

Procurement risk management in a petroleum refinery.

OLIVEIRA, F.S.

2023

© 2021 Decision Sciences Institute. This is the peer reviewed version of the following article: OLIVEIRA, F.S. 2023. Procurement risk management in a petroleum refinery. *Decision sciences* [online], 54(3), pages 277-296, which has been published in final form at <https://doi.org/10.1111/dec.12536>. This article may be used for non-commercial purposes in accordance with Wiley Terms and Conditions for Use of Self-Archived Versions. This article may not be enhanced, enriched or otherwise transformed into a derivative work, without express permission from Wiley or by statutory rights under applicable legislation. Copyright notices must not be removed, obscured or modified. The article must be linked to Wiley's version of record on Wiley Online Library and any embedding, framing or otherwise making available the article or pages thereof by third parties from platforms, services and websites other than Wiley Online Library must be prohibited.

Procurement Risk Management in a Petroleum Refinery

Fernando S. Oliveira

fernando.oliveira@auckland.ac.nz, Information Systems and Operations Management, The University of Auckland Business School, Sir Owen G Glenn Building, 12 Grafton Road, Auckland, 1010, New Zealand

We analyze a petroleum refinery's procurement strategy, explaining how risk management affects optimal sourcing from long-term, spot, and swap contracts. We use time series analysis to model the interaction between petroleum prices, transportation costs, and gross product worth. These models are then used to generate the scenarios incorporated in the stochastic program applied to compute the conditional value-at-risk. We prove the necessary and sufficient conditions for the optimal procurement and risk management strategies, and show that risk aversion can be better represented by the weighted average between expected profit and conditional value-at-risk, deriving the respective ISO-curves. We estimate that an increase in the degree of risk aversion decreases the use of swap contracts. Our model is applied to the analysis of a refinery based in Singapore. Using regression analysis, we show we cannot reject the hypothesis of a statistically significant relationship between the way Saudi Arabia prices the long-term contracts and the shape of the forward curve. We then study how risk aversion influences the procurement strategies, profitability and risk-exposure of the refinery. Finally, we analyze the pricing of long-term (forward) contracts by Saudi Arabia, and study how the country could benefit from a different pricing policy.

Subject classifications: Manufacturing; Stochastic Processes; Supply Chain Contracts and Incentives; Supply Chain Risk Management; Time Series Methods.

1. Introduction

The petroleum market is characterized by a very large supply chain based on the shipment of crude by sea and land, over long-distances, from producers to refineries in different regions of the planet. In this industry, long-term contracts between refineries and major producers are commonly used for several reasons: first, the relative sophistication of other financial products (such as swap contracts) that are associated with extreme risk exposures if not handled appropriately; second, the conservative nature of petroleum producer firms managed by government officials who may be made responsible for any perceived losses due to financial derivative contracts; and third, risk management is not a priority when such control may lead to the loss of the upper side associated with price volatility, which, in many instances, is created by coordinated curtailment of production.

In this article we discuss the different contracts available to a refinery and how these affect its expected profits and risk. In particular, we analyze the optimal proportion of crude oil to be procured from Saudi Arabia's long-term contracts, considering alternative suppliers in the spot market, and how crack spread swap contracts can be used by the refinery to increase profit and reduce risk.

Martinez-de-Albeniz and Simchi-Levi (2005), and Secomandi and Kekre (2014) analyzed a similar problem in which a firm uses long-term, spot, and option contracts. Our study differs from theirs in significant ways: a) we consider risk aversion using conditional value-at-risk in a static model; b) we focus on the interaction between long-term contracts (with an index price) and the spot market; c) using time series analysis, we study the properties of petroleum prices (and those of its products) and derive an equation to describe the pricing of Arab Light forward contracts; and finally, d) we analyze how alternative contract designs affect the producers' profits.

We focus our analysis on a Singapore based refinery, which can buy its crude from Saudi Arabia, from South East Asia (Indonesia in our case), and from West Africa (represented by Angola). As discussed in Zhang et al. (2014), when considering the price of raw materials, there are different types of contract, such as fixed price, cost reimbursement, procurement control, index-linked payment and relational. The Arab Light long-term contract is an index-linked payment, as the price for each shipment depends on the actual spot price of the Oman-Dubai crude. This contract represents financial hedging by Saudi Arabia, as they avoid being locked into a long-term contract with possible low prices. From the refinery's perspective, it reduces quantity risk as it ensures access to crude oil from a reliable source and product availability (without which the refinery cannot operate). For this reason, in the problem analyzed in this article, the buyer (the refinery) allocates proportions of its procurement to long-term and to spot contracts, based not on the speed of the supply (which will be faster from Indonesia, in our case) but solely on price risk. The ability of refineries to manage such risk is directly connected to the possibility of passing petroleum price increases onto the final consumers.

Therefore, long-term contract price indexing reduces risk for both parties. As, in this contract the income stream is not fixed, some would argue there is no price hedging. This is clearly true. However, by not fixing the price, the refinery and the producer avoid some of the potential losses associated with price risk: when the forward price is fixed the refinery suffers a financial loss if the spot price, at the delivery month, is less than the forward price, but receives a profit otherwise. By using a forward price indexed to the spot price, the potential for financial loss is no longer a concern for the refinery.

From a policy perspective, an important research question analyzed in this paper is the impact of different forward contract designs on the refinery's procurement strategy. Additionally, we also

study how contract design affects the producer’s profitability as, in order to establish the credibility of a given contract, we need to show that it is profitable for the producer as well.

The methodologies used in our analysis are twofold. First, we apply time series methods to analyze the pricing and refining operational data, showing that prices are non-stationary. We describe how to use this information in the parameterization of the procurement problem. Second, the refinery optimization model is solved using static stochastic programming, applying Monte Carlo simulation to generate scenarios. In this procurement problem, the forward contracts are held for 12 months and involve monthly spot decisions. The static model is an adequate approach as, each month, the refinery aims to have total purchases in the forward and spot contracts equal production requirements. For this reason, the spot decisions in one month have no impact on subsequent months and therefore each month starts from the same basis, i.e., the quantity bought forward in Arab Light contracts.

In summary, the main contribution of this article is to analyze how different types of petroleum sources (taking into consideration how quality influences production output), long-term and crack spread swap contracts, together with risk aversion, influence the procurement plan of a refinery. We analyze how the degree of risk aversion influences the refinery’s procurement strategy, using conditional value-at-risk (*CVaR*) in a single-period problem (a typical year). We use a static *CVaR* as the risk measure as it is better suited to representing the uncertainty associated with operational risk and it is a *coherent risk measure*, as defined by Artzner et al. (1999), whereas the variance or standard deviation are not. Moreover, we estimate empirically the indexing mechanism used by Saudi Arabia to price the Arab Light contracts, based on the perceived state of the market, and analyze how the producer benefits from using a different indexing mechanism. This is modeled using regression analysis based on Oman-Dubai petroleum price data.

This article is organized as follows: in section 2 we review the literature on procurement and risk management and justify the use of conditional value-at-risk; in section 3 we propose a stochastic programming model to analyze the refinery’s procurement decisions, including risk aversion, and apply duality theory to describe the properties of the refinery’s optimal procurement; in section 4 we develop a regression analysis of Saudi Arabia’s current pricing policy and describe the main computational results of our analysis as applied to a refinery based in Singapore; and section 5 concludes the article.

2. Risk Analysis in the Refinery’s Procurement Problem

Hong et al. (2018) recognized five major sources of risk (demand, price, yield, lead time, and disruption) and seven major future research directions (multi-item procurement, risk awareness, risk dependency, lead time risk, yield risk, reactive management, and risk hedging using financial

instruments). Moreover, Pournader et al. (2020) identified as major research topics the increased interest on behavioral and sustainability related risk management practices. In this article we address the procurement problem, considering price risk with financial hedging.

Given a set of spot and forward petroleum prices, a refinery's major decision is the procurement of the petroleum required to meet demand. Typically, the production structure of the refinery is optimized for a given type of crude, but it can also refine other types, although less efficiently. The decision about which type of petroleum to refine depends on the margins per barrel the refinery can make from the petroleum bought from different sources. This margin is a function of the *gross product worth* (GPW), price, refining and transportation cost of a barrel of petroleum, and is source dependent. The GPWs are the weighted average value of the products refined from a barrel of petroleum, in which the weights are equal to the proportion of the different products refined: they represent the requirement of adapting the operations of the refinery to capture the operational flexibility (and setup costs) of working with different types of petroleum.

The supply disruption literature (e.g., Vakharia and Yenipazarli 2009, Bilser et al. 2011, Kumar and Park 2019) considers different types of uncertainty, including supplier (e.g., Huang et al. 2016, Goldschmidt et al. 2020, Li et al. 2020), demand and cost (or price), all of which are important in the study of supply chain design as they influence the trade-off between cost and reliability, e.g., Yildiz et al. (2016). Demand uncertainty is associated with consumer behavior. In terms of cost uncertainty, transportation may be an important component. Obviously, transportation costs depend on the distance, volumes involved and speed used. However, in petroleum markets the transportation cost is, in general, a small part of the product cost, and it is the uncertainty associated with the petroleum price that is crucial. The literature tends to analyze how supply and cost uncertainty relate to supplier reliability. Supply uncertainty arises in situations in which the buyer is not certain that the full order will be delivered on time due to quality issues, delays, and disruptions.

The models used to capture uncertainty include: a) the random yield model (to capture operational risks), in which the final delivered quantity is a stochastic function of the order volume, e.g., Burke et al. (2007), Tajbakhsh et al. (2010), Yan et al. (2012); b) an all or nothing model of the supply disruptions preventing, with some small probability, the order from being met, typically due to external factors, and assuming the producers have different degrees of reliability, e.g., Costantino and Pellegrino (2010), Meena et al. (2011), Sawik (2014); and c) models that aim to capture both operational and disruption risks, e.g., Sawik (2011), Torabi et al. (2015).

The methodologies used to study supplier selection, accounting for uncertainty, are: a) recourse programs, in which there are two stages and in which the first stage decisions may be adjusted in the second stage, e.g., Xu and Nozick (2009), Torabi et al. (2015); b) chance constrained

models based on relaxation of the deterministic constraints that the policy maker is allowed to violate, with some (small) probability, e.g., Bilser et al. (2011), Scott et al. (2015); c) single-stage models that assume risk-neutrality, e.g., Burke et al. (2007), Yan et al. (2012), or risk aversion by incorporating volatility, e.g., Hong and Lee (2013); and d) conditional value-at-risk models that capture risk aversion in the buyer's behavior by allowing the buyer to minimize the expected loss that may happen with a very small probability in the presence of very unlikely, but extreme events, e.g., Sawik (2011, 2014), Nejad et al. (2014).

We focus on the CVaR, which has been used, for example, to solve the portfolio optimization problem (e.g., Rockafellar and Uryasev 2000, Staino and Russo 2020), to analyze production decisions by risk electricity producers (e.g., Conejo et al. 2010), to choose technology adoption in fleet management (e.g., Ansariipoor et al. 2016) and commodity and energy operations (Devalkar et al. 2018, Oliveira and Ruiz 2020), and to study the news-vendor problem (e.g., Chen et al. 2009, 2015, Yang et al. 2018). The value-at-risk (VaR) is a quantile that is a function of the proportion of observations in the tail, α , and actions, x . The CVaR is the expected profit conditional on being less than VaR (Definition 1).

DEFINITION 1 (CVAR). Let α be the percentage of observations in the tail of the loss function, such that the observed loss is larger than VaR; x stand for the decision variables and s for the scenario index; $p(s)$ be the probability of scenario s and $\pi(s, x)$ be the profit function: $CVaR_\alpha(x) = \frac{1}{\alpha} \int_{\pi(s, x) \leq VaR_\alpha(x)} \pi(s, x) p(s) ds$.

In the linear program (1), we summarize the formulation to compute the CVaR (and the VaR) of a profit function, based on Rockafellar and Uryasev (2000).

$$\underset{\mu, z_s, x}{\text{Maximize}} \quad \mu - \frac{\sum_s z_s}{S\alpha} \tag{1a}$$

subject to:

$$z_s \geq \mu - \pi(s, x) \quad \perp \quad \lambda_s \quad \forall s \tag{1b}$$

$$z_s \geq 0 \quad \forall s \tag{1c}$$

In the linear program (1): a scenario is represented by s ; S is the number of scenarios; α is the probability of the profit function being on the tail and z_s is the value of this tail in scenario s . The tail, which is the difference between the value at risk, μ , and the actual profit observed, $\pi(s, x)$, is always non-negative and, accordingly, the program is constrained to ensure this is always the case. For this reason, the constraints (1b – 1c) determine that for each scenario only the values in the tail of the possible profits less or equal to μ are considered. The λ_s associated to constraint (1b) is the respective dual variable (or Lagrangian multiplier). This formulation of the CVaR is basically

the same as described in Chen et al. (2009), Chen et al. (2015), and Yang et al. (2018), but all require continuity and monotonic profit functions. Additionally, Chen et al. (2009) need twice-differentiability of the profit function, so as to derive the analytical results. Chen et al. (2015) work with continuous bounded distributions, and assume the stochastic variable follows a uniform distribution in some of the analytical results in order to derive a closed-form solution. Moreover, in their examples, they also assume knowledge of the profit function distribution. The formulation in Yang et al. (2018), even though general, requires knowledge of the stochastic distribution function in order to be applied: in their examples, demand is assumed to be known and to follow a uniform distribution. Instead, as in Rockafellar and Uryasev (2000) and Ansariipoor et al. (2016), among others, we approximate the CVaR with scenarios, as it has been shown to converge to the optimal policy and does not require any knowledge about the profit distribution, which may be non-differentiable and not continuous. We show that, even in this case, we can still derive interesting analytical results.

