

Capacity Management in Agricultural Commodity Processing and Application in the Palm Industry

Author:

Boyabatlı, O; Nguyen, J; Wang, T

Publication details:

Manufacturing & Service Operations Management

v. 19

Chapter No. 4

pp. 551 - 567

1523-4614 (ISSN); 1526-5498 (ISSN)

Publication Date:

2017-10-01

Publisher DOI:

<https://doi.org/10.1287/msom.2017.0624>

License:

<https://creativecommons.org/licenses/by-nc-nd/4.0/>

Link to license to see what you are allowed to do with this resource.

Downloaded from http://hdl.handle.net/1959.4/unsworks_49801 in <https://unsworks.unsw.edu.au> on 2024-05-01



Manufacturing & Service Operations Management

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Capacity Management in Agricultural Commodity Processing and Application in the Palm Industry

Onur Boyabatlı, Jason Nguyen, Tong Wang

To cite this article:

Onur Boyabatlı, Jason Nguyen, Tong Wang (2017) Capacity Management in Agricultural Commodity Processing and Application in the Palm Industry. *Manufacturing & Service Operations Management* 19(4):551-567. <https://doi.org/10.1287/msom.2017.0624>

Full terms and conditions of use: <http://pubsonline.informs.org/page/terms-and-conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2017, INFORMS

Please scroll down for article—it is on subsequent pages



INFORMS is the largest professional society in the world for professionals in the fields of operations research, management science, and analytics.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Capacity Management in Agricultural Commodity Processing and Application in the Palm Industry

Onur Boyabatli,^a Jason Nguyen,^b Tong Wang^c

^a Lee Kong Chian School of Business, Singapore Management University, Singapore 178899; ^b Carlson School of Management, University of Minnesota, Minneapolis, Minnesota 55455; ^c NUS Business School, National University of Singapore, Singapore 119245

Contact: oboyabatli@smu.edu.sg,  <http://orcid.org/0000-0002-8374-8318> (OB); nguy1762@umn.edu (JN); tong.wang@nus.edu.sg (TW)

Received: October 24, 2011

Revised: December 4, 2015; July 24, 2016;
December 5, 2016; January 15, 2017

Accepted: January 16, 2017

Published Online in Articles in Advance:
August 10, 2017

<https://doi.org/10.1287/msom.2017.0624>

Copyright: © 2017 INFORMS

Abstract. This paper examines the capacity investment decisions of a processor that uses a commodity input to produce both a commodity output and a by-product in the context of agricultural industries. We employ a multiperiod model to study the optimal one-time processing and (output) storage capacity investment decisions—in addition to the periodic processing and inventory decisions—when both input and output spot prices as well as production yield are uncertain. We characterize the optimal decisions and perform sensitivity analysis to investigate how spot price uncertainty affects the processor's optimal capacity and profitability. Using a calibration based on the palm industry, we study (both numerically and analytically) the performance of a variety of heuristic capacity investment policies that can be used in practice. We find that if the yield uncertainty is ignored in capacity planning, then basing those plans on the *average* yield is preferable to basing them (as often occurs in practice) on the *maximum* yield. However, planning based on the average yield performs well only when the relative (processing-to-storage) capacity investment cost is high; otherwise, it leads to a significant loss of profit. We also find that ignoring spot price uncertainty in capacity planning results in a relatively small profit loss. In contrast, ignoring by-product revenue—which constitutes a small portion of total revenues—during capacity planning substantially reduces the processor's profit.

Funding: The work of the third author was supported by the Singapore Ministry of Education Academic Research Fund Tier 1 [Grants WBS R-314-000-098-112 and R-314-000-104-112].

Supplemental Material: The online appendix is available at <https://doi.org/10.1287/msom.2017.0624>.

Keywords: capacity management • multiproduct firm • commodity risk management • spot market • agriculture • dynamic programming • processing • storage

1. Introduction

In this paper, we study the capacity investment decisions of a processor that—in the context of agricultural industries—uses a primary commodity input to produce a commodity output as well as a by-product. In particular, we analyze investment decisions related to input processing capacity and output storage capacity. Our analysis is applicable to several agricultural industries, including the oilseed industry (e.g., palm, soybean, rapeseed, sunflower seed, coconut) and the grain industry (e.g., corn and wheat).

Consider, for example, the palm industry. In this industry, palm oil mills produce crude palm oil (CPO; a commodity output) and palm kernel (a by-product) from palm fresh fruit bunches (FFBs; a commodity input). As seen in Table 11 of the U.S. Department of Agriculture report on oilseeds (2015), palm is the largest oilseed industry globally, with 59.39 million metric tons of crude palm oil produced between 2013 and 2014, the estimated market value of which is more than US\$49 billion. In a palm oil mill, the palm fresh fruit bunches go through several processing stations

(receiving, sterilization, threshing, pressing, and centrifugation) to produce palm kernel and crude palm oil. Crude palm oil is transferred to storage tanks prior to dispatch from the mill. The processing volume of the palm fresh fruit bunches is constrained by the joint capacity of the processing stations, while crude palm oil production and inventory volume is constrained by storage tank capacity. It follows that choosing the optimal levels of processing and storage capacity is critical for a mill's profitability. Similar capacity investment decisions are of relevance to other oilseed processors, which produce crude vegetable oil and meal or cake, and grain processors, which produce biofuel and animal feed.

