CENTRE FOR APPLIED MACROECONOMICS AND COMMODITY PRICES (CAMP)

CAMP Working Paper Series
No 1/2021

OPEC's crude game

Strategic competition and regime-switching in global oil markets

Thomas Størdal Gundersen and Even Soltvedt Hvinden

BI

OPEC's crude game*

Strategic competition and regime-switching in global oil markets

Thomas Størdal Gundersen† Even Soltvedt Hvinden‡

January 25, 2021

Abstract

We develop a model of oligopolistic competition under imperfect monitoring and dynamic observable demand. Efficient symmetric equilibria feature disciplined cooperative regimes interrupted by rare but severe price wars. The model predicts that the frequency, duration, and supply schedule associated with each regime may persistently deviate from average behavior. We find evidence for the theoretical predictions of our model in historical Organization of Petroleum Exporting Countries (OPEC) output using a Markov-switching Bayesian vector autoregressive model of the global oil market. The evidence suggests that conventional models without regime-switching of oil supply underestimates the linkages between quantities supplied and oil prices.

JEL codes: Q31, C34, C73

Keywords: Regime-switching, OPEC, cartel, price war, crude oil demand and supply

*This work is part of the research activities at the Centre for Applied Macroeconomics and Commodity Prices (CAMP) at the BI Norwegian Business School. The theoretical framework presented in this paper has previously circulated in the working paper "OPEC's crude game: The supply curve in a dynamic, strategic environment". We thank Knut Are Aastveit, Jørgen Juel Andersen, Arthur van Benthem, Hilde C. Bjørnland, Jamie Cross, Bård Harstad, Martin Blomhoff Holm, Felix Kapfhammer, Plamen Nenov, Dario Sidhu, Francesco Ravazzolo, Leif Anders Thorsrud, and Ragnar Torvik, in addition to numerous seminar participants at the 28th Annual Symposium of the Society for Nonlinear Dynamics and Econometrics, VU Amsterdam, BI Norwegian Business School, the Wharton School at the University of Pennsylvania, and the University of Oslo for helpful comments and suggestions. The authors declare that they have no relevant or material financial interests that relate to the research described in this paper.

†Centre for Applied Macroeconomics and Commodity Prices (CAMP), BI Norwegian Business School. Corresponding author: thomas_gundersen@me.com

‡Centre for Applied Macroeconomics and Commodity Prices (CAMP), BI Norwegian Business School. E-mail: even.c.hvinden@gmail.com

1 Introduction

The origin and propagation of shocks to the price of crude oil and on to the wider economy has been a subject of long-standing interest and lively debate (Hamilton 1985, Kilian 2009, Baumeister and Hamilton 2019a). To develop empirical models of the global oil market an assumption has to be made, among others, on the market behavior of oil producers in the Organization of Petroleum Exporting Countries (OPEC). Members of OPEC are endowed with considerable market power, strive actively to increase their profits by coordinating output restraint, and face no legal constraints on doing so. Yet empirical evidence suggests output coordination has been only partially successful and that OPEC's conduct varies considerably over time.[1] In this paper we study the implication of strategic behavior in an infinitely repeated game of oligopolistic quantity competition in which players (OPEC members) coordinate on a symmetric profit-maximizing public equilibrium in an environment with variable current- and future expected demand, capacity constraints to output and imperfect monitoring. The model predicts that under symmetric efficient equilibria supply correspondences are subject to regime-switching and that observed behavior may potentially deviate persistently from the average. We then provide evidence of such behavior in the data and show that empirically accounting for regime-switching overturns the conclusion that OPEC is passive with respect to price developments.

The conventional approach in the empirical macroeconomic literature is to assume that global oil supply is well-approximated by a stationary, linear process, e.g. Kilian (2009) and Baumeister and Hamilton (2019a). While our model of dynamic quantity competition does admit such behavior, we show that optimal symmetric equilibria imply price-quantity relationships that are non-linear along two dimensions.

First, OPEC's aggregate production alternates between a reward- and punishment phases. In the reward phase output is restrained and prices are raised relative to the static Nash equilibrium benchmark. Conversely, in the punishment phase output is elevated and prices are lower than the static benchmark. Optimal equilibria are generally characterized by persistent and disciplined

[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).

reward phases that are disrupted by rare but severe output wars. Switches between reward- and punishment phases are manifested empirically by pro- and countercyclical supply regimes, where OPEC's output respectively moderates- and exacerbates price fluctuations.

Second, conditional on past- and current market conditions, shifts in expected future profitability cause persistent deviations from average behavior. The expectation of falling demand raises the value of deviating from jointly profit-maximizing production relative to the losses sustained during a future output war. OPEC's optimal response is then to reduce cartel discipline in the cooperative regime and increase the frequency of punishments. The reduction in anticipated future losses sustained under punishment may yield highly non-linear and even non-monotonic price-quantity relationships within regimes.[2]

Our theoretical analysis motivates two empirical research questions: Is there evidence in the data that OPEC behaves strategically, manifested by non-linear price-quantity relationships? Second, does the non-linearity affect structural inference of aggregate dynamic properties of supply and demand? In particular, do standard oil market models underestimate the link between quantities supplied and prices?

These questions are investigated with a Markov switching Bayesian vector autoregressive model (MS-BVAR) adapted from Kilian (2009). This model has three variables, OPEC output, a global demand indicator and the real price of oil and is estimated using monthly data from 1985–2019. A natural benchmark is provided by a BVAR without regime-switching that is otherwise identical. The dynamics in each regime are governed by a distinct covariance matrix. Contemporary structural oil market models are typically identified through restrictions on the contemporaneous responses of variables to structural shocks obtained by transformations of the (reduced form) covariance matrices. Strategic competition is predicted to yield substantial and persistent switches in covariances and may therefore have a considerable impact on structural inference.

We find robust evidence for the existence of pro- and countercyclical regimes in OPEC output. The empirical model suggests that OPEC's change in behavior moderates or exacerbates price fluctuations. In contrast, the single-regime BVAR suggests that OPEC is on average passive with respect to price developments. We corroborate our identified pro-and countercyclical regimes

[2]This property is reflected in the observation by Sadek Boussena, OPEC conference president 1989-1990, that "OPEC is strong when prices are weak, and weak when prices are strong". See Bret-Rouzaut and Favennec (2011).

with historical accounts, contemporaneous analysis by the International Energy Agency (IEA), and inferred price expectations from futures contracts. Our identified regimes are consistent with all three sources of evidence, suggesting that our model –albeit simple– is not merely capturing extraneous correlations.

Finally, our results suggests the fundamental dynamics of OPEC's behavior has not been constant throughout the sample. Like previous studies, we find that the 1985–1999 period is marked by low cartel discipline with frequent switches between regimes and low confidence in regime classification. The evidence suggests that post-1999 countercyclical regimes were longer and more disciplined while the procyclical regime became shorter and more intense.

This paper contributes to three strands of literature. We 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 current demand. Other notable contributions are Salant (1976), Hnyilicza and Pindyck (1976), Huppmann (2013), Nakov and Nuño (2013), Behar and Ritz (2017), and Jaakkola (2019). The common theme in these contributions is that a representative OPEC producer competes inter-temporally with a non-OPEC fringe. However, the non-cooperative aspects of OPEC members' interaction are not modeled and variation in cartel discipline is absent or exogenously imposed. In contrast we consider a model where OPEC's output is jointly and endogenously determined by current and future expected conditions in an explicitly non-cooperative setting.

Our analysis contributes to the literature on time-varying oil market dynamics, identification of global supply and demand shocks, and regime-switching techniques in econometrics. A closely related analysis is Almoguera, Douglas, and Herrera (2011) who test for regime shifts in OPEC output using a switching simultaneous equation model (SEM) adapted from Porter (1983). In contrast, our paper combines theoretical advances with a fully dynamic econometric methodology and yields novel results. Unlike their static SEM framework, the dynamic properties of our MS-BVAR readily inform the structural vector autoregression (SVAR) models currently used in the literature. Nevertheless, we reach similar conclusions for most of the 1985–1998 time period in which our studies overlap. However, in contrast to their findings, our analysis provides evidence that OPEC's dynamic pattern of behavior shifted post-1998. Another related contribution is Ratti and Vespignani (2015), who document a structural break in OPEC's response to market developments in the first quarter of 1997 using an SVAR model based on Kilian (2009). In a time-varying parameter (TVP) VAR framework, Baumeister and Peersman

(2013) finds that the increase- and decrease in oil price and production volatility may plausibly be explained by a substantial reduction in the demand and supply elasticities leading up to 1985, the start of our sample. Using high-frequency data, Känzig (*Forthcoming*) identifies an OPEC oil supply news shock and finds that these shocks have statistically and economically significant effects. Our paper complements this analysis by shedding light on how strategic competition between OPEC members shapes the underlying process generating these shocks in the data.

The evidence of time-varying cartel discipline in our findings informs a lively and ongoing debate on the identification of supply and demand shocks in the global oil market, see Hamilton (1985), Kilian (2009), Kilian and Murphy (2014), Caldara, Cavallo, and Iacoviello (2019), Baumeister and Hamilton (2019a) Baumeister and Hamilton (2019b), Bjørnland (2019), among others. Our results suggest that traditionally estimated, constant supply elasticities, by measuring average behavior over distinct pro- and countercyclical regimes, are biased towards zero. The traditional approach will generally find that OPEC is passive on average, a result that does not describe the observed pattern of persistent regime-contingent responses. Our analysis suggests that ignoring regime switching will therefore underestimate the impact of supply developments. Thus our paper joins a burgeoning literature applying regime switching econometrics in macroeconomic research, see for example, Lo and Piger (2005), Auerbach and Gorodnichenko (2012), Billio et al. (2016), and Bjørnland, Larsen, and Maih (2018).

We proceed as follows. In Section 2 we consider the stylized facts of OPEC's time-varying behavior and briefly review the salient properties of their strategic environment to motivate our model. Section 3 presents the model and solution concept. Empirical predictions on market dynamics under strategic competition are stated in Section 4. We introduce the MS-BVAR and evaluate evidence for our predictions in Sections 5 and 6, respectively. Section 7 presents robustness exercises of our empirical results. We conclude and discuss avenues for future research in Section 8.

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] The evidence reported in the literature suggests 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. 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, sample splitting, and time-varying parameter models.[4] All the aforementioned contributions find persistent time-variation in OPEC behavior.[5]

For a stylized illustration of OPEC's time-varying behavior, consider Figure 1, plotting twelve-month changes in monthly OPEC crude oil output and log real crude oil prices between January 1985 and December 2019.[6] We have highlighted by shaded bars four significant episodes of oil price drops: 1986, 1997-1998, 2008-2009, and 2014-16. Measured across the entire 1985-2019 sample, OPEC output and price growth are linearly uncorrelated. However, this masks significant and sign-varying correlation in sub-samples. Between 2002 and 2014, OPEC's output 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, with the correlation becoming negative.

We formalize OPEC's time-varying behavior in a model of oligopolistic quantity competition with imperfect monitoring, dynamic residual demand, and capacity constraints. The remainder of this section briefly motivates our

[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]See Almoguera, Douglas, and Herrera (2011), Barros, Gil-Alana, and Payne (2011), Baumeister and Peersman (2013), Kolodzeij and Kaufmann (2014), Huppmann and Holz (2012), Ratti and Vespignani (2015) respectively. See also Dees et al. (2007).

[5]See also Dvir and Rogoff (2009) and Dvir and Rogoff (2014) who consider very 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 we have deflated the U.S. refiner's acquisition cost of crude oil imports from the U.S. Energy Information Administration with the average all-item U.S. CPI from the Federal Reserve Bank of St. Louis data service (FRED).

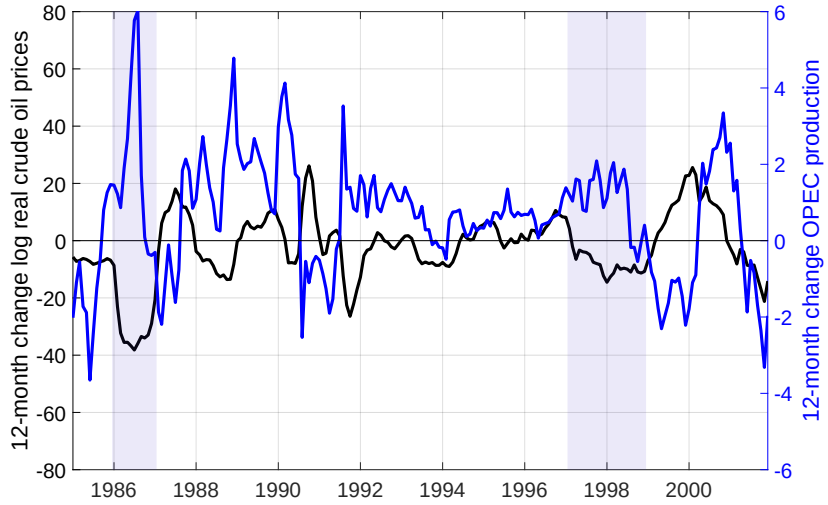

(a) 1985 - 2002

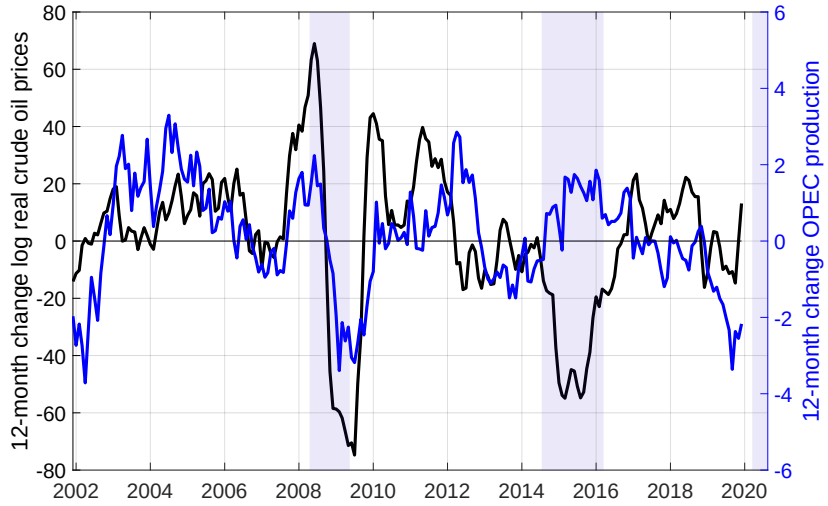

(b) 2002 - 2019

Figure 1: Log real oil prices. OPEC crude oil output, millions of barrels per day. Twelve-month change. January 1985 to December 2019. Highlighted historical episodes: 1986 and 1997 output wars, global financial crisis of 2008, and the 2014-2016 price fall. Source: International Energy Agency Monthly Oil Data Statistics, U.S. Energy Information Administration, Federal Reserve Bank of St. Louis (FRED).

model specification and equilibrium selection criterion.

Imperfect monitoring is viewed as the likely fundamental cause behind intentional market flooding by OPEC members. The idea is that OPEC agreements which successfully restrict total production create an incentive for individual producers to cheat.[7] 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.[8] The incidence of such punishments, or output wars, has directed attention in the literature to imperfect monitoring models as a salient framework capturing an important property of OPEC's strategic environment.[9][10]

The standard reference model of oligopolistic competition under imperfect monitoring is due to Green and Porter (1984). This framework features an environment with static observable demand and no capacity constraint to output. Yet persistent changes to commodity demand and fixed short-run production capacity are widely viewed as empirically important features of the oil market and have the theoretical potential to alter producer behavior.[11]

[7]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 company, is a plausible non-OPEC candidate for a dominant producer. It has been reported that the Russian government assisted in coordinating output cuts with OPEC following the 2014 price fall, leading to the so-called "OPEC+" format, see e.g. "Russia, Saudi Arabia agree OPEC+ format should be extended", accessed August 21 2018 from www.reuters.com.

[8]This is a general result in the theory of repeated games under imperfect monitoring. See for instance Mailath and Samuelson 2006 pp. 233.

[9]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 statement 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 "Satellites Aid the Chase for Better Information on Oil Supplies", accessed October 12 2018 from www.nytimes.com.

[10]Market analysts and historians have argued that the steep price declines in 1986 and 1997 were explicitly due to intentional market flooding by leading OPEC producers aimed to punish quota violations by other OPEC members. See the accounts in Noreng (2006), Downey (2008), Yergin (2011), and also Coll (2012).

[11]See for instance the Energy Information Agency: What drives crude oil prices? or the many econometric analyses of the crude oil market cited above. The impact of variation in current- and future expected demand on cartel discipline is studied in Rotemberg and

We therefore impose capacity constraints and dynamic residual demand in our model, giving it a clear short- to medium-run interpretation. Non-OPEC output is not explicitly modeled, and is interpreted as competitive and subsumed in the dynamic residual demand function.

Repeated games admit a multiplicity of equilibria with distinct behavior. Empirical prediction are conditional on an equilibrium selection argument. For instance, our model admits equilibria where with positive probability the stage-game Nash quantity is played forever on the equilibrium path. However, equilibria with such grim trigger punishments are generally sub-optimal even *ex-ante* in the reward phase (Abreu, Pearce, and Stacchetti 1986). In studying efficient symmetric equilibria we appeal to the notion that OPEC members individually have an incentive to coordinate on behavior that maximizes their expected profits.

3 Model of oligopolistic quantity competition

We briefly present our model of oligopolistic quantity competition. The setup is standard, satisfying the key assumptions in Abreu, Pearce, and Stacchetti (1990) but augmented with a Markov chain for residual demand. Parametric restrictions are detailed in Appendix A.1.

3.1 The stage game

In each stage game G, two symmetric, dominant producers compete in homogeneous quantities.[12] Each producer $i \in \{1, 2\}$ chooses an output level q from a finite action set $Q \subset \mathbb{R}_{\geq 0}$ and receives an expected payoff $\pi : Q^2 \times X \to \mathbb{R}$, where π is strictly continuous and concave in q and where $X \subset \mathbb{R}_{\geq 0}$ is the residual demand space. Producers' can render their output unobserved at no extra cost. The expected profits of player i are given by $\pi(\boldsymbol{q}, x)_i = p(\boldsymbol{q}, x)q_i - c(q_i)$ where $x \in X$, $\boldsymbol{q} \in Q^2$, and inverse demand $p(\boldsymbol{q}, x)$ is inelastic with respect to quantities, implying elastic demand. Realized prices $p(\boldsymbol{q}, x, \tilde{\theta}) = \tilde{\theta}p(\boldsymbol{q}, x)$, and hence also realized profits, depend on a unobserved stochastic variable $\tilde{\theta}$, log-normally distributed $\ln \tilde{\theta} \sim N\left(-\sigma_\theta^2/2, \sigma_\theta^2\right)$. We further assume that the inverse elasticity is weakly increasing in quantities, a standard property. Finally, residual demand levels are such that for every $x \in X$ the stage-game Nash equilibrium- and jointly profit-maximizing quantities q^n, q^m are in Q.

Saloner (1986) and Haltiwanger and Harrington Jr (1991).

[12] Due to symmetry the game generalizes straightforwardly to any number $n > 2$ of players.

3.2 The repeated game

The repeated game $G^{\infty}(\delta)$ is played over an infinite horizon, with time indexed by t. The common discount factor is $\delta \in (0,1)$. Now gather demand levels in a d-dimensional vector $\boldsymbol{x} = (x^1, \ldots, x^d) \subset X$ where $0 \leq x^1 \leq \ldots \leq x^d < \infty$. Demand evolves as a Markov chain over $\boldsymbol{x}$ with transition matrix M which is stationary and irreducible. Let $D = \{1, \ldots, d\}$ be the index set over states and normalize M to be right-stochastic so the elements of each row $\boldsymbol{m}'_j$, $j \in D$ sum to unity, $\sum_{m_{js} \in \boldsymbol{m}_j}^{d} m_{js} = 1$. In the following, state-j- and time-t values of endogenous variables are denoted by a $j \in D$ superscript and t subscript. The stage game proceeds as follows:

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

3.3 Optimal equilibrium

The model has a unique, optimal symmetric equilibrium, formally derived in Appendix A.2. For each demand state $x \in X$ producers' play either the reward or punishment quantity, $\overline{q}(x)$ and $\underline{q}(x)$. Transitions between phases are endogenously determined by observable actions and the price signal. Due to imperfect monitoring, punishments occur on the equilibrium path and there are therefore $2d$ possible states, governed by the $2d \times 2d$ transition matrix T. The key property of optimal equilibria is that they are necessarily extreme in the sense that reward- and punishment phases yield respectively the highest- and lowest payoffs that may be supported as equilibrium outcomes. The optimal reward phase is generally persistent and features restrained output, whereas the punishment phases are short and feature elevated output. We refer to this distinct pattern of behavior as high-powered incentive creation. We now turn to the salient empirical predictions of high-powered equilibria on the dynamic patterns of quantities supplied.

4 Market dynamics under strategic competition

This section derives empirical predictions on the joint behavior of prices and quantities in the optimal equilibrium of our model. Recall that structural in-

ference in macroeconomic models relies on linear transformations of the covariance matrix of the reduced form VAR. Consequently, structural responses are identified under the assumption that relationships between endogenous variables are well-approximated by a linear stochastic process. Our focus is thus on the impact of strategic competition on the reduced form covariances between prices and quantities. We highlight two sources of non-linearity in the relationship between prices and quantities. First, there is regime-switching across supply correspondences of the reward- and punishment phases. Second, there may be persistent deviations from average behavior, both in the pattern of switching and output within regimes. The strength of these non-linear effects is generally increasing in incentive power. The consequence is that average representations of dynamic supply behavior in linear models may be rendered uninformative.

To fix ideas we consider an analytical expression of the covariance between quantities and prices. Recall that the equilibrium features $2d$ states. Gather reward- and punishment regimes in state $1, \ldots, d$ and $d+1, \ldots, 2d$ respectively and let r_{ij} denote the transition probability from state i to j. The covariance is thus given by

$$\mathrm{Cov}(p, q) = \sum_{i=1}^{2d} \sum_{j \geq i}^{2d} r_{i,j}(q^i - q^j)(p^i - p^j) \tag{1}$$

with $r_{i,j} = (\lambda_i t_{ij} + \lambda_j t_{ji})$ where t_{ij} is an element i, j in the $2d \times 2d$ transition matrix T (see Equation (13) in Appendix A.2) and λ_i is element i in the corresponding stationary distribution of the Markov process. Notice that the covariance identity is straightforwardly transformed to an autocovariance by taking the appropriate time lags.

We decompose the covariance as

$$\begin{aligned}\mathrm{Cov}(p,q) = & \underbrace{\sum_{i=1}^{d} r_{i,i+d}(q^i - q^{i+d})(p^i - p^{i+d})}_{\text{Transition between reward and punishment phases, fixed demand}} \\ & + \underbrace{\sum_{i=1}^{d}\sum_{j\geq i}^{d} r_{i,j}(q^i - q^j)(p^i - p^j)}_{\text{Demand variation, reward phase}} + \underbrace{\sum_{i=d+1}^{d}\sum_{j\geq i}^{2d} r_{i,j}(q^i - q^j)(p^i - p^j)}_{\text{Demand variation, punishment phase}} \\ & + \underbrace{\sum_{i=1}^{d}\sum_{j\geq d-i+1}^{2d-i} r_{i,j}(q^i - q^j)(p^i - p^j)}_{\text{Transition between reward and punishment phases, variable demand}}\end{aligned} \tag{2}$$

where our predictions on the between- and within regime variation in covariances concern the first- and second row of Equation (2), respectively. The following proposition states that the first, second and third rows have arguments with sign that is generally weakly negative, positive, and indeterminate.

Proposition 1. ***Covariance.*** *Let $\delta \to 1$, $\sigma_\theta \to 0$ and $\boldsymbol{x} : \boldsymbol{\pi}^m \geq \boldsymbol{\pi}^n$. Then*

$$\begin{aligned}(\underline{q}^i - \underline{q}^j)(\underline{p}^i - \underline{p}^j) &\geq 0 \\ (\overline{q}^i - \overline{q}^j)(\overline{p}^i - \overline{p}^j) &\geq 0 \\ (\overline{q}^i - \underline{q}^i)(\overline{p}^i - \underline{p}^i) &\leq 0\end{aligned}$$

for every $j \geq i \in \{1, \ldots, d\}$ and with strict inequality for some i.