Rockafellar and Uryasev (2000) discuss in detail the conditions under which (1) ensures that the proportion of observations in the tail is α . An alternative way of understanding this result, is to analyze the dual of the linear program (1), which we prove in the Appendix to be the program (2). For a simple introduction to duality theory in linear programming see, e.g., Bradley et al. (1977, ch. 4). In this problem the objective function is now dependent on λ_s only and, most interestingly, it is a simple *minimization of the expected profit* (as the λ_s add up to 1), subject to the constraint that the weight (λ_s) place on each scenario s cannot exceed $\frac{1}{S\alpha}$. For this reason, by minimizing the expected profit, the program places all its weight on the $S\alpha$ scenarios with the lowest profit. It is worth noting that due to the transformation of the objective function, the dual problem is, in general, non-linear. For this reason, the less intuitive, but linear primal, is used instead to compute the CVaR. Moreover, in general α is an upper-bound on the proportion of scenarios in the tail, as the constraint may be trivially met with fewer scenarios on the tail.

$$\underset{\lambda_s}{\text{Minimize}} \quad \sum_s \pi(s, \lambda_s) \lambda_s \quad (2a)$$

subject to:

$$\sum_s \lambda_s = 1 \quad (2b)$$

$$\lambda_s \leq \frac{1}{S\alpha} \quad \forall s \quad (2c)$$

$$\lambda_s \geq 0 \quad \forall s \quad (2d)$$

As the CVaR is a coherent risk measure, as defined by Artzner et al. (1999), i.e., it satisfies the axioms for: convexity, monotonicity, translation equivariance, and positive homogeneity. Note that

the variance and the standard deviation of the profit are not coherent. The CVaR is monotonic as higher losses lead to higher risk; this is not the case when using the variance or the standard deviation. For example, a large but certain loss has zero variance (and zero standard deviation). The CVaR has the property of translation equivariance, as if the loss increases (decreases) by a constant value, then the risk increases (decreases) by the same amount. This is not the case with the variance, which therefore does not have the property of translation equivariance. The CVaR is a positive homogeneous risk measure as the risk increases in direct proportion to risk exposure while the variance is not positive homogeneous, but the standard deviation is. Finally, the CVaR is convex, i.e., the linear combination of two variables decreases risk, as is also the case for the variance and the standard deviation.

Another important advantage of using CVaR as a risk measure is that, when following the approach proposed in this article, it is constructed from the bottom-up without having to impose a statistical distribution on the profit function, or having any knowledge of its nature. We study the random variables affecting the profit function and model them. From the interaction between these variables we then obtain, at the same time as solving the risk-management problem, the final description of the utility function, which includes both the expected profit and the risk components. Alternatively, when using the mean-variance approach, we are required to know the different components of the profit function (the expected profit and variance for each of them) and how they interact (the profit covariance matrix). The mean-variance model uses a top-down approach, which is much more demanding in terms of information required, and better suited for simpler problems.

Moreover, another advantage of the CVaR is that it can be readily used when uncertainty is not symmetrical, and when the underlying risk distribution is unknown, because, as used in this article, this technique is scenario based. Whereas the VaR is solved using chance-constrained programming (e.g., Charnes and Cooper 1959, Chen et al. 2010) and therefore is possibly non-convex, the CVaR, on the other hand, is solved by convex programming (e.g., Shapiro et al. 2009).

3. The Refinery's Procurement Problem

We start the analysis of the refinery's procurement problem in section 3.1 by describing the profit function and how to incorporate the CVaR into the objective function. Then we provide the linear programming formulation to solve the problem. In section 3.2 we apply duality theory to derive the analytical results on the impact of risk aversion on the optimal procurement policy. In section 3.3 we conclude the analytical results by studying the relationship between risk aversion and procurement from spot, long-term and crack spread swap contracts.

3.1. Integrating Expected Profit Maximization with Risk Aversion

The refinery's procurement problem has three components: a) a long-term (forward) contract covering the 12 months in the planning horizon; b) the spot contracts that take into account the expected demand and spot prices for each month; and c) the crack spread swap contracts.

The planning horizon includes T months (t is the index for a given month). Let j stand for the petroleum source index. The stochastic parameters are g_{jts} , w_{jts} , and η_{jts} , representing, respectively, the gross product worth, the petroleum price, and the transportation cost from source j , in month t , and scenario s . The refining marginal cost is r per barrel. The GPWs are a function of the actual technical setup of the refinery (optimized to operate a given type of petroleum).

The refinery decides the amount of petroleum (q_{jts}) to purchase from each source j , in month t , and scenario s . There are two possible sources: the spot market, in which the quantity traded changes every month depending on the requirements of the refinery, and the long-term contract, which is represented by q_a , the quantity to be acquired every month during the duration of the contract. The price paid in any given month for this forward contract is indexed to the spot price.

Moreover, the refinery also decides how many crack spread swap contracts to purchase (k_h), all of which are financially settled. Let g_h and w_h , respectively, stand for the gross product worth and price of the benchmark crude h , e.g. Oman-Dubai, in which these swap contracts are traded. The refinery has the possibility of fixing the crack spreads 12 months ahead to hedge the risk of exposure to product and crude price uncertainties. All hedging decisions are taken when the contract is signed and cannot be modified during its life (one year).

The refinery's profit function in scenario s , is represented in equation (3).

$$\begin{aligned} \pi(s, q_{jts}, q_a, k_h) = & \sum_{j \neq a, t} (g_{jts} - w_{jts} - r - \eta_{jts}) q_{jts} + \sum_t (g_{ats} - w_{ats} - r - \eta_{ats}) q_a \\ & + \sum_t ((g_h - g_{hts}) - (w_h - w_{hts})) k_h \end{aligned} \quad (3)$$

Equation (3) has three main components: the first two represent the profits gained from refining activities, $(g_{jts} - w_{jts} - r - \eta_{jts}) q_{jts}$ and $\sum_t (g_{ats} - w_{ats} - r - \eta_{ats}) q_a$, which include the refining margins (gross product worth minus the petroleum price, minus the refining cost per barrel, r), to which we subtract the transportation cost per barrel from the spot market (η_{jts}) and long-term contracts (η_{ats}), respectively; the third stands for the profit received from the hedging actions $((g_h - g_{hts}) - (w_h - w_{hts})) k_h$, which include the gains made by fixing the product prices minus any costs from fixing the petroleum prices. An alternative representation of the swap profit component is $((g_h - w_h) - (g_{hts} - w_{hts})) k_h$, which represents the difference between the swap crack spread and the observed crack spread in month t and scenario s . Note that the refinery registers a profit in

the swap contracts when the crack spread in month t and scenario s , $g_{hts} - w_{hts}$, *decreases*, and has a loss when this same crack spread *increases*.

The problem is solved numerically using the stochastic linear program (4). The objective of the refinery is to maximize ϕ , the weighted average of the expected profit and CVaR, as represented by function (4a), where β is the weight allocated to the expected profit and $1 - \beta$ is the CVaR weight.

$$\underset{\mu, z_s, q_{jts}, q_a, k_h}{\text{Maximize}} \quad \phi = \beta \sum_s \frac{\pi(s, q_{jts}, q_a, k_h)}{S} + (1 - \beta) \left(\mu - \frac{\sum_s z_s}{S\alpha} \right) \quad (4a)$$

subject to:

$$\sum_{j \neq a} q_{jts} + q_a = \zeta \quad \forall ts \quad \xi_{ts} \quad (4b)$$

$$k_h - \zeta \leq 0 \quad \perp \quad \gamma \quad (4c)$$

$$-z_s + \mu - \pi(s, q_{jts}, q_a, k_h) \leq 0 \quad \forall s \quad \perp \quad \lambda_s \quad (4d)$$

$$-z_s \leq 0 \quad \forall s \quad \perp \quad \delta_s \quad (4e)$$

$$-q_{jts} \leq 0 \quad \forall j \neq a, \forall ts \quad \perp \quad \sigma_{jts} \quad (4f)$$

$$-q_a \leq 0 \quad \perp \quad \chi \quad (4g)$$

$$-k_h \leq 0 \quad \perp \quad \theta \quad (4h)$$

For every month t in scenario s , the refinery contracts all its requirements guaranteeing $\sum_{j \neq a} q_{jts} + q_a = \zeta$, constraint (4b), to ensure all the quantities are, in effect, proportions of the purchase requirements at any given time. The ζ is the refinery's available capacity, which we standardize to 1. (We assume that the refinery is expecting to take the full load. This assumption represents how the refinery works in reality, as it has enough storage capacity to always take the full load of petroleum purchased.) All procured quantities are non-negative, as declared by constraints (4f) and (4g). The proportion of capacity covered by the swap contracts is k_h , where $0 \leq k_h \leq \zeta$, as imposed by constraints (4c) and (4h).

Constraints (4d) and (4e) are used to ensure the tails, z_s , are non-negative. We use these two constraints to linearize the maximization operator, $z_s = \max(\mu - \pi(s, q_{jts}, q_a, k_h), 0)$, which imposes that only non-negative tails are accounted for computing the CVaR, following Rockafellar and Uryasev (2000). The z_s represents the level of profit in the tail, for each scenario s . In order to get an accurate measure of the CVaR, a very high number of scenarios needs to be generated. As the CVaR is assessing the cost associated with events that are very unlikely. This formulation of the problem has been used in Shapiro et al. (2009, p. 272) and Ansariipoor et al. (2016), Oliveira and Ruiz (2020), among others, to model different degrees of risk aversion: by changing the β from 0

to 1 you move from risk aversion to risk-neutrality. The level of risk aversion can also be set using α , as analyzed in Chen et al. (2015), Devalkar et al. (2018) and Yang et al. (2018). Chen et al. (2015, p. 80) discuss how by considering a negative α , we can model risk-seeking behavior. Equally, the formulation (4) can be modified to address risk-seeking behavior by changing the tail from the left to the right of the profit function, i.e., (4d) will become $z_s - \mu - \pi(s, q_{jts}, q_a, k_h) \leq 0$ and the objective function will be $\phi = \beta \sum_s \frac{\pi(s, q_{jts}, q_a, k_h)}{S} + (1 - \beta) \left(\frac{\sum_s z_s}{S\alpha} - \mu \right)$. However, even though easy to model, such behavior would not represent how a refinery behaves in the real world, as discussed in the Introduction.

Finally, in the linear program (4), the \perp represents the complementarity condition, which means that when the respective dual variable is positive the constraint is binding (and if the constraint is not binding, then the dual variable is equal to zero). See for example Bradley et al. (1977, ch. 4) for a general introduction to duality in linear programming. ξ_{ts} is the dual variable of the equality constraint (4b). The dual variables of the constraints (4c), (4d), (4e), (4f), (4g), and (4h) are, respectively, γ , λ_s , δ_s , σ_{jts} , χ and θ , all of which are discussed in sections 3.2 and 3.3.

3.2. How Risk Aversion Affects Optimal Procurement

In this section we derive analytical results on the impact of risk aversion on the refinery's optimal policy by using duality theory. The necessary and sufficient conditions for the uniqueness of the optimal solution are summarized in Proposition 1. The proofs of all the propositions are provided in the Appendix. The interpretation of Proposition 1 is the topic of the rest of this section (conditions A) and B)) and section 3.3 (conditions C), D) and E).) Conditions F) and G) were analyzed in section 3.1. Conditions H) to L) are the complementarity constraints of problem (4).

PROPOSITION 1 (Necessary and Sufficient Conditions).

$$A: \sum_s \lambda_s = 1 - \beta$$

$$B: \lambda_s + \delta_s = \frac{1 - \beta}{S\alpha} \quad \forall s$$

$$C: \left(\frac{\beta}{S} + \lambda_s \right) (g_{jts} - w_{jts} - r - \eta_{jts}) - \xi_{ts} + \sigma_{jts} = 0 \quad \forall j \neq a, \forall t, \forall s$$

$$D: \sum_{ts} \left(\frac{\beta}{S} + \lambda_s \right) (g_{ats} - w_{ats} - r - \eta_{ats}) - \sum_{ts} \xi_{ts} + \chi = 0$$

$$E: \sum_{ts} \left(\frac{\beta}{S} + \lambda_s \right) ((g_h - g_{hts}) - (w_h - w_{hts})) - \gamma + \theta = 0$$

$$F: \sum_{j \neq a} q_{jts} + q_a = \zeta \quad \forall t, \forall s$$

$$G: 0 \leq z_s - \mu + \pi(s, q_{jts}, q_a, k_h) \perp \lambda_s \geq 0 \quad \forall s$$

$$H: 0 \leq z_s \perp \delta_s \geq 0 \quad \forall s$$

$$I: 0 \leq q_{jts} \perp \sigma_{jts} \geq 0 \quad \forall jts$$

$$J: 0 \leq q_a \perp \chi \geq 0$$

$$K: 0 \leq k_h \perp \theta \geq 0$$

$$L: 0 \leq \zeta - k_h \perp \gamma \geq 0$$

The dual variables λ_s and δ_s are essential to the understanding of the optimality conditions. From *Proposition 1.A*), $\sum_s \lambda_s = 1 - \beta$, it is evident that the sum of the dual variables associated with non-negative tails, i.e., to scenarios labeled as posing a significant risk, is equal to the weight put on CVaR. Consequently, these dual variables allow us identify the riskier scenarios. Therefore, β influences the weight put on CVaR and on the different scenarios.

From *Proposition 1.B*), $\lambda_s + \delta_s = \frac{1-\beta}{S\alpha}$, it follows that in a scenario s where the tail is positive, i.e., $z_s > 0$, the $\delta_s = 0$ and $\lambda_s > 0$, for this reason, only λ_s is relevant for calculation of the CVaR. On the other hand, in a scenario s in which the profit equals the value-at-risk, i.e., $\mu = \pi(s, q_{jts}, q_a, k_h)$, the tail is zero, i.e., $z_s = 0$, δ_s and λ_s are both non-negative (it may happen that they are both positive or zero, simultaneously).