In the operations management (OM) literature there is a vast amount of research that studies capacity investment decisions in processing environments (for a review, see Van Mieghem 2003), but a very limited amount of this research in the context of agricultural industries. Among the few papers that focus on agricultural industries (e.g., Allen and Schuster 2004), there is no work that considers the processing of a commodity product—a common characteristic of the

majority of agricultural products in practice. To this end, the literature most relevant to our paper is the OM research on commodity processing. The papers in this field examine operating decisions (e.g., processing and inventory) of a commodity processor in a variety of models. These studies capture the idiosyncratic features of different commodity markets, including those for electricity (Zhou et al. 2014), electronic equipment (Pei et al. 2011), metals (Plambeck and Taylor 2013), natural gas (Secomandi 2010a, b; Lai et al. 2011), petroleum (Dong et al. 2014), and semiconductors (Kleindorfer and Wu 2003), as well as commodity markets associated with such agricultural industries as beef (Boyabatlı et al. 2011), citrus fruit (Kazaz and Webster 2011), cocoa (Boyabatlı 2015), corn (Goel and Tanrisever 2017), olives (Kazaz 2004), processed food (Mehrotra et al. 2011), and soybeans (Devalkar et al. 2011, 2017). Because the focus of these papers is on operating decisions, they either assume (often implicitly) abundant processing and storage resources or consider fixed capacity levels for these resources. In summary, there is no work that studies the joint processing and storage capacity investment decisions of commodity processors in agricultural industries. In this paper, we attempt to fill this void.


Processors in agricultural industries feature unique characteristics that present challenges for capacity management. *First*, since both the input and the output are commodities, there exist regional exchange or “spot” markets (Devalkar et al. 2011). So in buying and selling these commodities, processors are exposed to prevailing spot prices. The input and output spot prices are closely linked and exhibit considerable variability, as shown for the palm industry in Figure 1(a). The uncertainty in spot prices may affect capacity investment decisions because the profit from processing depends on those prices. Moreover, the processor

can hold inventory for sale at a later date to benefit from fluctuations in output spot price, where that inventory can be sourced from in-house production (Fackler and Livingston 2002) and the spot market (Kouvelis et al. 2013). *Second*, there is some uncertainty also in the production yield (extraction rate) from each input, as plotted (again for the palm industry) in Figure 1(b). This uncertainty is driven by several factors that include weather conditions and the extent of pests and diseases during the input’s growing period (Boyabatlı et al. 2016), the harvest timing of the input, and the processing technology used (Chang et al. 2003). The uncertainty in production yield may affect capacity investment decisions because profits from processing depend also on that yield.

Given these characteristics, our first objective is to study how the processor should determine the optimal levels of investment in processing and storage capacity. Because there is substantial variability in input and output spot prices observed in practice, our second objective is to investigate how spot price uncertainty affects the processor’s optimal capacity and profitability. Our final research objective is to examine the performance of heuristic capacity investment policies that are already used—or that can be used—in practice in comparison with the optimal capacity investment policy. Because these heuristic policies ignore some operational factors (e.g., production yield, by-product revenue) during capacity planning, this performance comparison is instrumental in understanding the criticality of these operational factors for capacity investment to generate valuable managerial insights.

To achieve these objectives, we model the processor’s decisions as a multiperiod optimization problem in which the firm (i) procures an input commodity, where the marginal procurement cost equals the commodity’s spot price; (ii) sells an output commodity,

Figure 1. (Color online) Characteristics of Fresh Fruit Palm Bunches and Crude Palm Oil in the Malaysian Peninsula for the Period January 2006 to December 2013, as Reported by the Malaysian Palm Oil Board



Note. Prices are reported in Malaysian ringgits (RM) per metric ton. Production yields (extraction rates) are reported as percentages.

where the marginal sales revenue equals this commodity's spot price; and (iii) sells a by-product that has a fixed marginal sales revenue. The output can also be procured from the spot market for storage and speculative sale with the marginal procurement cost equal to the output spot price. The firm maximizes its expected total profit over a finite planning horizon. At the beginning of this horizon, the firm chooses input processing and output storage capacity levels. In the rest of the planning horizon, constrained by these capacity levels, the firm periodically makes decisions about the processing volume and output inventory. More specifically, in each period, the processing volume is chosen with respect to production yield uncertainty, and the output inventory level is chosen after this uncertainty is realized.

We characterize the optimal levels of investment in processing and storage capacity (as well as the periodic processing and inventory decisions) in closed form. We distinguish two optimal capacity investment strategies based on the investment cost of processing capacity relative to storage capacity. When that relative cost is sufficiently high, the firm invests in a *storage-dominating portfolio*, where the storage capacity is strictly greater than what is required for production (with full utilization of processing capacity) under all yield realizations. When that relative cost is sufficiently low, the firm invests in a *high yield-balanced portfolio (HYBP)*, where the processing capacity is at the level required for production (with full utilization of storage capacity) under the maximum yield. We complement our structural analysis with numerical analysis by calibrating our model to represent a typical palm oil mill. We use publicly available data from the Malaysian Palm Oil Board (MPOB) as well as publicly available and proprietary data from palm oil mills located in Malaysia. Our main findings and their contribution can be summarized as follows.

(1) We conduct sensitivity analyses, both analytically and numerically, to investigate the effects of correlation between input and output spot prices and their respective volatility. We find that the processor always benefits from a lower correlation but benefits from a lower input or output price volatility when this volatility is low; otherwise, a higher volatility is beneficial. These results are reminiscent of the sensitivity results of Plambeck and Taylor (2013), who examine the effect of spot price uncertainty on a clean-tech manufacturer's profitability, and Dong et al. (2014), who examine the effect of spot price uncertainty on the value of operational flexibility in the context of an oil refinery. In both papers, only a single period is modeled. We extend the sensitivity analyses to a multiperiod setting. In addition, we examine the effect of spot price uncertainty on capacity investment decisions of a typical

processor. We find that the optimal processing capacity decreases with an increase in price correlation. The optimal processing capacity also decreases with an increase in output or input price volatility, but only when this volatility is low; otherwise, it increases with an increase in that volatility. In contrast, the optimal storage capacity increases with an increase in output price volatility but it is not affected otherwise. These results showcase the significant differences in how spot price uncertainty affects each capacity type.