The proof is in Appendix A.3. The intuition is as follows. The limiting conditions ensure that supply schedules in the reward phase are monotonic in price-quantity space. Given upward-sloping supply schedules it follows straightforwardly that demand variation induces a weakly positive covariance between prices and quantities. Likewise, transitions across regimes along the demand correspondence induce a weakly negative covariance between prices and quantities. Two corollaries follow immediately.

Corollary 1. *If $\overline{\boldsymbol{q}} = \underline{\boldsymbol{q}} = \boldsymbol{q}^n$ then* $\mathrm{Cov}(p,q) \geq 0$.

If the equilibrium features no intertemporal incentives then stage-game equilibria are played every period. Without transitions between reward- and punishment phases, the first- and second rows of the covariance identity in Equation (2) equal zero and the supply schedule induced by the stage-game equilibrium is monotonic in price-quantity space.

Corollary 2. *Transitions* $(\overline{q}^i - \underline{q}^j)(\overline{p}^i - \underline{p}^j)$ *and* $(\underline{q}^i - \overline{q}^j)(\underline{p}^i - \overline{p}^j)$ *have an indeterminate sign.*

When transitions across regimes are combined with shifts in demand the linear covariance depends on relative magnitudes that are in general ambiguous.

The implication of regime-switching is that the relationship between prices and quantities within regimes – the second row of Equation (2) – may not be well-approximated by a regime-independent and linear supply schedule. In particular the autocovariances will be biased towards zero if strategically motivated changes to price and quantities – the first row of Equation (2) – are substantial and persistent.

Prediction 1. ***Distinct pro- and countercyclical regimes.*** *Strategic competition induces a sign change in the autocovariances between prices and quantities with* $\mathrm{Cov}(p, q) \leq 0$ *and* $\mathrm{Cov}(p, q) \geq 0$ *denoted pro- and countercyclical, respectively.*

We show in Appendix A.4 that this bias is generally stronger for higher-powered equilibria. Optimal strategic competition generates persistent, disciplined reward phases disrupted by rare but severe output wars. Lower-powered, and hence sub-optimal, equilibria feature less disciplined, persistent cooperative phases with more frequent, shorter or less intense output wars. The strategically induced change in market dynamics (autocovariances) will therefore be starker under higher-powered equilibria with greater and more persistent changes to quantities supplied. The upshot is that under higher-powered competition a single representative supply schedule becomes an increasingly ill-suited representation of quantities supplied.

We now turn to variation in behavior within regimes. The power of equilibrium incentives is jointly and endogenously determined by fundamental properties of the strategic environment. Appendices A.5-A.7 consider the comparative statics of exogenous changes to monitoring quality σ_θ, demand level x, and transition probabilities M on incentive power. The key insight is that producers' ability to restrain output is jointly determined by relative current- and future expected market conditions. The implication is that evolving expectations over future profitability exert a force on quantities supplied independent of past- and present market conditions. This effect may in principle be arbitrarily powerful: Appendix A.8 demonstrates that for a given level of current demand x, there exist parameters (expectations) that induce any output level in $\{q^m(x), \ldots, q^n(x)\}$ and $\{q^n(x), \ldots, q^*\}$ in the reward and punishment phase, respectively. Lasting changes to expected profitability induce

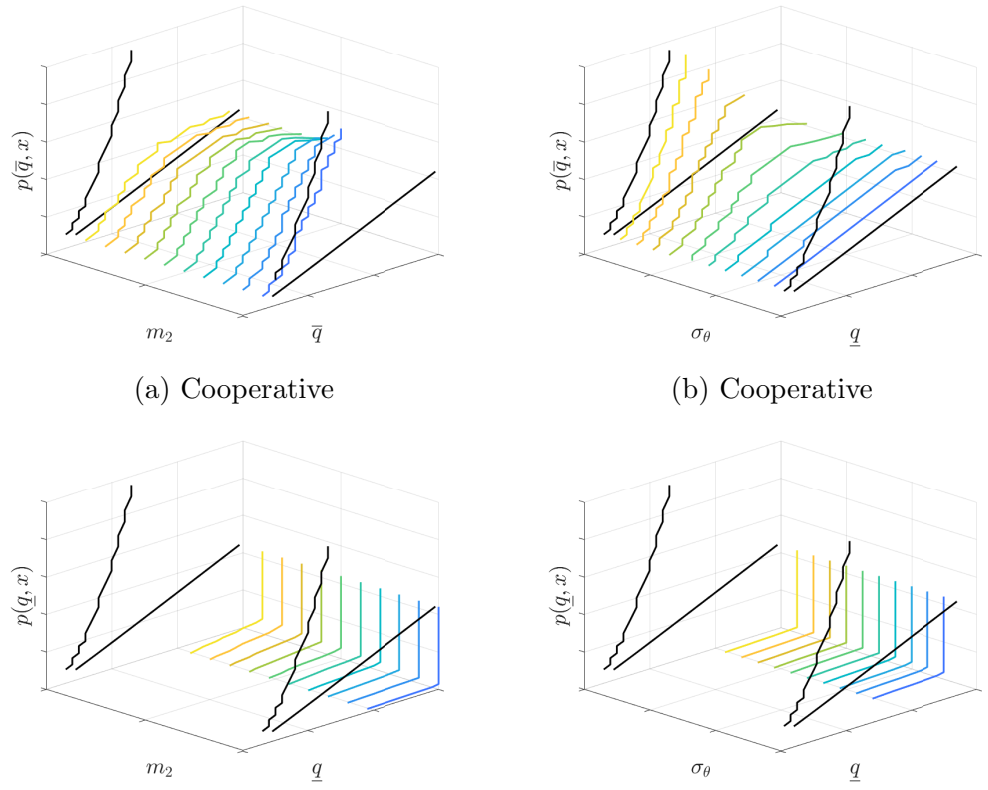

(a) Cooperative

(b) Cooperative

(c) Punishment

(d) Punishment

Figure 2: Equilibrium prices $p(q, x)$ vs. quantities q for a range of demand, persistence- and signal noise parameters x_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.

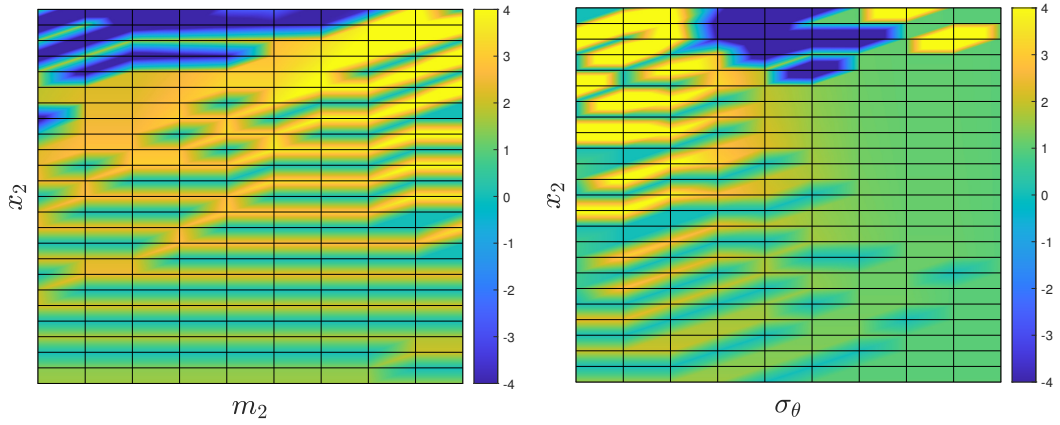

(a) Persistence (b) Signal noise

Figure 3: Heat map plots of approximate, reward-phase price elasticities of supply vs demand level x_2, persistence m_2, and signal noise σ_θ. The color bar is censored at -4, and 4.

substantial variation not only in the frequency and duration of regime-switches (Appendix A.4) but also the slope and level of supply correspondences, manifested through changes in the sign and magnitude of within-regime autocovariances, the first row of Equation (2). Thus conditional on the past- and present state of the market, constraints on incentive power yield more frequent, less severe punishments and shorter, but less disciplined reward phases. Under evolving market conditions we expect both patterns of regime-switching and within-regime behavior to deviate persistently from the data sample average.

Prediction 2. ***Deviation from average behavior.*** *Strategic competition may induce variation in the frequency and duration of switches across regimes $r_{i,j}$ and the autocovariances* $\mathrm{Cov}(p, q)$ *between prices and quantities within regime, yielding persistent deviations from average covariance.*

The potential for strategic competition to induce substantial changes to within-regime behavior increases in a high-powered equilibrium as the larger potential range of incentive power supported on the equilibrium path allows the supply correspondence to trace out a greater range of price-quantity space.

To illustrate and help build intuition we present numerically solved equilibrium of a game under $d = 2$ states labeled 1 and 2 with corresponding demand levels satisfying $0 < x^1 < x^2 < \infty$. See Appendix B for table of parameter values. The numerical solution algorithm is detailed in online Appendix G. Figure 2 plots the corresponding supply correspondences, that is, equilibrium

prices $p(q, x)$ vs. total quantities q in reward- and punishment phases. The limiting competitive (stage-game equilibrium) and profit-maximizing (monopoly) actions are plotted in black and the equilibrium action in color, with lighter colors indicating more constrained incentive creation. The (implicit) equilibrium price elasticity of supply is plotted in Figure 3 and provides a unit-free measure of the supply schedule. Figure 4 plots the stationary distribution over the four states $\boldsymbol{\mu}$, satisfying $T'\boldsymbol{\mu} = \boldsymbol{\mu}$ where T is the transition matrix over phases and states. It shows that the share of time spent in the punishment state increases in demand, and more so when monitoring is poor or expected profitability is low. Finally, Figure 12 in Appendix C plots equilibrium incentive power for a range of persistence and monitoring quality parameters, illustrating how incentive power is non-monotonic in demand but increasing in monitoring quality and future expected profitability.

We conclude by briefly discussing non-linear supply correspondences. The non-linearity results from the interaction of multiple non-linear effects and is therefore challenging to characterize analytically. Intuition may be provided by the special case with constant marginal cost- and elasticity of demand. In that case Collie et al. (2004) shows (numerically) that supporting monopoly quantities in equilibrium requires a higher discount factor when the inverse elasticity of demand is greater, given a grim-trigger strategy with absorbing punishment. But by Assumption 1 the equilibrium inverse elasticity is weakly increasing in quantities. Because $q^m(x)$ is weakly increasing in x it follows that the price elasticity of demand is falling in x. Thus jointly profit-maximizing quantities require greater incentive power to sustain at higher demand levels. Yet if incentives are sufficiently constrained by a combination of poor monitoring environment, falling profitability, and capacity constraints to punishment intensity (Appendix A.5-A.8) it may be optimal to increase cooperative quantities $\overline{q}$ to sustain cooperation by reducing the one-shot deviation profit, a point first made by Rotemberg and Saloner (1986). But an increase in $\overline{q}$ will reduce profits $\overline{\pi}$ which again reduces incentive power, an effect itself increasing in the inverse elasticity of demand and thus accentuated at high prices. The implication is that it not possible in general to rule out highly non-linear and even non-monotonic supply correspondences in price-quantity space.

5 A Regime-Switching Oil Market VAR Model

We evaluate evidence of regime-switching and persistent deviations from average behavior using an MS-BVAR model. To provide a natural benchmark, we estimate a standard BVAR model using the same data and specification,

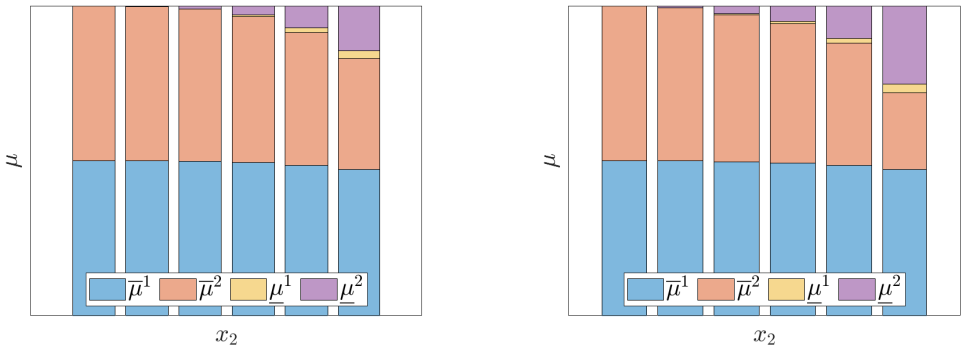

(a) Low signal noise

(b) High signal noise

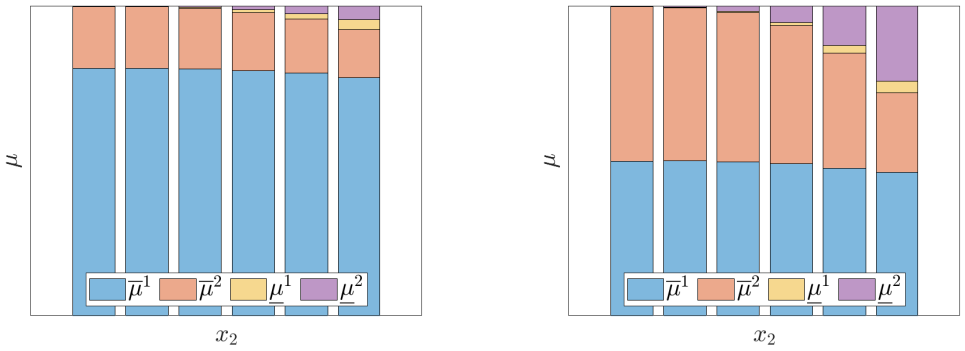

(c) Low persistence

(d) High persistence

Figure 4: Stationary distribution $\boldsymbol{\mu}$ of the transition matrix T, Equation 13, satisfying $T'\boldsymbol{\mu} = \boldsymbol{\mu}$ for $d = 2$, against x_2 the level of residual demand in the high-demand state. The elements $\mu_1 = \overline{\mu}^1$, $\mu_2 = \overline{\mu}^2$ the reward phase and $\mu_3 = \underline{\mu}^1$, and $\mu_4 = \underline{\mu}^2$ the punishment phase, for low- and high demand respectively.

but without switching and with uninformative priors.

5.1 The Model

Consider the following reduced form VAR model

$$\boldsymbol{y}_t = \boldsymbol{\mu}(s_t) + \sum_{\ell=1}^{L} A_\ell(s_t)\boldsymbol{y}_{t-\ell} + \boldsymbol{e}_t \qquad \mathbb{E}\left[\boldsymbol{e}_t\boldsymbol{e}_t' \mid s_t\right] = \Sigma_{\boldsymbol{e}}(s_t) \tag{3}$$

where $\boldsymbol{y}_t = \left[\,\Delta q_t,\, \Delta x_t,\, \Delta \ln p_t\,\right]'$ is the vector of endogenous variables, namely the twelve-month change in OPEC production, a measure of real economic activity and the twelve-month change in the log real price of oil. The vector $\boldsymbol{\mu}$ contains intercepts and $\boldsymbol{e}_t$ are the reduced form errors with a positive semi-definite and symmetric covariance matrix $\Sigma_{\boldsymbol{e}}$. Our sample runs over the 1985:M01–2019:M12 time period. See Figure 15 in Appendix C for a plot of the standardized raw data. A description of the data and sources can be found in Appendix D.

Our model differs from Kilian (2009) in three respects. Firstly, the estimated parameters $\boldsymbol{\mu}$, $\{A_\ell\}_{\ell=1}^{L}$ and $\Sigma_{\boldsymbol{e}}$ are allowed to change discretely across time with the state variable $s_t \in \{\mathrm{PC}, \mathrm{CC}\}$ denoting the pro- and counter-cyclical regime, respectively. The regime indicator s_t evolves according to a Markov chain with transition probability matrix P. Secondly, we substitute aggregate global crude oil production with that of just the OPEC member countries. Third and finally we employ the OECD+6 index of industrial production provided by Baumeister and Hamilton (2019a) rather than the Kilian (2009) index as a measure of real economic activity.[13]

In order to keep our model parsimonious and to facilitate a clear interpretation of regimes, we partition the model so that only the parameters in the OPEC oil supply equation are functions of s_t

$$\begin{bmatrix}\Delta q_t \\ \boldsymbol{y}_{\circ t}\end{bmatrix} = \begin{bmatrix}\mu_\bullet(s_t) \\ \boldsymbol{\mu}_\circ\end{bmatrix} + \begin{bmatrix}\sum_{\ell=1}^{L} A_{\bullet\ell}(s_t)\boldsymbol{y}_{t-\ell} \\ \sum_{\ell=1}^{L} A_{\circ\ell}\boldsymbol{y}_{t-\ell}\end{bmatrix} + \begin{bmatrix}e_{\bullet t} \\ \boldsymbol{e}_{\circ t}\end{bmatrix}, \tag{4}$$

where the $\bullet$ and $\circ$ subscripts denote regime-switching and regime-fixed blocks respectively. These blocks can be estimated independently and subsequently

[13] We detrend the OECD+6 index by taking twelve-month growth rates. This is equivalent to extracting the cyclical component of the series as the forecast error of a random walk model with a twelve months ahead forecast horizon, as recommended by Hamilton (2018). Hence, we achieve a consistent twelve-month change transformation across all variables included in the model. The Kilian index on the other hand, constructed as a deviation from a linear trend, should not be transformed (Kilian and Zhou, 2018).

transformed to obtain the VAR representation (Hamilton, 1994; Hamilton, 2016). While the regime-fixed component is estimated as a standard VAR model, the OPEC equation is estimated as

$$\Delta q_t = \mu_\bullet(s_t) + B_\bullet(s_t)\boldsymbol{y}_{\circ t} + \sum_{\ell=1}^{L} C_{\bullet\ell}(s_t)\boldsymbol{y}_{t-\ell} + v_t \qquad v_t \sim N(0, \sigma_v^2(s_t)) \tag{5}$$

where $B_\bullet(s_t)$ is the vector of parameters describing the contemporaneous relationship between Δq_t and $\boldsymbol{y}_{\circ t} = [\,\Delta x_t,\, \Delta \ln p_t\,]'$.

5.2 Priors, estimation and structural inference

To estimate the model we follow Hamilton (2016) and employ the Gibbs sampler, a Markov Chain Monte Carlo (MCMC) algorithm.[14] We refer to Appendix E for a detailed overview of the prior distributions selected and to online Appendix H for details about the estimation procedure. Following Hamilton and Herrera (2004) we estimate our model with 24 lags of the endogenous variables. For the regime-fixed component, we employ uninformative natural conjugate normal inverse-Wishart priors so that the posterior distributions will have mean and variance corresponding to the ordinary least-squares (OLS) estimates.

For Equation (5) we employ independent normal inverse-Wishart priors and set the prior means of parameters $\beta_2(s_t) \in B_\bullet(s_t)$ to be -0.2 and 0.2 in the pro- and countercyclical regime respectively. Two considerations motivate our choice of prior means for $\beta_2(s_t)$. First, having sufficiently different priors across regimes reduces the probability that our estimation algorithm runs into degeneracy, i.e. that s_t takes on only a single value for all t.[15] Second, given the theoretical predictions discussed above, we hold prior beliefs that this parameter will change sign with s_t as OPEC exhibits pro- and countercyclical behavior.[16]

The remaining VAR parameters have prior means of zero as they are growth rates and specified with sizable variance. Finally we prescribe Dirichlet pri-

[14] We make 200,000 draws from the sampler and discard the first 100,000.

[15] This is the label-switching problem in Bayesian estimation of latent Markov models (see among others Celeux, Hurn, and Robert, 2000; Jasra, Holmes, and Stephens, 2005; Geweke, 2007) where identical marginal posterior distributions across subsamples render the regimes unidentifiable.

[16] It is worth noting that the main results remain robust to different choices of the prior means so long as they are sufficiently different. We report below that the posterior means of $\beta_2(s_t)$ are larger in magnitude than the prior means, suggesting that the data favor opposite signs for $\beta_2(s_t)$.

ors for the columns of transition probabilities matrix P. Our choice of shape parameters however, result in prior distributions that are approximately uniformly distributed on $[0, 1]$ so the data is allowed to inform regime persistence.

The structural shocks are identified by Cholesky decomposition of the variance-covariance matrix. This method is tractable, transparent, without excess computational burden and has been widely applied to study oil market dynamics in small VAR models (see among many others Kilian, 2009; Ratti and Vespignani, 2015; Gundersen, 2020). Moreover, it makes the comparison to the regime-invariant BVAR model readily feasible. The widespread use of this identification scheme also facilitates a direct comparison of estimated dynamic behavior with existing literature. Online Appendix I demonstrates that the sign of unrestricted responses are determined by the sign of estimated covariances when identifying the model using Cholesky decomposition. The identification scheme assumes a recursive structure on the timing of the impact of structural shocks on the endogenous variables.[17] More specifically, the reduced form errors can be decomposed as $\boldsymbol{e}_t = S(s_t)\boldsymbol{\varepsilon}_t$, or

$$\begin{bmatrix} e^q \\ e^x \\ e^p \end{bmatrix}_t = \begin{bmatrix} \varsigma_{11}(s_t) & 0 & 0 \\ \varsigma_{21}(s_t) & \varsigma_{22}(s_t) & 0 \\ \varsigma_{31}(s_t) & \varsigma_{32}(s_t) & \varsigma_{33}(s_t) \end{bmatrix} \begin{bmatrix} \varepsilon^q \\ \varepsilon^x \\ \varepsilon^p \end{bmatrix}_t \tag{6}$$

with matrix $S(s_t)$ being the lower triangular component of the Cholesky decomposition of $\Sigma_{\boldsymbol{e}}(s_t)$ and $\boldsymbol{\varepsilon}_t$ the vector of structural uncorrelated shocks, $\mathbb{E}[\boldsymbol{\varepsilon}_t\boldsymbol{\varepsilon}_t'] = I$. Because $\Sigma_{\boldsymbol{e}}(s_t)$ varies with s_t, we will get two sets of structural parameters and shocks.

Following Kilian (2009) we order the supply variable at the top, followed by global activity and finally the real price of oil. An OPEC supply shock in the countercyclical regime is interpreted as unexpected shortfalls or unanticipated increases in the crude oil output of OPEC member countries. This is the classic interpretation of an oil supply shock (Hamilton, 1985; Kilian, 2009). In the procyclical regime where there is either a flooding or balancing of the market, an OPEC supply shock will reflect output wars of unusual intensity or unusually large withdrawals of crude from the market.[18] By ordering OPEC crude oil production as the first equation we impose a vertical

[17] For cases where such zero-restrictions cannot be justified, popular but computationally demanding alternatives such as sign-restrictions (Rubio-Ramirez, Waggoner, and Zha, 2010) and the more general Baumeister-Hamilton approach (Baumeister and Hamilton, 2015) should be applied.

[18] It is important to note that the shocks themselves must be interpreted as occurring within a given regime and are independent of each other. We do not model a relationship between structural shocks and regime switches.

within-month supply curve. Our identifying assumption is that OPEC cannot respond within a month to demand- or price shocks. The typical argument is that shifting the production schedule is costly and producers observe oil consumption at a low frequency (Kilian 2009). An aggregate demand shock that increase global real economic activity is interpreted as a sudden change to the demand for industrial commodities whereas shocks to the real price of oil are called oil-specific demand shocks. The latter may be interpreted as shifts in precautionary demand for oil caused by uncertainty about future availability of crude oil. Our final identifying assumption is that the response of global real activity lags by at least one month to oil-specific demand shocks. As oil prices change, oil consumers are slow at adjusting their activity level.

For structural inference, we follow Karamé (2010) and compute exact impulse response functions (EIRF) defined as

$$\phi_h^j(y,\varepsilon) = \mathbb{E}\left[\frac{\partial y_h}{\partial \varepsilon_0}\middle| s_0 = j, P\right],$$

that is, the expected value of endogenous variable $y \in \boldsymbol{y}$, h periods after a structural shock $\varepsilon \in \boldsymbol{\varepsilon}$ from the stationary solution of state $s_0 = j \in \{\text{PC}, \text{CC}\}$, taking the expectation over all possible paths of the state variable (regime indicator) from s_0 to s_H. The H-horizon EIRF is obtained by taking a probability-weighted sum over all conditional responses given by the 2^H possible sequences of the state variable from s_1 to s_H.[19] Hence we emphasize that the dynamic properties of the EIRF are conditional on the transition matrix P. To isolate the regime-contingent behavior, we compute the regime-dependent impulse response functions (RDIRF)

$$\psi_h^j(y,\varepsilon) = \mathbb{E}\left[\frac{\partial y_h}{\partial \varepsilon_0}\middle| s_h = j \;\forall h \in [0,H]\right]$$

which conditions on remaining in the same regime throughout. Finally we compute the conventional impulse response function (IRF)

$$\varphi_h(y,\varepsilon) = \mathbb{E}\left[\frac{\partial y_h}{\partial \varepsilon_0}\right]$$

in the case of the fixed-regime BVAR.