Moreover, in order to fully understand the implications of risk aversion for the way the refinery optimizes its decisions on the spot, long-term and swap contracts, we need to explain the meaning of the $\frac{\beta}{S} + \lambda_s$ term appearing in each one of the necessary and sufficient conditions associated with these contracts. In Proposition 2 we prove that $\sum_s (\frac{\beta}{S} + \lambda_s) = 1$. Each one of the terms $\frac{\beta}{S} + \lambda_s$ represents the risk-adjusted probability a refinery applies to compute the risk-adjusted expected profits of the different contracts. The $\frac{1}{S}$ is the risk-neutral probability that is transformed by being multiplied by β and added to λ (the shadow variable of binding profit tails) to represent the risk-adjusted (subjective) probability. Whereas the risk-neutral refinery assigns a probability $\frac{1}{S}$ to each scenario s , the risk-averse refinery assigns probability $\frac{\beta}{S} + \lambda_s$ to scenarios in which $z_s > 0$. Nonetheless, as proved in Proposition 2, these risk-adjusted probabilities also add up to 1.

PROPOSITION 2. *For a risk-averse refinery, with $\beta < 1$, for each scenario s , the risk-adjusted probability $\frac{\beta}{S} + \lambda_s$ is such that $\sum_s (\frac{\beta}{S} + \lambda_s) = 1$.*

It is worth emphasizing that $\frac{\beta}{S} + \lambda_s$ are subjective probabilities, which depend on the risk profile of the refinery, as captured both by the β and α (which together determine λ). Consequently, in order to understand the determinants of the objective function, we need to better explain the interaction between α and β in determining λ , as described by the Iso- λ curves derived in Proposition 3. (Iso- δ curves can be derived similarly for all scenarios s such that $z_s = 0$.)

PROPOSITION 3. *Let $0 < \alpha, \alpha^0 \leq 1$ and $0 \leq \beta \leq 1$. Assume that a refinery's risk profile is completely defined by a $\beta^0 = 0$ and an α^0 . Let state s be such that $z_s > 0$ and therefore $\lambda_s^0 = \frac{1}{S\alpha^0}$ and $\delta_s^0 = 0$. a) Every λ_s^0 can be alternatively obtained by the Iso- λ curve $\alpha = \frac{1}{S\lambda_s^0} - \frac{1}{S\lambda_s^0}\beta$. b) Additionally let $1 \geq \alpha > 1 - \beta \geq 0$, then there is a $\lambda = \frac{1-\beta}{S\alpha}$ that cannot be replicated by a refinery with $\beta^0 = 0$.*

A first interesting insight from Proposition 3 is that λ is independent of the specific scenario s , and it is a linear function of α and β , as represented in the Iso- λ curves. Moreover, it follows

directly from Proposition 3 that when $\alpha > 1 - \beta$, we can still obtain a solution for problem (4), but such a situation cannot be represented by using only α . This would represent half of the domain of the Iso- λ curves depicted in Figure 1, in which the number associated with each ISO-curve is the respective λ .

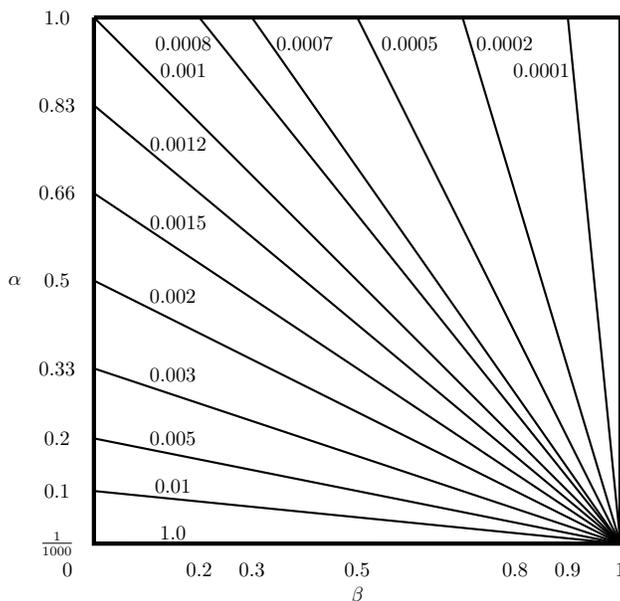


Figure 1 Iso- λ Curves – Example Assuming $S = 1000$.

In Figure 1, λ increases counter-clockwise from $\lambda = 0.0001$ in the line with coordinates $(\beta = 1, \alpha = 0) - (\beta = 0.9, \alpha = 1)$, to $\lambda = 1.0$ in the line with coordinates $(\beta = 1, \alpha = 0) - (\beta = 1, \alpha = \frac{1}{1000})$. In this example, the ISO-curves in the range $0.001 > \lambda > 0.0$ are only representable by simultaneously using α and β , as proved in Proposition 3.

An interesting property of this contour map is that all the Iso- λ curves pass through the same point $(\beta = 1, \alpha = 0)$, in which λ assumes all possible values from infinitesimally small to infinitely high. This point is a singularity as α cannot be equal to zero, except when the refinery is risk-neutral ($\beta = 1$), as in this case the actual value of α and λ no longer feature in the objective function.

3.3. The Interaction between Spot, Long-Term and Crack Spread Swap Contracts

We now consider the behavior of a refinery when actively trading long-term contracts and explain the interaction between the three profit sources: spot, long-term and swap contracts, as described in Proposition 1, conditions C) to E).

It follows from Proposition 2 that: a) the risk-adjusted expected profit of the spot contract (for all triples jts , such that $j \neq a$) is $(\frac{\beta}{S} + \lambda_s)(g_{jts} - w_{jts} - r - \eta_{jts})$; b) the risk-adjusted expected

profit of the long-term contract is $\sum_{ts} \left(\frac{\beta}{S} + \lambda_s\right) (g_{ats} - w_{ats} - r - \eta_{ats})$; and c) the risk-adjusted expected profit of the crack spread swap contract is $\sum_{ts} \left(\frac{\beta}{S} + \lambda_s\right) ((g_h - g_{hts}) - (w_h - w_{hts}))$. Each one of these profit functions is a weighted average between the expected profit and the CVaR of the respective contract. In fact these three conditions decompose the the overall objective function (4a) in its three components, imposing the optimality conditions on each of them.

In *Proposition 1.C*), for the spot contract j at time t and scenario s , equation $\left(\frac{\beta}{S} + \lambda_s\right) (g_{jts} - w_{jts} - r - \eta_{jts}) - \xi_{ts} + \sigma_{jts} = 0$ describes the optimality condition. The term $g_{jts} - w_{jts} - r - \eta_{jts}$ is the profit per barrel. ξ_{ts} and σ_{jts} are dual variables. When the refinery purchases from source j at time t and scenario s , the ξ_{ts} stands for the risk-adjusted profit margin per barrel and $\sigma_{jts} = 0$. When the refinery does not purchase from source j at time t and scenario s , $q_{jts} = 0$ and $\sigma_{jts} = -\left(\frac{\beta}{S} + \lambda_s\right) (g_{jts} - w_{jts} - r - \eta_{jts}) > 0$ is the *risk-adjusted loss* per barrel if purchased from source j .

In *Proposition 1.D*), for the long-term contract, the optimality condition is described by equation $\sum_{ts} \left(\frac{\beta}{S} + \lambda_s\right) (g_{ats} - w_{ats} - r - \eta_{ats}) - \sum_{ts} \xi_{ts} + \chi = 0$. The term $\sum_{ts} \left(\frac{\beta}{S} + \lambda_s\right) (g_{ats} - w_{ats} - r - \eta_{ats})$ is the risk-adjusted expected profit per barrel, which is equal to the sum of the dual variables $\sum_{ts} \xi_{ts}$, as $\chi = 0$ when the refinery purchases long-term contracts. This means that $\sum_{ts} \xi_{ts} = \sum_{ts} \left(\frac{\beta}{S} + \lambda_s\right) (g_{ats} - w_{ats} - r - \eta_{ats})$ stands for the weighted average of profit per barrel and CVaR (from long-term contracts) a risk-averse refinery expects to earn (in which the weights are subjectively determined by the degree of risk aversion as discussed in section 3.2).

In *Proposition 1.E*), the optimality condition for the crack spread swap contract is $\sum_{ts} \left(\frac{\beta}{S} + \lambda_s\right) ((g_h - g_{hts}) - (w_h - w_{hts})) - \gamma + \theta = 0$. The *risk-adjusted expected profit per barrel* from swap contracts is described by the expression $\sum_{ts} \left(\frac{\beta}{S} + \lambda_s\right) ((g_h - g_{hts}) - (w_h - w_{hts}))$, in which the weights are the same as in the long-term contracts. When $k_h = \zeta$ and $\theta = 0$, $\gamma = \sum_{ts} \left(\frac{\beta}{S} + \lambda_s\right) ((g_h - g_{hts}) - (w_h - w_{hts}))$ is the risk-adjusted expected profit per barrel from swap contracts the refinery would receive if it was to purchase more contracts than its refining capacity. On the other hand, when $k_h = 0$ and $\gamma = 0$, $\theta = -\sum_{ts} \left(\frac{\beta}{S} + \lambda_s\right) ((g_h - g_{hts}) - (w_h - w_{hts})) > 0$ stands for the resulting risk-adjusted expected loss per barrel for the refinery trading swap contracts.

In Proposition 4 we describe the interaction between spot and long-term contracts.

PROPOSITION 4. *Let a refinery procure some of the petroleum using long-term contracts and therefore $q_a > 0$ and $\chi = 0$. Under these conditions $\sum_{ts} \left(\frac{\beta}{S} + \lambda_s\right) (g_{ats} - w_{ats} - r - \eta_{ats}) - \frac{\sum_{j \neq a, ts} \left\{ \left(\frac{\beta}{S} + \lambda_s\right) (g_{jts} - w_{jts} - r - \eta_{jts}) \right\}}{J-1} = \frac{\sum_{j \neq a, ts} \sigma_{jts}}{J-1}$.*

From Proposition 4 it follows that the average value of the dual variables associated with procurement in the spot market (σ_{jts} , for $j \neq a$) is equal to the difference between the risk-adjusted expected marginal profit from long-term contracts and the average of the risk-adjusted expected

marginal profit from the spot contracts over the entire planning horizon. When $\frac{\sum_{j \neq a, ts} \sigma_{jts}}{J-1} = 0$ the refinery procures simultaneously in all spot markets and the risk-adjusted expected profit per barrel from long-term contracts equals the average of the risk-averse expected profit per barrel from spot contracts.

On the other hand, when $\frac{\sum_{j \neq a, ts} \sigma_{jts}}{J-1} > 0$, this term is equal to the marginal profit a refinery would receive if it was allowed to sell in the spot market the petroleum bought using long-term contracts. In practice this means that there are some arbitrage opportunities between the long-term and the spot contracts from which the refinery would be able to profit, at the expense of the long-term contract seller. This means that the refinery would be able to pay petroleum at a discount using long-term contracts and sell it in the spot markets elsewhere in the world, at a profit. Consequently, in order to prevent this type of arbitrage, either the long-term contract should include a clause preventing the refinery from selling the petroleum in the spot market, and (or) the long-term contract price should always be higher than the spot prices. This premium is analyzed in detail in section 4.1, equation (5).

Similarly, by using dual variables, we can analyze how a refinery optimally purchases swap contracts, as summarized in Figure 2.

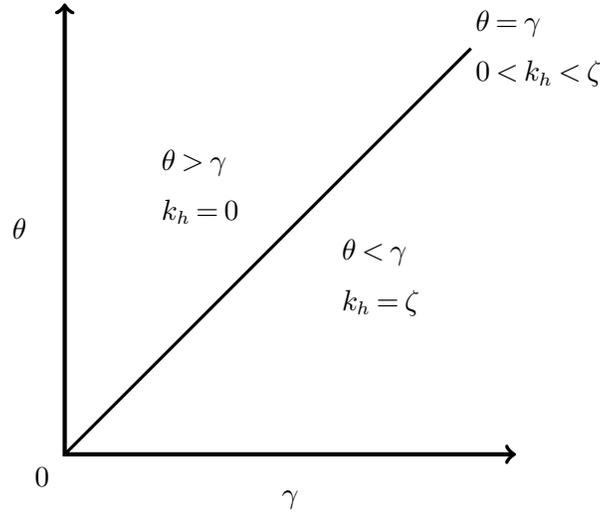


Figure 2 Policy Implications of the Relationship between θ and γ .

From Proposition 1 we know that for the crack spread swap contracts $\sum_{ts} \left(\frac{\beta}{S} + \lambda_s \right) ((g_h - g_{hts}) - (w_h - w_{hts})) = \gamma - \theta$, from which it is clear that the volume of swap trading depends on the degree of risk aversion, represented by β and λ_s (because it depends on α and β), and on the difference between crack spread fixed by using swap contracts ($g_h - w_h$) and the risk-adjusted expected crack spreads in the spot market, $\sum_{ts} \left(\frac{\beta}{S} + \lambda_s \right) (g_{hts} - w_{hts})$. If the crack

spread fixed by using swap contracts ($g_h - w_h$) is higher than the risk-adjusted expectation of the crack spread in spot market prices, $\sum_{ts} (\frac{\beta}{S} + \lambda_s) (g_{hts} - w_{hts})$, then the refinery hedges 100% of production capacity using swap contracts. In this case, $\sum_{ts} (\frac{\beta}{S} + \lambda_s) ((g_h - g_{hts}) - (w_h - w_{hts})) = \gamma$ and $\theta = 0$. Otherwise, if the risk-adjusted expected crack spreads are higher in the spot markets, i.e., $g_h - w_h < \sum_{ts} (\frac{\beta}{S} + \lambda_s) (g_{hts} - w_{hts})$, then $-\sum_{ts} (\frac{\beta}{S} + \lambda_s) ((g_h - g_{hts}) - (w_h - w_{hts})) = \theta$ and $\gamma = 0$ (the refinery does not purchase any swap contracts). Finally, if the expected crack spreads implicit in the swap contracts and in the spot prices are equal, i.e., $g_h - w_h = \sum_{ts} (\frac{\beta}{S} + \lambda_s) (g_{hts} - w_{hts})$, then $\sum_{ts} (\frac{\beta}{S} + \lambda_s) ((g_h - g_{hts}) - (w_h - w_{hts})) = 0$, $\theta = \gamma = 0$, and the proportion of swap contracts to be purchased is undetermined, but not zero and less than 100% of capacity. In Proposition 5 we analyze the conditions under which the refinery purchases swap contracts.