(2) We study the performance of a variety of heuristic capacity investment policies in comparison with the optimal capacity investment policy. To this end, we numerically compute the profit loss due to employing the heuristic policy and also provide analytical bounds on this profit loss. We find that should the production yield uncertainty be ignored in capacity planning, rather than making the planning based on the maximum yield, as is often done in practice in the palm industry, it is better to plan based on the average yield. However, planning based on the average yield performs well only when the relative (processing-to-storage) capacity investment cost is sufficiently high; otherwise, it leads to considerable profit loss (an average profit loss of 14.53% in the numerical instances considered). We also find that ignoring spot price uncertainty in capacity planning results in a relatively small profit loss (an average profit loss of 5.87% in the numerical instances considered). Another finding of interest is that the processor's profits are substantially reduced if its capacity planning ignores by-product revenue (a minimum profit loss of 61.86% in the numerical instances considered)—even though that revenue accounts for just a small portion of the firm's total revenue. Based on our theoretical analysis, as a heuristic policy, we propose setting storage capacity at the level required for production (with full utilization of processing capacity) under the maximum yield. Because this policy assumes a particular relationship between processing capacity and storage capacity, it provides an operational simplification in making capacity investment decisions. We show that this heuristic policy is nearly optimal (an average profit loss of only 0.57% in the numerical instances considered). These results contribute to the OM research on commodity processing. The papers in this field (e.g., Devalkar et al. 2011, Lai et al. 2011, Zhou et al. 2014) compare the performance of heuristic operating policies with the performance of the optimal policy and propose near-optimal decision rules for making operating decisions. Our focus is similar but applied to capacity investment decisions. Our results offer insights with potential practical relevance for both optimal and approximate capacity investment decisions in agricultural processing.

The rest of this paper proceeds as follows. Section 2 describes the model and the basis for our assumptions,

and Section 3 derives the optimal strategy. Section 4 structurally examines the effects of spot price uncertainty on the optimal capacity investment policy and on firm profitability. It also compares the optimal policy's performance with that of heuristic policies. Section 5 conducts the same analysis numerically using a model calibration that represents a typical processor in the context of the palm industry. Section 6 concludes with a discussion of the limitations of our analysis. All proofs are relegated to the online technical appendix.

2. Model Description

We use the following notation and conventions throughout this paper. A realization of the random variable \tilde{y} is denoted by y . The expectation operator is denoted by \mathbb{E} . We use $(u)^+ = \max(u, 0)$. Boldface letters represent row vectors. The monotonic relations (increasing, decreasing) are used in the weak sense. Subscript t denotes period t . Superscript I denotes input-related parameters and decision variables, while superscript O (B) denotes the parameters and variables related to the output (by-product).

We consider a firm that uses a commodity input to produce a commodity output and a by-product and that seeks to maximize its expected total (discounted) profit over a finite planning horizon. The firm operates under capacity constraints related to input processing and output storage. At the beginning of the planning horizon, the firm chooses its level of investment in each type of capacity. In the rest of the horizon, the firm makes periodic decisions about processing volume and output inventory subject to its chosen capacity levels.

Let $\mathbf{K} = (K^I, K^O)$ denote the firm's capacity investment portfolio, where K^I is the input processing capacity and K^O is the output storage capacity. In agricultural industries, the processors are located near the plantations where the input originates; because the available land is scarce, the marginal capacity investment cost is increasing in the capacity level. In line with this observation, we assume that the capacity investment cost $C(\mathbf{K})$ is convex increasing in \mathbf{K} .¹ Specifically, we consider a quadratic cost with parameters β^I and β^O : $C(\mathbf{K}) \doteq \beta^I (K^I)^2 + \beta^O (K^O)^2$. The structural analysis of Section 3 holds also for a general convex increasing $C(\mathbf{K})$ (except for the closed-form characterization of the capacity portfolio, which requires a specific functional form).

We assume the marginal procurement cost of the commodity input to be given by that commodity's spot price. In practice, this case is relevant when the input is procured through an exchange (spot) market or when the input is procured through bilateral contracts under which the unit price is benchmarked to the exchange market price. Similarly, we assume that the marginal sales revenue of the commodity output is given by its spot price. Input and output spot prices are assumed

to follow correlated Markovian stochastic processes; in other words, current spot price realizations are sufficient to characterize the distribution of future spot prices. We defer the specification of these stochastic processes to Section 4 because the structural analysis in Section 3 is not affected thereby.² The firm may also procure output from the spot market for the purpose of storage and speculative sale. We assume that the marginal procurement cost is given by the spot price. Output storage incurs a per-period unit holding cost of h .

We consider a per-period unit processing cost $c > 0$. For each unit of the processed input, the production yield of the output (by-product) is given by a (a^B). We assume that $a^B \in (0, 1)$ is constant and that the by-product is not stored, but sold at a fixed unit price p^B . Hence, $c \doteq c - a^B p^B$ is the effective processing cost, which can be negative if the by-product revenue is sufficiently high. The output production yield \tilde{a} is uncertain and independent and identically distributed across periods. We assume that \tilde{a} is statistically independent of the spot price processes, which is a reasonable assumption in the palm industry (as verified empirically in Section 5). We consider a Bernoulli distribution for \tilde{a} : $a = a^I$ with probability $q \in [0, 1]$ and $a = a^H$ with probability $1 - q$ for $0 < a^I < a^H \leq 1 - a^B$, where the last inequality follows because the overall production yield cannot exceed 1. Let $\bar{a} = qa^I + (1 - q)a^H$ denote the average production yield. The structural analysis in Section 3 will also hold for a general discrete distribution of \tilde{a} with more than two realizations.