[19]The number of paths increases exponentially in H. For example there are 262,144 possible sequences for $H = 18$ which must be computed for each draw from the Gibbs sampler. To ease the computational burden, we only consider sequences that have more than 10^{-6} probability of occurring. The excluded paths account for about 3–5% of the total probability mass.

6 Results

We present theoretical predictions in terms of empirical model properties and the corresponding evidence for these found in the data. We evaluate the model output against contemporary and historical narrative evidence and price expectations inferred from futures contracts.

From the discussion above, we had Prediction 1 which states that we expect regime-switching between distinct pro- and countercyclical regimes, with the difference in response across regimes increasing in incentive power. We will now postulate empirical predictions in terms of empirical model parameters.

Prediction 3. ***Sign reversal.*** *Under strategic behavior, the price semi-elasticity of supply* $\alpha_{13}(s_t) \in A_{\bullet 1}$ *in Equation (3) will have opposite signs* $\alpha_{13}(PC) < 0 < \alpha_{13}(CC)$ *for all elements of the 68% posterior credible set.*

This prediction follows from Proposition 1, stating that the first row of Equation (2) has negative sign. Notice that Prediction 3 concerns properties of reduced form parameters $A_{\bullet 1}$ and covariance matrix $\Sigma_{\boldsymbol{e}}$ which jointly govern the EIRF. To present the predicted sign-reversal in terms of the EIRF, we ease notation by setting $\phi_h^j = \phi_h^j(\Delta q, \Delta \ln p)$, i.e. the expected response in OPEC output change Δq to a shock in log price changes $\Delta \ln p$ from state j at horizon h.

Prediction 4. ***Pro- and countercyclical responses.*** *Under strategic behavior we expect distinct pro- and countercyclical responses of OPEC output to price shocks on impact* $\phi_0^{PC} \leq 0 \leq \phi_0^{CC}$.

To evaluate whether regime-switching yields distinct inference on aggregate dynamic behavior we compare the EIRF ϕ and RDIRF ψ of the regime-switching model with the non-switching IRF φ of the regime-fixed model.

Prediction 5. ***Distinct structural inference.*** *Under strategic behavior the regime-contingent responses satisfy* $\psi_h^{PC} \leq \phi_h^{PC} \leq \varphi_h \leq \phi_h^{CC} \leq \psi_h^{CC}$ *for all* h *and elements in the 68% posterior credible set.*

Recall from earlier that Prediction 2 states that persisting changes to the market environment may induce persistent deviations from average behavior. A lasting decrease in incentive power yields more frequent but less severe punishments and shorter but less disciplined reward phases. Conversely, an increase in incentive power yields longer-lasting reward phases with greater cartel output restraint and less frequent but more severe punishments. Formally we expect such structural breaks to yield episodes of behavior that is

unlikely given average relationships or indeterminate with respect to regime classification.

Let $\underline{S}_T$, $\overline{S}_T$ denote the longest unbroken sequences of pro- and countercyclical regimes respectively, observed in a sample of length T. Let $\underline{F}(X) = \Pr(\underline{S}_T \geq X|P,T)$ and $\overline{F}(X) = \Pr(\overline{S}_T \geq X|P,T)$ be the probability of observing an unbroken sequence of pro- and countercyclical regime lasting at least X periods, conditional on transition matrix P and sample length T. Finally, let CC^* and PC^* be the longest unbroken sequence of pro- and countercyclical regimes observed for a draw of the model.

Prediction 6. ***Break in dynamic behavior.*** *Under strategic competition we expect evolving market conditions to yield* $\lim_{T\to\infty} \underline{F}(CC^*) \to 0$ *and* $\lim_{T\to\infty} \overline{F}(PC^*) \to 0$.

We expect that low cartel discipline and frequent switching between less distinct regimes will lead to a more uncertain regime classification.

Prediction 7. ***Indeterminate regime classification.*** *Under strategic competition we expect evolving market conditions to yield periods with a posterior mean of the regime indicator sequence* s_t *to be* $\Pr(s_t = CC) \approxeq \frac{1}{2}$.

Finally, recall also from Section 4 that we may not exclude non-linear or even non-monotonic price-quantity relationships when incentive power is constrained and OPEC optimally reduces cartel discipline. We expect also such a local reduction in covariance between prices and quantities to be manifested in regime indeterminacy.

6.1 Prediction 1: Evidence of pro- and countercyclical output regimes

We assess evidence for distinct pro- and countercyclical regimes, see Predictions, 3, 4 and 5 above. Consider the impulse-response functions plotted in Figure 5. The solid lines denote the EIRF, dashed lines the RDIRF, and green lines denote standard IRF from a (one-regime) BVAR. Shaded areas report the 68% credible sets, i.e. the 16- and 84th percentiles of the posterior distributions. The OPEC response to a precautionary demand shock is given by the first row and third column panel and clearly demonstrates distinct pro- and countercyclicality. Moreover, the response in the countercyclical regime is more persistent, with the credible set containing 0 after approximately 15 months, compared to 7 under procyclical behavior. The fixed-regime BVAR on the other hand exhibits a passive, acyclical response. This difference in

responses between the regime-switching- and fixed model highlights how assuming a regime-independent supply schedule leads the estimated response to be biased towards zero, see the discussion in Section 4. Hence we conclude that the data support empirical Predictions 4 and 5.[20] Finally notice that the price response to a supply shock is also regime-contingent. In this case, the fixed-regime BVAR gives an IRF that has a large credible set, but is at least initially in agreement with the procyclical regime response. This may be due to the larger persistence of the output response to an OPEC supply shock. The remaining responses are qualitatively similar to those found in Kilian (2009).

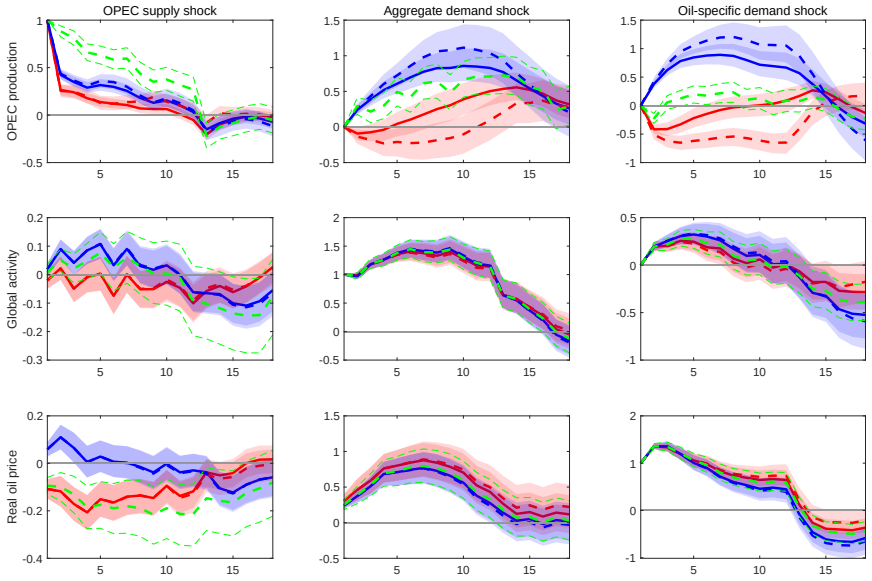

Figure 5: Exact impulse response functions (EIRF). The pro- and countercyclical regimes are plotted in red and blue, respectively. The dashed lines show regime-dependent response (RDIRF) and shaded areas 68% credible sets. The impulse-response functions (IRF) of the one-regime BVAR are reported in green together with the 68% credible sets.

[20]To evaluate Prediction 3 consider Figure 13 in Appendix C, plotting the implied prior- and posterior distributions over the first lag of the MS-BVAR coefficients. The first row reports posteriors for the pro- and counter-cyclical regime in red and blue respectively, with priors in dashed lines. Note the sign change in posterior mean and median for the coefficients of interest $\partial \Delta q_t(s_t)/\partial \Delta p_{t-1} = \alpha(s_t)_{\bullet 13} \in A(s_t)$, Equation 3, with $\alpha(\mathrm{PC})_{\bullet 13} < 0 < \alpha(\mathrm{CC})_{\bullet 13}$. The final two rows report posteriors of the non-switching block in gray and priors in green.

We now turn to the inferred path of output regimes by comparing model estimates with contemporary and historical narrative evidence. Consider Figure 7, which reports the posterior mean (blue) and median (black) of the distribution over regime indicator sequences. A mean posterior value near PC or CC indicates high confidence of pro- or counter-cyclical behavior respectively. Indeterminate values imply lower confidence. To fix ideas on how to interpret the regime indicator sequence, consider Figure 6 which partitions $\Delta \log p_t$, Δq_t space into four quadrants of which I, IV and II, III are associated with pro- and countercyclical behavior, $s_t = PC$ and $s_t = CC$, respectively. We emphasize that only the observations with market flooding, $\Delta \ln p < 0$ and $\Delta q_t > 0$, or quadrant IV in Figure 6, are to be interpreted as an output war. Observations in quadrant I represents cooperative, procyclical behavior where oil is withdrawn from the market. Hence, observations in quadrants I, II and III have an interpretation as a cooperative or reward phase actions.

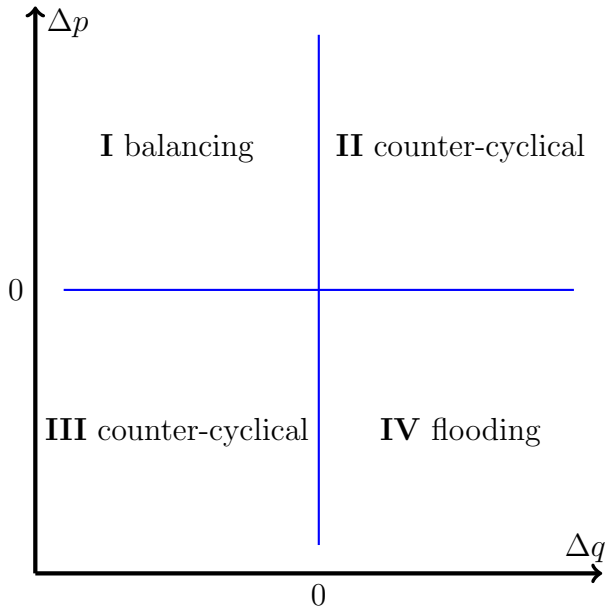

Figure 6: Implied regimes in the Δq, Δp space. Observations from pro- and countercyclical regimes are in the quadrants I, IV and II, III respectively. The punishment stage correspond to quadrant IV, the reward phase to quadrants I, II and III. The origin is marked by the intersecting blue lines.

We define output wars as periods in which we have high confidence that OPEC is flooding the market, i.e. observations t such that $\Delta \ln p_t < 0 < \Delta q_t$ (quadrant IV in Figure 6) and the posterior mean regime probability

$\Pr(s_t = \text{PC}) > .75$. They are indicated with light red shading. The darkly shaded red- and blue areas denote months in which the International Energy Agency's monthly Oil Market Report (OMR) finds that OPEC's actions are substantially raising or decreasing the price of oil, see Appendix D for details on the construction of this time series.

Notice that because the procyclical regime includes cooperative behavior, their duration serves as a lower bound on the expected length of output wars. Second, the models' classification of procyclical behavior is consistent with judgment in the OMR, excepting two OPEC meetings in 1995 which are classified as indeterminate, i.e. reflecting low cartel discipline. Furthermore, the output wars we identify in 1986 and 1998 are consistent with historical narrative accounts of OPEC's actions and motives at the time, see for instance Noreng (2006), Yetiv (2010) and Yergin (2011).

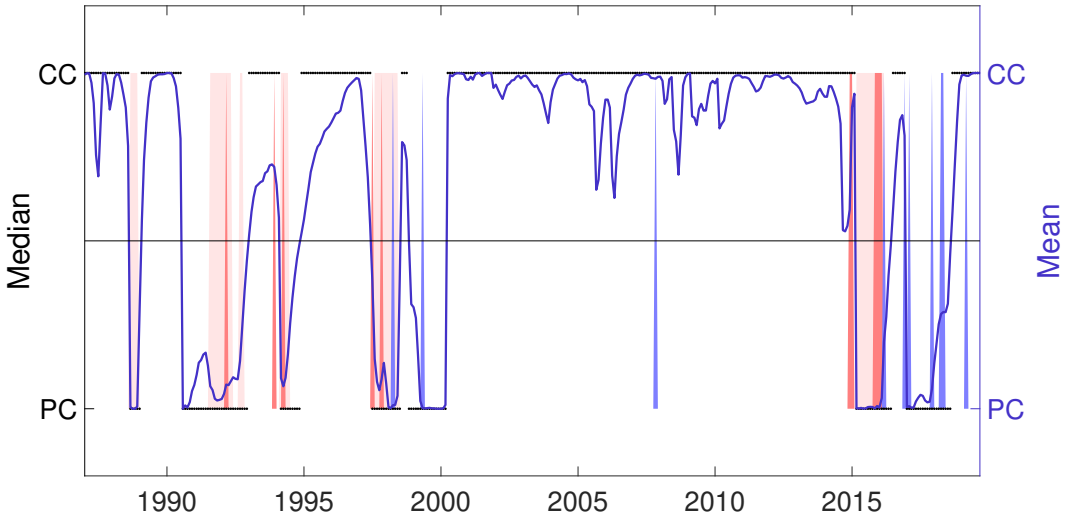

Figure 7: Posterior regime indicator against time. Pro- and counter-cyclical regimes denoted by PC and CC. Median, black points and mean, blue. The lightly shaded areas denote output wars with high confidence of market flooding behavior, $\Delta \ln p_t < 0 < \Delta q_t$ and $\Pr(s_t = \text{PC}) > .75$. The darkly shaded red and blue areas denote months in which the International Energy Agency's Oil Market Report (OMR) finds that OPEC's actions are substantially decreasing or increasing the price of oil.

6.2 Prediction 2: Deviations from average behavior

We now turn to evidence that OPEC has exhibited persistent deviations from average behavior, see Predictions 6 and 7. The estimated regime sequence in Figure 7 suggests a qualitative shift in the dynamic patterns of behavior around 1999. The rapid regime shifts and prolonged periods of indeterminacy, such as in the 1986-1998 period, indicate frequent switching between less distinct regimes or less steeply sloped supply schedules, i.e. that the observed acyclical quantities supplied are given by a convex combination of pro- and countercyclical supply (Prediction 7).

Notice that from approximately the year 2000 the countercyclical behavior appears to become more persistent. We formally consider the likelihood of observing this countercyclical episode, given average behavior. For each stored draw from the Gibbs sampler, we identify CC^*, the longest unbroken sequence of countercyclical behavior. We then store the date at which it begins. The posterior distributions over the dates and duration of the longest sequences are reported in Figure 8.

The distribution over dates is remarkably concentrated at the mode of April 1998, with a long right tail pulling the median to June 2000. In contrast, the distribution over duration is multi-modal. The evidence suggests that the longest countercyclical period began between 1998 and 1999, but the duration of this regime depends on whether the indeterminate episodes between 1998 and 2014 are classified as pro- or countercyclical. Following Prediction 6, we quantify the likelihood of observing a sequence at least as long as CC^*, i.e. $\overline{F}(\text{CC}^*)$, as follows: For each draw we simulate the probability of observing at least once a countercyclical regime of equal or greater length than the greatest observed sequence. The mean and median probabilities are approximately 30% and 22.5%, respectively. Based on this metric the observed duration of the longest countercyclical regime period is consistent with the average predictions of the model.

Yet from Figure 7 we see that the posterior distribution suggests it is unlikely that any output wars occurred between mid-1998 and late 2014, with observed procyclical episodes all firmly in the cooperative quadrant I of Figure 6. We examine whether contemporary narrative evidence of exogenous supply disruptions can explain isolated incidents of regime indeterminacy. Consider Figure 9 which plots periods in which the OMR reports substantial exogenous disruptions to OPEC output. Notice in particular the 2006-2007 period, where OPEC production is falling and prices are increasing. The OMR ascribes these price movements respectively to international sanctions on Iranian oil exports and unrest in Nigeria, coinciding with the modes at 62 and 100 months in

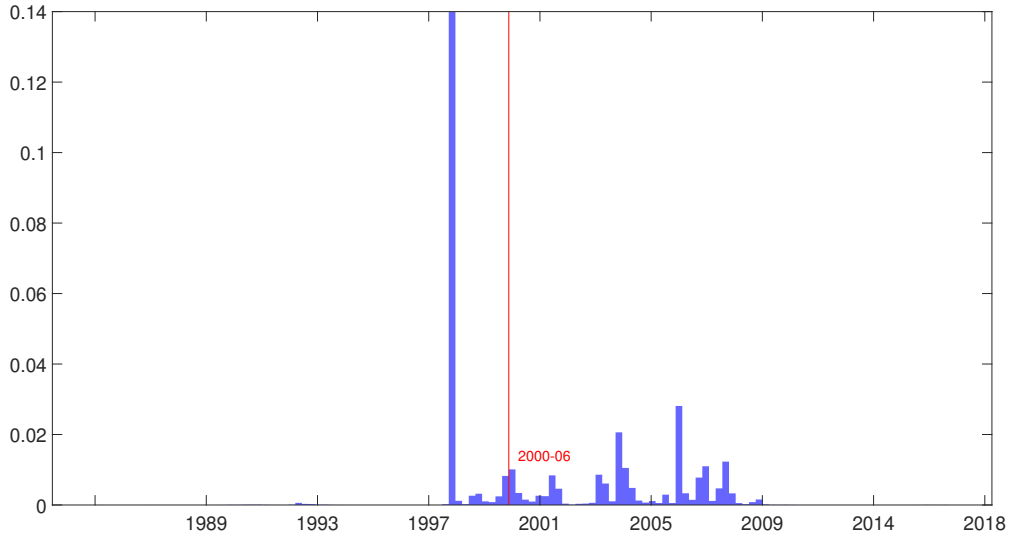

(a) Month

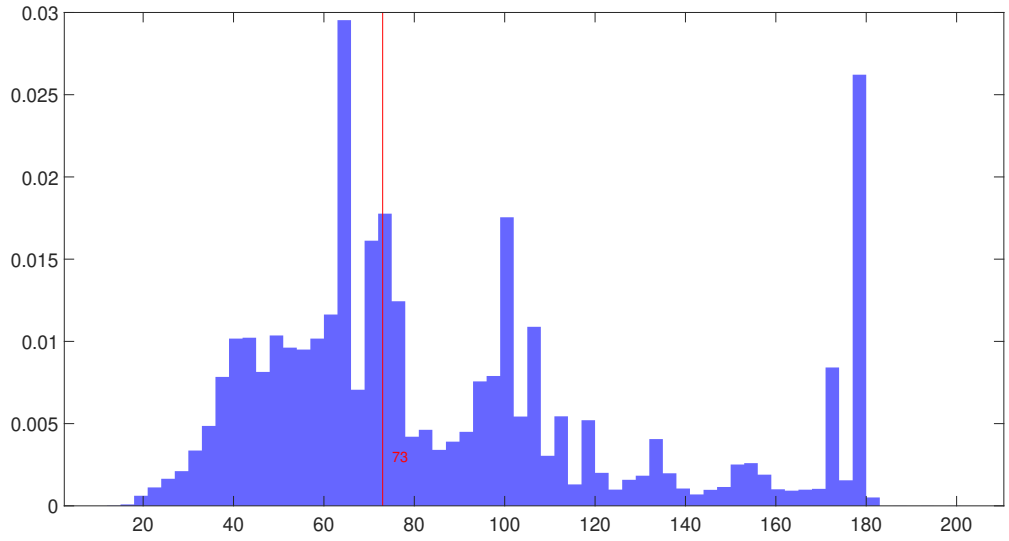

(b) Number of months

Figure 8: Histogram over the starting month and duration of the longest unbroken sequence of a countercyclical regime across Gibbs sampler draws.

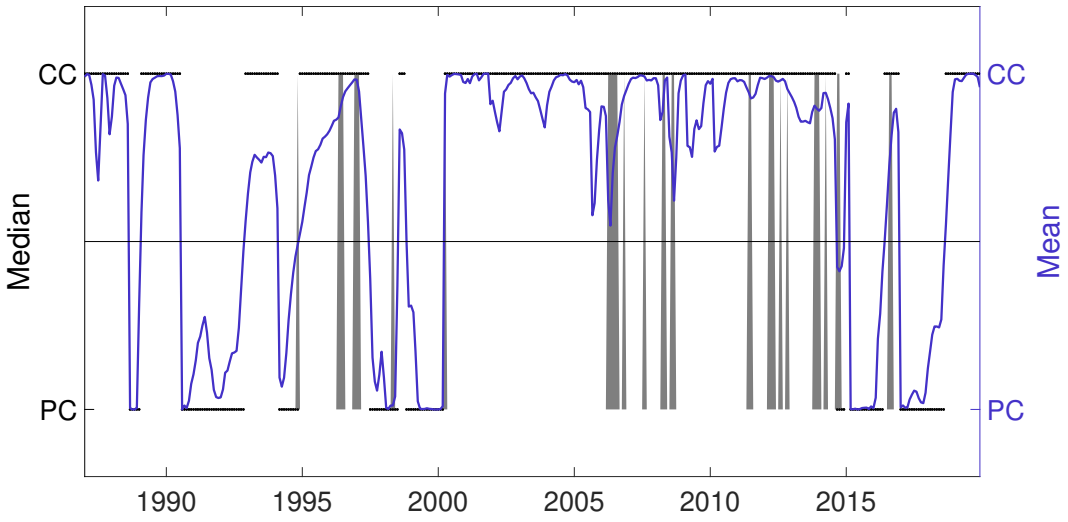

Figure 9: Posterior regime indicator against time. Pro- and counter-cyclical regimes denoted by PC and CC. The shaded areas denote months in which the International Energy Agency's Oil Market Report (OMR) finds that there are substantial exogenous disruptions to OPEC's oil production.

Figure 8.

We therefore repeat the exercise but counting only the pro-cyclical regimes that coincides with increases in OPEC output, i.e. the red-shaded areas in Figure 7. This exercise amounts to assuming that OPEC's behavior would be countercyclical in absence of these disruptions, accepting the OMR description of unintended production outages. The longest countercyclical regime period is now 200 months and the probability mean of randomly observing this realization is 1.1%, with a median of only 0.04%. Under the assumption of otherwise countercyclical behavior the data suggest that there was a change in the underlying data generating process in mid-1998. Our interpretation is that OPEC's cartel discipline was low during the 1986–1999 period but increased markedly post-1998. This finding is consistent with the results in Almoguera, Douglas, and Herrera (2011) and Ratti and Vespignani (2015) but contrasts with the claim of Baumeister and Kilian (2016) who state that OPEC permanently collapsed in 1986, never again attempting to influence market outcomes.

6.3 Corroborative evidence from futures prices

As a final exercise we consider if price expectations inferred by futures contracts are consistent with regime-switching.[21] Let $\Delta \ln p_t^F = \ln p_t^{12} - \ln p_t^0$ be the time-t difference in log WTI prices for delivery in twelve months and within month. Figure 10 plots the posterior mean of the regime indicator together with an indicator for when $\Delta \ln p_t^F$ exceeds its 90th percentile. Notice that steeply sloped futures curves coincide with periods in which either our model, the IEA or both classify procyclical behavior. This pattern of futures prices is consistent with market expectations of heavily but temporarily depressed oil prices, exactly as predicted by the theory. The only two exceptions are during October 2008 to November 2009 global financial crisis, a large transitory demand shock, and November 2006, coinciding with OMR reports of geopolitical tensions in Iran.

7 Robustness

Our results hold up to numerous robustness checks and extensions. We estimate the model with different lag orders, priors and variables as well as perform MCMC convergence diagnostics. We present a brief summary of these exercise here, relegating the details to Appendix F.