PROPOSITION 5. *Let $0 < \alpha \leq 1$ and $0 \leq \beta \leq 1$, a refinery purchases swap contracts if and only if $(g_h - w_h) \geq \frac{1}{T} \sum_{ts} (\frac{\beta}{S} + \lambda_s) (g_{hts} - w_{hts})$.*

We end this formal analysis by considering how β affects the refinery's use of crack spread swap contracts. However, before proceeding, we need to establish the relationship between crack spread contracts and extreme losses, in the petroleum refining business, which is also illustrated numerically in Section 4, Tables 6 and 8: the proportion of contracts is directly associated with extreme losses (i.e., lower CVaR). This is a well known property of financial derivatives see, e.g., Hull (2012, ch. 33) for an introduction energy and commodity derivatives.

In the case study presented in section 4, in Table 6, it is overwhelmingly evident that k_h increases with β , i.e., the closer the risk profile is to neutrality, the higher the proportion of swap contracts purchased, and the higher the exposure to extreme risk. In order to better understand this empirical evidence, Proposition 6 analyzes the conditions under which a refinery buys swap contracts. To include a valid representation of reality in our analysis, Postulate 1 declares that due to the leverage effect of swap contracts, i.e., as all extreme losses (and profits) are caused by swap contracts and there is a one-to-one correspondence between the total losses and the total swap contract losses, the scenarios in the tail of the total profit and swap contracts profit distributions are the same.

POSTULATE 1. *Leverage effect. For every scenario s such that $z_s \geq 0$ and $\mu - \pi(s, q_{jts}, q_a, k_h) \geq 0$, the corresponding swap contract losses $\sum_t ((g_h - g_{hts}) - (w_h - w_{hts})) k_h$ are also in the $(1 - \beta)$ lowest profits.*

PROPOSITION 6. *Let $0 < \alpha < 1$ and $0 \leq \beta < 1$. The lower the β the lower the optimal proportion of swap contracts purchased, k_h^* , on average.*

For the risk-averse refinery ($0 < \alpha < 1$ and $\beta < 1$) its risk-adjusted expectation is that, during the lifetime of the contract, the crack spreads bought in the spot market will *increase* and the

refinery will lock in a *loss*. For this reason, the more risk averse the refinery is (β closer to zero), the less likely is that it will purchase swap contracts. This is a surprising result as our expectation was that the risk-averse refinery would be the one using swap contracts to fix the refining margin and therefore decrease risk. However, this is not what we proved in Proposition 6, and observed in the simulations reported in section 4.2, Table 6.

Proposition 6 explains why, from the perspective of a risk-averse refinery, fixing the crack spread is not risk avoidance, but instead is risk-seeking behavior. A risk-averse refinery prefers to avoid purchasing financial derivatives, even if fixing the crack spread, because in the tail scenarios of the profit distribution, the possible losses arising from crack spread swap contracts are *higher*.

4. Procurement Risk Management by a Singapore Based Refinery

In this section we analyze how a Singapore based refinery uses long-term, spot and swap contracts in procurement risk management. We start in section 4.1 by describing the statistical properties of petroleum prices, GPWs, and transportation costs, and how they are correlated. This analysis is extended to estimate how Saudi Arabia prices the Arab Light long-term contracts. With this information, we can then generate scenarios for the price trajectories of the three different crude prices considered, 12 months ahead. Finally, we study how the refinery optimizes its procurement strategy, calculating how much to purchase from different sources of petroleum and how many crack spread swap contracts to use for risk hedging. We analyze how procurement depends on risk aversion (section 4.2) and is a function of alternative long-term contract pricing policies (section 4.3).

4.1. An Empirical Study

The relationship between the state of the market, backwardation vs. contango, and petroleum prices (e.g., Bodie et al. 2003, p. 762) is essential to understanding the current pricing policy of Saudi Arabia. It is said a market is in normal backwardation (Keynes 1930) when the series of futures prices for contracts with different duration converges from below to the expected spot price. The classic explanation for this market state is that the risk-averse producers of petroleum are willing to offer an expected profit to the refineries who adopt a long position. A different explanation for the normal backwardation in petroleum markets is that the physical purchase of petroleum provides a convenience value to the refinery that the acquisition of a forward contract does not. This value, when expressed as a rate, has been described as a convenience yield, e.g., Routledge et al. (2000). This concept was first introduced in the theory of storage (i.e., Kaldor 1939, Working 1948, 1949, Telser 1958), whereby the convenience yield arises from the possibility of consuming (or selling to a consumer) the petroleum now, or keeping it in storage to be consumed (or sold to consumers) in the future. A market is said to be in contango if the futures prices for a

sequence of contracts of different duration converge to the expected spot price from above, as the contracts get closer to the delivery date. In this case, the risk hedgers are the refineries as they are willing to pay a premium over the expected spot price in order to lock in a futures price before delivery.

We start by estimating the stochastic processes describing the petroleum prices and the GPWs in the different geographical areas considered: Arab Light (Saudi Arabia), Attaka (Indonesia), Cabinda (Angola), and Oman-Dubai vs. Brent (the benchmark crude) prices. The GPWs reflect the operational flexibility of the refinery, i.e., its ability to transform the petroleum into final products, considering both the input and output price uncertainties. Whereas Dong et al. (2014) explicitly modeled the blending of petroleum products with different degrees of efficiency, we use the GPWs to account for this same effect, as these are obtained from real-world data and already take into account how the different sources of petroleum are used by the refinery. The data sources used in this section are British Petroleum (bp.com) and Platts market data (spglobal.com/platts).

All the prices are highly correlated with each other and non-stationary. We use this correlation, together with the random walk properties of the different time series, to generate the scenarios in section 4.2. There are three major reasons for using the random-walk hypothesis in analyzing the behavior of petroleum prices and GPWs: first, we tested it using time series methods and found the hypothesis could not be rejected; second, the swap contracts are priced using this hypothesis, so the time series results and the pricing simulations are internally consistent; and third, the random-walk hypothesis assumes that all relevant information is already known at the decision time, and that there is no information to be gained by mining the data. Therefore, it avoids over-fitting the data by capturing any patterns that only are observable after the event has occurred.

As we have used Brent petroleum as the benchmark crude for both the generation of prices and GPWs, we need to estimate the monthly rates of change of the different prices (and GPWs) and the respective correlation with the Brent prices, as summarized in Table 1. There is no price reported for Arab Light as this crude is priced based on Oman-Dubai; all the prices and GPWs are highly correlated with the Brent price (the correlations are all above 90%).

Table 1 Standard Deviation, and Coefficient of Correlation to Brent of the Arab Light, Attaka, Brent, Cabinda and Oman-Dubai Monthly Rate of Change (Δ) of Prices and GPWs, in %.

		Arab Light	Attaka	Brent	Cabinda	Oman-Dubai
Δ Prices	S.D.		8.2	8.0	8.4	7.5
	Corr. to Brent		0.957	1.0	0.984	0.967
Δ GPWs	S.D.	7.9	7.9	8.2	7.4	7.7
	Corr. to Brent	0.945	0.939	0.945	0.928	0.945

In Table 2 we summarize the transportation costs from the three sources (Arab Light, Attaka and Cabinda) to Singapore as a percentage of the respective petroleum prices. On average, Arab Light

transportation costs are about 1.12% of the petroleum price (s.d. of 0.55%), whereas transportation from Attaka and Cabinda is more expensive at 2.28% (s.d. 0.91%) and 2.55% (s.d. 0.97%) of the petroleum price, respectively.

Table 2 Sample Mean and Standard Deviation for the Transportation Costs in % of Petroleum Price.

	Arab Light	Attaka	Cabinda
Mean	1.12	2.28	2.55
S.D.	0.55	0.91	0.97

In Figure 3, to better describe the problem faced by the refinery, we plot the average refining margins per month in \$/bbl. These decrease with increasing petroleum prices. In many of the months analyzed, the refining margins are negative, especially for higher petroleum prices.

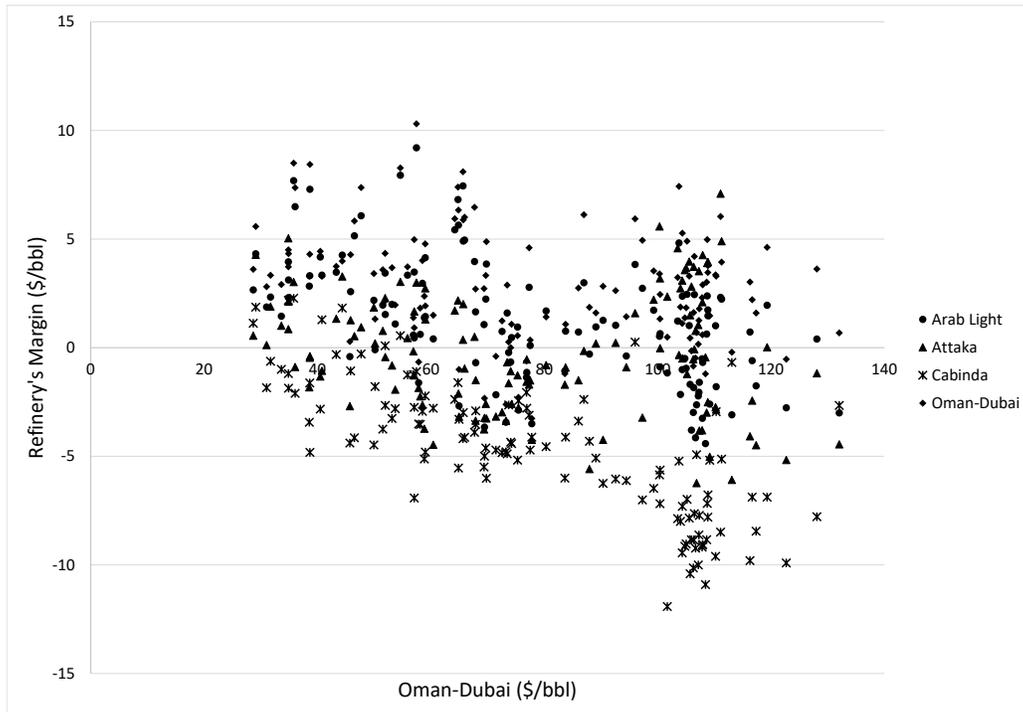


Figure 3 Average Refining Margins vs. Oman-Dubai Petroleum Price.

To understand how the long-term contracts are priced in Saudi Arabia, we interviewed expert traders with very long experience in the industry about how the pricing is done. First, we were told that Saudi Aramco uses an adjustment factor based on the Oman-Dubai spot prices. Second, this adjustment would be dependent on the state of the market, i.e., the slope of the forward curve. Third, the adjustment would also be affected by the difference in GPWs between Arab Light and Oman-Dubai petroleum. Our working hypothesis was that this adjustment factor is dependent on

the level of backwardation of the petroleum prices (as represented in Figure 4), the past values of the adjustment factor f_{t-1} , and the difference between the gross product worth of Arab Light and Oman-Dubai crude in a typical Singapore refinery. We tested all these variables in our analysis.

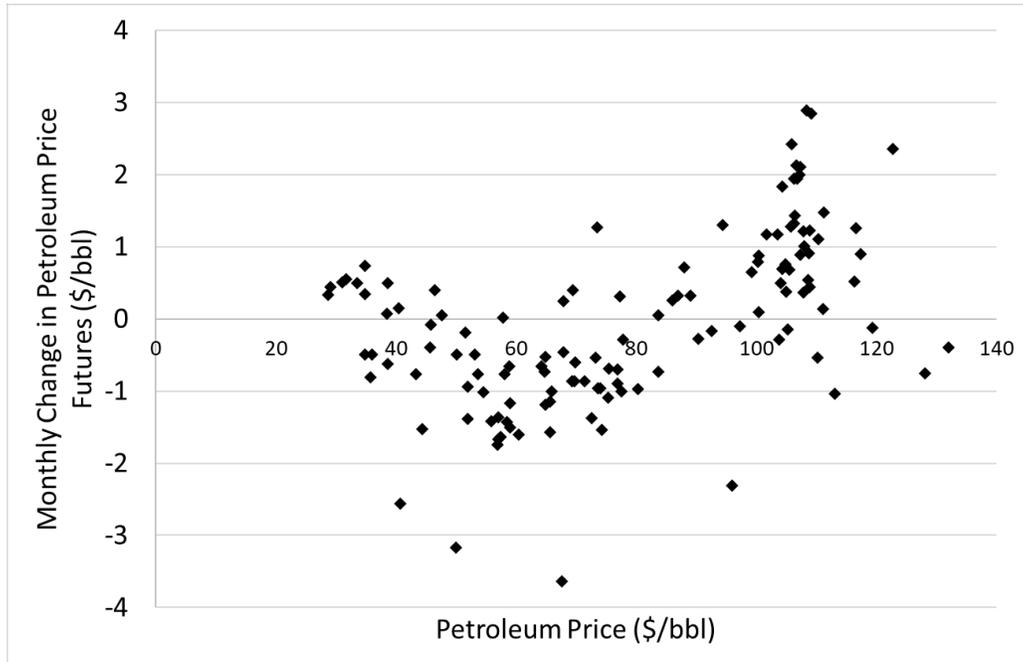


Figure 4 Oman-Dubai Petroleum Price vs. the Forward Curve. The Monthly Change in Petroleum price is Equal to Difference between the Futures Price 1-Month and 2-Months Ahead of the Spot Market.

As a result of these interviews, we then hypothesized that Saudi Arabia prices its crude to take into account the state of the market (forward curve) and the differences in quality between its own production and the benchmark crude in each market (difference in GPWs). This adjustment factor, called the Saudi Arabia offset, is computed specifically for each regional market. We focused on trying to describe how the Arab Light adjustment factor is computed in relation to Oman-Dubai crude prices. This is an index-linked payment, as analyzed in Zhang et al. (2014). These were the basic ideas we tested in order to improve scenario generation for the Arab Light long-term contract prices. When first tested, this hypothesis seemed correct and to be supported by the statistical model. However, upon analysis of the residuals, we realized these were heteroskedastic.