The storage capacity K^O affects processing activities because output is placed in the storage facility before being dispatched from the plant. Profitability will decline if the output yield from processing exceeds the available storage capacity. This reduction in profitability can result from the cost associated with process interruption due to retrieving the excess output from the facility or the cost of using temporary storage tanks to handle the excess output. The reduction in profitability can also be driven by the decline in the marginal sales revenue due to the output's inferior quality as a result of improper storage conditions. In the oilseeds industries, for example, the quality of crude vegetable oil decreases not only with metal (e.g., iron) contamination, which occurs when the storage facility is not lined with suitable protective coating, but also with solidification and fractionation, which occur if the storage facility cannot maintain a specific temperature. In the palm industry, palm oil mills in practice minimize quality issues related to storage conditions by planning operations in such a way that the entire volume of output (crude palm oil) goes through the storage facility after processing. In line with this observation, we assume that the firm adopts a policy of no excess production; that is, processing volume is chosen

in such a way that the available output storage capacity is sufficient under all yield realizations.³

We formulate the firm's problem as a finite-horizon stochastic dynamic program. The per-period processing capacity K^I (in units of, say, metric tons of input per day) and the storage capacity K^O (metric tons of output) are determined at period $t = 0$, and they are fixed in all subsequent periods. In each period $t \in [1, T]$, the sequence of events is as follows:

1. At the beginning of period t , the firm observes the input and the output spot prices $\mathbf{P}_t = (p_t^I, p_t^O)$ as well as the output inventory level s_{t-1} (carried from period $t - 1$). The firm then decides on the input processing volume z_t within the processing capacity level K^I , and considering the available output storage capacity $K^O - s_{t-1}$ due to policy of no excess production.

2. The production yield a is realized, which determines the available output volume, and the firm decides on the output inventory level s_t within the storage capacity K^O . Required inventory that is not provided by the available output volume is procured from the spot market; output volume that is not stored is sold to the spot market.

The firm's immediate payoff in period $t \in [1, T]$ is given by

$$L(z_t, s_t | s_{t-1}, \mathbf{P}_t) \\ \doteq -p_t^I z_t - c z_t - h s_t + \mathbb{E}_{\tilde{a}}[-p_t^O(s_t - (s_{t-1} + \tilde{a} z_t))^+ \\ + p_t^O(s_{t-1} + \tilde{a} z_t - s_t)^+]. \quad (1)$$

In (1), the first two terms capture the effective processing and procurement cost, the third term denotes the inventory holding cost, and the last term expresses the expected cash flows resulting from the realized production yield. The first term in these expected cash flows denotes the spot procurement cost for the inventory level beyond the available output volume, and the second term denotes the spot sale revenue for the available output volume that is not stored.

Let $V_t(s_{t-1}, \mathbf{P}_t)$ for $t \in [1, T]$ be the optimal value function from period t onward given s_{t-1} and \mathbf{P}_t ; this function satisfies

$$V_t(s_{t-1}, \mathbf{P}_t) \\ = \max_{z_t \geq 0, s_t \geq 0} \{L(z_t, s_t | s_{t-1}, \mathbf{P}_t) + \delta \mathbb{E}_t[V_{t+1}(s_t, \tilde{\mathbf{P}}_{t+1})]\} \\ \text{s.t. } z_t \leq \min\left(K^I, \frac{K^O - s_{t-1}}{a^h}\right), \quad s_t \leq K^O, \quad (2)$$

with boundary condition $V_{T+1}(s_T, \mathbf{P}_{T+1}) = 0$ and initial inventory level $s_0 = 0$, where $\delta \in [0, 1]$ is the discounting factor and $\mathbb{E}_t[\cdot]$ is our shorthand notation for $\mathbb{E}[\cdot | \mathbf{P}_t]$. In (2), the constraint $z_t \leq (K^O - s_{t-1})/a^h$ captures the no excess production assumption, where $(K^O - s_{t-1})/a^h$ denotes the input volume required to fill the available output storage capacity under the high yield realization.

At period $t = 0$, the firm observes \mathbf{P}_0 and chooses $\mathbf{K} = (K^I, K^O)$, thereby incurring the capacity investment cost $C(\mathbf{K}) = \beta^I (K^I)^2 + \beta^O (K^O)^2$. The firm's optimal expected total (discounted) profit over the planning horizon is given by $\Pi^* = \max_{\mathbf{K} \geq 0} \delta \mathbb{E}_0[V_1(0, \tilde{\mathbf{P}}_1)] - C(\mathbf{K})$.

3. Characterization of the Optimal Strategy

In this section, we describe the firm's optimal strategy. In particular, we first characterize the periodic input processing and output inventory decisions (Section 3.1) and then characterize the optimal capacity investment decisions (Section 3.2).


To facilitate the analysis, we make the following observation. It costs the same to source inventory from the output available in-house (i.e., the realized output yield after processing combined with the inventory carried over from the previous period) as it does from the output spot market: in both cases, the cost is the prevailing output spot price. For the in-house case, this cost is the opportunity cost of not selling to the spot market (spot sale revenue); when sourced from outside, it is the spot procurement cost. Therefore, the firm's immediate payoff in period $t \in [1, T]$, as given by (1), can be decoupled into two components:

$$L_{\text{pr}}(z_t | s_{t-1}, \mathbf{P}_t) \doteq (-p_t^I - c + \tilde{a} p_t^O) z_t + p_t^O s_{t-1}, \\ L_{\text{sc}}(s_t | \mathbf{P}_t) \doteq (p_t^O + h) s_t,$$

where the subscripts "pr" and "sc" refer, respectively, to "processing return" and "storage cost." This decoupling suggests that one can view same-period processing and inventory decisions as being independent. Thus, the firm first decides on the processing volume to sell to the output spot market (together with the inventory carried over from the previous period), which generates the processing return $L_{\text{pr}}(z_t | s_{t-1}, \mathbf{P}_t)$, and then chooses the output inventory level s_t to source from the spot market incurring the storage cost $L_{\text{sc}}(s_t | \mathbf{P}_t)$. The processing decision does not affect any other decisions, and the inventory decision affects only the subsequent period's processing decision through limiting the available storage capacity. Therefore, the optimization problem given by (2) can be written in terms of independent two-stage optimization problems by grouping the inventory decision in period $t - 1$ with the processing decision in period t ; see Figure 2.