First, we re-estimate the model with reduced lag-order. Changing the lag order to e.g. 12 or 18 does not qualitatively overturn our results or conclusions. We assess prior sensitivity in two ways. First, we increase the variance of prior distributions for all parameters. Second, for $\beta_2(s_t)$ (governing the relationship between contemporaneous prices on quantities), we move the prior means closer to zero and in line with the priors of other VAR parameters. The results remain unchanged, demonstrating that the data are highly informative about the parameter values. However, the sensitivity analysis shows that the prior means of $\beta_2(s_t)$ must be sufficiently different in order to identify separate regimes.

We extend our analysis and estimate our baseline model with a total of three variable substitutions. First, we replace the OECD+6 index of indus-

[21] The price of futures contracts is widely used as a measure of oil price expectations. For example, p.7 in the Norges Bank' Monetary Policy Report 4/14 plainly states that "[projections] in this report are based on the assumption that oil prices move in line with futures prices [...]". For storeable commodities such as oil the slope of the futures curve is typically interpreted as representing the equilibrium cost of storage (Reichsfeld, Roache, et al. 2011). If the cost of storing oil increases convexly in the quantity stored, a steeply positively sloped curve implies a supply glut.

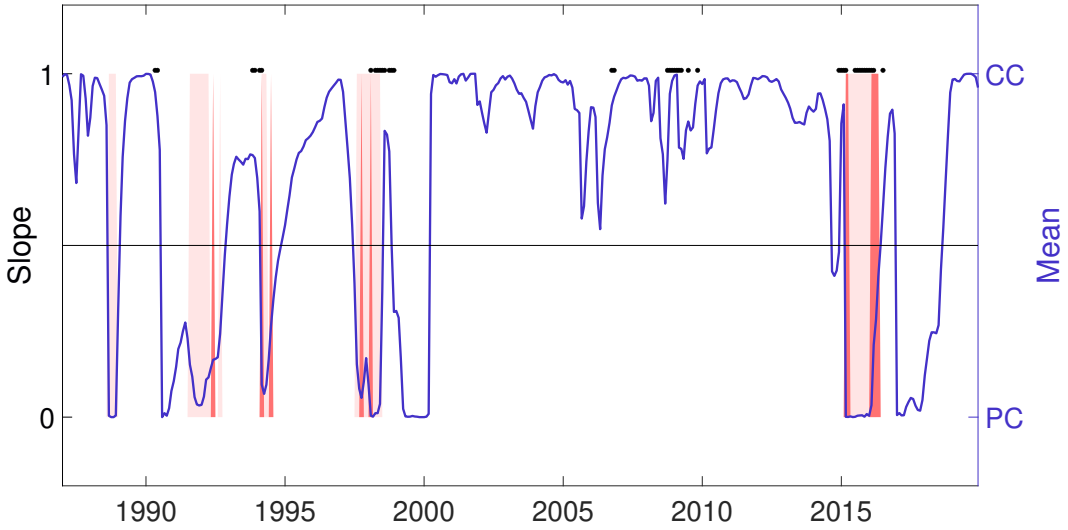

Figure 10: Posterior mean regime indicator against time in blue. Pro- and counter-cyclical regimes denoted by PC and CC. Black points indicate when the twelve-month log WTI future curve slope $\ln p_t^{12} - \ln p_t^0$ exceeds its 90th percentile. The lightly shaded areas denote output wars with high confidence of market flooding behavior, $\Delta \ln p_t < 0 < \Delta q_t$ and $\Pr(s_t = s^1) > .75$. The darkly shaded red areas bars denote months in which the International Energy Agency's Oil Market Report (OMR) finds that OPEC's actions are substantially reducing the price of oil.

trial production with the Kilian (2009) freight index to measure global real economic activity. The results are qualitatively similar to our baseline specification, albeit with less precisely estimated regimes. Second, we switch OPEC crude oil production with non-OPEC output. Our estimation procedure mechanically identifies pro- and countercyclical regimes. However, the 68% credible sets of impulse response functions of the real price of oil to non-OPEC supply shocks contain zero for all horizons and hardly differ across regimes. Moreover, regime persistence is low, with many periods classified as indeterminate and little evidence of prolonged countercyclical output. Third and similarly to the previous point, we substitute OPEC production with global aggregate output. In this case, we find qualitatively similar results as in the baseline case except that non-OPEC output shortages (e.g. coincidentally with Hurricane Katrina) alter the regime classification in the 2000s. This suggests that OPEC output variation indeed does dominate the global oil supply response.

Finally, we perform simple MCMC convergence diagnostics. We compute Geweke inefficiency factors and do difference-in-means tests for all estimated parameters. Both metrics are within conventionally accepted limits and we conclude that the MCMC sampler converged.

8 Conclusion

A long-standing literature has employed empirical macroeconomic models of the oil market to investigate the origin and propagation of shocks to the price of crude oil. Structural inference is conducted under the assumption that global oil supply is well-approximated by a stationary, linear process. In this paper, we argue that strategic competition between OPEC members has induced non-linear price quantity relationships and that this non-linearity has rendered conventional structural inference uninformative.

We derived the symmetric profit-maximizing public equilibrium in a model of oligopolistic quantity competition with evolving observable demand. Efficient, symmetric equilibria induce switching between cooperative and punishment regimes, persistent deviations from average behavior, and potentially non-linear supply correspondences. The model was used to derive two empirical predictions. First, strategic competition induces switching between distinct pro- and countercyclical regimes. Second, under evolving market conditions we expect persistent deviations from average behavior, manifested through switching frequencies and supply elasticity.

Adapting the model of Kilian (2009) to allow for regime-switching in crude

oil supply, we find robust evidence in the data supporting both hypotheses. We also find that allowing for regime-shifting substantially alters inferred behavior. Inference in traditional oil market models generally find that oil producers are passive with respect to price developments. Our analysis suggests that this acyclical behavior may be due to an unwarranted conflation of pro- and countercyclical regimes. We therefore suggest that a stationary, linear supply correspondence is not an appropriate approximation of OPEC's actual output behavior.

We conclude by suggesting two avenues for future research. First, our theoretical analysis assumes symmetry among producers. More research is needed to identify the salient dimensions of heterogeneity among OPEC producers and how they interact with optimizing behavior in a dynamic environment, e.g. the domestic political economy of OPEC countries or their costs of production. Empirically one could explore such heterogeneity through a panel of OPEC oil producers. Second, future research may study methods of structural inference under regime-switching that are computationally feasible and allow for a relaxation of the zero-restrictions employed in this paper. A natural candidate would be an extension of the Baumeister and Hamilton (2015) methodology.

References

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.

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.

Auerbach, Alan J and Yuriy Gorodnichenko (2012). "Measuring the output responses to fiscal policy". In: *American Economic Journal: Economic Policy* 4.2, pp. 1–27.

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 (2015). "Sign restrictions, structural vector autoregressions, and useful prior information". In: *Econometrica* 83.5, pp. 1963–1999.

— (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, UC San Diego*, pp. 139–160.

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.

Billio, Monica et al. (2016). "Interconnections between eurozone and US booms and busts using a Bayesian panel Markov-switching VAR model". In: *Journal of Applied Econometrics* 31.7, pp. 1352–1370.

Bjørnland, Hilde C (2019). "Supply flexibility in the shale patch: Facts, no fiction". In: *manuscript, BI Norwegian Business School*.

Bjørnland, Hilde C, Vegard H Larsen, and Junior Maih (2018). "Oil and macroeconomic (in) stability". In: *American Economic Journal: Macroeconomics* 10.4, pp. 128–51.

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

Caldara, Dario, Michele Cavallo, and Matteo Iacoviello (2019). "Oil price elasticities and oil price fluctuations". In: *Journal of Monetary Economics* 103, pp. 1–20.

Celeux, Gilles, Merrilee Hurn, and Christian P Robert (2000). "Computational and inferential difficulties with mixture posterior distributions". In: *Journal of the American Statistical Association* 95.451, pp. 957–970.

Chan, Joshua CC (2020). "Large Bayesian vector autoregressions". In: *Macroeconomic Forecasting in the Era of Big Data*. Springer, pp. 95–125.

Coll, Steve (2012). *Private empire: ExxonMobil and American power*. Penguin.

Collie, David et al. (2004). "Collusion and the elasticity of demand". In: *Economics Bulletin* 12.3, pp. 1–6.

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*. 1st ed. 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.

Fattouh, Bassam and Lavan Mahadeva (2013). "OPEC: What Difference Has It Made?" In: *Annual Review of Resoure Economics* 5.1, pp. 427–443.

Geweke, John (2007). "Interpretation and inference in mixture models: Simple MCMC works". In: *Computational Statistics & Data Analysis* 51.7, pp. 3529–3550.

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.

Griffin, James M (1985). "OPEC behavior: a test of alternative hypotheses". In: *The American Economic Review* 75.5, pp. 954–963.

Gundersen, Thomas S. (2020). "The Impact of US Supply Shocks on the Global Oil Price". In: *The Energy Journal* 41.1, pp. 151–174.

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.

Hamilton, James D (1985). "Historical causes of postwar oil shocks and recessions". In: *The Energy Journal* 6.1.

— (2016). "Macroeconomic regimes and regime shifts". In: *Handbook of Macroeconomics*. Vol. 2. Elsevier, pp. 163–201.

— (1994). *Time Series Analysis*. Princeton University Press.

— (2018). "Why you should never use the Hodrick-Prescott filter". In: *Review of Economics and Statistics* 100.5, pp. 831–843.

Hamilton, James D and Ana María Herrera (2004). "Comment: Oil shocks and aggregate macroeconomic behavior: The role of monetary policy". In: *Journal of Money, credit and Banking*, pp. 265–286.

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.

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.

Jaakkola, Niko (2019). "Carbon taxation, OPEC and the end of oil". In: *Journal of Environmental Economics and Management* 94, pp. 101–117.

Jasra, Ajay, Chris C. Holmes, and David A. Stephens (2005). "Markov Chain Monte Carlo methods and the Label Switching Problem in Bayesian Mixture Modeling". In: *Statistical Science* 20.1, pp. 50–67.

Känzig, Diego R (*Forthcoming*). "The Macroeconomic Effects of Oil Supply News: Evidence from OPEC Announcements". In: *American Economic Review*.

Karamé, Frédéric (2012). "An algorithm for generalized impulse-response functions in Markov-switching structural VAR". In: *Economics Letters* 117.1, pp. 230–234.

— (2010). "Impulse-response functions in Markov-switching structural vector autoregressions: A step further". In: *Economics Letters* 106.3, pp. 162–165.

Kilian, Lutz (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 Helmut Lütkepohl (2017). *Structural vector autoregressive analysis*. Cambridge University Press.

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 (2018). "Modeling fluctuations in the global demand for commodities". In: *Journal of International Money and Finance* 88, pp. 54–78.

Kolodzeij, Marek and Robert K Kaufmann (2014). "Oil demand shocks reconsidered: a cointegrated vector autoregression". In: *Energy Economics* 41, pp. 33–40.

Koop, Gary, M Hashem Pesaran, and Simon M Potter (1996). "Impulse response analysis in nonlinear multivariate models". In: *Journal of econometrics* 74.1, pp. 119–147.

Lo, Ming Chien and Jeremy Piger (2005). "Is the response of output to monetary policy asymmetric? Evidence from a regime-switching coefficients model". In: *Journal of Money, credit and Banking*, pp. 865–886.

Mailath, George J and Larry Samuelson (2006). *Repeated games and reputations: long-run relationships*. Oxford university press.

Mas-Colell, Andreu, Michael Dennis Whinston, Jerry R Green, et al. (1995). *Microeconomic theory*. Vol. 1. Oxford university press New York.

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.

Porter, Robert H (1983). "A study of cartel stability: the Joint Executive Committee, 1880-1886". In: *The Bell Journal of Economics*, pp. 301–314.

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*.

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.

Reichsfeld, David A, Shaun K Roache, et al. (2011). "Do commodity futures help forecast spot prices?" In: *IMF Working Paper 11/254*.

Rotemberg, Julio J and Garth Saloner (1986). "A supergame-theoretic model of price wars during booms". In: *The American Economic Review* 76.3, pp. 390–407.

Rubio-Ramirez, Juan F, Daniel F Waggoner, and Tao Zha (2010). "Structural vector autoregressions: Theory of identification and algorithms for inference". In: *The Review of Economic Studies* 77.2, pp. 665–696.

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.

Yergin, Daniel (2011). *The Prize: The epic quest for oil, money & power*. New York: Simon and Schuster.

Yetiv, Steve A (2010). *Crude awakenings: global oil security and American foreign policy*. Cornell university press.

A Appendix: Theory

A.1 Model of oligopolistic quantity competition

The stage game is described in detail. Producer $i \in \{1,2\}$ chooses an observable and unobservable quantity q_i and $\tilde{h}_i$, both in $Q := \{0, \epsilon, ..., q^*\} \subset \mathbb{R}$ an evenly ϵ-spaced grid, and such that total output respects the constraint $q_i + \tilde{h}_i \leq q^*$. We denote variables that are not publicly observed with a tilde, e.g. the total production $\tilde{q}_i := q_i + \tilde{h}_i$ of player i is not publicly observed. We will omit the tilde when it is irrelevant whether an action is observable or not, i.e. when conditioning on actions. Let $\boldsymbol{q} \in Q^2$ be the action profile and $q := q_1 + q_2 \in Q + Q = \{0, \epsilon, \ldots, 2q^*\}$ the total output.[22] To easily relate action profiles to total production we introduce the vector $\boldsymbol{\iota} = (1,1)$ so $\boldsymbol{\iota}'\boldsymbol{q} = q$, used interchangeably. The inverse demand function is $p(q,x)$, where $x \in \mathbb{R}_{\geq 0}$ is interpreted as the market demand for the dominant producers' product net of production from a non-strategic competitive fringe. The price $p(q,x) = xp(q)$ is continuous in all arguments, strictly decreasing and weakly convex in output $\partial p(q)/\partial q < 0$, $\partial^2 p/\partial^2 q \geq 0$ and strictly positive $p(q,x) > 0$ for all q. The players observe

$$p(q, x, \tilde{\theta}) = \tilde{\theta} p(q, x) \tag{7}$$

where the multiplicative term $\tilde{\theta}$ is an unobserved log-normally distributed random variable $\ln \tilde{\theta} \sim \mathrm{N}\left(-\sigma_\theta^2/2, \sigma_\theta^2\right)$ with independent realizations over t. The observed prices are then conditionally log-normally distributed

$$\ln p(x,q) \sim \mathrm{N}(\ln p(q,x) - {}^{\sigma_\theta^2}/_2, \sigma_\theta^2) \tag{8}$$

with distribution $F_p(\cdot|q,x)$ and density $f_p(\cdot|q,x)$. The parameter σ_θ governs a mean-preserving spread of the distribution and has a natural interpretation as monitoring quality. Notice that because prices are strictly positive for all levels of production the producers may never infer their opponents action with certainty from the price realization. Let $c : \mathbb{R} \to \mathbb{R}$ be a weakly convex and strictly increasing cost function, $\partial c(q)/\partial q > 0$, $\partial^2 c(q)/\partial^2 q^2 > 0$. Conditional on actions, the realized- and ex-ante expected profits in each stage game are

$$\begin{aligned} \pi(\boldsymbol{q}, x, \tilde{\theta})_i &= \tilde{\theta} x p(\boldsymbol{\iota}'\boldsymbol{q}) q_i - c(q_i) \\ \implies \mathrm{E}_{\tilde{\theta}} \pi(\boldsymbol{q}, x, \tilde{\theta})_i &= x p(\boldsymbol{\iota}'\boldsymbol{q}) q_i - c(q_i) =: \pi(\boldsymbol{q}, x)_i \end{aligned} \tag{9}$$

[22] We will denote matrices, vectors, and their scalar elements respectively by uppercase, lowercase bold- and standard symbols, e.g. a matrix Z, vector $\boldsymbol{z}$, and scalar elements $z \in Z$, $z \in \boldsymbol{z}$.

for $i \in \{1, 2\}$.

We make two assumptions that are satisfied, among others, for standard elastic constant elasticity of demand functions or an appropriately scaled linear function. First, suppressing player subscripts, we assume that the expected profits are strictly concave in own output

$$\frac{\partial^2 \pi(\boldsymbol{q}), x}{\partial^2 q} < 0$$

for all q so that a best-response always exists. In the limiting case of constant marginal costs this condition holds for all demand functions such that the first order effect dominates

$$q\frac{\partial^2 p(\boldsymbol{\iota}'\boldsymbol{q})}{\partial^2 q^2} < -2\frac{\partial p(\boldsymbol{\iota}'\boldsymbol{q})}{\partial q}$$

a standard property. Second, we assume that (inverse) demand is everywhere (inelastic) elastic and the elasticity (increases) decreases in output q.

Assumption 1. ***Elastic demand.*** *Let the inverse price elasticity of demand be inelastic and increasing in quantities*

$$-\frac{\partial p(q, x)}{\partial q}\frac{q}{p(q, x)} = -\frac{\partial p(q)}{\partial q}\frac{q}{p(q)} \leq 1$$

$$\frac{\partial}{\partial q}\frac{\partial p(q)}{\partial q}\frac{q}{p(q)} \geq 0$$

for all q.

The natural interpretation of Assumption 1 is that the overall supply, including the competitive fringe, is sufficient to push consumer demand away from the inelastic region. Notice that while there is no direct effect of demand x on the inverse price elasticity, there is an indirect effect through an increase in equilibrium output. For example, the (inverse) elasticity will (increase) decrease in demand x when evaluated at $q^m(x)$.

Define $\boldsymbol{q}^m(x) := \text{argmax}_{\boldsymbol{q} \in Q^2} \pi_1(x, q_1, q_2) + \pi_2(x, q_1, q_2)$ the monopoly or jointly profit-maximizing quantity, symmetric when $\epsilon \to 0$ by the assumption of concave profits. Next, denote the conditional best-response function as $q^{\text{br}}(q', x) := \text{argmax}_{q \in Q} \pi(q, q', x)$, with q' denoting the opponents' action. Let the symmetric, pure strategy, stage-game Nash equilibrium such that $\boldsymbol{q}^n(x) := \boldsymbol{q} : q^{\text{br}}(q', x) = q'$ and $q_1 = q_2$ in $\boldsymbol{q}$. The existence of $\boldsymbol{q}^n(x)$ is not guaranteed for every $x \in \mathbb{R}_+$ because actions are discrete, $\epsilon > 0$. In the following we restrict attention to sufficiently fine action sets such that jointly profit-maximizing

quantities are symmetric $Q^2 : q = q' \, \forall \, q, q' \in \boldsymbol{q}^m(x)$ and demand levels x such that pure strategy symmetric stage-game Nash equilibria exist $X := \{x : \boldsymbol{q}^n(x) \in Q^2\}$. The set X exists because the best-response under symmetric actions in the continuous action limit are increasing in demand.[23]

A.2 Optimal equilibrium

We derive the strongly symmetric, perfect public equilibrium (PPE, Mailath and Samuelson 2006) in our oligopoly model when augmented by a Markov process on the residual demand level. We show that the results from Abreu, Pearce, and Stacchetti (1990) apply straightforwardly so there exists a unique, efficient strongly symmetric PPE. Each period the dominant producers observe the observable quantities q_1 and q_2, the demand level x, and the realized price $\tilde{p}$. Recall that the producers do not observe or the signal noise $\tilde{\theta}$ or the hidden quantities $\tilde{h}_1$ and $\tilde{h}_2$. The signal space of the stage game is therefore $S := Q_1 \times Q_2 \times \mathbb{R}_+$. Let a history $h_t = \{x_s, p_s, \boldsymbol{q}_s\}_{s=1}^{t-1}$ be the information set at the outset of stage t with $h_1 = \emptyset$ and $H_t = S^t$ the space of period-t histories. A perfect public strategy is a map from the signal history and current demand state to actions $\sigma : H \times X \to Q$. A profile of strategies is strongly symmetric if $\sigma_1(h_t) = \sigma_2(h_t)$ for all $h_t \in H$. Let Σ the set of all strongly symmetric public strategies, non-empty by the existence of a symmetric stage-game Nash equilibrium $\boldsymbol{q}^n(x)$ in pure strategies. A profile $\sigma \in \Sigma$ which for every $h_t \in H^t$ and t are a Nash equilibrium of the repeated game is a strongly symmetric PPE. Let $v(\sigma)$ the pay-off induced by $\sigma \in \Sigma$ and $V := \{v(\sigma), \sigma \in \Sigma\} \subset \mathbb{R}^d$. Note that V is non-empty by the non-emptiness of Σ and bounded, above by repeated play of $\boldsymbol{q}^m(y^j)$ and below, through individual rationality, by a pay-off of 0. Let $\boldsymbol{q}^j(\sigma)$ the equilibrium action profile and $E_{\Omega|\boldsymbol{q}} v^j$ the continuation

[23] Let $\epsilon \to 0$ derive the first-order condition for the best-response

$$\frac{\partial \pi(x, \boldsymbol{q})}{\partial q} | q' = 0$$

and apply the implicit function theorem, and solve for the best-response under symmetric actions to yield

$$\frac{\partial q^{\mathrm{br}}}{\partial x} = - \left(\frac{\partial^2 \pi(\boldsymbol{q}, x)}{\partial^2 q^2} \right)^{-1} \left(\frac{\partial^2 \pi(\boldsymbol{q}, x)}{\partial q \partial x} + \frac{\partial^2 \pi(\boldsymbol{q}, x)}{\partial q \partial q'} \right).$$

By symmetry $\partial q / \partial x = \partial q' / \partial x$. But if $\partial q' / \partial x < 0$, the right-hand-side is strictly positive and $\partial q / \partial x > 0$, a contradiction. Thus $q^{\mathrm{br}}(q, x)$ crosses q once in x, $q^{br}(q, x)$ space. But then there exists $x : \lim_{\epsilon \to 0} q^{\mathrm{br}}(q) = q \in \boldsymbol{q}$ for every a symmetric action profile $\boldsymbol{q} \in Q$ and $\epsilon > 0$.

payoffs conditional on state j, with expectations taken over the signal space when conditioning on actions $\boldsymbol{q}$. The discounted average state-j payoff may be generically decomposed into current and future payoffs

$$v(\sigma, x^j) = (1-\delta)\pi(x^j, \boldsymbol{q}^j(\sigma)) + \delta \sum_{m_{js} \in \boldsymbol{m}_j} m_{js} E_{\Omega|\boldsymbol{q}} v(\sigma, x^s)$$

conditional on a history. Suppose for the moment that V is compact, and consider the following property:

Definition 1. ***Bang-bang property.*** *A strategy profile σ such that after any history $h_t \in H^t$ and t the continuation values are $v(\sigma) \in \text{ext}V$ is said to be bang-bang.*

Abreu, Pearce, and Stacchetti (1990) demonstrate that any efficient strongly symmetric equilibrium with $d = 1$ (absorbing demand state) are necessarily bang-bang in the sense of Definition 1, see in particular Theorem 7 and pp. 1056-1057. We now show that this result generalizes to environments where payoffs follow a stationary and irreducible higher-dimensional Markov process. We derive a contraction mapping B of which the equilibrium payoffs are a fixed point, numerically computed by recursively applying B to a superset of the equilibrium payoffs $W \supset V$,

$$V = \lim_{n \to \infty} B^n(W) = \underbrace{B(B(\dots B(B(W))\dots))}_{n}.$$

The algorithm for computing equilibrium values is detailed in Appendix G. Before stating the result, it is worth considering the intuition for why the bang-bang result of Abreu, Pearce, and Stacchetti (1990) generalizes through the notion of a lottery (6.B.1, Mas-Colell, Whinston, Green, et al. 1995). In optimal equilibria of the game with static demand the players face a lottery over extreme payoffs, with transition probabilities parameterized by the action profile. In the game with Markov demand producers face a compound lottery over extreme state-contingent payoffs while retaining exactly the recursive structure of the game with static demand. But the risk neutral players are indifferent between simple and compound lotteries with the same expected value. Hence the equilibria with Markov demand may be "collapsed" statewise by convex combinations into an incentive-equivalent game with one state and the APS algorithm applied. Because the convex combination is linear, the extreme one-dimensional continuation value corresponds uniquely to the convex combination of extrema across states. Hence the bang-bang property applies in every state.