We then took the logarithms of the prices so as to reduce the heteroskedasticity issue and tested again in order to obtain statistical evidence on the way Saudi Arabia prices the long-term contracts. However, working with the logarithms of prices and GPWs led us to conclude that the difference

in GPWs between Arab Light and Oman-Dubai petroleum prices does contribute to explain Arab Light offsets (i.e., the respective parameter was not statistically significant).

Let w_{jt}^m stand for the price of the petroleum bought from j at time t , and to be delivered in m months (with $m = 0$ for the spot market). In this equation we compare two different producers: Oman-Dubai ($j = o$) and Arab Light from Saudi Arabia ($j = a$). $f_t = w_{at}^0 - w_{ot}^0$ represents the behavior of Saudi Arabia when pricing its crude as a function of the Oman-Dubai spot prices, as described in equation (5), in which the log represents the natural logarithm. In this equation, $\log(f_t)$ stands for the difference between the logarithm of the price of Oman-Dubai and the logarithm of the price of Arab Light. The residuals of equation (5) are not autocorrelated and are normal distributed.

$$\mathbf{E}(\log(f_t)) = 0.0073 + 0.639(\log(w_{o(t-1)}^1) - \log(w_{o(t-1)}^2)) + 0.445\log(f_{t-1}) \quad (5)$$

Equation (5) describes the expected adjustment factor for Arab Light when compared to Oman-Dubai prices, as a function of two main components, as follows. (A) Backwardation vs Contango: this is expressed in the term $\log(w_{o(t-1)}^1) - \log(w_{o(t-1)}^2)$, which has an associated parameter of 0.639, meaning that for each 1% increase in the price difference between the Oman-Dubai contract for one or two months, the adjustment factor increases by 0.639%. This means that when the inventory is tighter (abundant), Saudi Arabia increases (decreases) the adjustment factor. (B) The adjustment factor is autoregressive with an associated parameter of 0.445, which means that for each 1% increase in the Arab Light offset in the previous month, on average Saudi Arabia charges 0.445% more in the current month.

From this equation it is evident that Saudi Arabia uses an adjustment process, based on reinforcement learning, to change the offsets at any given month. This learning mechanism is based on the proportion of the adjustment in the previous month, and on the difference between the Oman-Dubai prices for delivery within 1 and 2 months.

Table 3 describes the regression equation explaining how Arab Light is priced in comparison to $(Oman - Dubai)_{t-1}$. In the first column, Oman-Dubai stands for $\log(w_{o(t-1)}^1) - \log(w_{o(t-1)}^2)$. All the parameters are statistically significant. The Adj. R-squared is 79.9%. This means there is a *strong positive correlation* between the Oman-Dubai price at time $t - 1$ and the Arab Light offsets. However, it does not mean there is causality, even though this is part of the hypothesis raised based on the interviews with the energy experts. Nonetheless, in order to improve the modeling of the prices of Arab Light long-term contracts, *we only need correlation* between the variables; we do not need causality. So, we can use this equation to approximate statistically how we expect the Arab Light to be priced, based on the Oman-Dubai price in the previous month. This simulation is

grounded on significant statistical evidence. In Figure 5 we plot the relationship between the *Actual vs. Expected* Arab Light offsets in natural logs, in order to visualize the quality of the regression analysis. It is clear that the points are clustered around an imaginary line that goes through the coordinate origin and has a 45 degree angle, as should be the case in a good quality statistical analysis.

Table 3 Estimation of the Arab Light Offsets.

Dependent Variable: $\log(f_t)$

Included observations: 126 after adjustments

Variable	Coefficient	Std. Error	t-Statistic	Prob.
C	0.0073	0.00008	8.94	0.000
$Oman - Dubai_{t-1}$	0.6398	0.063	10.05	0.000
$\log(f_{t-1})$	0.445	0.052	8.5	0.000

R-squared	0.805	Adj R-squared	0.799
S.E. of regress	0.799	Sum sq. resid	0.003
F-statistic	247.4	Prob(F-statistic)	0.000

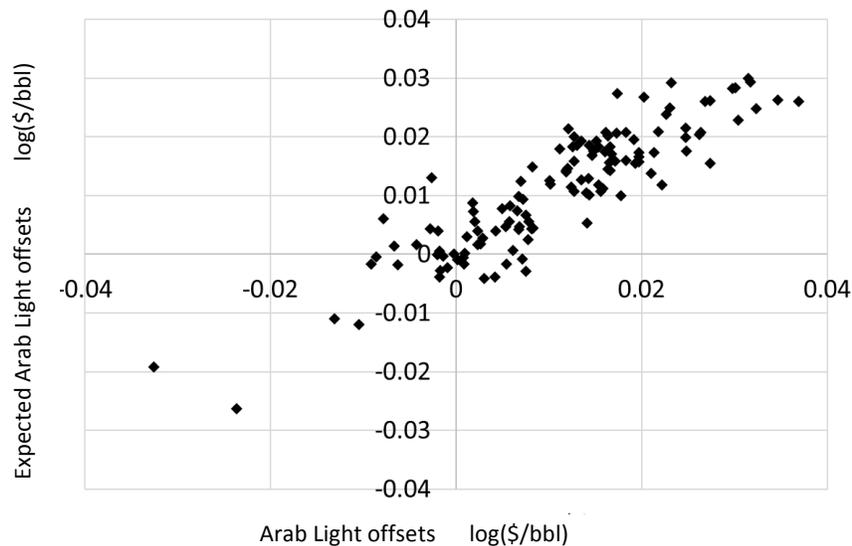


Figure 5 Actual Arab Light Offsets vs. Expected Arab Light Offsets, in Natural Logs.

4.2. The Refinery's Optimal Procurement Policy

In this section we summarize the numerical results on the refinery's procurement problem. In Table 4 we describe the procedure used to generate the scenarios for the prices, GPWs and transportation costs required to parameterize the profit function (3), explaining how to go from time series analysis to stochastic programming. In points 1 and 2 we estimate the models for the petroleum prices, GPWs and Arab Light offsets as described in section 4.1. Then in point 3.a), using the models estimated and based on Tables 1 and 2, we use Monte Carlo simulation (e.g., Kroese et al. 2011, pp. 143-146; Brandimarte 2014, pp. 278-282) to generate the scenarios for the petroleum prices, GPWs, and transportation costs from the three different sources, together with the Oman-Dubai prices as these are the base for Arab Light pricing. Note that we have 126 moving windows (the sample size). For each one of these, we generate 5000 paths (each one 12 months in duration) representing the behavior of the random variables. Then in point 3.b), having generated the scenarios required to compute the profit equation (3), we solve problem (4), for each time window, using the stochastic variables generated for each of the 5000 scenario paths. This means that by using a scenario based CVaR, we avoid over-fitting the data as we are not optimizing for a given sample, but rather for a very wide variety of scenarios.

Table 4 Procedure for Price Scenario Generation and Solution of the Stochastic Model.

1. Model for Petroleum Prices, GPWs and Transportation Costs.
2. Estimation of the Arab Light offsets.
3. Generate Scenarios for the Petroleum Prices, GPWs and Transportation Costs (Arab Light, Attaka, Cabinda and Oman-Dubai).
For each month b in the sample:
(a) For each month $t = 1$ to 12 in the planning horizon:
(i) Generate the petroleum prices, GPWs and transportation costs for S scenarios;
(ii) Compute the Arab Light offsets, or alternative pricing policies, for S scenarios.
(b) Solve the stochastic program (4) obtaining:
(i) Optimal procurement strategy (petroleum sources and swap contracts);
(ii) CVaR;
(iii) Expected Profit.

Moreover, another advantage of this methodology is that the scenarios are generated for the stochastic variables only, without having to compute the profit distribution and covariance matrix for the different policy decisions, as would be the case if using standard deviation or variance as risk-measures. Consequently, for problems in which the profit function is not trivial to compute, or in which the decisions may impact the profit from each source, scenario-based CVaR optimization is a simpler and better approach.

Additionally, still regarding the possible issues with over-fitting the optimal policy to the sample data, it should be reinforced that the use of scenario paths is necessary but not sufficient to avoid over-fitting the policy to a specific set of observations; it is also necessary to ensure that only the information available to the refinery at the start of the planning horizon is used in generating the scenarios. For this reason, we use the random-walk hypothesis, as otherwise the scenario generation would be biased by knowledge the refinery, or the Saudi producers in the case of the policy for pricing long-term contracts, would not have when making their decisions.

For each one of the 126 months in the sample, we analyze the optimal procurement spot market from alternative producers (in South Asia and West Africa), including the possible purchasing of swap contracts to fix crack spreads in advance, and the use of a one-year forward contract at a price indexed to Oman-Dubai petroleum, plus an offset announced monthly. The procurement problem is solved using a stochastic programming model accounting both for expected profit and the degree of risk aversion. For the percentage of observations in the tails we set $\alpha = 0.05$. For each scenario, the petroleum prices (w_{jts}), GPWs (g_{jts}), and transportation costs (η_{jts}) are simulated as cointegrated random walks. This method is consistent with the theory of efficient markets and the non-arbitrage argument for pricing swap contracts.

We start in Table 5 by analyzing the impact of β on the optimal procurement strategy. The results depicted in this section, including Table 5, are the averages collected from the simulations for all 126 time windows in our sample. In these experiments, the average Arab Light offset to Oman-Dubai crude oil prices was 1.01\$/bbl. Independently of the β , the refinery always buys most of the petroleum from Saudi Arabia as the optimal proportion for long-term contracts remains in the 50 – 60% range. Why there appears to be no clear trade-off between the β and the proportion of long-term contracts? To answer this question, in Table 6 we summarize the proportion of crack spread swap contracts purchased by the refinery. The impact of the β on the trading of these contracts is evident. These results are the averages, for each β , of the proportions of swap contracts bought, in the 126 time windows.

Table 5 Proportion Procured from Long-Term vs. Spot Contracts, in %.

β	1.0	0.9	0.75	0.5	0.25	0.0
Arab Light (Long-Term)	52	54	57	58	56	54
Attaka (Spot)	33	33	29	28	30	35
Cabinda (Spot)	15	15	14	14	15	11

First, the risk-neutral refinery ($\beta = 1$) bought enough swap contracts to fix the crack spread of about 85% of its production capacity. The proportion of swap contracts bought decreases directly with the degree of risk aversion. The refinery with a $\beta = 0.75$ buys a large enough number of swap

Table 6 Swap Contracts Trading as Proportion of Refining Capacity, in %.

β	1.0	0.9	0.75	0.5	0.25	0.0
Swap Contracts	85	61	20	20	19	18

contracts to cover 20% of its capacity, in line our analysis in Proposition 6. The main reason for this behavior is that a refinery with a risk profile closer to neutral (β approximating one) purchases swap contracts (at time zero) in a higher number of time windows than their more risk-averse counterparts (β closer to zero), even when the non-arbitrage principle holds.

It also follows from the results depicted in Tables 5 and 6 that the degree of risk aversion mainly affects the proportion of swap contracts purchased by the refinery. On the other hand, the proportion of long-term and spot trading remains comparatively unaffected by the degree of risk aversion. The explanation for this result is the very high correlation between the long-term and spot trading prices (given Saudi Arabia’s price indexing policy), and the low correlation between refining and swap contract profits.

Second, we still need to explain why the long-term contracts remain in the 50 – 60% range. As analyzed in Figure 3, the margins the Singaporean refinery makes from Arab Light are higher, on average, than the margins received from the other types of crude. This occurs for two reasons: the transportation costs are lower (see Table 2) and the GPWs are higher (as the refinery’s production process has been optimized to work with Arab Light). Consequently, on average, Arab Light is the most profitable of all the crudes for this specific refinery.

Nonetheless, this explanation is still incomplete as, at any given time, the risk-neutral refinery should choose only the cheapest source, and not diversify. This is also the case in these experiments. When the refinery is risk-neutral ($\beta = 1$), everything is bought from Saudi Arabia in 52% of the moving windows, hence the result. It is also true that, for any given time window, procurement diversification only happens with $\beta < 1$.

By analyzing the refinery’s profitability as a function of risk aversion (β), we can better appreciate how this parameter influences the optimal policy. The expected profit, decomposed into refining and swap trading, is presented in Table 7. The risk-neutral refinery ($\beta = 1.0$) has a positive expected total profit, almost all of which is due to refining. The total profit decreases with the degree of risk aversion (with the decrease in β), going into loss for a β of 0.25 or less. The swap contracts show a zero expected profit when the refinery has a β less than or equal to 0.75 due to a trade-off between expected profit and CVaR, as described in Table 8.

The main reason the risk-averse refinery ($\beta = 0$) has expected losses is due to the very difficult market in which it was operating during the time covered by our data set, as the refining margins are negative in many months (see Figure 3). There are several reasons for this outcome. Margins

Table 7 Decomposition of the Refinery's Expected Profits per Activity, in \$/bbl.

β	1.0	0.9	0.75	0.5	0.25	0.0
Refining	0.46	0.46	0.30	0.04	-0.14	-1.29
Swaps	0.06	0.05	0.00	0.00	0.00	0.00
Total	0.52	0.51	0.30	0.04	-0.14	-1.29

are very small and therefore an increase in petroleum price may not be passed to the consumer fast enough, especially as petroleum products are highly regulated in the region. Additionally, for many companies refining is just a small part of their supply chain operations, with production and retail being more profitable, thus justifying the maintenance of the refinery, even if operating at a loss.

Moreover, it is also evident from Table 8, that there is a trade-off between maximization of expected profit and CVaR. The CVaR of a risk-neutral refinery ($\beta = 1.0$) is about 3.5 times higher than the CVaR of a refinery with a $\beta = 0.75$.

Table 8 Relationship between Expected Profit and CVaR, in \$/bbl.

β	1.0	0.9	0.75	0.5	0.25	0.0
Expected Profit	0.52	0.51	0.30	0.04	-0.14	-1.29
CVaR	-27.0	-10.6	-7.8	-7.4	-7.2	-7.2

On the other hand, it is also obvious that extreme risk aversion ($\beta = 0.0$) is counterproductive as the refinery avoids very large losses but it locks in an expected loss. For this reason, we focus the rest of our analysis on the refinery with a $\beta = 0.75$, as the expected profit is still positive (\$0.3/bbl) and most of the CVaR decrease is already captured by the optimal policy.

