Macroeconomics and Commodity Prices (CAMP)
will bring together economists working on applied macroeconomic issues, with special emphasis on petroleum economics.

BI Norwegian Business School
Centre for Applied Macro - Petroleum economics (CAMP)
N-0442 Oslo

www.bi.no/camp

CAMP Working Paper Series
ISSN: 1892-2198