To compactly state our result we present the generic strategy, introduce the necessary notation, and state the result. Symmetric strategies prescribe a reward and punishing action for each state $j \in (1, \ldots, d)$ for a total of $2d$ states. In state j of the regular phase play $\overline{\boldsymbol{q}}^j \in Q^2$ in publicly observable quantities. If an opponent observably deviates, $q \notin \overline{Q}^j$, or the signal realization $p_t \in \overline{P}^j \subset \mathbb{R}$, switch to the punishment phase in the next period and play $\underline{\boldsymbol{q}}_j \in Q^2$. The punishment phase continues if $q \notin \underline{Q}^j$ or $p_t \in \underline{P}_j \subset \mathbb{R}$, else switch to the reward phase in the next period. Begin in the regular phase at $t = 1$. Let $\tau : H \times X \times Q^2 \to [0, 1]$ the transition probability implied by the trigger regions $\overline{P}_j, \underline{P}_j$, action profiles $\underline{\boldsymbol{q}}^j$, $\overline{\boldsymbol{q}}^j$ and demand level $x \in \boldsymbol{x}$. Equilibrium path actions are thus determined by the collection $\varsigma = \{\overline{\boldsymbol{q}}^j, \underline{\boldsymbol{q}}^j, \overline{\tau}^j, \underline{\tau}^j\}_{j=1}^d$. Let $W \supset V$ a superset of the equilibrium payoff set with $\max W = \overline{\boldsymbol{w}}$, $\min W = \underline{\boldsymbol{w}}$. The reward- and punishment phase payoffs are thus given by

$$\begin{aligned} \overline{v}^j(\boldsymbol{q}, \overline{\tau}^j, W) &:= (1-\delta)\pi(x^j, \boldsymbol{q}) + \delta(\overline{\tau}^j(\boldsymbol{q})\boldsymbol{m}_j'\overline{\boldsymbol{w}} + (1-\overline{\tau}^j(\boldsymbol{q}))\boldsymbol{m}_j'\underline{\boldsymbol{w}}) \\ \underline{v}^j(\boldsymbol{q}, \underline{\tau}^j, W) &:= (1-\delta)\pi(x^j, \boldsymbol{q}) + \delta(\underline{\tau}^j(\boldsymbol{q})\boldsymbol{m}_j'\overline{\boldsymbol{w}} + (1-\underline{\tau}^j(\boldsymbol{q}))\boldsymbol{m}_j'\underline{\boldsymbol{w}}) \end{aligned} \tag{10}$$

given a bang-bang strategy with exogenously fixed extreme continuation values.

Proposition 2. ***Optimal equilibrium.*** *There exists a unique, efficient symmetric public perfect equilibria in bang-bang strategies. The equilibrium may be computed as the fixed-point of a monotone set-valued contraction mapping*

$$B(W) = \{B_j(W)\}_{j \in D} : \mathbb{R}^d \to \mathbb{R}^d \tag{11}$$

where

$$B_j(W) = \text{co} \left(\begin{array}{c} \max \\ \boldsymbol{q} \in Q^2, \overline{p}^j \in R_+ \end{array} \overline{v}^j(\boldsymbol{q}, \overline{\tau}^j(\boldsymbol{\iota}'\boldsymbol{q}, \overline{p}^j), W) \right)$$
$$\left(\begin{array}{c} \min \\ \boldsymbol{q} \in X^2, \underline{\tau}^j \in [0,1] \end{array} \underline{v}^j(\boldsymbol{q}, \underline{\tau}^j(\boldsymbol{q}), W) \right)$$

such that

$$\overline{v}^j(\boldsymbol{q}, \overline{\tau}^j(\boldsymbol{\iota}'\boldsymbol{q}, \overline{p}^j), W) \geq \overline{v}^j(\boldsymbol{q}^\dagger, \overline{\tau}^j(\boldsymbol{\iota}'\boldsymbol{q}', \overline{p}^j), W) \; \forall \; \boldsymbol{q}^\dagger \in Q^2$$

and

$$\underline{v}^j(\boldsymbol{q}, \underline{\tau}^j(\boldsymbol{q}), W) \geq \underline{v}^j(\boldsymbol{q}^\dagger, \underline{\tau}^j(\boldsymbol{q}^\dagger), W) \; \forall \; \boldsymbol{q}^\dagger \in Q^2$$

where the value functions are given by (10), $\boldsymbol{q}^\dagger$ *are action profiles where the player's output is free and the opponent plays equilibrium actions,* $\underline{\tau}^j(\boldsymbol{q}) = (1 - \mathbb{1}(\boldsymbol{\iota}'\boldsymbol{q} \neq \boldsymbol{\iota}'\underline{\boldsymbol{q}}^j))\underline{\tau}^j$ *and*

$$\overline{\tau}^j(\boldsymbol{\iota}'\boldsymbol{q}, \overline{p}^j) = \left(1 - F_\theta\left(\frac{\overline{p}^j}{p(\boldsymbol{\iota}'\boldsymbol{q}, x^j)}\right)\right)(1 - \mathbb{1}(\boldsymbol{\iota}'\boldsymbol{q} \neq \boldsymbol{\iota}'\overline{\boldsymbol{q}}^j)).$$

Proof. The proof proceeds in two steps. We first show that the results in Abreu, Pearce, and Stacchetti (1990), hereafter APS, apply such that a unique, optimal equilibrium exists. We then derive the operator B in terms of primitive variables.

Existence of optimal equilibrium

We follow the notation of APS exactly and proceed in three steps. We first verify that restrictions on the stage game in APS hold. Second we generalize definitions of the primitive objects employed by APS to the model with Markov demand. Third, we show that the resulting equilibrium can be decomposed into d equilibria isomorphic with APS and the equilibrium may be computed by applying the APS algorithm independently to each state and updating the value set.

The five assumptions on the stage game structure in APS are satisfied for the model presented 3. Action spaces are finite (A1). The signal is continuously distributed with support independent of actions (A2, A3). Stage-game pay-offs are continuous in the signal (A4) and a pure-strategy Nash equilibrium exists in the stage game (A5).

The strongly symmetric equilibrium payoffs $V \subset \mathbb{R}^d$, non-empty by the existence of pure strategy stage-game Nash equilibria, follow a d-state Markov process with right-stochastic transition matrix M, stationary and irreducible. Let $\boldsymbol{m}_j$ the jth row of M a standard d-simplex. The continuation payoffs in state j are in $V_j := \{\boldsymbol{m}_j' v : v \in V\} \subset \mathbb{R}$. Recall that APS define $L(\Omega, \mathbb{R}^N)$ the set of all equivalence classes of essentially bounded Lebesgue measurable functions u from signal space Ω to the N-dimensional reals. I now define the analogous set for strongly symmetric equilibria in a Markov setting. By strong symmetry each state has a one-dimensional payoff set so $N = 1$ for states $j \in D$. Now consider $W \supset V \in \mathbb{R}^d$ containing the equilibrium payoffs. Let $w_j \subset \mathbb{R}$ the jth dimension of W and $W_j := \{\boldsymbol{m}_j' w : w \in W\} \subset \mathbb{R}$. Define the set

$$L_j(\Omega, W) := \left\{ \sum_{m_{jk} \in \boldsymbol{m}_j} m_{jk} u_k : u_k \in L(\Omega, w_k) \right\}$$

containing all convex combinations of Lebesgue-measurable functions $u_k \in L(\Omega, w_k)$. Let $L_M(\Omega, W) := \{L_j(\Omega, W)\}_{j=1}^d$ and

$$v(\boldsymbol{q}, x^j | u(\cdot|\boldsymbol{q})) = (1-\delta)\pi(x^j, \boldsymbol{q}) + \delta \sum_{m_{jk} \in \boldsymbol{M}_j} m_{jk} \int_\Omega u_k(z|\boldsymbol{q}) dz$$

the payoff induced by an action profile $\boldsymbol{q}$ and $u \in L_j(\Omega, W)$. Let $\boldsymbol{q}' \in Q^2$ denote an action profile where the players actions $q \in Q$ varies freely and the opponents' action $\boldsymbol{q}'$ fixed. Then $\boldsymbol{q}$ is admissible with respect to u if

$$v(\boldsymbol{q}, y^j | u_d(\cdot|\boldsymbol{q})) \geq v(\boldsymbol{q}', x^j | u(\cdot|\boldsymbol{q}'))$$

for all $\boldsymbol{q}' \in Q^2$. Let $B_j(W)$ be the set of all $v(\boldsymbol{q}, x^j | u_j(\cdot|\boldsymbol{q}))$ such that $\boldsymbol{q} \in Q^2$ is admissible with respect to $u_j \in L_j(\Omega, W)$. APS show that there exists a unique $u^* \in L(\Omega, \text{ext}W_j)$ that renders $\text{ext}B(W_j)$ admissible with respect to W_j. The convex combination is strictly increasing so extrema of the range $\text{ext}W_j$ are uniquely given as a function of extreme points in the domain $\{\text{ext}w_j\}_{j=1}^d$. Hence there exist a unique set of functions $\{u_k^* \in L(\Omega, w_k)\}_{k=1}^d$ in $L_j(\Omega, W)$ such that $\sum_{m_{jk} \in \boldsymbol{m}_j} m_{jk} u_k^*(\omega) = u^*(\omega)$ for all $\omega \in \Omega$ and $\text{ext}B(W_j) \in B_j(W)$.

Suppose there existed a $u \in L(\Omega, W) \setminus L(\Omega, \text{ext})$ non-extreme that supported payoffs $\underline{w} \leq \min B(W_j)$ and $\overline{w} \geq \max B(W_j)$. But the convex combination of Lebesgue-measurable functions $u \in L(\Omega, W)$ is itself a measurable function from Ω to $\boldsymbol{m}_j' w$. For every $u_j \in L_j(\Omega, W)$ there exists a $u \in L(\Omega, W_j)$ such that $u(\omega) = u_j(\omega)$ for all $\omega \in \Omega$. This is a contradiction. Thus $\text{ext}B_j(W) = \text{ext}B(W_j)$ rendered admissible by a unique $u^* \in L(\Omega, W_j)$ and $u_j^* \in L_j(\Omega, W)$ such that $u^*(\omega) = u_j^*(\omega)$ for all $\omega \in \Omega$. The notion of self-generation generalizes directly to $B_M(W) := \{B_j(W)\}_{j=1}^d$ as $B_1(W), \ldots, B_d(W)$ are evaluated independently.

Deriving the optimal equilibrium

The generic, symmetric bang-bang strategy prescribes a reward and punishing action for each state $j \in (1, \ldots, d)$ for a total of $2d$ states. We now state the players' strategy: In state j of the regular phase play $\overline{\boldsymbol{q}}^j \in Q^2$ in publicly observable quantities. If an opponent observably deviates, $q \notin \overline{Q}^j$, or the signal realization $p_t \in \overline{P}^j \subset \mathbb{R}$, switch to the punishment phase in the next period and play $\underline{\boldsymbol{q}}_j \in Q^2$. The punishment phase continues if $q \notin \underline{Q}^j$ or $p_t \in \underline{P}_j \subset \mathbb{R}$, else switch to the reward phase in the next period. Begin in the regular phase at $t = 1$. Let $\tau : H \times X \times Q^2 \to [0, 1]$ the transition probability implied by the trigger regions $\overline{P}_j, \underline{P}_j$, action profiles $\underline{\boldsymbol{q}}^j$, $\overline{\boldsymbol{q}}^j$ and demand level $x \in \boldsymbol{x}$. Equilibrium path actions are thus determined by the collection $\sigma_{bb} = \{\overline{\boldsymbol{q}}^j, \underline{\boldsymbol{q}}^j, \overline{\tau}^j, \underline{\tau}^j\}_{j=1}^d$. Average pay-offs under this strategy satisfy the following stationary system

$$\begin{aligned} \overline{\boldsymbol{v}} &= (1-\delta)\overline{\boldsymbol{\pi}} + \delta(\overline{\boldsymbol{t}} \circ M\overline{\boldsymbol{v}} + (\boldsymbol{\iota} - \overline{\boldsymbol{t}}) \circ M\underline{\boldsymbol{v}}) \\ \underline{\boldsymbol{v}} &= (1-\delta)\underline{\boldsymbol{\pi}} + \delta(\underline{\boldsymbol{t}} \circ M\overline{\boldsymbol{v}} + (\boldsymbol{\iota} - \underline{\boldsymbol{t}}) \circ M\underline{\boldsymbol{v}}) \end{aligned} \tag{12}$$

where $\circ$ denotes element-wise multiplication and $\boldsymbol{v}$, $\boldsymbol{\pi}$, $\boldsymbol{t}$ are d-dimensional vectors respectively stacking state-wise present-valued payoffs, stage-game profits, and transition probabilities for the regular- and punishing phases. If we define the first $1, \ldots, d$ of the $2d$ total states as regular and the last $d+1, \ldots, 2d$ as punishing the system is governed in equilibrium by the $2d \times 2d$ transition matrix

$$T := \begin{pmatrix} \overline{T} \circ M & (\iota - \overline{T}) \circ M \\ \underline{T} \circ M & (\iota - \underline{T}) \circ M \end{pmatrix} \tag{13}$$

where $\overline{T} := (\overline{\boldsymbol{t}}, \overline{\boldsymbol{t}})$, $\underline{T} = (\underline{\boldsymbol{t}}, \underline{\boldsymbol{t}})$ and $\bar{\iota} := (\boldsymbol{\iota}, \boldsymbol{\iota})$. An example of (13) for $d = 2$ is illustrated in Figure 11, where t_{ij} is the element in row i and column j of T. Recall that PPE are sequential equilibria (9.C.4, Mas-Colell, Whinston, Green, et al. 1995) so the continuation values equal equilibrium payoffs under any deviation. Let $\overline{\boldsymbol{q}}'$ and $\underline{\boldsymbol{q}}'$ denote action profiles where the opponent plays

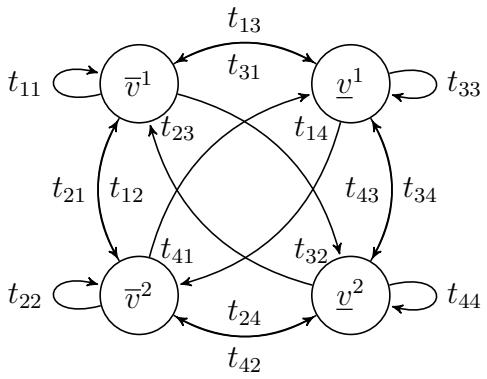

Figure 11: Equilibrium states and transition probabilities for T in Equation 13 for the case with two demand states, $d = 2$.

equilibrium actions and the player's action is free. For generic state-j equilibria with action profile $\boldsymbol{q}$ let the payoff from playing $\boldsymbol{q}'$ be

$$v^j(q, q^j) = (1-\delta)\pi(\boldsymbol{q}', x^j) + \delta \sum_{m_{js} \in \boldsymbol{m}_j} m_{js} \left(\tau(\boldsymbol{q}')\overline{v}^s + (1-\tau(\boldsymbol{q}')\underline{v}^s\right)$$

so the equilibrium condition for state j may be generically expressed as $v^j(\boldsymbol{q}^j) \geq v^j(\boldsymbol{q}')$ for all $j \in D$ and $\tilde{q}_i \in Q$. To ease notation let $\Delta\overline{\boldsymbol{\pi}} := \boldsymbol{\pi}(\overline{\boldsymbol{q}}') - \boldsymbol{\pi}(\overline{\boldsymbol{q}})$, $\Delta\underline{\boldsymbol{\pi}} := \boldsymbol{\pi}(\overline{\boldsymbol{q}}') - \boldsymbol{\pi}(\overline{\boldsymbol{q}})$ and $\Delta\bar{\boldsymbol{t}} := \bar{\boldsymbol{t}}(\overline{\boldsymbol{q}}) - \bar{\boldsymbol{t}}(\overline{\boldsymbol{q}}')$, $\Delta\underline{\boldsymbol{t}} := \underline{\boldsymbol{t}}(\underline{\boldsymbol{q}}) - \underline{\boldsymbol{t}}(\underline{\boldsymbol{q}}')$ vectors of relative payoffs- and transition probabilities under a deviation in the reward- and punishment phase, respectively. Let $\Delta\boldsymbol{v} := \overline{\boldsymbol{v}} - \underline{\boldsymbol{v}} \geq 0$ the value function differential. The equilibrium conditions may thus be compactly denoted

$$\begin{aligned} (1-\delta)\Delta\overline{\boldsymbol{\pi}} &\leq \delta\Delta\bar{\boldsymbol{t}} \circ M\Delta\boldsymbol{v} \\ (1-\delta)\Delta\underline{\boldsymbol{\pi}} &\leq \delta\Delta\underline{\boldsymbol{t}} \circ M\Delta\boldsymbol{v} \end{aligned} \tag{14}$$

for all $\boldsymbol{q}' \in Q^2$. Equation (14) 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 M. Thus, higher-powered incentives in any state enforce lower transition probabilities and greater one-shot deviation pay-offs in all states. Notice that by Proposition 2 the unique optimal pure strategy equilibrium maximizes the value of $\Delta\boldsymbol{v}$. It is useful to solve for $\Delta\boldsymbol{v}$ in terms of primitives. Let $\Delta\boldsymbol{\pi} := \overline{\boldsymbol{\pi}} - \underline{\boldsymbol{\pi}}$ and $\Delta\boldsymbol{t} := \bar{\boldsymbol{t}} - \underline{\boldsymbol{t}}$ the relative payoff and duration of the reward

phase. Incentive power may then be expressed as

$$\Delta \boldsymbol{v} = (1-\delta)\Delta\boldsymbol{\pi} + \delta(\Delta \boldsymbol{t} \circ M\Delta \boldsymbol{v}).$$

Let $M_{\Delta\tau} = (\Delta\tau_i m_{ij})_{i,j\in D}$ a $d \times d$ matrix such that $M_{\Delta\tau}\Delta\boldsymbol{v} = \Delta\boldsymbol{t} \circ M\Delta\boldsymbol{v}$, so

$$\Delta \boldsymbol{v} = (I - \delta M_{\Delta\tau})^{-1}(1-\delta)\Delta\boldsymbol{\pi} \tag{15}$$

expresses the value function differential in terms of primitives.[24] Notice that the marginal effect of relative profits $\Delta\boldsymbol{\pi}$ and duration $\Delta\boldsymbol{t}$ on incentive power are complementary and that there are decreasing returns to punishment duration on $\Delta\boldsymbol{v}$ due to discounting.

We now characterize the transition functions. We begin by arguing that the equilibrium action profiles satisfy

$$0 \leq \boldsymbol{\iota}'\boldsymbol{q}^m(x^j) \leq \boldsymbol{\iota}'\overline{\boldsymbol{q}}^j \leq \boldsymbol{\iota}'\boldsymbol{q}^n(x^j) \leq \boldsymbol{\iota}'\underline{\boldsymbol{q}}^j \tag{16}$$

as all reward-phase action profiles $\boldsymbol{\iota}'\boldsymbol{q}^m(x^j) > \boldsymbol{\iota}'\overline{\boldsymbol{q}}$ and $\boldsymbol{\iota}'\overline{\boldsymbol{q}} > \boldsymbol{\iota}'\boldsymbol{q}^n(x^j)$ violate individual rationality and punishing actions $\boldsymbol{\iota}'\underline{\boldsymbol{q}} < \boldsymbol{\iota}'\boldsymbol{q}^n(x^j)$ trivially do not satisfy the bang-bang criterion. It follows that there is no myopic incentive to decrease (increase) output in the reward (punishing) phase and the incentive constraints (14) hold trivially for such deviations. In deriving transition probabilities, attention may be restricted to deviations which increase (decrease) output in the reward (punishing) phase. Recall that equilibrium output must be in observable quantities. Fixing continuation values, We argue that the most severe punishment is achieved by

$$\underline{\tau}^j(\boldsymbol{q}) = (1 - \mathbb{1}(\boldsymbol{\iota}'\boldsymbol{q} \neq \boldsymbol{\iota}'\underline{\boldsymbol{q}}^j))\underline{\tau}^j \tag{17}$$

where $\mathbb{1}(\cdot)$ the indicator function and $\underline{\tau}^j \in [0,1]$ governs the stochastic length of the punishment. Note that (17) demands that the action profiles be in observable quantities only, so any downward deviation is immediately detected and there is no information asymmetry. In equilibrium $\underline{\tau}^j$ must satisfy

$$\pi(\underline{\boldsymbol{q}}', x^j) - \pi(\underline{\boldsymbol{q}}^j, x^j) = \frac{\delta}{1-\delta}(\underline{\tau}^j - 0) \sum_{m_{js} \in \boldsymbol{M}_j} m_{js}(\overline{v}^s - \underline{v}^s) \tag{18}$$

[24] To verify that $(I - \delta M_{\Delta\tau})^{-1}$ is well-defined let $\lambda_1, \ldots, \lambda_d$ the eigenvalues of $\delta M_{\Delta\tau}$ so the matrix $I - \delta M_{\Delta\tau}$ has eigenvalues $\rho_j = 1 - \lambda_j$ and the inverse exists if $|\rho_j| < 1$ for all $j \in D$. But this holds by the following argument: Let $M_{\Delta\tau}\boldsymbol{v} = \lambda\boldsymbol{v}$ for some eigenvalue λ. Let v_k the largest entry in the eigenvector such that $|v_k| \geq |v_s|$ for all k, s in D. Then $\sum_{j=1}^d m_{kj}v_j = \frac{\lambda}{\delta\Delta\tau_k}v_k$. But then $|\frac{\lambda}{\delta\Delta\tau_k}||v_k| = \sum_{j=1}^d m_{kj}|v_j| \leq \sum_{j=1}^d m_{kj}|v_k| = |v_k|$ so $|\frac{\lambda}{\delta\Delta\tau_k}| \leq 1$. Since $\delta < 1$ and $|\Delta\tau_k| \leq 1$ we conclude that $\lambda < 1$ implying $0 < \rho_j < 1$ and thus $(I - \delta M_{\Delta\tau})^{-1}$ exists.

with equality for all j. If not, strictly lower pay-offs exist and the pay-off is not extreme. Notice that, by implication, $\underline{\tau} = 0$ if the punishment is in stage-game actions $\boldsymbol{q}^n$. Turn to the regular phase, and fix some $\overline{\boldsymbol{q}}^j$ to be enforced. The optimal transition function $\overline{\tau}$ minimizes inefficient equilibrium transitions (size, $\overline{\tau}$) while maintaining incentive compatibility (power, $\overline{\tau}(\boldsymbol{q}') - \overline{\tau}(\boldsymbol{q})$). The trade-off between size and power is optimized by restricting the transition probability to the convex region of F_θ. We show that the regular-phase transition probability is given by

$$\overline{\tau}^j(\boldsymbol{\iota}'\boldsymbol{q}, \overline{p}^j) = \left(1 - F_\theta\left(\frac{\overline{p}^j}{p(\boldsymbol{\iota}'\tilde{\boldsymbol{q}}, x^j)}\right)\right)(1 - \mathbb{1}(\boldsymbol{\iota}'\boldsymbol{q} \neq \boldsymbol{\iota}'\overline{\boldsymbol{q}}^j)) \tag{19}$$

where $\mathbb{1}(\boldsymbol{\iota}'\boldsymbol{q} \neq \boldsymbol{\iota}'\overline{\boldsymbol{q}}^j)$ is in an indicator for a public defection and the trigger price is given by the smallest

$$\overline{p}^j \in [0, \exp(-\sigma_\theta^2 3/2) p(y^j, \boldsymbol{\iota}'\boldsymbol{q})]$$

satisfying the incentive conditions (14) for every $\boldsymbol{q} \in Q^2$.