We proceed by analyzing the refinery's procurement policy. Figure 6 depicts the proportion of crude purchased by a refinery with a $\beta = 0.75$ from the three different sources, as a function of the Oman-Dubai petroleum price. This refinery exhibits purchasing behavior in line with what we observe in practice. The Arab Light long-term contract is still preferred, but when market prices are above average, a significantly higher proportion is bought from South East Asia and West Africa, in the spot market.

In Figure 7 we depict the relationship between expected profit (\$/bbl) and CVaR (\$/bbl) for a refinery with a $\beta = 0.75$. There are a couple of contracts in which the CVaR is positive, meaning the expected tail profit is still positive. In the large majority of contracts, the refinery is profitable. However, due to the small margins, the smaller number of contracts with losses end-up destroying an important portion of the value created by the positive results. The risk faced by the refinery is still very important, with the expected profit being a small fraction of the possible losses represented by the CVaR.

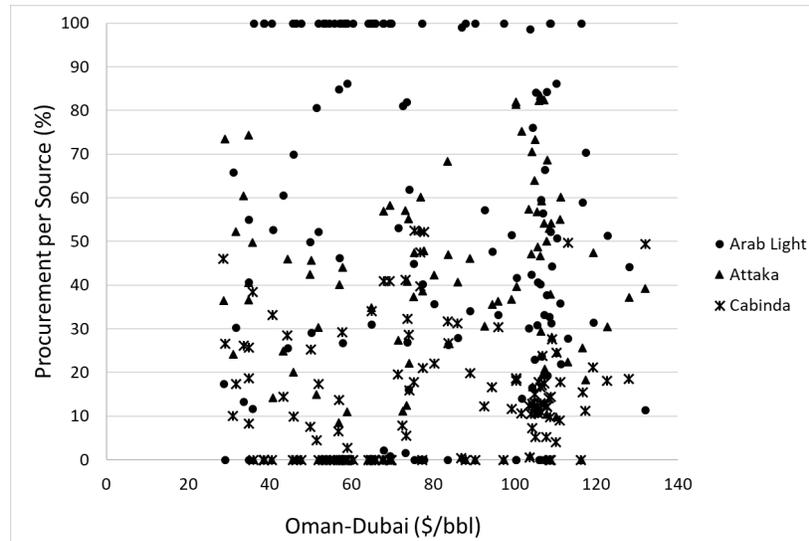


Figure 6 Refinery with a $\beta = 0.75$. Proportion Bought from the Different Producers.

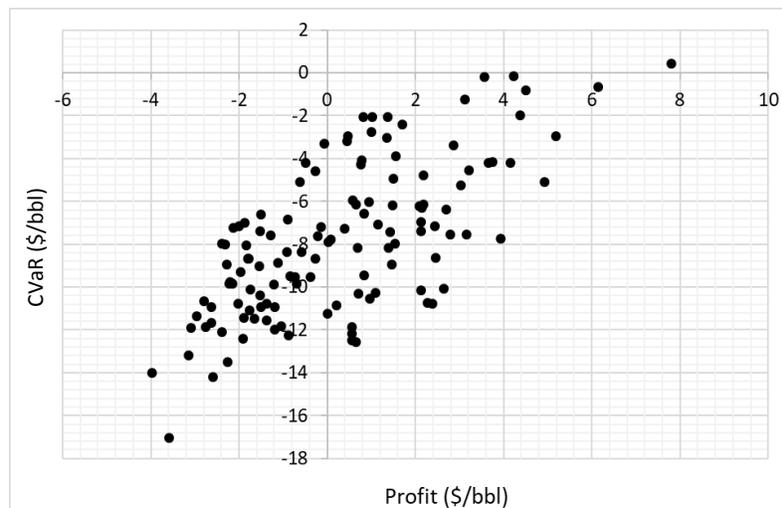


Figure 7 Refinery with a $\beta = 0.75$. Relationship between CVaR and Expected Profit.

Going back to the profit decomposition between refining and swap contracts, as represented in Table 7, we would like to better explain the relationship between these two activities. To understand the role of the swap contracts as part of the refinery's procurement plan, in Figure 8 we plot the relationship between the expected profit from refining and from swap contracts for the refinery with $\beta = 0.75$. It is clear there is no statistically significant relationship between refining and swap trading profitability. Moreover, the distribution of trading profits show a high volatility in comparison with its expected profit of about zero. This reinforces the idea that no value is created from the trading

activities when the crack spread contracts are priced so that no arbitrage opportunities are present, as would be expected if the contracts were correctly priced in equilibrium, and in accordance with finance theory on asset pricing, e.g., Bodie et al. (2003). Consequently, at any point in time, and independently of the profitability of the refining activities, the expected profit from swap contracts is zero, from which we infer that the profitability of these activities is not correlated. Nonetheless, it is well known from portfolio theory that holding onto non-correlated assets reduces the volatility of portfolio profit.

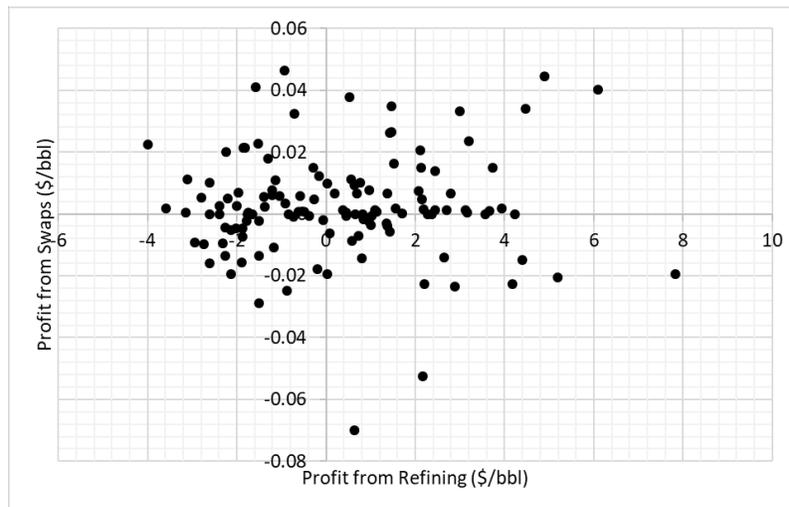


Figure 8 Refinery with a $\beta = 0.75$. Relationship between the Refining and Swap Trading Profits.

4.3. Alternative Design for the Arab Light Long-Term Contracts

In this subsection we explore alternative pricing policies to the current Arab Light offsets. This is an important analysis as contract design is essential to determining the incentives, and the risk, associated with the long-term contracts. Moreover, in order to establish the credibility of the different design settings, we need to assess the impact of the different contracts. This will allow us to rule out any design that decreases the profitability of petroleum production.

In Table 9 we depict the average producers' revenues in \$/bbl. Each column represents a different pricing policy for Arab Light. The Offsets column is calculated using equation (5). Attaka, Brent and Oman-Dubai columns apply the prices of the respective types of crude oil as the basis for Arab Light pricing, without using any offsets.

Of the four different pricing methods analyzed, the best for Saudi Arabia is to set the price of Arab Light equal to the Oman-Dubai price, which increases the average revenue by 9\$/bbl when compared to the Arab Light offsets policy (55\$/bbl under Oman-Dubai pricing minus 46\$/bbl

Table 9 Refinery with a $\beta = 0.75$. Average Producers' Revenues in \$/bbl as a Function of Different Pricing Policies for Arab Light.

Petroleum	Offsets	Attaka	Brent	Oman-Dubai
Arab Light	46	2	24	55
Attaka	26	55	42	19
Cabinda	11	24	16	8

under the current policy). The main explanation for the benefits of using Oman-Dubai as the price for Arab Light is that, currently, the Arab Light offsets are set too high when the market is overheated (decreasing demand). For this reason, the average quantity sold from Arab Light to the refinery is 52% in the current policy, and 69% when using an Oman-Dubai index, instead.

Another interesting insight is related to the way the producers' profits are affected by Saudi Arabia's indexing: an Attaka index would lead the refinery to buy most of its petroleum in the spot markets (as for risk mitigation it would be the only price not correlated with the other two). A Brent-based index would also be beneficial for Attaka as the risk mitigation from buying Saudi Arabia's petroleum would be reduced. Overall, from the perspective of Saudi Arabia, the Arab Light offset policy, and the Oman-Dubai price index are the most beneficial.

5. Conclusions and Discussion

In this article we have analyzed the impact of long-term petroleum contracts on a refinery's procurement policy, taking into account operational flexibility (considering the gross product worth of the different types of petroleum), price indexation, and degree of risk aversion based on CVaR, for the one-year planning horizon. The refinery's problem is deciding the proportion of contracts to buy forward, in the different spot markets, and additionally, the proportion of crack spread swap contracts to purchase.

We have proved the necessary and sufficient conditions for the optimal solution to be unique, analyzing it and interpreting its meaning. We have proved that risk aversion can be better captured by using an objective function that is the weighted average between the risk-neutral expected profit and CVaR, and shown that such an objective function is equivalent to the risk-adjusted expected profit, in which the probabilities associated with each scenario are transformed to incorporate the subjective degree of risk aversion.

We have proved that when the non-arbitrage condition holds, independently of the degree of risk aversion, a refinery buys an undetermined number (but strictly positive and less than expected production) of crack spread contracts. We have also assessed how the proportion of swap trading is affected by the degree of risk aversion, concluding that, surprisingly, an *increase in the degree of risk aversion leads to the purchase of fewer swap contracts*. This result was proved both analytically (Proposition 6) and using computer simulations parameterized based on real-world data (Table 6).

We analyzed the data on refining margins in our sample to empirically show that, on average, these are negatively correlated with petroleum prices, and were negative in more than half of the observations in our data (see Figure 3). By analyzing the relationship between forward and spot petroleum prices (Figure 4), we estimated a positive correlation between these variables (which is stronger when using the logarithm of the prices so as to correct for the heteroskedasticity effect). We were then able to estimate a statistically significant relationship, using regression analysis (equation 5), between the adjustment factor used by Saudi Arabia to price the Arab Light when compared to Oman-Dubai prices (79.9% R-squared Adjusted). In the analyzed data, Saudi Arabia increases the price of Arab Light by 0.639% for an increase of 1% in the Oman-Dubai petroleum price. This analysis also showed that there is statistically significant support for the hypothesis that Saudi Arabia uses reinforcement learning when adjusting the Arab Light prices.

We then tested the model, and the theory, in computing the refinery's procurement policy. Our first conclusion from the numerical analysis is that even just a 10% weight on the CVaR (90% on the expected value) significantly reduces risk exposure without compromising profitability, as depicted in Table 8. Moreover, the CVaR is able to capture the trade-off between risk and expected profit, i.e., in order to decrease exposure to high losses, the refinery accepts receiving a lower expected profit. In the case under analysis, a refinery with a 25% CVaR weighting captures almost all possible risk reduction and still has a positive expected profit. However, a refinery that only aims to reduce risk exposure ends up locking in a negative expected profit. This is irrational, as in the long-term, without additional external capital, it will not be able to survive.

The low profitability of the refinery business model is actually not unusual in the industry, as summarized by the profit decomposition depicted in Table 7. We know from the historical data that refining was not profitable in many of the months (see the negative margins in Figure 3). The total profits per barrel were \$0.46 for the risk-neutral refinery, and a *loss* of \$1.29 for the risk-averse refinery. This means that in order to break even, the risk-neutral refinery may need to buy petroleum from riskier sources, whereas the risk-averse refinery diversifies its sources, resulting in an actual loss, but avoiding a potentially very costly exposure. The swap contracts produced no significant profits, as they are priced so that there are no arbitrage opportunities.

In exploring the pricing policy of Saudi Arabia, we analyzed several alternatives to the setting of Arab Light offsets. We found that Saudi Arabia would increase its expected profit by setting the price for Arab Light equal to Oman-Dubai crude prices, without using offsets.

References

- Artzner, P., Delbaen, F., Eber, J., & Heath, D. (1999). Coherent measures of risk. *Mathematical Finance*, 9, 203–228.

- Ansariipoor, A. H., Oliveira, F. S., & Liret, A. (2016). Recursive expected conditional value at risk in the fleet renewal problem with alternative fuel vehicles. *Transp. Res. Part C: Emerging Technologies*, 65, 156–171.
- Bilsel, R. U., & Ravindran, A. (2011). A multiobjective chance constrained programming model for supplier selection under uncertainty. *Transportation Research Part B: Methodological*, 45 (8), 1284-1300.
- Bodie, Z., Markus, A. J., & Kane, A. (2003). *Investments*. Irwin/McGraw-Hill.
- Bradley, S. P., Hax, A. C., & Magnanti, T. L. (1977). *Applied Mathematical Programming*. Reading, MA: Addison-Wesley.
- Brandimarte, P. (2014). *Handbook in Monte Carlo simulation: Applications in financial engineering, risk management, and economics*. John Wiley & Sons.
- Burke, G. J., Carrillo, J. E., & Vakharia, A. J. (2007). Single versus multiple supplier sourcing strategies. *European Journal of Operational Research*, 182(1), 95–112.
- Charnes, A., & Cooper, W. W. (1959). Chance-constrained programming. *Management Science*, 6(1), 73–79.
- Chen, W., Sim, M. , Sun, J., & Teo, C.-P. (2010). From CVaR to uncertainty Set: implications in joint chance-constrained optimization. *Operations Research*, 58(2), 470–485.
- Chen, Y., Xu, M., & Zhang, Z. G. (2009). Technical note - A risk-averse Newsvendor model under the CVaR criterion. *Operations Research*, 57(4), 1040–1044.
- Chen, L. C., Long, D. Z., & Perakis, G. (2015). The impact of a target on Newsvendor decisions. *Manufacturing & Service Operations Management*, 17(1), 78–86.
- Conejo, A. J., Carrión, M., & Morales, J. M. (2010). *Decision making under uncertainty in electricity markets*. International Series in Operations Research & Management Science, Springer, New York.
- Costantino, N., & Pellegrino, R. (2010). Choosing between single and multiple sourcing based on supplier default risk: A real options approach. *Journal of Purchasing and Supply Management*, 16(1), 27–40.
- Devalkar, S. K., Anupindi, R., & Sinha, A. (2018). Dynamic risk management of commodity operations: model and analysis. *Manufacturing & Service Operations Management*, 20(2), 317–332.
- Dong, L., Kouvelis, P., & Wu, X. (2014). The value of operational flexibility in the presence of input and output price uncertainties with oil refining applications. *Management Science*, 60(12), 2908–2926.
- Goldschmidt, K., Kremer, M., Thomas, D.J., & Craighead, C.W. (2020). Strategic Sourcing Under Severe Disruption Risk: Learning Failures and Under-Diversification Bias. *Manufacturing & Service Operations Management*, <https://doi.org/10.1287/msom.2020.0907>.
- Hong, Z., & Lee, C. (2013). A decision support system for procurement risk management in the presence of spot market. *Decision Support Systems*, 55(1), 67–78.
- Hong, Z., Lee, C. K. M., & Zhang, L. (2018). Procurement risk management under uncertainty: a review. *Industrial Management & Data Systems*, 118(7), 1547–1574.