Because inventory is not needed in period T , $L_{\text{sc}}(\cdot | \mathbf{P}_T) = 0$, and so the optimal value function in period $t \in [1, T - 1]$ can be written as

$$V_t(s_{t-1}, \mathbf{P}_t) = \max_{0 \leq z_t \leq \min(K^I, (K^O - s_{t-1})/a^h)} \{L_{\text{pr}}(z_t | s_{t-1}, \mathbf{P}_t)\} \\ + \sum_{\tau=t}^{T-1} \delta^{\tau-t} \mathbb{E}_t[G_{\tau}(\tilde{\mathbf{P}}_{\tau})], \quad (3)$$

Figure 2. Schematic Representation of the Formulations in (2) and (3)

where the optimal expected profit $G_t(\mathbf{P}_t)$ for the two-stage problem in period t is given by

$$G_t(\mathbf{P}_t) \doteq \max_{0 \leq s_t \leq K^O} \left\{ -L_{sc}(s_t | \mathbf{P}_t) + \delta \mathbb{E}_t \left[\max_{0 \leq z_{t+1} \leq \min(K^I, (K^O - s_t)/a^h)} L_{pr}(z_{t+1} | s_t, \tilde{\mathbf{P}}_{t+1}) \right] \right\}, \quad (4)$$

with $L_{sc}(s_t | \mathbf{P}_t) = (p_t^O + h)s_t$, and $L_{pr}(z_{t+1} | s_t, \mathbf{P}_{t+1}) = (-p_{t+1}^I - c + \bar{a}p_{t+1}^O)z_{t+1} + p_{t+1}^O s_t$.

3.1. Periodic Input Processing and Output Inventory Decisions

We now derive the optimal solution for (4). We first characterize the optimal processing volume $z_{t+1}^*(s_t, \mathbf{P}_{t+1})$ for a given output inventory level s_t and then characterize the optimal output inventory level $s_t^*(\mathbf{P}_t)$.

Lemma 1. The optimal processing volume $z_{t+1}^*(s_t, \mathbf{P}_{t+1})$ is given by

$$z_{t+1}^*(s_t, \mathbf{P}_{t+1}) = \begin{cases} 0 & \text{if } -p_{t+1}^I - c + \bar{a}p_{t+1}^O \leq 0, \\ \min((K^O - s_t)/a^h, K^I) & \text{if } -p_{t+1}^I - c + \bar{a}p_{t+1}^O > 0. \end{cases}$$

Here, $-p_{t+1}^I - c + \bar{a}p_{t+1}^O$, the difference between the output spot sale revenue per expected yield and the sum of input spot procurement and unit processing costs, denotes the *processing margin per input*. If this margin is not positive, then it is not profitable to process. Otherwise, the firm optimally processes up to $(K^O - s_t)/a^h$ unless constrained by the processing capacity K^I .

Using Lemma 1 in (4), the optimal inventory decision becomes the solution to

$$\max_{0 \leq s_t \leq K^O} \left\{ (-p_t^O - h + \delta \mathbb{E}_t[\tilde{p}_{t+1}^O])s_t + \delta \mathbb{E}_t[(-\tilde{p}_{t+1}^I - c + \bar{a}\tilde{p}_{t+1}^O)^+] \min(K^I, (K^O - s_t)/a^h) \right\}. \quad (5)$$

Proposition 1. The optimal output inventory level $s_t^*(\mathbf{P}_t)$ is characterized by

$$s_t^*(\mathbf{P}_t) = \begin{cases} 0 & \text{if } -p_t^O - h + \delta \mathbb{E}_t[\tilde{p}_{t+1}^O] \leq 0, \\ (K^O - a^h K^I)^+ & \text{if } 0 < -p_t^O - h + \delta \mathbb{E}_t[\tilde{p}_{t+1}^O] \leq (\delta/a^h) \mathbb{E}_t[(-\tilde{p}_{t+1}^I - c + \bar{a}\tilde{p}_{t+1}^O)^+], \\ K^O & \text{if } -p_t^O - h + \delta \mathbb{E}_t[\tilde{p}_{t+1}^O] > (\delta/a^h) \mathbb{E}_t[(-\tilde{p}_{t+1}^I - c + \bar{a}\tilde{p}_{t+1}^O)^+]. \end{cases} \quad (6)$$

Here, $-p_t^O - h + \delta \mathbb{E}_t[\tilde{p}_{t+1}^O]$, the difference between the discounted expected spot sale revenue in the subsequent period and the storage cost (the sum of output spot procurement and holding costs), denotes the *storage margin per output*. If the storage margin is not positive, then it is not profitable to hold inventory. Otherwise, it is profitable to hold inventory, and $s_t^*(\mathbf{P}_t)$ is determined by the trade-off between the storage margin and the opportunity cost of holding inventory (i.e., since doing so limits the subsequent period's processing volume because then there is less unoccupied storage capacity). In particular, $(\delta/a^h) \mathbb{E}_t[(-\tilde{p}_{t+1}^I - c + \bar{a}\tilde{p}_{t+1}^O)^+]$, the discounted expected positive part of subsequent period's processing margin *per output* (and thus, scaled by the high yield realization), denotes this opportunity cost. If the opportunity cost is higher than the storage margin, then the firm stores only up to $(K^O - a^h K^I)^+$ so that the subsequent period's processing volume is not limited by the unoccupied storage capacity under any yield realization. If the opportunity cost is lower than the storage margin, then the firm stores up to the full storage capacity K^O .

It is clear from Proposition 1 that the firm's rationale for holding inventory is to benefit from output spot price fluctuations across periods. In particular, instead of selling the output to the spot market in period t , the firm stores the output for later sale to the spot market at what is expected to be a higher price. That benefit does not exist if, for instance, the discounted output spot price follows a Martingale process. In this case, since $\delta \mathbb{E}_t[\tilde{p}_{t+1}^O] = p_t^O$, it follows that the storage margin is negative, and so, by Proposition 1, the firm does not hold inventory.