A defection in hidden quantities is inferred from realizations of conditionally log-normally distributed prices, $\ln p(\tilde{\boldsymbol{q}}, x, \tilde{\theta})$ where we recall that the distribution $F_p(\cdot|\boldsymbol{q})$ is parameterized by the action profile $\boldsymbol{q}$. Notice that the players' inference problem corresponds to a goodness of fit test across $\{\ln p(\tilde{\boldsymbol{q}}, x, \tilde{\theta}) : \boldsymbol{q} \in Q^2\}$. The likelihood-ratio test of the hypothesis $\tilde{\boldsymbol{q}} > \overline{\boldsymbol{q}}$ 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

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

for $\boldsymbol{\iota}'\hat{\boldsymbol{q}} > \boldsymbol{\iota}'\boldsymbol{q}$, so the likelihood of a deviation is monotonically increasing in the realized price level. A tail test of observed prices $\overline{p}^j \leq p(\boldsymbol{q}, x, \tilde{\theta})$ is a sufficient statistic for the likelihood ratio. Thus the functional form of the transition function is

$$\Pr(p(\tilde{\boldsymbol{q}}, x, \tilde{\theta}) \geq \overline{p}^j) = \Pr\left(\tilde{\theta} \leq \frac{\overline{p}^j}{p(\boldsymbol{q}, x)}\right) = 1 - F_\theta\left(\frac{\overline{p}^j}{p(\boldsymbol{q}, x)}\right).$$

To determine the upper bound $\overline{p}$, notice that incentives are provided by the conditional difference (power) $\overline{\tau}^j(\boldsymbol{q}') - \overline{\tau}^j(\overline{\boldsymbol{q}}^j)$, $\boldsymbol{q}' \neq \overline{\boldsymbol{q}}$ not the level $\overline{\tau}^j(\overline{\boldsymbol{q}}^j)$ or false positive rate. It is never optimal for the trigger price $\overline{p}^j$ to locate $F_\theta(\overline{p}^j/p(\boldsymbol{q}), x)$ in the concave region, as the same power can be achieved with

a lower probability of false positives, increasing pay-offs while maintaining incentives. The second derivative F''_θ changes sign at the mode,

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

where replacing z with $\overline{p}^j/p(\boldsymbol{q},x)$ in the final expression above yields the proposed bound, restricting F_θ to the efficient convex region. Finally, and again for fixed actions $\overline{\boldsymbol{q}}$, the trigger $\overline{p}^j$ is optimally set to the lowest level such that all incentive compatibility constraints hold, minimizing false positives. Notice that the transition probability under punishment $\underline{\tau}$ requires the output to be in observable quantities. If $\underline{X}^j$ were non-observable the transition probability would be governed by a tail test and imply efficiency losses. □

A.3 Proposition 1. Covariance.

Proof. We show that

$$\lim_{\sigma_\theta \to 0} \lim_{\delta \to 1} \overline{\boldsymbol{q}} - \underline{\boldsymbol{q}} = \boldsymbol{q}^m - \boldsymbol{q}^*$$

such that the stated result follows directly from Equation (16) in Appendix (A.2) and that jointly profit-maximizing quantities $\boldsymbol{q}^m(x)$ clearly are weakly increasing in x.

By the assumption $\boldsymbol{x} : \boldsymbol{\pi}^m \geq \boldsymbol{\pi}^n$ there are returns to coordinating output in at least one demand state so $M\Delta\boldsymbol{\pi} > 0$ if $\boldsymbol{q}^m$ may be enforced. But in the limit $\boldsymbol{q}^m$ are enforceable: By the bang-bang property optimal punishment is no less severe than permanent reversion to stage-game Nash equilibria. It is therefore sufficient to show that $\boldsymbol{q}^m$ may be enforced under that punishment. Combine the equilibrium conditions (14) with the value function differential identity (15) to yield the reward-phase equilibrium condition

$$\Delta\overline{\boldsymbol{\pi}} \leq \delta\Delta\overline{\boldsymbol{t}} \circ M(I - \delta M_{\Delta\tau})^{-1}\Delta\boldsymbol{\pi} \tag{20}$$

where the equilibrium condition under punishment holds trivially under permanent reversion to stage-game equilibria and is omitted. But then

$$\lim_{\sigma_\theta \to 0} \lim_{\delta \to 1} \delta M(I - \delta M)^{-1}\Delta\boldsymbol{\pi} = \boldsymbol{\infty}$$

as $\lim_{\sigma_\theta \to 0} \Delta \boldsymbol{t} = \mathbf{1} - \mathbf{0} = \mathbf{1}$ and $\max \text{eig} M = 1$ as M is a stochastic matrix. Thus $\overline{\boldsymbol{\pi}} = \boldsymbol{\pi}^m$ may be supported and incentives may become arbitrarily powerful.

Recall from Proposition 2 that the optimal equilibrium demands the most severe punishment, minimizing $\underline{v}$, or equivalently, maximizing $\overline{v}$. We show that greater punishment quantities have a direct positive effects on incentives via greater per-period losses and a negative indirect effect through reduced punishment duration, but that the former dominates under the stated limiting conditions. To begin consider the equilibrium condition under punishment

$$\Delta \underline{\boldsymbol{\pi}} \leq \delta \Delta \underline{\boldsymbol{t}} \circ M(I - \delta M_{\Delta\tau})^{-1} \Delta \boldsymbol{\pi} \tag{21}$$

where we again combine the equilibrium condition (14) and the incentive power identity (15). Notice that raising the punishment quantity $\underline{q}$ increases the per-period loss $\Delta\pi$ and thus the power of incentives

$$\Delta \boldsymbol{v} = (I - \delta M_{\Delta\tau})^{-1}(1 - \delta) \Delta \boldsymbol{\pi}$$

directly. But the one-shot deviation payoff $\Delta \underline{\boldsymbol{\pi}}$ will also increase. Clearly $\underline{\boldsymbol{t}} = \mathbf{0}$ under the absorbing punishment. It follows immediately from the equilibrium condition (21) that any punishment with $\underline{q} > q^n$ may not be absorbing, so $\underline{\boldsymbol{t}}$ in (14) must increase with $\underline{q}$. But raising $\Delta \underline{\boldsymbol{t}}$ reduces $\Delta \boldsymbol{t}$ which indirectly lowers $\Delta \boldsymbol{v}$ through a reduction in $M_{\Delta\tau}$. But notice that magnitudes of the direct- and secondary effects are respectively increasing- and decreasing in δ. For sufficiently patient players the direct effect dominates and punishments are t □

A.4 Proposition 3: Incentive power.

Proposition 3. ***Incentive power.*** *Consider an unanticipated one-time shift in incentive power from $\overline{\boldsymbol{v}}$, $\underline{\boldsymbol{v}}$ to $\overline{\boldsymbol{v}}'$, $\underline{\boldsymbol{v}}'$ such that $\Delta \boldsymbol{v}' \leq \Delta \boldsymbol{v}$. The corresponding optimal reward-phase output $\overline{\boldsymbol{q}}' \geq \overline{\boldsymbol{q}}$ and transition probabilities $\overline{\boldsymbol{t}}' \geq \overline{\boldsymbol{t}}$, $\underline{\boldsymbol{t}}' \geq \underline{\boldsymbol{t}}$ and punishment quantities $\underline{\boldsymbol{q}}' \leq \underline{\boldsymbol{q}}$ weakly increase- and decrease, respectively.*

Proof. In optimal equilibria the equilibrium conditions will hold with equality for at least one violation, see the proof of Proposition 2 in Appendix A.2. The equilibrium conditions will therefore be violated for a sufficient reduction in incentive power such that initial action profiles and transition functions are no longer enforceable. Applying $B(\Delta \boldsymbol{V}')$ from (2) it may be verified directly that $\Delta \overline{\boldsymbol{\pi}}' \leq \overline{\boldsymbol{\pi}}$, $\Delta \underline{\boldsymbol{\pi}}' \leq \underline{\boldsymbol{\pi}}$ and $\Delta \overline{\boldsymbol{t}}' \geq \Delta \overline{\boldsymbol{t}}$, $\Delta \underline{\boldsymbol{t}}' \geq \Delta \underline{\boldsymbol{t}}$. The required reduction in deviation payoffs implies $\overline{\boldsymbol{q}}' \geq \overline{\boldsymbol{q}}$ and $\underline{\boldsymbol{q}}' \leq \underline{\boldsymbol{q}}$ by the quantity ranking in Equation (16). □

A.5 Proposition 4. Monitoring quality.

Proposition 4. ***Monitoring quality.*** *Incentives $\Delta \boldsymbol{v}$ are decreasing in signal noise σ_θ.*

Proof. A marginal increase σ_θ reduces the slope of $F_\theta(\cdot)$ if

$$\frac{\partial \Delta \bar{t}}{\partial \sigma_\theta} = \frac{\partial}{\partial \sigma_\theta}\left(F_\theta\left(\frac{\overline{p}}{p(\overline{\boldsymbol{q}}', x)}\right) - F_\theta\left(\frac{\overline{p}}{p(\overline{\boldsymbol{q}}, x)}\right)\right) < 0$$

is falling for all deviations $\overline{\boldsymbol{q}}'$. To verify recast the cumulative density function in terms of the error function

$$F_\theta\left(\frac{\overline{p}}{p(\overline{\boldsymbol{q}}, x)}\right) = \frac{1}{2} + \frac{1}{2}\text{erf}\left(\frac{\ln\left(\frac{\overline{p}}{p(\overline{\boldsymbol{q}}, x)}\right) - \frac{\sigma_\theta^2}{2}}{2\sigma_\theta}\right).$$

But then

$$F_\theta\left(\frac{\overline{p}}{p(\overline{\boldsymbol{q}}', x)}\right) - F_\theta\left(\frac{\overline{p}}{p(\overline{\boldsymbol{q}}, x)}\right) = \text{erf}\left(\frac{\ln\left(\frac{\overline{p}}{p(\overline{\boldsymbol{q}}', x)}\right)}{2\sigma_\theta} - \frac{\sigma_\theta}{2^{3/2}}\right) - \text{erf}\left(\frac{\ln\left(\frac{\overline{p}}{p(\overline{\boldsymbol{q}}, x)}\right)}{2\sigma_\theta} - \frac{\sigma_\theta}{2^{3/2}}\right)$$

with both arguments convexly decreasing in σ_θ. But the error function is strictly increasing and

$$\overline{\boldsymbol{q}}' > \overline{\boldsymbol{q}} \implies p(\overline{\boldsymbol{q}}', x) < p(\overline{\boldsymbol{q}}, x) \implies \ln\left(\frac{\overline{p}}{p(\overline{\boldsymbol{q}}', x)}\right) > \ln\left(\frac{\overline{p}}{p(\overline{\boldsymbol{q}}, x)}\right)$$

and hence the slope $\Delta \bar{t}$ decreases. By the proof of Proposition 2, at least one incentive compatibility constraint binds under regular play, so the increase in σ_θ will leave the equilibrium conditions violated upon impact. In response, quantities $\overline{\boldsymbol{q}}$ or trigger price $\overline{p}$ must increase, decreasing $\overline{\boldsymbol{v}}$, in turn increasing $\underline{\boldsymbol{v}}$, and thus also $\Delta \boldsymbol{v}$. □

A.6 Proposition 5. Dynamic demand.

Proposition 5. ***Demand dynamics and incentive power.*** *Let j, s and $k \neq s$ indices in $D = \{1, \ldots, d\}$ and consider transition matrices M, $\hat{M}$ where $\hat{m}_{js} > m_{js}$, $m_{jk}/(1 - m_{js}) = \hat{m}_{jk}/(1 - \hat{m}_{js})$ and $\boldsymbol{M}_k = \hat{\boldsymbol{M}}_k$ for all $k \in D \setminus j$. Let V and V' the set of equilibrium payoffs induced by M and M'. Then $V \subset \hat{V}$ if $\sum_{m_{jk} \in \boldsymbol{M}_j \setminus m_{js}} m_{jk} \Delta v^k < \Delta v^s$.*

Proof. Consider $W = B(V|\hat{M})$, where $\{w^k \in W : k \in D \setminus j\} = \{v^k \in V : k \in D \setminus j\}$. Let $\text{ext}W_j = \{\overline{w}^j, \underline{w}^j\}$ the extrema in dimension j and note that

$$\begin{aligned}
\Delta w^j &= (1-\delta)\Delta\pi^j + \delta\Delta\tau \sum_{\hat{m}_{jk}\in\hat{\boldsymbol{M}}_j} m_{jk}\delta v^k =: \overline{w}^j - \underline{w}^j \\
&= (1-\delta)\Delta\pi^j + \delta\tau \left(\hat{m}_{js}\Delta v^s + (1-\hat{m}_{js}) \sum_{\hat{m}_{jk}\in\hat{\boldsymbol{M}}_j\setminus\hat{m}_{js}} \frac{\hat{m}_{jk}}{1-\hat{m}_{js}}\Delta v^k \right) \\
&= (1-\delta)\Delta\pi^j + \delta\tau \left(\hat{m}_{js}\Delta v^s + (1-\hat{m}_{js}) \sum_{m_{jk}\in\boldsymbol{M}_j\setminus m_{js}} \frac{m_{jk}}{1-m_{js}}\Delta v^k \right) \\
&\Rightarrow \Delta w^j - \Delta v^j = (\hat{m}_{js} - m_{js}) \left(\Delta v^s - \sum_{m_{jk}\in\boldsymbol{M}_j\setminus m_{js}} m_{jk}\Delta v^k \right) \geq 0
\end{aligned}$$

by construction. Greater incentive power in state j yields weakly more extreme outcomes so $B(V|\hat{M})$ is a proper superset of V and $V \subset \hat{V}$. □

A.7 Proposition 6: Non-monotonic incentive power.

Proposition 6. ***Demand level and incentive power.*** *Let $\sigma_\theta \to 0$, $\delta \to 1$ and $m_{jj} \to 1$. Then there exist demand levels $x^j \in \{x', x'', x'''\} \subset X$ increasing $x' < x'' < x'''$ such that incentive-power $\Delta\boldsymbol{v}(x') < \Delta\boldsymbol{v}(x'') > \Delta\boldsymbol{v}(x''')$.*

Proof. Let $\Delta\pi^c(x) := \pi^m(x) - \pi^n(x)$ the difference between symmetric maximum- and stage-game Nash profits and $\Delta\pi(x) \in \overline{\pi}(x) - \underline{\pi}(x)$ the difference between unconstrained reward- and punishment profits. Note that because $M > 0$ incentive power is complementary across states and $\Delta v^j(x') \leq \Delta v^j(x'') \geq \Delta v^j(x''')$ implies $\Delta\boldsymbol{v}(y') \leq \Delta\boldsymbol{v}(x'') \geq \Delta\boldsymbol{v}(x''')$. By Proposition 2 incentive power is locally increasing if $\Delta\pi^c(x'') \geq \Delta\pi^c(x')$. Let $x' := x : \pi(\epsilon, \epsilon, x') = 0$. The unconstrained reward phase payoff $\lim_{m_{jj}\to 1} \Delta\overline{v}^j(x') = 0$ and by individual rationality $\underline{v}^j(x') \geq 0$. Hence $\lim_{m_{jj}\to 1} \Delta v^j(x') = 0$. Next let $x'' \in x : q^m(y'') < q^n(x'')$, implying and $\Delta\pi^c(x'') > 0$ so the unconstrained $\Delta v(y'') > 0$. Finally let $x''' : q^n(y) = x^*$ capacity constrained implying that $\underline{x} = x^*$ so $\Delta\pi^c(y''') \geq \Delta\pi$. But $\pi(\boldsymbol{q}^*, y)$ increases linearly at rate $p(\boldsymbol{q}^*)q^*$ in y whereas the jointly profit-maximizing payoff $\pi(q^{\boldsymbol{q}^m}, x)$ increases at most by $p(\boldsymbol{q}^m(x))q^m(x)$. By Assumption (1) the term $p(\boldsymbol{q})q$ is increasing everywhere in q so $p(\boldsymbol{q}^*)q^* \geq p(\boldsymbol{q}^m)q^m$. But then $\Delta\pi^c(y)$ decreases in y for all $y : q^n = q^*$ and $\lim_{x\to\infty} \Delta v(x) = 0$.

□

A.8 Proposition 7: Intertemporal incentives

Proposition 7. ***Intertemporal incentives.*** *Let x' such that $q^m(x') < q^n(x') < q^*$. There exists parameters $\boldsymbol{x}$, M, $\delta < 1$ and σ_θ such that $\overline{q}(^i) = q^m(^i)$, $\underline{q}(^i) = q^*$, and $\overline{q}(x^j) = q^n(x^j) = \underline{q}(x^j)$ for identical demand levels $x^i = x^j = x'$.*

Proof. Proof is by construction. Let M such that $\boldsymbol{m}_i \Delta \boldsymbol{v} > 0$. By the proof of Proposition 1 there exist $\delta < 1$ and $\sigma_\theta > 0$ such that $\overline{q}(^i) = q^m(^i)$, $\underline{q}(^i) = q^*$. Let $0 = x^k \in\in \boldsymbol{x}$ and fix δ and σ_θ. Then $\lim_{m_{jk} \to 1} \boldsymbol{m}_j \Delta \boldsymbol{v} = 0$ and $\overline{q}(x^j) = q^n(x^j) = \underline{q}(x^j)$. □

B Parameter values for numerical solution

The inverse demand function is $p(y,q) = y(\xi + \beta q)^{-1}$, with ξ and β governing the inverse elasticity of demand. The reference parameters are held constant when evaluating a range of the alternative parameter.

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

Table 1: Parameter values for numerical solutions.

C Appendix: Figures

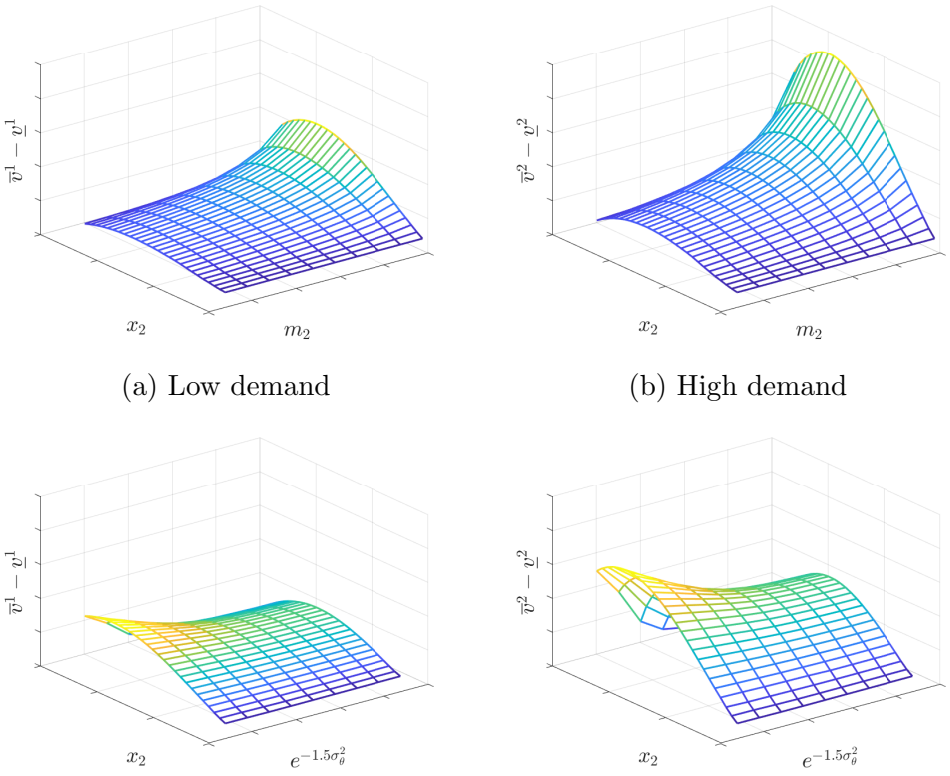

(a) Low demand (b) High demand

(c) Low demand (d) High demand

Figure 12: Incentive power $\Delta \boldsymbol{v}$ against the level x_2 and persistence m_2 of the high demand state and the signal noise distributions' mode $\exp(-1.5\sigma_\theta^2)$ (see Appendix A.2). The incentive power is evaluated in the low (a,c) and high (b,d) demand states x^1 and x^2, respectively.

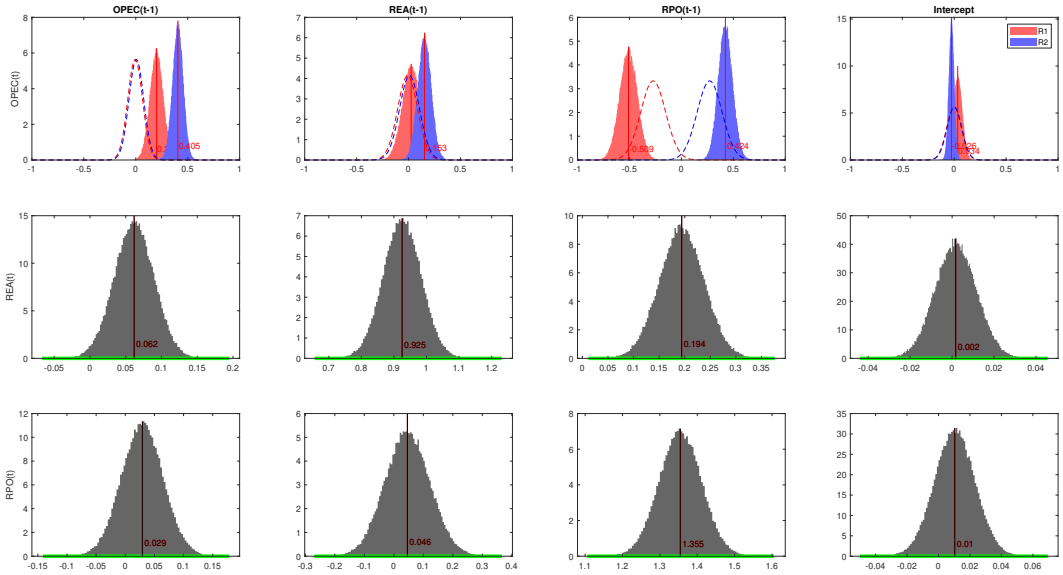

Figure 13: Implied prior and realized posterior distributions over VAR coefficients. The first two rows reports the non-switching block, with priors in green and posterior in gray. The final row shows the switching equation with priors in dashed lines.

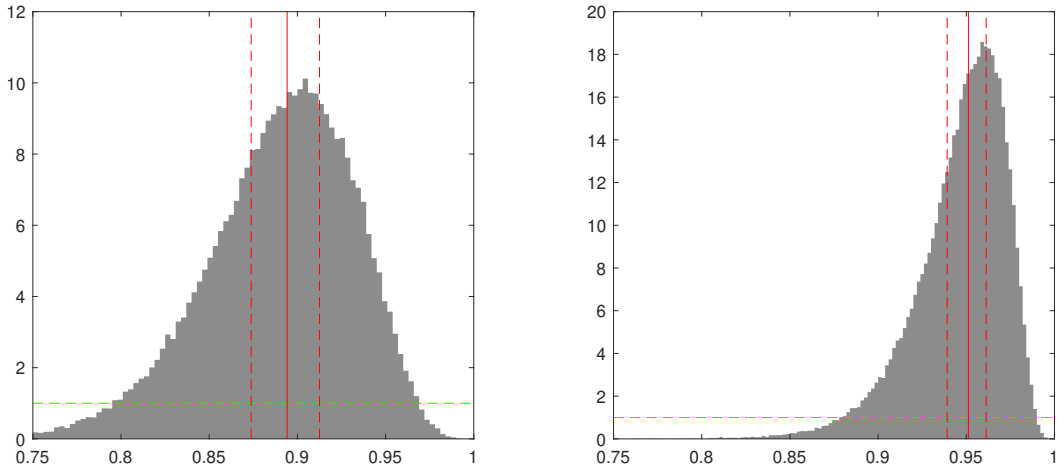

Figure 14: The prior (green) and posterior (gray) distributions over transition probabilities for the pro- and countercyclical regimes, $\Pr(s_{t+1} = \text{PC}|s_t = \text{PC})$ (left-hand-side) and $\Pr(s_{t+1} = \text{CC}|s_t = \text{CC})$ (right-hand-side) with median and 68% credible sets are reported in red.

D Appendix: Data

D.1 Data description, sample and source

The data are described in Table 2. The normalized, endogenous variables of the MS-BVAR are plotted in Figure 15.