-
- Huang, H., Shen, X., & Xu, H. (2016). Procurement contracts in the presence of endogenous risk. *Decision Sciences*, 47(3), 437–472.
- Hull, J. C. (2012). *Options, Futures, and Other Derivatives*. Eighth edition. Prentice Hall.
- Kaldor, N. (1939). Speculation and economic stability. *Review of Economic Studies*, 7, 1–27.
- Keynes, J. M. (1930). *A treatise on money*. Harcourt Brace, New York.
- Kroese, D. P., Taimre, T., & Botev, Z. I. (2011). *Handbook of Monte Carlo methods*. John Wiley & Sons.
- Kumar, R. L., & Park, S. (2019). A portfolio approach to supply chain risk management. *Decision Sciences*, 50(2), 210–244.
- Li, G., Liu, M., Bian, Y., & Sethi, S. P. (2020). Guarding against disruption risk by contracting under information asymmetry. *Decision Sciences*, 51(6), 1521–1559.
- Martinez-de-Albeniz, V., & Simchi-Levi, D. (2005). A portfolio approach to procurement contracts. *Production and Operations Management*, 14(1), 90–114.
- Meena, P., Sarmah, S., & Sarkar, A. (2011). Sourcing decisions under risks of catastrophic event disruptions. *Transportation Research Part E: Logistics and Transportation Review*, 47(6), 1058–1074.
- Nejad A. E., Niroomand, I., & Kuzgunkaya, O. (2014). Responsive contingency planning in supply risk management by considering congestion effects. *Omega*, 48, 19–35.
- Oliveira, F.S., & Ruiz, C. (2020). Analysis of futures and spot electricity markets under risk aversion. *European Journal of Operational Research*, <https://doi.org/10.1016/j.ejor.2020.10.005>.
- Pournader, M., Kach, A., & Talluri, S. (2020). A Review of the Existing and Emerging Topics in the Supply Chain Risk Management Literature. *Decision Sciences*, 51(4), 867–919.
- Rockafellar, R. T., & Uryasev, S. (2000). Optimization of conditional value-at-risk. *Journal of Risk*, 2, 21–42.
- Routledge, B. R., Seppi, D. J., & Spatt, C. S. (2000). Equilibrium forward curves for commodities. *Journal of Finance*, LV(3), 1297–1338.
- Sawik, T. (2011). Selection of supply portfolio under disruption risks. *Omega*, 39, 194–208.
- Sawik, T. (2014). Joint supplier selection and scheduling of customer orders under disruption risks: single vs. dual sourcing. *Omega*, 43, 83–95.
- Scott, J., Ho, W., Dey, P. K., & Talluri, S. (2015). A decision support system for supplier selection and order allocation in stochastic, multi-stakeholder and multi-criteria environment. *International Journal of Production Economics*, 166, 226–237.
- Shapiro, A., Dentcheva, D., & Ruszczyński, A. (2009). *Lectures on stochastic programming: modeling and theory*. MPS-SIAM Series on Optimization, Philadelphia PA.
- Staino, A., & Russo, E. (2020). Nested Conditional Value-at-Risk portfolio selection: A model with temporal dependence driven by market-index volatility. *European Journal of Operational Research*, 280, 741–753.

- Secomandi, N., & Kekre, S. (2014). Optimal Energy Procurement in Spot and Forward Markets. *Manufacturing & Service Operations Management*, 16(2), 270–282.
- Tajbakhsh, M. M., Lee, C.-G., & Zolfaghari, S. (2010). Sole sourcing in eoq models with binomial yield. *Journal of Purchasing and Supply Management*, 16(3), 163–170.
- Telser, L. G. (1958). Futures trading and the storage of cotton and wheat. *Journal of Political Economy*, 66, 233–255.
- Torabi, S., Baghersad, M., & Mansouri, S. (2015). Resilient supplier selection and order allocation under operational and disruption risks. *Transportation Research Part E: Logistics and Transportation Review*, 79, 22–48.
- Vakharia, A. J., & Yenipazarli, A. (2009). *Managing supply chain disruptions. Foundations and Trends in Technology, Information and Operations Management*, 2(4), 243–325.
- Working, H. (1948). Theory of inverse carrying charge in future markets. *Journal of Farm Economics*, 30, 1–28.
- Working, H. (1949). The theory of the price of storage. *American Economic Review*, 39, 1254–1262.
- Xu, N., & Nozick, L. (2009). Modeling supplier selection and the use of option contracts for global supply chain design. *Computers & Operations Research*, 36(10), 2786–2800.
- Yan, X., Ji, Y., & Wang, Y. (2012). Supplier diversification under random yield. *International Journal of Production Economics*, 139(1), 302–311.
- Yang, L., Cai, G., & Chen, J. (2018). Push, pull, and supply chain risk-averse attitude. *Production and Operations Management*, 27(8), 1534–1552.
- Yildiz, H., Yoon, J., Talluri, S., & Ho, W. (2016). Reliable supply chain network design. *Decision Sciences*, 47(4), 661–698.
- Zhang, W., Zhou, D., & Liu, L. (2014). Contracts for changing times: sourcing with raw material price volatility and information asymmetry. *Manufacturing & Service Operations Management*, 16(1), 133–148.

Appendix

We start by summarizing in Table A1 the notation used in the article.

Table A1 Notation.

Indexes
a : source with a long-term contract, Arab Light, Saudi Arabia
h : a benchmark crude in which swap contracts are traded
j : a petroleum source
m : delivery time in the futures market, in months (0 stands for spot market)
o : Oman-Dubai petroleum market
s : scenario
t : a given month in the planning horizon
Objective function
L : Lagrangian objective function
$\pi(s, x)$; $\pi(s, \lambda_s)$; $\pi(s, q_{jts}, q_a, k_h)$: profit function in scenario s
ϕ : the weighted average of the expected profit and CVaR
Decision variables
k_h : proportion of capacity covered by the swap contract h
q_a : quantity of petroleum purchased per month using the long-term contract
q_{jts} : quantity of petroleum purchased from source j , in month t , and scenario s
z_s : the level of profit in the tail, for each scenario s
μ : value-at-risk
Dual variables
ξ_{ts} : dual variable of the equality constraint (4b)
γ : dual variable of constraint (4c)
λ_s : dual variable of constraint (4d)
δ_s : dual variable of constraint (4e)
σ_{jts} : dual variable of constraint (4f)
χ : dual variable of constraint (4g)
θ : dual variable of constraint (4h)
Parameters
g_h : the gross product worth of the benchmark crude h
g_{jts} : gross product worth from source j , in month t , and scenario s
r : marginal cost of refining per barrel
T : number of months in the planning horizon
w_h : the price of the benchmark crude h
w_{jts} : spot petroleum price from source j , in month t , and scenario s
w_{jt}^m : futures petroleum price, from j , at time t , to be delivered in m months
α : proportion of observation in the tail
β : expected profit weight ($1 - \beta$ is the CVaR weight)
ζ : expected available capacity, standardize to 1
η_{jts} : transportation cost from source j , in month t , and scenario s

Dual Problem (2) - Proof: Let $\lambda_s \geq 0$ be the dual variable of constraint (1b). For the primal problem (1), let the dual function be defined as $\psi = \inf \left(u - \frac{\sum_s z_s}{S\alpha} + \sum_s \lambda_s (z_s - u + \pi(s, x)) + \sum_s \delta_s z_s \right)$. Then, by taking the partial derivatives: from $\frac{\partial \psi}{\partial u} = 0$ we obtain $\sum_s \lambda_s = 1$; and from $\frac{\partial \psi}{\partial z_s} = 0$ it follows $\lambda_s + \delta_s = \frac{1}{S\alpha}$, which as $\delta_s \geq 0$, implies $\lambda_s \leq \frac{1}{S\alpha}$. And from $\frac{\partial \psi}{\partial x} = 0$ it follows $\sum_s \lambda_s \frac{\partial \pi(s, x)}{\partial x} = 0$. This function is then used to express x as a function of λ_s and to obtain the profit function as a function of λ_s as well: $\pi(s, \lambda_s)$. Therefore, the dual function simplifies to $\psi = \sum_s \pi(s, \lambda_s) \lambda_s$, under constraints $\sum_s \lambda_s = 1$, $\lambda_s \leq \frac{1}{S\alpha}$ and $\lambda_s \geq 0$, i.e., dual problem (2). \square

Let L in equation (A1) be the Lagrangian of linear program (4). From the Lagrangian, and given the convexity of the objective function ϕ , we derive the necessary and sufficient conditions for uniqueness of the optimal solution, as summarized in Proposition 1.

$$\begin{aligned} L(\mu, z_s, q_{jts}, q_a, k_h, \xi_{ts}, \gamma, \lambda_s, \delta_s, \sigma_{jts}, \chi, \theta) = & \phi + \sum_{ts} \xi_{ts} \left(\zeta - \sum_{j \neq a} q_{jts} - q_a \right) - \gamma (k_h - \zeta) \\ & - \sum_s \lambda_s (-z_s + \mu - \pi(s, q_{jts}, q_a, k_h)) - \sum_s \delta_s (-z_s) - \sum_{jts} \sigma_{jts} (-q_{jts}) - \chi (-q_a) - \theta (-k_h) \end{aligned} \quad (\text{A1})$$

Proof of Proposition 1: The necessary and sufficient conditions for a policy to be a global optimum of the linear program (4) are derived from the Lagrangian (A1): $\frac{\partial L}{\partial \mu} = 0$, which is equal to $\frac{\partial \phi}{\partial \mu} - \sum_s \lambda_s = 0$, and as $\frac{\partial \phi}{\partial \mu} = 1 - \beta$, it then follows that $\sum_s \lambda_s = 1 - \beta$. For all scenarios s : $\frac{\partial L}{\partial z_s} = 0$, which is equivalent to $\frac{\partial \phi}{\partial z_s} + \lambda_s + \delta_s = 0$ and as $\frac{\partial \phi}{\partial z_s} = -\frac{1-\beta}{S\alpha}$, we obtain $\lambda_s + \delta_s = \frac{1-\beta}{S\alpha}$. For all triples jts , such that $j \neq a$: $\frac{\partial L}{\partial q_{jts}} = 0$, this is equivalent to $\frac{\partial \phi}{\partial q_{jts}} - \xi_{ts} + \lambda_s \frac{\partial \pi(s, q_{jts}, q_a, k_h)}{\partial q_{jts}} + \sigma_{jts} = 0$ and, as $\frac{\partial \phi}{\partial q_{jts}} = \frac{\beta}{S} \frac{\partial \pi(s, q_{jts}, q_a, k_h)}{\partial q_{jts}}$ and $\frac{\partial \pi(s, q_{jts}, q_a, k_h)}{\partial q_{jts}} = g_{jts} - w_{jts} - r - \eta_{jts}$, it follows that, $(\frac{\beta}{S} + \lambda_s) (g_{jts} - w_{jts} - r - \eta_{jts}) - \xi_{ts} + \sigma_{jts} = 0$. Similarly, for the long-term contract supplier, a : $\frac{\partial L}{\partial q_a} = 0$ and this is equivalent to $\frac{\partial \phi}{\partial q_a} - \sum_{ts} \xi_{ts} + \sum_s \lambda_s \frac{\partial \pi(s, q_{jts}, q_a, k_h)}{\partial q_a} + \chi = 0$, and as $\frac{\partial \phi}{\partial q_a} = \frac{\beta}{S} \sum_s \frac{\partial \pi(s, q_{jts}, q_a, k_h)}{\partial q_a}$, we then obtain $\frac{\beta}{S} \sum_s \frac{\partial \pi(s, q_{jts}, q_a, k_h)}{\partial q_a} - \sum_{ts} \xi_{ts} + \sum_s \lambda_s \frac{\partial \pi(s, q_{jts}, q_a, k_h)}{\partial q_a} + \chi = 0$, which simplifies to $\sum_s (\frac{\beta}{S} + \lambda_s) \frac{\partial \pi(s, q_{jts}, q_a, k_h)}{\partial q_a} - \sum_{ts} \xi_{ts} + \chi = 0$. Consequently, as $\frac{\partial \pi(s, q_{jts}, q_a, k_h)}{\partial q_a} = \sum_t (g_{ats} - w_{ats} - r - \eta_{ats})$, it follows that: $\sum_{ts} (\frac{\beta}{S} + \lambda_s) (g_{ats} - w_{ats} - r - \eta_{ats}) - \sum_{ts} \xi_{ts} + \chi = 0$. For the crack spread swap contracts, from the necessary condition $\frac{\partial L}{\partial k_h} = 0$, it follows that $\frac{\partial \phi}{\partial k_h} + \sum_s \lambda_s \frac{\partial \pi(s, q_{jts}, q_a, k_h)}{\partial k_h} - \gamma + \theta = 0$, and as $\frac{\partial \phi}{\partial k_h} = \frac{\beta}{S} \sum_s \frac{\partial \pi(s, q_{jts}, q_a, k_h)}{\partial k_h}$, we get $\frac{\beta}{S} \sum_s \frac{\partial \pi(s, q_{jts}, q_a, k_h)}{\partial k_h} + \sum_s \lambda_s \frac{\partial \pi(s, q_{jts}, q_a, k_h)}{\partial k_h} - \gamma + \theta = 0$, which is equivalent to $\sum_s (\frac{\beta}{S} + \lambda_s) \frac{\partial \pi(s, q_{jts}, q_a, k_h)}{\partial k_h} - \gamma + \theta = 0$. Subsequently, as $\frac{\partial \pi(s, q_{jts}, q_a, k_h)}{\partial k_h} = \sum_t ((g_h - g_{hts}) - (w_h - w_{hts}))$ it follows that $\sum_{ts} (\frac{\beta}{S} + \lambda_s) ((g_h - g_{hts}) - (w_h - w_{hts})) - \gamma + \theta = 0$. From the equality constraint, for all t and for all s , we have $\sum_{j \neq a} q_{jts} + q_a = \zeta$. Finally, we require the complementarity constraints: for all s , $0 \leq z_s - \mu + \pi(s, q_{jts}, q_a, k_h) \perp \lambda_s \geq 0$ and $0 \leq z_s \perp \delta_s \geq 0$; for all triples jts , $0 \leq q_{jts} \perp \sigma_{jts} \geq 0$; $0 \leq q_a \perp \chi \geq 0$; $0 \leq \zeta - k_h \perp \gamma \geq 0$ and $0 \leq k_h = 0 \perp \theta \geq 0$. \square