Substituting the optimal inventory level $s_t^*(\cdot)$ in the optimal processing volume $z_{t+1}^*(s_t, \mathbf{P}_{t+1})$ for a given inventory level s_t , as characterized by Lemma 1, we observe that when the processing margin $-p_{t+1}^I - c + \bar{a}p_{t+1}^O$ is strictly positive, $z_{t+1}^*(s_t^*, \mathbf{P}_{t+1}) = \min(K^I, K^O/a^h)$ unless $s_t^* = K^O$ (in which case the firm optimally does not process because there is no available storage capacity for the output to be placed after processing). Using this observation, the optimal expected profit $G_t(\mathbf{P}_t)$ for

the two-stage problem in period t , as given in (4), can be written as

$$G_t(\mathbf{P}_t) = \max \left(-p_t^O - h + \delta \mathbb{E}_t[\tilde{p}_{t+1}^O], \right. \\ \left. \delta / a^h \mathbb{E}_t[(-\tilde{p}_{t+1}^I - c + \bar{a} \tilde{p}_{t+1}^O)^+] \right) \min(a^h K^I, K^O) \\ + (-p_t^O - h + \delta \mathbb{E}_t[\tilde{p}_{t+1}^O])^+ (K^O - a^h K^I)^+. \quad (7)$$

For the first $a^h K^I$ units of the storage capacity K^O , the firm faces the trade-off between holding inventory this period, which has a unit profit of storage margin per output, and processing in the subsequent period, which has a unit expected profit of discounted processing margin (when it is profitable to process) scaled by the high yield realization. Therefore, the marginal revenue of these capacity units is given by the maximum profit from the two options. For the remaining $(K^O - a^h K^I)$ units of the storage capacity, holding inventory this period does not limit the subsequent period's processing volume. Therefore, the marginal revenue of these capacity units is given by the storage margin (when it is profitable to hold inventory).

3.2. Capacity Investment Decisions

Next we solve for the firm's optimal capacity investment decision. At period $t = 0$, the firm observes \mathbb{P}_0 and chooses the capacity portfolio $\mathbf{K} = (K^I, K^O)$, while incurring the capacity investment cost $C(\mathbf{K}) = \beta^I (K^I)^2 + \beta^O (K^O)^2$, to maximize its expected total (discounted) profit over the entire planning horizon: $\max_{\mathbf{K} \geq 0} V(\mathbf{K}) - C(\mathbf{K})$, where $V(\mathbf{K}) \doteq \delta \mathbb{E}_0[V_1(0, \tilde{\mathbf{P}}_1)]$ signifies the expected profit for a given capacity portfolio \mathbf{K} . It follows from (3) that

$$V_1(0, \mathbf{P}_1) = \max_{0 \leq z_1 \leq \min(K^I, K^O/a^h)} \{(-p_1^I - c + \bar{a} p_1^O) z_1\} \\ + \sum_{\tau=1}^{T-1} \delta^{\tau-1} \mathbb{E}_t[G_\tau(\tilde{\mathbf{P}}_\tau)],$$

and using Lemma 1 and the characterization of $G_t(\mathbf{P}_t)$ given in (7), the expected total (discounted) profit over the planning horizon for a given \mathbf{K} can therefore be written as

$$\Pi(\mathbf{K}) = M_1 \min(a^h K^I, K^O) + M_2 (K^O - a^h K^I)^+ \\ - \beta^I (K^I)^2 - \beta^O (K^O)^2, \quad (8)$$

where

$$M_1 \doteq \frac{\delta}{a^h} \mathbb{E}_0[(-\tilde{p}_1^I - c + \bar{a} \tilde{p}_1^O)^+] \\ + \mathbb{E}_0 \left[\sum_{t=1}^{T-1} \delta^t \max \left(-\tilde{p}_t^O - h + \delta \mathbb{E}_t[\tilde{p}_{t+1}^O], \right. \right. \\ \left. \left. \frac{\delta}{a^h} \mathbb{E}_t[(-\tilde{p}_{t+1}^I - c + \bar{a} \tilde{p}_{t+1}^O)^+] \right) \right], \quad (9)$$

$$M_2 \doteq \mathbb{E}_0 \left[\sum_{t=1}^{T-1} \delta^t (-\tilde{p}_t^O - h + \delta \mathbb{E}_t[\tilde{p}_{t+1}^O])^+ \right].$$

In (9), M_2 denotes the total expected storage profit over the entire planning horizon. This term is relevant for the storage capacity units $(K^O - a^h K^I)^+$, which have no effect on the processing activities. The term M_1 denotes the expected marginal revenue of the first $a^h K^I$ units of the storage capacity K^O . For these capacity units, because holding inventory in each period limits the subsequent period's processing volume, the expected marginal revenue is given by the maximum of the storage margin and the expected processing benefit per output—that is, expected discounted processing margin (when it is profitable to process) scaled by the high yield realization. Because storage is empty at the beginning of the planning horizon, the first period's processing volume is not constrained by inventory, and so only the processing benefit per output is relevant.

Proposition 2 characterizes the optimal solution for the firm's capacity investment decision.

Proposition 2. The optimal capacity investment portfolio $\mathbf{K}^* = (K^{I*}, K^{O*})$ is characterized by

$$(K^{I*}, K^{O*}) = \begin{cases} \left(\frac{a^h (M_1 - M_2)}{2\beta^I}, \frac{M_2}{2\beta^O} \right) & \text{if } \beta \in \Omega_1 = \{\beta: \beta^I/\beta^O > (a^h)^2 (M_1/M_2 - 1)\}, \\ \left(\frac{a^h M_1}{2\beta^I + 2(a^h)^2 \beta^O}, \frac{(a^h)^2 M_1}{2\beta^I + 2(a^h)^2 \beta^O} \right) & \text{if } \beta \in \Omega_2 = \{\beta: \beta^I/\beta^O \leq (a^h)^2 (M_1/M_2 - 1)\}, \end{cases}$$

where $\beta = (\beta^I, \beta^O)$ and M_i for $i = 1, 2$ is as given in (9). The optimal expected profit is given by

$$\Pi^* = \begin{cases} \frac{(a^h (M_1 - M_2))^2}{4\beta^I} + \frac{(M_2)^2}{4\beta^O} & \text{if } \beta \in \Omega_1, \\ \frac{(a^h M_1)^2}{4(\beta^I + \beta^O (a^h)^2)} & \text{if } \beta \in \Omega_2. \end{cases}$$