Data	Description	Sample	Source
Δq_t	The twelve-month change in total OPEC crude oil production in thousands of barrels per day. The time series is standardized prior to estimation.	1984:M01–2019:M12	International Energy Agency Monthly Oil Data Service (IEA MODS)
Δx_t	The twelve-month log change in the OECD+6 industrial production index provided by Baumeister and Hamilton (2019a). This detrending procedure is in line with the recommendations of Hamilton (2018) and is equivalent to the forecast error from a random walk model with a twelve-month forecast horizon.	1958:M01–2019:M12	`https://sites.google.com/site/cjsbaumeister/research`
	The Kilian (2009) index of demand for industrial commodities.	1968:M01–2020:M04	`https://sites.google.com/site/lkilian2019/research/data-sets`
$\Delta \log p_t$	The twelve-month log change in the real price of oil. The underlying time series is the refiner's acquisition cost of imported crude oil.	1974:M01–2019:M12	U.S. Energy Information Administration (EIA).
	We deflate the price by the U.S. CPI.		`CPIAUCSL` at the FRED Data Service, Federal Reserve Bank of St. Louis.
$p^f_{t,t+h}$	The end of month West Texas Intermediate futures contract price at time t with delivery h months in the future.	1985:M01–2019:M12	Thomson Reuters Datastream: `NCLC.h`

Table 2: Variable Summary and Data Sources

D.2 Narrative evidence

We downloaded the 2001–2019 Oil Market Report (OMR) from `iea.org`. The 1990–2000 reports are archived on `https://archive.org/web`. Each report features an executive summary of the key forces driving oil price change since last report. We construct an indicator series for when IEA reports that OPEC's actions increased or decreased prices and for when there was an exogenous disruption, yielding three series in total. The coding involves some

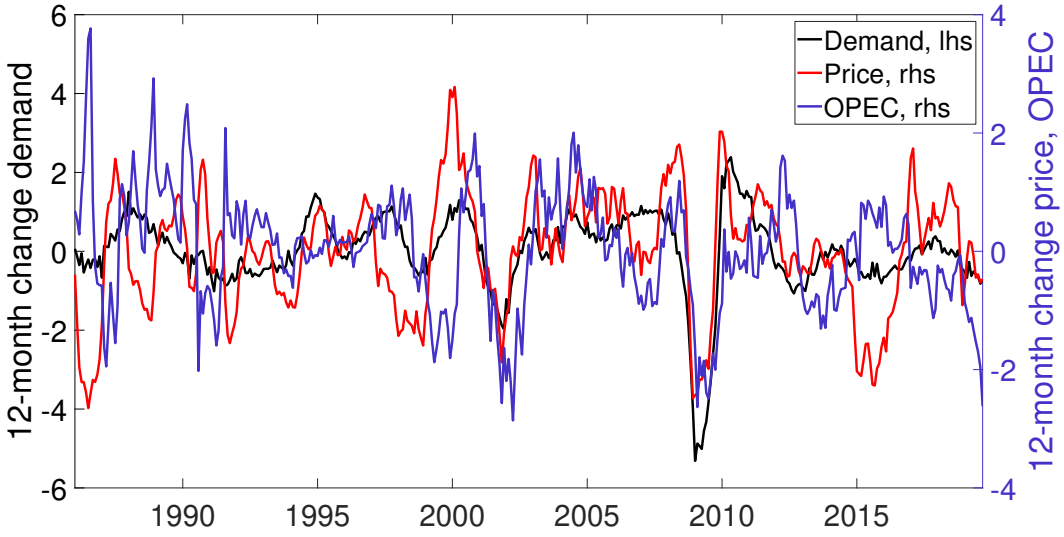

Figure 15: Normalized raw data. January 1985 - September 2019. The demand indicator is plotted the left-hand-side, twelve-month changes in log real oil prices and OPEC output on the right-hand-side.

judgment. The full set of executive summaries and our coding procedure are available upon request. We provide an example for each series below.

- **OPEC actions reducing prices:** December 2014: "Oil prices continued to plunge in November and into early December. The selloff gained pace after OPEC on 27 November decided to keep its output target unchanged."
- **OPEC actions increasing prices:** December 2017: "Benchmark crude prices rose by $4-5/bbl on average in November and traded at their highest level in more than two years in early December. The extension of the OPEC/non-OPEC output cuts and, latterly, the closure of the Forties pipeline system were factors."
- **Exogenous event affects OPEC production:** January 2006: "NYMEX WTI averaged $59.45/bbl in December and pushed above $64/bbl in early January on strong gasoline prices, Nigerian outages and uncertainty about Iran's nuclear program."

Prior distributions for the switching equation

Parameter(s)	Variable(s)	Density	Mean	St. dev.
$\mu_\bullet(s^i)$	Intercept	Normal	0	0.05
$\beta_1(s^i) \in B_\bullet$	Contemporaneous	Normal	0	0.071
$\beta_2(s^i) \in B_\bullet$	Contemporaneous	Normal	$(-1)^i \times 0.2$	0.071
$\forall\, \zeta(s^i) \in C_{\bullet 1}(s^i)$	1st lag	Normal	0	0.071
p_{11}	Probability	Dirichlet	0.50	0.29
p_{22}	Probability	Dirichlet	0.50	0.29
ρ_j	Probability	Dirichlet	0.50	0.22

Table 3: Prior distributions for Equation 5, the transition matrix P and the initial probabilities needed to initialize the Hamilton filter ρ_j. $s^1 = PC$ and $s^2 = CC$.

E Appendix: Prior selection

Priors on coefficients of switching Equation 5 and transition matrix P are listed in Table 3. Parameters in $B_\bullet$ and $C_{\bullet\ell}$ denote coefficients on contemporaneous and lagged terms respectively. The columns p_{i1} and p_{i2} in P are Dirichlet distributed $\text{Dir}(\boldsymbol{\alpha}_j)$ with $\boldsymbol{\alpha}_1 = [\,1\;1\,]$ and $\boldsymbol{\alpha}_2 = [\,1\;1\,]$. This choice of shape parameters results in uniform (flat) prior distributions for the transition probabilities. For ρ_j, we specify Dirichlet prior distribution $\boldsymbol{\alpha}_j = [\,2\;2\,]$ as to put equal probability of being in a given regime at $t = 1$.[25]

The implied prior distributions of the VAR system will for the unrestricted equation be transformations of the distributions listed above and the diffuse priors from the non-switching equations that have Gaussian inverse-Wishart posterior distributions. As explained by Hamilton (2016, p. 175), the coefficients for the full VAR system can be recovered by computing $\hat{\Sigma}_{12}(s_t) = \hat{B}_\bullet(s_t)\hat{\Sigma}_{22}$ and then $\hat{A}_\bullet(s_t) = \hat{C}_\bullet(s_t) + \hat{\Sigma}_{12}(s_t)\hat{\Sigma}_{22}^{-1}\hat{A}_\circ = \hat{C}_\bullet(s_t) + \hat{B}_\bullet(s_t)\hat{A}_\circ$.[26] If we evaluate these relationships at the prior means, noting that the prior means of $A_\circ$ is the OLS estimates, and assume for for simplicity that we have

[25] In our dual regime case, the Dirichlet distribution has only two inputs of shape parameters. For this special case, its probability density function is identical to that of the beta distribution.

[26] With the notation used in the main text, these computations can be done directly if we suppress the intercept term and assume one lag. A more compact notation is used to describe our model and the estimation in Appendix H from which these relationships can be computed directly.

only one lag, then we are left with

$$\bar{A}_{\bullet}(s_t) = \begin{bmatrix} 0 & \bar{\beta}_2(s_t) \end{bmatrix} \begin{bmatrix} \hat{\alpha}_{21} & \hat{\alpha}_{22} & \hat{\alpha}_{23} \\ \hat{\alpha}_{31} & \hat{\alpha}_{32} & \hat{\alpha}_{33} \end{bmatrix}$$

The imputed prior means for our parameters of interest $\alpha_{13}(s_t)$ is then -0.273 and 0.273 for PC and CC, respectively. These imputed priors can be seen in Figure 13 for the first lag.

F Robustness and extensions

This section describe a variety of extensions and robustness checks. Apart from the specification of the baseline model, we also look at extensions of the model where we substitute baseline model variables with other measures of oil demand and oil supply. Finally, we report a summary of MCMC covergence diagnostics.

F.1 Model specification and prior sensitivity

Our choice to specify a model of 24 lags is motivated by the literature. Among others Hamilton and Herrera (2004), Kilian and Murphy (2014) argue that a long lag order is required to capture the full transmission of oil price shocks and to study the dynamics of business cycles in commodity markets. Kilian and Lütkepohl (2017) advocate an *ex ante* choice of lag order rather than relying on information criteria. We have estimated our model with 12 and 18 lags as well. The main results are all similar in substance, with the exception that regime classification becomes less precise during the 2000s.

Our baseline specification prescribes non-informative priors on the non-switching equation and zero-mean normal priors for coefficients in the switching equation. Zero means are motivated by the data being in growth rates and not in levels. Estimating the model with looser priors for the switching block does not change our conclusions. For the contemporaneous term $\beta_2(s_t) \in B_{\bullet}$, a prior mean of zero is not robustly feasible because it causes degeneracy in expectation when the full VAR representation is computed (see Appendix E for details on this transformation). Neither do we hold prior beliefs that the contemporaneous correlation between quantities produced and prices is zero. Our sensitivity analysis also suggests that the prior means of $\beta_2(PC)$ and $\beta_2(CC)$ must be sufficiently different (approximately ± 0.05) in order to reliably identify the regimes.

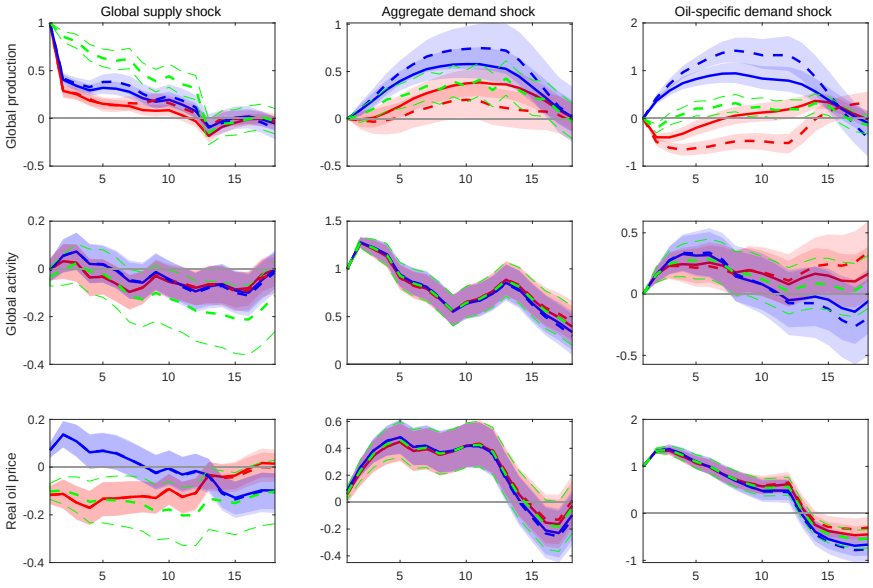

Figure 16: Exact impulse response functions (EIRF) using the specification with the Kilian (2009) indicator for real economic activity. The pro- and countercyclical regimes are plotted in red and blue, respectively. The dashed lines show regime-dependent response (RDIRF) and shaded areas 68% credible sets. The impulse-response function (IRF) of the one-regime BVAR reported in green together with the 68% credible set.

F.2 The Kilian index of real economic activity

We re-estimate our baseline model with one simple change. In place of the twelve-month growth rate of the OECD+6 index of industrial production (Baumeister and Hamilton, 2019a), we include the Kilian index of global real economic activity (Kilian, 2009). As can be seen in Figures 16 and 17, the exact impulse response functions and posterior distribution over regime indicator variables do not substantially change when this substitution is made. This exercise highlights that our findings are driven mainly by switching in the price-quantity relationship and not the choice of demand indicator. The main difference lies again in poorer regime classification as reflected by an indeterminate posterior mean of the regime indicator sequence.

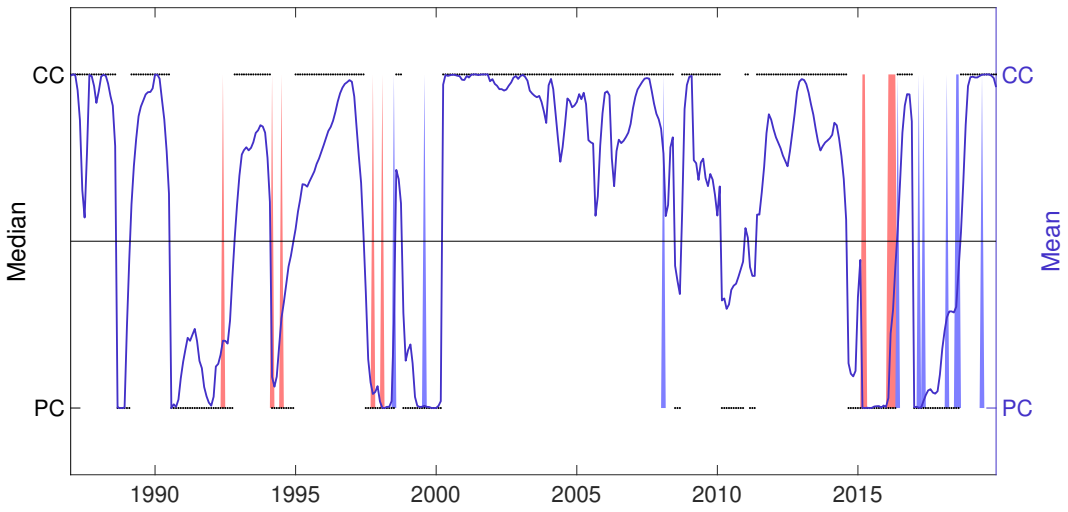

Figure 17: Posterior regime indicator against time using the specification with the Kilian (2009) indicator for global demand. Pro- and counter-cyclical regimes are denoted by PC and CC. Median, black points and mean, blue. The lightly shaded areas denote output wars with high confidence of market flooding behavior, $\Delta \ln p_t < 0 < \Delta q_t$ and $\Pr(s_t = \text{PC}) > .75$. The darkly shaded red and blue areas denote months in which the International Energy Agency's Oil Market Report (OMR) finds that OPEC's actions are substantially decreasing or increasing the price of oil.

F.3 Alternative oil supply variables

We run two additional estimations of our baseline model to assess whether our results are spurious. Instead of estimating the model with OPEC crude oil production, we include total non-OPEC production as well as global crude oil production. The estimation procedure mechanically decomposes both non-OPEC and total world output into pro- and countercyclical regimes. See Figures 18a and 18b for the impulse response functions.

Notice that there is no response of real oil prices to non-OPEC output shocks while the aggregate response is nearly identical. For the specification employing total world production, the posterior regime classifications overlap during output wars, showing that OPEC's actions affect the aggregate outcome in a substantively similar way, see Figure 19b. Unsurprisingly, the specification with total output registers several instances of procyclical behavior not present in the baseline model with OPEC output. For instance, the procyclical production in August and September of 2005 likely reflects the effects of Hurricane Katrina. Finally and importantly, there is no evidence of non-OPEC output wars, see Figure 19a. Interestingly, the posterior mean of the regime indicator sequence indicates that non-OPEC output is largely passive prior to 2013, after which it becomes more strongly pro-cyclical. This is consistent with evidence in Gundersen (2020) of increased importance of U.S oil production following the shale revolution.

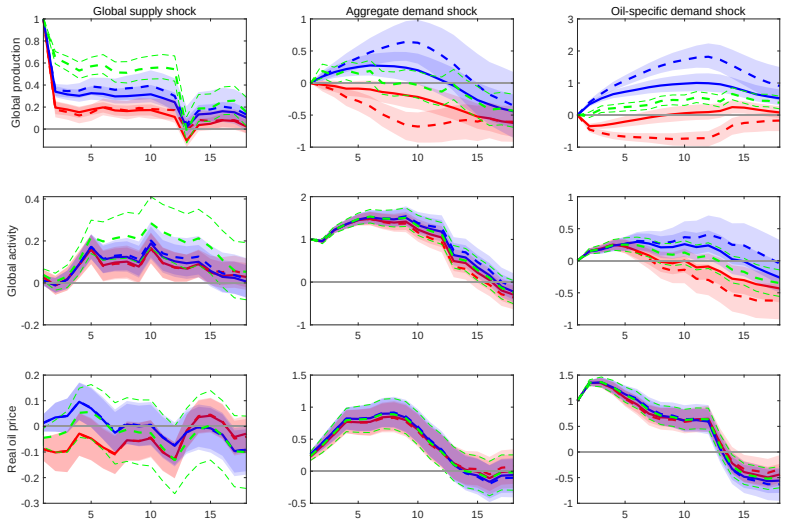

(a) Non-OPEC

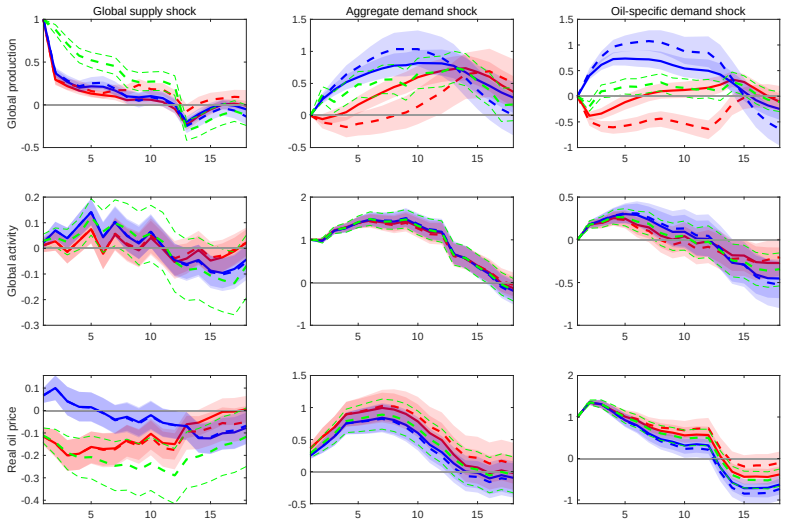

(b) Total

Figure 18: Exact impulse response functions (EIRF) using the specification with Non-OPEC and total global output. The pro- and countercyclical regimes are plotted in red and blue, respectively. The dashed lines show regime-dependent response (RDIRF) and shaded areas 68% credible sets. The impulse-response function (IRF) of the one-regime BVAR reported in green together with the 68% credible set.

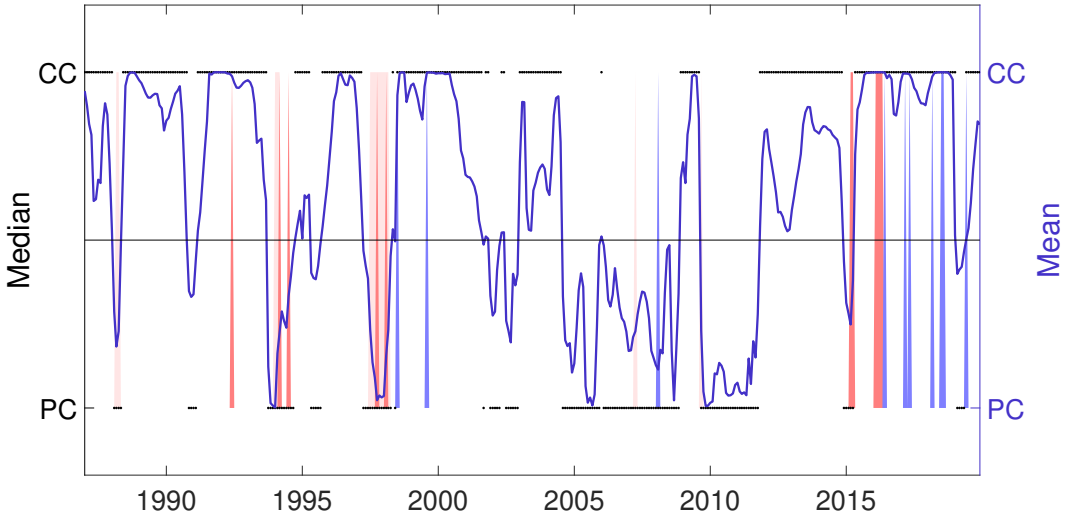

(a) Non-OPEC

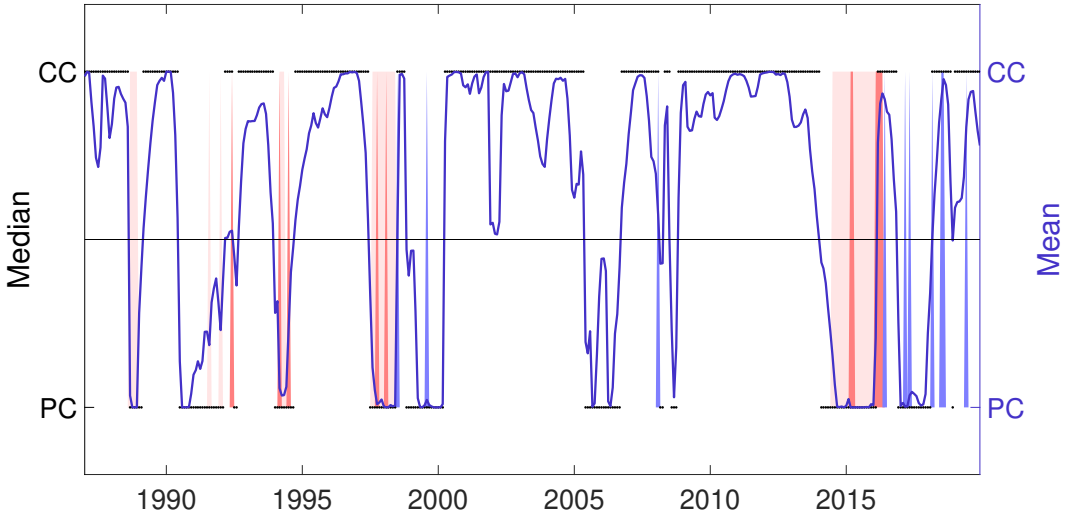

(b) Total

Figure 19: Posterior regime indicator against time using the specification with non-OPEC (top) and total (bottom) output. Pro- and counter-cyclical regimes are denoted by PC and CC. Median, black points and mean, blue. The lightly shaded areas denote output wars with high confidence of market flooding behavior, $\Delta \ln p_t < 0 < \Delta q_t$ and $\Pr(s_t = \text{PC}) > .75$. The darkly shaded red and blue areas denote months in which the International Energy Agency's Oil Market Report (OMR) finds that OPEC's actions are substantially decreasing or increasing the price of oil.

F.4 MCMC convergence diagnostics

We perform simple diagnostics to assess convergence of the MCMC sampler. First, we compute Geweke inefficiency factors for each parameter. An efficiency factor of 5 means that we will need 50,000 draws from the Gibbs sampler to obtain the same efficiency as 10,000 ideally obtained IID draws. With our 100,000 draws, the inefficiency factors of most parameters range between 0.5 and 2. A handful of parameters have larger inefficiency factors than this, but none are above 17. A rule of thumb is that inefficiency factors below 20 suggest convergence.[27] See Figure 20 for the distribution of inefficiency factors.

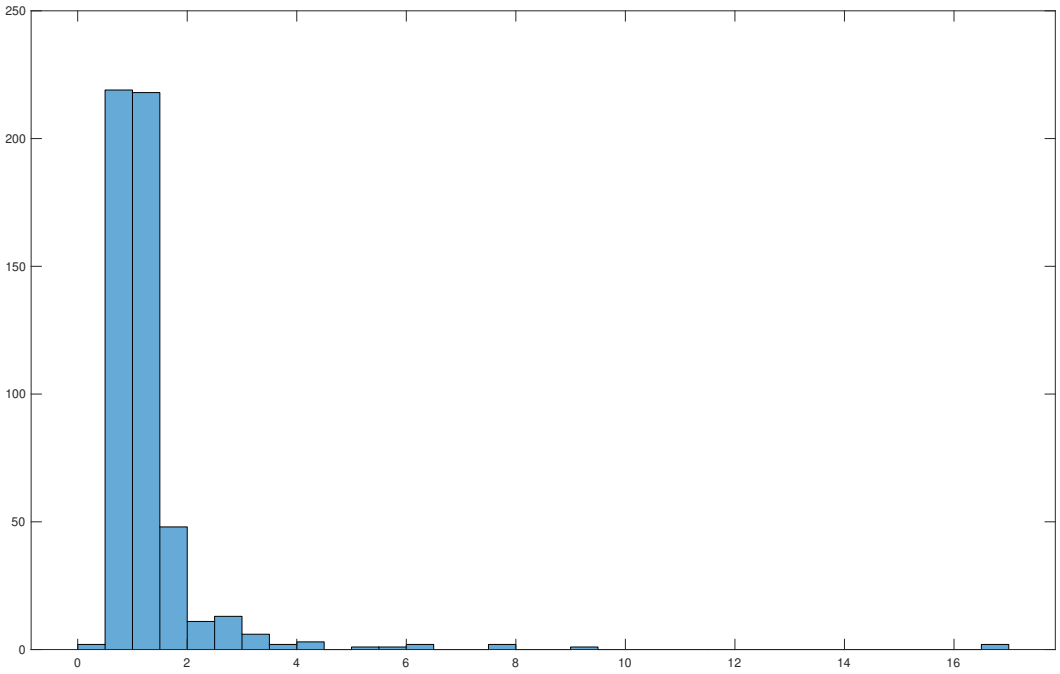

Figure 20: Distribution of inefficiency factors for all parameters. An inefficiency factor below 20 is conventionally deemed acceptable.