Proof of Proposition 2: By summing $\frac{\beta}{S} + \lambda_s$ for all s , we get $\sum_s (\frac{\beta}{S} + \lambda_s) = \beta + \sum_s \lambda_s$. Then as from Proposition 1 we know that $\sum_s \lambda_s = 1 - \beta$, we equivalently obtain $\sum_s (\frac{\beta}{S} + \lambda_s) = \beta + 1 - \beta$ and therefore $\sum_s (\frac{\beta}{S} + \lambda_s) = 1$. \square

Proof of Proposition 3: From Proposition 1 we know that: $\sum_s \lambda_s = 1 - \beta$; for all s , $\lambda_s + \delta_s = \frac{1-\beta}{S\alpha}$; for all s , $0 \leq z_s - \mu + \pi(s, q_{jts}, q_a, k_h) \perp \lambda_s \geq 0$ and $0 \leq z_s \perp \delta_s \geq 0$. We then need to understand which α and β pairs lead to the same solutions. As we are analyzing states on the tail of the profit function, such that

$z_s > 0$ and $\delta_s = 0$, as for refinery 0 the risk profile is described by $\beta^0 = 0$, then $\lambda_s^0 + \delta_s^0 = \frac{1-\beta^0}{S\alpha}$ is equal to $\lambda_s^0 = \frac{1}{S\alpha^0}$. a) To derive the Iso- λ curve, we need to find an α and a β such that $\lambda_s^0 = \frac{1-\beta}{S\alpha}$, which is equivalent to $\alpha = \frac{1}{S\lambda_s^0} - \frac{1}{S\lambda_s^0}\beta$. b) As $\lambda = \frac{1-\beta}{S\alpha}$, refinery 0 with $\beta^0 = 0$ can replicate this λ if and only if $\frac{1}{S\alpha^0} = \frac{1-\beta}{S\alpha}$, which is equivalent to $\alpha^0 = \frac{\alpha}{1-\beta}$. However, as $1 \geq \alpha > 1 - \beta \geq 0$, it follows that $\alpha^0 > 1$, creating a contradiction with the initial condition $0 < \alpha^0 \leq 1$. \square

Proof of proposition 4: From Proposition 1 we have the following results. For all triples jts , such that $j \neq a$, $(\frac{\beta}{S} + \lambda_s)(g_{jts} - w_{jts} - r - \eta_{jts}) - \xi_{ts} + \sigma_{jts} = 0$, which is equivalent to $(\frac{\beta}{S} + \lambda_s)(g_{jts} - w_{jts} - r - \eta_{jts}) + \sigma_{jts} = \xi_{ts}$. Therefore, by summing over all ts for a source $j \neq a$, we have $\sum_{ts} \{(\frac{\beta}{S} + \lambda_s)(g_{jts} - w_{jts} - r - \eta_{jts}) + \sigma_{jts}\} = \sum_{ts} \xi_{ts}$, and by summing over the $J - 1$ sources that are not a , we derive $\frac{\sum_{j \neq a, ts} \{(\frac{\beta}{S} + \lambda_s)(g_{jts} - w_{jts} - r - \eta_{jts}) + \sigma_{jts}\}}{J-1} = \sum_{ts} \xi_{ts}$. For the long-term contract supplier, a : $\sum_{ts} (\frac{\beta}{S} + \lambda_s)(g_{ats} - w_{ats} - r - \eta_{ats}) - \sum_{ts} \xi_{ts} + \chi = 0$. As $\chi = 0$, by plugging in the optimal condition for the sum over the sources $j \neq a$ into the long-term contract equation, we get $\sum_{ts} (\frac{\beta}{S} + \lambda_s)(g_{ats} - w_{ats} - r - \eta_{ats}) = \frac{\sum_{j \neq a, ts} \{(\frac{\beta}{S} + \lambda_s)(g_{jts} - w_{jts} - r - \eta_{jts}) + \sigma_{jts}\}}{J-1}$, which is equivalent to $\sum_{ts} (\frac{\beta}{S} + \lambda_s)(g_{ats} - w_{ats} - r - \eta_{ats}) - \frac{\sum_{j \neq a, ts} \{(\frac{\beta}{S} + \lambda_s)(g_{jts} - w_{jts} - r - \eta_{jts})\}}{J-1} = \frac{\sum_{j \neq a, ts} \sigma_{jts}}{J-1}$. \square

Proof of proposition 5: From Proposition 1.E) we know that for the crack spread swap contracts, $\sum_{ts} (\frac{\beta}{S} + \lambda_s)((g_h - g_{hts}) - (w_h - w_{hts})) = \gamma - \theta$. As explained in Figure 2, we know that a refinery trades crack spread swap contracts if and only if $\gamma - \theta \geq 0$, which is equivalent to $\sum_{ts} (\frac{\beta}{S} + \lambda_s)((g_h - g_{hts}) - (w_h - w_{hts})) \geq 0$ and to $\sum_{ts} (\frac{\beta}{S} + \lambda_s)(g_h - w_h) \geq \sum_{ts} (\frac{\beta}{S} + \lambda_s)(g_{hts} - w_{hts})$, from which we get $(g_h - w_h) \sum_{ts} (\frac{\beta}{S} + \lambda_s) \geq \sum_{ts} (\frac{\beta}{S} + \lambda_s)(g_{hts} - w_{hts})$. As from Proposition 2 we know that $\sum_s (\frac{\beta}{S} + \lambda_s) = 1$, it therefore follows that $(g_h - w_h) \geq \frac{1}{T} \sum_{ts} (\frac{\beta}{S} + \lambda_s)(g_{hts} - w_{hts})$. \square

Proof of Proposition 6: From Proposition 5 we know that $k_h^* > 0$ if and only if $(g_h - w_h) \geq \frac{1}{T} \sum_{ts} (\frac{\beta}{S} + \lambda_s)(g_{hts} - w_{hts})$. Moreover, given the one-to-one correspondence between the extreme total and swap contract losses, summarized in Postulate 1, this condition can be decomposed into three components: $(g_h - w_h) \geq \frac{1}{T} \sum_{t,s:z_s>0} (\frac{\beta}{S} + \lambda_s)(g_{hts} - w_{hts}) + \frac{1}{T} \sum_{t,s:z_s=0,\mu-\pi_s=0} (\frac{\beta}{S} + \lambda_s)(g_{hts} - w_{hts}) + \frac{1}{T} \sum_{t,s:z_s=0,\mu-\pi_s<0} (\frac{\beta}{S} + \lambda_s)(g_{hts} - w_{hts})$.

From Proposition 1.B) we know that $\lambda_s + \delta_s = \frac{1-\beta}{S\alpha}$, and from Proposition 2 we know that the probability associated with a state s is equal to $\frac{\beta}{S} + \lambda_s$. There are three possible types of risk-adjusted probabilities, for any given s : $z_s > 0$ from which we obtain $\lambda_s = \frac{1-\beta}{S\alpha}$ and $\frac{\beta}{S} + \lambda_s = \frac{1-(1-\alpha)\beta}{S\alpha}$; $z_s = 0$ and $\mu - \pi(s, q_{jts}, q_a, k_h) < 0$, which implies, from Proposition 1, $\lambda_s = 0$, $\delta_s = \frac{1-\beta}{S\alpha}$ and the weight put on s is equal to $\frac{\beta}{S} + \lambda_s = \frac{\beta}{S}$; $z_s = 0$ and $\mu - \pi(s, q_{jts}, q_a, k_h) = 0$, which implies, from Proposition 1.B), $\lambda_s + \delta_s = \frac{1-\beta}{S\alpha}$ and $\frac{\beta}{S} + \lambda_s = \frac{1-(1-\alpha)\beta}{S\alpha} - \delta_s$ and, consequently, $\frac{\beta}{S} + \lambda_s \leq \frac{1-(1-\alpha)\beta}{S\alpha}$.

Therefore, to explain how the purchasing of swap contracts depends on β we need to analyze the change in $\frac{\beta}{S} + \lambda_s$, for a risk averse refinery ($\beta < 1$ and $\alpha < 1$), as β converges to zero. a) When $z_s > 0$, $\frac{\beta}{S} + \lambda_s = \frac{1-(1-\alpha)\beta}{S\alpha}$ and $\lim_{\beta \rightarrow 0} \frac{1-(1-\alpha)\beta}{S\alpha} = \frac{1}{S\alpha}$. As $\alpha < 1$, then $\frac{1}{S\alpha} > \frac{1}{S}$, and $\frac{\beta}{S} + \lambda_s$ increases from $\frac{1}{S}$ (when $\beta = 1$ and $\alpha = 1$, risk-neutral refinery) to $\frac{1}{S\alpha}$ as β converges to 0 (risk-averse refinery). Intuitively, this means that the probabilities get more concentrated in a small group of scenarios with larger z_s as the refinery becomes more risk-averse. b) When $z_s = 0$ and $\mu - \pi(s, q_{jts}, q_a, k_h) < 0$, $\frac{\beta}{S} + \lambda_s = \frac{\beta}{S}$ and the $\lim_{\beta \rightarrow 0} \frac{\beta}{S} = 0$. This means that the risk-averse refinery places less weight on the states that are not in the profit distribution tail.

Then, as when $z_s = 0$ and $\mu - \pi(s, q_{jts}, q_a, k_h) = 0$, $\frac{\beta}{S} + \lambda_s \leq \frac{1-(1-\alpha)\beta}{S\alpha}$, a stronger condition is obtained by using only the two major components: $(g_h - w_h) \geq \frac{1}{T} \sum_{t,s:z_s \geq 0, \mu - \pi_s \geq 0} \left(\frac{\beta}{S} + \lambda_s\right) (g_{hts} - w_{hts}) + \frac{1}{T} \sum_{t,s:z_s=0, \mu - \pi_s < 0} \left(\frac{\beta}{S} + \lambda_s\right) (g_{hts} - w_{hts})$, with $\sum_s \lambda_s = 1 - \beta$.

From a) and b), this is equivalent to $(g_h - w_h) \geq \frac{1}{T} \sum_{t,s:z_s \geq 0, \mu - \pi_s \geq 0} \frac{1-(1-\alpha)\beta}{S\alpha} (g_{hts} - w_{hts}) + \frac{1}{T} \sum_{t,s:z_s=0, \mu - \pi_s < 0} \frac{\beta}{S} (g_{hts} - w_{hts})$. Then, as $\lim_{\beta \rightarrow 0} \frac{1-(1-\alpha)\beta}{S\alpha} = \frac{1}{S\alpha}$ and $\lim_{\beta \rightarrow 0} \frac{\beta}{S} = 0$, we obtain, $\lim_{\beta \rightarrow 0} (g_h - w_h) \geq \lim_{\beta \rightarrow 0} \frac{1}{T} \sum_{t,s:z_s \geq 0, \mu - \pi_s \geq 0} \frac{1-(1-\alpha)\beta}{S\alpha} (g_{hts} - w_{hts}) + \lim_{\beta \rightarrow 0} \frac{1}{T} \sum_{t,s:z_s=0, \mu - \pi_s < 0} \frac{\beta}{S} (g_{hts} - w_{hts})$ and equivalently $(g_h - w_h) \geq \frac{1}{T} \sum_{t,s:z_s \geq 0, \mu - \pi_s \geq 0} \frac{1}{S\alpha} (g_{hts} - w_{hts})$. Consequently, as β converges to 0, the extreme losses $\frac{1}{T} \sum_{t,s:z_s \geq 0, \mu - \pi_s \geq 0} \frac{1-(1-\alpha)\beta}{S\alpha} (g_{hts} - w_{hts})$ receive a higher weight, and the profits (or lower losses), i.e., $\frac{1}{T} \sum_{t,s:z_s=0, \mu - \pi_s < 0} \frac{\beta}{S} (g_{hts} - w_{hts})$ get a lower weight (and eventually no weight when $\beta = 0$). For this reason, the condition for $k_h^* > 0$, i.e., $(g_h - w_h) \geq \frac{1}{T} \sum_{ts} \left(\frac{\beta}{S} + \lambda_s\right) (g_{hts} - w_{hts})$, becomes harder to meet as β converges to 0, and the average proportion of swap contracts purchased decreases. \square

Acknowledgments

I would like to acknowledge the financial support by KAPSARC, Saudi Arabia. I am deeply grateful to the editor Xenophon Koufteros for his support and to the department editor Nicholas Petruzzi, also to an anonymous associate editor and referees for their constructive comments. I would also like to thank Jennifer Considine, Tava Olsen, and Axel Pierru for their comments on earlier versions of the manuscript, and to Ziyad Fawzan for his help with data collection. All content is my responsibility alone.