The optimal processing and storage capacity levels are characterized by the ratio of the expected marginal revenue of an additional capacity unit to its marginal investment cost. The marginal investment cost of each capacity type is given by $2\beta^j$ for $j \in \{I, O\}$ if $\beta \in \Omega_1$, and by $2\beta^I + 2(a^h)^2 \beta^O$ if $\beta \in \Omega_2$, in which case $K^{O*} = a^h K^{I*}$. The expected marginal revenue of each capacity type takes different forms based on the capacity investment costs $\beta = (\beta^I, \beta^O)$. When the processing capacity cost relative to the storage capacity cost is sufficiently high (i.e., $\beta \in \Omega_1$), there is excess storage capacity; that is, this capacity is strictly larger than what is required for production (with full utilization of processing capacity) under both yield realizations ($K^{O*} > a^h K^{I*}$). We denote the optimal capacity portfolio in this case as the *storage-dominating portfolio*. Because there is no production benefit to having additional storage capacity, its

marginal revenue is given by the total expected storage profit M_2 . In contrast, an additional unit of processing capacity can be used for production (because there is excess storage capacity); therefore, its marginal revenue is given by the additional benefit of processing margin over the storage margin $M_1 - M_2$ per input (and thus, scaled by the high yield realization a^h). When the relative (processing-to-storage) capacity cost is sufficiently low (i.e., $\beta \in \Omega_2$), there is no excess storage capacity in the optimal solution ($K^{O*} = a^h K^{I*}$). We denote the optimal capacity portfolio in this case as the *high yield-balanced portfolio* because the processing capacity is at the level required for production (with full utilization of storage capacity) under high yield realization. In this case, because there is no excess storage capacity, M_2 has no effect on the expected marginal revenue of either capacity type.

4. Comparative Statics and Heuristics

In this section, we study the effects of spot price uncertainty on the processor's optimal capacity investment policy as well as on its profitability (Section 4.1), and the performance of the optimal capacity investment policy in comparison with a variety of heuristic policies (Section 4.2).

Throughout this section, we make two additional assumptions for the sake of tractability. First, we assume that the processing margin is nonnegative for all price realizations: $-p_t^I - c + \bar{a}p_t^O \geq 0$. In Section 5.1, we employ the same data used for calibrating our numerical experiments to verify that this is a reasonable assumption in the palm industry. Based on this assumption, when characterizing the optimal capacity investment portfolio and the optimal expected profit given in Proposition 2, we replace $\mathbb{E}_t[(-\bar{p}_{t+1}^I - c + \bar{a}\bar{p}_{t+1}^O)^+]$ for $t \in [0, T-1]$ with $\mathbb{E}_t[-\bar{p}_{t+1}^I - c + \bar{a}\bar{p}_{t+1}^O]$ in M_1 . Second, to study the effects of spot price uncertainty, we impose additional structure on our model of the spot price process. In particular, we use a single-factor, bivariate, mean-reverting price process to describe how both the input and output spot prices evolve.⁴ Thus, input and output spot prices at time τ , $\mathbf{P}_\tau = (p_\tau^I, p_\tau^O)$, are now modeled in as follows:

$$\begin{aligned} dp_\tau^I &= \theta^I(\bar{p}^I - p_\tau^I) d\tau + \sigma^I d\tilde{W}_\tau^I, \\ dp_\tau^O &= \theta^O(\bar{p}^O - p_\tau^O) d\tau + \sigma^O d\tilde{W}_\tau^O, \end{aligned} \quad (10)$$

where $\theta^j > 0$ is the mean-reversion parameter, \bar{p}^j is the long-term price level, and σ^j is the volatility for $j \in \{I, O\}$. We use $(d\tilde{W}_\tau^I, d\tilde{W}_\tau^O)$ to denote the increment of a standard bivariate Brownian motion with correlation ρ . We assume $\rho > 0$ throughout our analysis. This is a reasonable assumption in the palm industry as we empirically demonstrate in Section 5.1. Because the capacity investment and operating (processing and

storage) decisions are made at discrete time periods $t \in [0, T]$, although the price process in (10) evolves on a continuous time τ , we only need to focus on the price evolution at these discrete time periods. We assume that τ and t are in the same time units (which we consider to be a weekday for our model calibration in Section 5.1). This price model implies that, at period \hat{t} and with realized spot prices $\mathbf{P}_{\hat{t}} = (p_{\hat{t}}^I, p_{\hat{t}}^O)$, the spot prices $\tilde{\mathbf{P}}_t = (\tilde{p}_t^I, \tilde{p}_t^O)$ at a future period $t > \hat{t}$ follow a bivariate normal distribution with

$$\begin{aligned} \mathbb{E}[\tilde{p}_t^j | \mathbf{P}_{\hat{t}}] &= e^{-\theta^j(t-\hat{t})} p_{\hat{t}}^j + (1 - e^{-\theta^j(t-\hat{t})}) \bar{p}^j, \\ \text{VAR}[\tilde{p}_t^j | \mathbf{P}_{\hat{t}}] &= \frac{1 - e^{-2\theta^j(t-\hat{t})}}{2\theta^j} (\sigma^j)^2, \\ \text{COV}[\tilde{p}_t^I, \tilde{p}_t^O | \mathbf{P}_{\hat{t}}] &= \frac{1 - e^{-(\theta^I + \theta^O)(t-\hat{t})}}{\theta^I + \theta^O} \rho \sigma^I \sigma^O, \end{aligned}$$

where VAR and COV denote variance and covariance, respectively.