Second, for each parameter, we perform a difference in means test where the means are computed for the first 10% and the last 50% of the draws. Apart from a few exceptions, we cannot reject the null hypotheses of no differences in means at the 1% or 5% levels.[28] A plot of the MCMC draws of the parameter

[27]The expression for the inefficiency factor is given by $1 + 2\sum_{k=1}^{\infty}\gamma(k)$, where $\gamma(k)$ is the autocorrelation function for the trace of a parameter at horizon k. Hence, an inefficiency factor of 1 means that there is no autocorrelation among the draws.

[28]It is well known and can be shown through simulation that the p-values of a test statistic is uniformly distributed if the null hypothesis is true. With hundreds of parameters to test,

β_2(CC) can be seen in Figure 21. This plot is representative of those of the other parameters and do indeed resemble white noise.

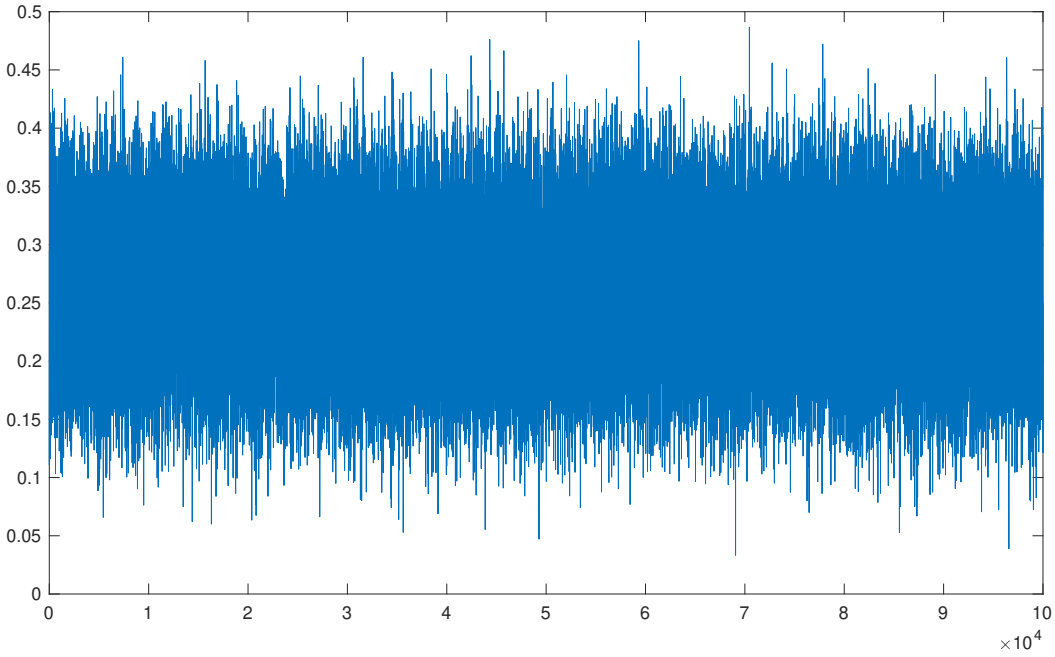

Figure 21: The trace of β_2(CC) from the Gibbs sampler. It is representative of other parameter draws in resembling white noise.

we are bound to observe p-values that are smaller than the conventional significance levels by pure chance.

ONLINE APPENDIX
— For Online Publication Only —

G Algorithm to solve for efficient equilibrium

We compute the equilibrium pay-offs V by numerically iterating the operator B, defined by equations (2) and (11), on a set of initial values W^* satisfying $V \subset W^*$. We describe the initialization procedure, then give an overview of the main steps in the computation, and finally detail the exact zero-finding algorithm.

Block A - Initialize

Select a tolerance ζ and parameters, see Table 1. Let

$$\overline{w}_0^j := \max_{\boldsymbol{q} \in Q^2} \pi(\boldsymbol{q}, x^j) + \pi(\boldsymbol{q}, x^j)$$
$$\underline{w}_0^j := \min_{\boldsymbol{q} \in Q} \pi(\boldsymbol{q}, x^j)$$

be initial values in W^* for all $j \in D$. Let the corresponding action profiles by $\overline{\boldsymbol{q}}_0^j$ and $\underline{\boldsymbol{q}}_0^j$. Let payoffs be $\overline{v}^j(\boldsymbol{q}, \overline{p}^j, W_t)$ and $\underline{v}^j(\boldsymbol{q}, \underline{\tau}^j, W_t)$ under action profile $\boldsymbol{q}$, continuation values W_t, trigger price $\overline{p}^j$, and transition probability $\underline{\tau}$. The set of feasible and in individually rational deviations are $\{\overline{q}^j + \epsilon, ..., q^{max}\}$ and $\{0, \epsilon, ..., \underline{q}^j - \epsilon\}$ in the reward- and punishment state respectively.

Block B - Iteration

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

1. **Compute candidate actions, trigger prices, and transition probabilities.** Search for each $j \in D$ for every candidate $\underline{\boldsymbol{q}}$, $\overline{\boldsymbol{q}}$ in Q^2 for a trigger $\overline{p}^j(\boldsymbol{q})$ and transition probability $\underline{\tau}^j$ satisfying incentive compatibility for all deviations with equality for at least one deviation

$$\overline{v}(\overline{\boldsymbol{q}}, \overline{p}^j, W_t) = \overline{v}(\overline{\boldsymbol{q}}', \overline{p}^j, W_t) \tag{22}$$

$$\underline{v}^j(\underline{\boldsymbol{q}}, \underline{\tau}^j, W_t) = \underline{v}^j(\underline{\boldsymbol{q}}', \underline{\tau}^j, W_t) \tag{23}$$

and for one $\underline{\boldsymbol{q}}$, $\overline{\boldsymbol{q}}$ in Q^2

$$\overline{v}(\overline{\boldsymbol{q}}, \overline{p}^j, W_t) \geq \overline{v}(\overline{\boldsymbol{q}}'', \overline{p}^j, W_t)$$

$$\underline{v}^j(\underline{\boldsymbol{q}}, \underline{\tau}^j, W_t) \geq \underline{v}^j(\underline{\boldsymbol{q}}'', \underline{\tau}^j, W_t)$$

for all other $\underline{\boldsymbol{q}}'' \neq \underline{\boldsymbol{q}}'$, $\overline{\boldsymbol{q}}'' \neq \overline{\boldsymbol{q}}'$ in Q^2. Gather candidate trigger prices, transition probabilities, and quantities in $\overline{\boldsymbol{p}}_t^j$, $\underline{\boldsymbol{\tau}}_j$, $\overline{Q}_t^j$, and $\underline{Q}_t^j$.

2. **Compute extreme continuation values.** Evaluate continuation values for every combination $(\overline{\boldsymbol{q}}_t^j, \overline{p}_t^j) \in \{\overline{Q}_t^j, \overline{\boldsymbol{p}}_t^j\}$ and $(\underline{\boldsymbol{q}}_t^j, \underline{\tau}_t^j) \in \{\underline{Q}_t^j, \underline{\boldsymbol{\tau}}_t^j\}$ and select the extreme continuation values:

$$\overline{w}_{t+1}^j = \max_{(\overline{\boldsymbol{q}}_t^j, \overline{p}_t^j) \in \{\overline{Q}_t^j, \overline{\boldsymbol{P}}_t^j\}} \overline{v}(\overline{\boldsymbol{q}}_t, \overline{p}_t^j, W_t) \tag{24}$$

$$\underline{w}_{t+1}^j = \min_{(\underline{\boldsymbol{q}}_t^j, \underline{\tau}_t^j) \in \{\underline{Q}_j, \underline{\boldsymbol{T}}_j\}} \underline{v}^j(\underline{\boldsymbol{q}}_t, \underline{\tau}_t^j, W_t) \tag{25}$$

Due to discounting, we have $\overline{w}_{t+1}^j \leq \overline{w}_t^j$ and $\underline{w}_{t+1}^j \geq \underline{w}_t^j$.

3. **Compute candidate transition probabilities.** Gather the $\overline{w}_{t+1}^j$ and $\underline{w}_{t+1}^j$ in d-dimensional vectors $\overline{w}_{t+1}$, $\underline{w}_{t+1}$ and define $W_{t+1} = \{\underline{w}_{t+1}, \overline{w}_{t+1}\}$.

H Appendix: Gibbs sampler

Our estimation procedure largely follows Hamilton (2016, p. 181). However, we include independent parallel sampling from two separate blocks in order to allow parts of the VAR system to remain constant across regimes. We use subscript $j \in \{1, 2\}$ to designate the two different regimes, but the sampler generalizes to allow for a any number of regimes $N \geq 2$. In our baseline specification, we do not allow the covariance matrix of the regime-switching block to be governed directly by the Markov chain. Rather it is conditional on the switching of the VAR coefficients.

For convenience, the two blocks are to be estimated are repeated below.

$$\Delta q_t = \mu_\bullet(s_t) + B_\bullet(s_t)\boldsymbol{y}_{\circ t} + \sum_{\ell=1}^{L} C_{\bullet\ell}(s_t)\boldsymbol{y}_{t-\ell} + v_t, \qquad v_t \sim N(0, \sigma_v^2(s_t))$$

$$\boldsymbol{y}_{\circ t} = \mu_\circ + \sum_{\ell=1}^{L} A_{\circ\ell}\boldsymbol{y}_{t-\ell} + \boldsymbol{e}_{\circ t}, \qquad \boldsymbol{e}_{\circ t} \sim N(0, \Sigma_\circ)$$

To further simplify the exposition we rewrite the equations into

$$\Delta q_t = \boldsymbol{G}_\bullet(s_t)\boldsymbol{z}_t + v_t$$
$$\boldsymbol{y}_{\circ t} = \boldsymbol{A}_\circ \boldsymbol{x}_{t-1} + \boldsymbol{e}_{\circ t}$$

where:

$$\boldsymbol{G}_\bullet(s_t) = \begin{bmatrix} B_\bullet(s_t) & \boldsymbol{C}_\bullet(s_t) \end{bmatrix}$$
$$\boldsymbol{z}_t = \left[\boldsymbol{y}'_{\circ t}, \boldsymbol{x}'_{t-1}\right]'$$
$$\boldsymbol{x}_{t-1} = \left[\boldsymbol{y}'_{t-1}, \boldsymbol{y}'_{t-2}, ..., \boldsymbol{y}'_{t-L}, 1\right]'$$

The collection of objects that we are interested in are the covariances $\boldsymbol{\sigma}_\bullet = \{\sigma_v^2\}$ and $\boldsymbol{\sigma}_\circ = \Sigma_\circ$, coefficients $\boldsymbol{\varphi}_\bullet = \{\boldsymbol{G}_{\bullet 1}, \boldsymbol{G}_{\bullet 2}\}$ and $\boldsymbol{\varphi}_\circ = \boldsymbol{A}_\circ$, Markov probabilities $p = \{\rho_1, \rho_2, p_{11}, p_{12}, p_{21}, p_{22}\}$ and finally the sequence, one for each date, of regime indicators $\mathcal{S} = \{s_1, ..., s_T\}$. The ρ_j are initial probabilities of being in regime j at $t = 1$ and are needed to initialize the Hamilton filter. We now describe how the conditional posterior distributions of these objects are obtained.

Initial conditions

To start the sampler we provide an initial value for each parameter. Given some $\mathcal{S}^0$ and our priors, the remaining objects may be drawn from the appropriate distributions as described below. In blocks A and B, we employ the OLS estimates of the parameters conditional on $\mathcal{S}^0$ to draw initial parameter values from these distributions. In theory the estimated posterior distribution will converge for any initial $\mathcal{S}^0$ that does not imply degeneracy, e.g. $\mathcal{S}^0 = \{j, j, ..., j\}$. Faster convergence may be achieved by providing an appropriate initial $\mathcal{S}^0$. To obtain $\mathcal{S}^0$, we estimate a simple single-equation Markov-switching model of the following form

$$\Delta \ln p_t = \mu(s_t) + \varphi(s_t)\Delta x_t + \varepsilon_t, \qquad \varepsilon_t \sim N(0, \sigma^2) \tag{26}$$

by Maximum Likelihood as described in Hamilton (2016). The data used to estimate this model is identical as in our VAR model framework. To generate $\mathcal{S}^0$, we compute the smoothed regime probabilities $\Pr(s_t = j|\Omega_T; \boldsymbol{\lambda})$ for each t. We classify each $s_t \in \mathcal{S}^0$ with the decision rule that if $\Pr(s_t = j|\Omega_T; \boldsymbol{\lambda}) \geq 0.5$, then $s_t = j$. To be confident that we will obtain the same posterior distribution irrespective of the initial conditions, we have run the Gibbs sampler using randomly generated $\mathcal{S}^0$ and obtain the same distribution.

Block A — $\sigma_v^2|\Omega_T, p, \boldsymbol{\varphi}_\bullet, \mathcal{S}$ and $\Sigma_\circ|\Omega_T, \boldsymbol{\varphi}_\circ$

Suppose that we have draws from a previous iteration of the sampler (or initial conditions) ℓ at hand and want to now draw sample $\ell + 1$. We first draw the variance σ_v^2 from an inverse-Wishart distribution with scale matrix $(I + H_\bullet^\ell)$ and $T + \eta$ degrees of freedom where

$$H_\bullet^\ell = \sum_{t=1}^{T} (\Delta q_t - \delta_{1t}^\ell \boldsymbol{G}_{\bullet 1}^\ell \boldsymbol{z}_t - \delta_{2t}^\ell \boldsymbol{G}_{\bullet 2}^\ell \boldsymbol{z}_t)(\Delta q_t - \delta_{1t}^\ell \boldsymbol{G}_{\bullet 1}^\ell \boldsymbol{z}_t - \delta_{2t}^\ell \boldsymbol{G}_{\bullet 2}^\ell \boldsymbol{z}_t)'$$

is the sum of residual outer products and δ_{jt} an indexing variable such that $\delta_{jt} = 1$ if $s_t = j$ and zero otherwise.[29] We then sample the regime-independent block covariance matrix $\Sigma_\circ$ from an inverse-Wishart distribution with scale parameter $(I + H_\circ^\ell)$ but where

$$H_\circ^\ell = \sum_{t=1}^{T} (\boldsymbol{y}_{\circ t} - A_\circ^\ell \boldsymbol{x}_{t-1})(\boldsymbol{y}_{\circ t} - A_\circ^\ell \boldsymbol{x}_{t-1})'.$$

At the conclusion of this step we have the variances $\boldsymbol{\sigma}_\bullet^{\ell+1}$ and $\boldsymbol{\sigma}_\circ^{\ell+1}$ of both blocks.

Block B — $\boldsymbol{G}_{\bullet j}|\Omega_T, p, \boldsymbol{\sigma}_\bullet, \mathcal{S}$ and $A_\circ|\Omega_T, \boldsymbol{\sigma}_\circ$

To draw the coefficients for the regime-switching block, we sample from a normal distribution given a precision matrix computed using the draw $\boldsymbol{\sigma}_\bullet^{\ell+1}$. In particular, ${\boldsymbol{g}_\bullet}_j^{\ell+1} \sim N(\hat{\boldsymbol{g}}_{\bullet j}, \boldsymbol{K}_g^{-1})$ where $\boldsymbol{K}_g = \boldsymbol{V}_g^{-1} + \delta_{jt}^\ell \boldsymbol{z}_t'(\sigma_v^{-2(\ell+1)})\boldsymbol{z}_t$, $\hat{\boldsymbol{g}}_{\bullet j} = \boldsymbol{K}_g^{-1}(\boldsymbol{V}_g^{-1}\boldsymbol{g}_{0j} + \delta_{jt}^\ell \boldsymbol{z}_t'(\sigma_v^{-2(\ell+1)})\Delta q_t)$, $\boldsymbol{V}_g$ is the prior covariance and $\boldsymbol{g}_{0j}$ the prior mean for regime j. For computational efficiency, we implement Algorithm

[29] As the inverse-Wishart distribution is multivariate analogue to the inverse-gamma distribution, we are in effect sampling from the inverse-gamma distribution as the regime-switching block only contains one equation.

1 described in Chan (2020). For the regime-independent block, we specify diffuse priors and sample $\boldsymbol{\alpha}_{\circ}^{\ell+1} \sim N(\hat{\boldsymbol{\alpha}}^{\ell+1}, \Sigma_{\circ}^{\ell+1} \otimes (\boldsymbol{x}_{t-1}\boldsymbol{x}_{t-1}')^{-1})$ where $\hat{\boldsymbol{\alpha}}^{\ell+1} = \text{vec}[\hat{A}_{\circ}^{\prime\ell+1}]$ which is obtained through OLS.

To obtain the VAR representations described in Equation 4, one for each regime, we apply the transformation described by Hamilton (2016, p. 175). As the final step, we evaluate the eigenvalues of the companion forms to ensure that we have obtained stationary systems. If at least one of the eigenvalues lies outside the unit circle, we abandon the draws for $\ell + 1$ and roll back the sampler to start over from Block A.

Block C — $p|\Omega_T, \boldsymbol{\sigma}_\bullet, \boldsymbol{\varphi}_\bullet, \mathcal{S}$

In this block, we draw the initial regime probabilities necessary to initialize the Hamilton filter $\{\rho_j\}_{j=1}^2$ and the columns of the transition probability matrix P, $\{p_{ij}\}_{j=1}^2$. The conditional posterior distribution for ρ_j is sampled from $D(\kappa_1 + \delta_{11}^\ell, \kappa_2 + \delta_{21}^\ell)$. We then proceed to count the number of jumps between regimes in the previous iteration as well as the number of times the regime did not change, $T_{ij}^\ell = \sum_{t=2}^T \delta_{it}^\ell \delta_{jt-1}^\ell$. The columns of P are then drawn from $D(\kappa_{i1} + T_{i1}^\ell, \kappa_{i2} + T_{i2}^\ell)$ for $i \in \{1, 2\}$. $D(\cdot)$ in this case refers to the Dirichlet distribution and the values separated by commas are the shape parameters.

Block D — $\mathcal{S}|\Omega_T, p, \boldsymbol{\sigma}_\bullet, \boldsymbol{\varphi}_\bullet$

In this final step, we apply the Hamilton filter (Hamilton 2016, p. 172) to obtain the sequences of probabilities $\Pr(s_t = j|\Omega_t, p, \boldsymbol{\sigma}_\bullet, \boldsymbol{\varphi}_\bullet)$. To obtain $\mathcal{S}^{\ell+1} = \{s_1, ..., s_T\}$, we begin iterating backwards by first drawing a number from $U(0, 1)$. If this draw is smaller than $\Pr(s_T = 1|\Omega_T, p, \boldsymbol{\sigma}_\bullet, \boldsymbol{\varphi}_\bullet)$ we set $s_T^{\ell+1} = 1$ and $s_T^{\ell+1} = 2$ otherwise. To continue for $T - 1$, we first compute the joint probability $\Pr(s_{T-1} = i, s_T = s_T^{\ell+1}|\Omega_T, p, \boldsymbol{\sigma}_\bullet, \boldsymbol{\varphi}_\bullet)$[30] with which we obtain the conditional probability $\Pr(s_{T-1} = i|s_T = s_T^{\ell+1}, \Omega_T, p, \boldsymbol{\sigma}_\bullet, \boldsymbol{\varphi}_\bullet)$. Using this probability, we make a new draw from $U(0, 1)$ and assign a new value for $s_{T-1}^{\ell+1}$ using the same rule as before. We repeat these steps for $t = T - 2, ..., 1$ to obtain the complete a draw from the conditional posterior distribution of $\mathcal{S}^{\ell+1}$. $\mathcal{S}^{\ell+1}$ will determine the value of δ_{jt} in iteration $\ell + 2$.

[30] See Hamilton (2016, p. 173) for details.

I Appendix: Identification

We show that the sign of covariance between quantities produced and prices determines the sign of the impact effect of OPEC supply shocks on the real price of oil. For any positive-definite matrix Σ, there exists a lower triangular matrix L such that $\Sigma = LL'$, the Cholesky decomposition of Σ. Suppose that Σ is a covariance matrix that is symmetric with positive entries on its diagonal.

$$\begin{aligned}\begin{pmatrix}\sigma_1^2 & \sigma_{12} & \sigma_{13}\\ \sigma_{21} & \sigma_2^2 & \sigma_{23}\\ \sigma_{31} & \sigma_{23} & \sigma_3^2\end{pmatrix} &= \begin{pmatrix}\ell_{11} & 0 & 0\\ \ell_{21} & \ell_{22} & 0\\ \ell_{31} & \ell_{32} & \ell_{33}\end{pmatrix}\begin{pmatrix}\ell_{11} & \ell_{21} & \ell_{31}\\ 0 & \ell_{22} & \ell_{32}\\ 0 & 0 & \ell_{33}\end{pmatrix}\\ &= \begin{pmatrix}\ell_{11}^2 & \ell_{11}\ell_{21} & \ell_{11}\ell_{31}\\ \ell_{11}\ell_{21} & \ell_{21}^2+\ell_{22}^2 & \ell_{21}\ell_{31}+\ell_{22}\ell_{32}\\ \ell_{31}\ell_{11} & \ell_{31}\ell_{21}+\ell_{32}\ell_{22} & \ell_{31}^2+\ell_{32}^2+\ell_{33}^2\end{pmatrix}\end{aligned}$$

From this relationship the elements of L may be expressed as functions of the variances and covariances. Three of the solutions are trivial:

$$\begin{aligned}\sigma_1^2 = \ell_{11}^2 &\Rightarrow \ell_{11} = \sigma_1,\\ \sigma_{21} = \ell_{11}\ell_{21} &\Rightarrow \ell_{21} = \frac{\sigma_{21}}{\sigma_1},\\ \sigma_{31} = \ell_{11}\ell_{31} &\Rightarrow \ell_{31} = \frac{\sigma_{31}}{\sigma_1}\end{aligned}$$

We can obtain the other three by substitution:

$$\begin{aligned}\sigma_2^2 = \ell_{21}^2+\ell_{22}^2 &\Rightarrow \ell_{22} = \sqrt{\sigma_2^2 - \frac{\sigma_{12}^2}{\sigma_1^2}},\\ \sigma_{23} = \ell_{21}\ell_{31}+\ell_{22}\ell_{32} &\Rightarrow \ell_{32} = \frac{1}{\sqrt{\sigma_2^2 - \frac{\sigma_{12}^2}{\sigma_1^2}}}\left(\sigma_{23} - \frac{\sigma_{12}\sigma_{13}}{\sigma_1^2}\right),\\ \sigma_3^2 = \ell_{31}^2+\ell_{32}^2+\ell_{33}^2 &\Rightarrow \ell_{33} = \sqrt{\sigma_3^2 - \frac{\sigma_{13}^2}{\sigma_1^2} - \frac{1}{\sigma_2^2 - \frac{\sigma_{12}^2}{\sigma_1^2}}\left(\sigma_{23} - \frac{\sigma_{12}\sigma_{13}}{\sigma_1^2}\right)^2}\end{aligned}$$

In the estimated reduced form model, the covariance between quantities produced and prices, σ_{31} is respectively negative- and positive in the pro- and countercyclical regime. In turn the impact responses, ℓ_{31}, of oil prices to OPEC supply shocks are respectively negative and positive. Changing the variable ordering, e.g. with OPEC production at the bottom will scale this covariance by a different factor.

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