4.1. Effects of Spot Price Uncertainty

In this section we conduct sensitivity analyses to study the effects of spot price correlation (ρ) and of input and output spot price volatility (σ^I and σ^O , respectively) on the firm's optimal capacity investment portfolio \mathbf{K}^* and optimal expected profit Π^* . The key observation from Proposition 2 is that price correlation and price volatility affect \mathbf{K}^* and Π^* through their impacts on the expected marginal revenue terms M_2 (which, in each period $t \in [1, T-1]$, depends on the positive part of the storage margin $-\bar{p}_t^O - h + \delta \mathbb{E}_t[\bar{p}_{t+1}^O]$) and M_1 (which, in each period, depends on the maximum of the storage margin and the discounted expected processing margin per output $(\delta/a^h) \mathbb{E}_t[-\bar{p}_{t+1}^I - c + \bar{a}\bar{p}_{t+1}^O]$ and that, in turn, is characterized by the processing margin $-\bar{p}_t^I - c + \bar{a}\bar{p}_t^O$ in that period).

Proposition 3 (Price Correlation ρ). *Inequalities $\partial K^{I*}/\partial \rho < 0$, $\partial K^{O*}/\partial \rho \leq 0$, and $\partial \Pi^*/\partial \rho < 0$ hold, where $\partial K^{O*}/\partial \rho = 0$ only when \mathbf{K}^* is given by the "storage-dominating" portfolio.*

Because the storage margin depends only on the univariate distribution of the output spot price, M_2 is independent of the spot price correlation ρ . Therefore, the effect of ρ on \mathbf{K}^* (and on Π^*) is characterized by how it affects M_1 , which, in each period, depends on the maximum of the storage and the processing margins. A lower ρ makes it more likely that when the input spot price is low (high), the output spot price will be high (low). Therefore, because the processing margin depends on the difference between output and input spot prices, a lower ρ increases the variability of the processing margin while the storage margin remains unaffected. With increasing variability of the processing margin, the maximum of the storage and

processing margins increases: a higher processing margin increases this maximum, whereas a lower processing margin does not decrease it because the maximum value is given by the storage margin. Therefore, M_1 increases. Thus, a lower ρ increases both \mathbf{K}^* and Π^* .

Proposition 4 (Input Price Volatility σ^I). *There exist values $\underline{\sigma}^I < \bar{\sigma}^I$ such that $\partial K^{I*}/\partial \sigma^I < 0$, $\partial \Pi^*/\partial \sigma^I < 0$ for $\sigma^I < \underline{\sigma}^I$, and $\partial K^{I*}/\partial \sigma^I > 0$, $\partial \Pi^*/\partial \sigma^I > 0$ for $\sigma^I > \bar{\sigma}^I$. If \mathbf{K}^* is given by the storage-dominating portfolio, then $\partial K^{O*}/\partial \sigma^I = 0$; otherwise, $\partial K^{O*}/\partial \sigma^I < 0$ for $\sigma^I < \underline{\sigma}^I$ and $\partial K^{O*}/\partial \sigma^I > 0$ for $\sigma^I > \bar{\sigma}^I$.*

Much as with the effect of ρ , M_2 is independent of the input price volatility σ^I , and thus the effect of σ^I on \mathbf{K}^* (and on Π^*) is characterized by how it influences the maximum of the storage and the processing margins in each period. Since $\rho > 0$, a higher σ^I decreases (increases) the processing margin variability when σ^I is low (high). Therefore, there exists a unique $\hat{\sigma}_t^I$ threshold in period $t \in [1, T-1]$ such that the maximum of the storage and the processing margins decreases in σ^I when σ^I is lower than this threshold, and increases in σ^I otherwise. Because the threshold is period dependent, we can identify the effect of σ^I on M_1 only when σ^I is either sufficiently low ($\sigma^I < \underline{\sigma}^I \doteq \min_t \{\hat{\sigma}_t^I\}$) or sufficiently high ($\sigma^I > \bar{\sigma}^I \doteq \max_t \{\hat{\sigma}_t^I\}$).

Proposition 5 (Output Price Volatility σ^O). *There exist $\underline{\sigma}^O < \bar{\sigma}^O$ such that*

(i) *if \mathbf{K}^* is given by the high yield-balanced portfolio, then $\partial K^{I*}/\partial \sigma^O < 0$, $\partial K^{O*}/\partial \sigma^O < 0$, and $\partial \Pi^*/\partial \sigma^O < 0$ for $\sigma^O < \underline{\sigma}^O$, whereas $\partial K^{I*}/\partial \sigma^O > 0$, $\partial K^{O*}/\partial \sigma^O > 0$, and $\partial \Pi^*/\partial \sigma^O > 0$ for $\sigma^O > \bar{\sigma}^O$;*

(ii) *if \mathbf{K}^* is given by the “storage-dominating” portfolio, then $\partial K^{I*}/\partial \sigma^O < 0$ for $\sigma^O < \underline{\sigma}^O$, and $\partial K^{O*}/\partial \sigma^O > 0$.*

The influence of σ^O on \mathbf{K}^* (and Π^*) is determined by its effect on both M_1 and M_2 . Recall that in each period M_2 depends on the positive part of the storage margin. A higher σ^O increases the storage margin variability. While a high storage margin is beneficial, a low storage margin is less consequential because the firm optimally chooses not to hold inventory. Therefore, M_2 is increasing in σ^O . The impact of σ^O on M_1 parallels the σ^I effect. In particular, there exists a unique $\hat{\sigma}_t^O$ threshold in period t such that the maximum of the storage and the processing margins decreases in σ^O when σ^O is lower than this threshold, and increases in σ^O otherwise. This threshold, too, is period dependent, so the effect of σ^O can only be partially characterized: M_1 decreases (increases) with σ^O if $\sigma^O < \underline{\sigma}^O \doteq \min_t \{\hat{\sigma}_t^O\} \forall t$ (if $\sigma^O > \bar{\sigma}^O \doteq \max_t \{\hat{\sigma}_t^O\} \forall t$). When \mathbf{K}^* is given by the high yield-balanced portfolio, the impact of σ^O is characterized by its effect on M_1 ; otherwise, its effect on M_2 is also relevant. In this latter case, K^{O*} increases in σ^O because M_2 increases, whereas K^{I*} decreases in σ^O when σ^O is low because then M_1 decreases and M_2 increases.

4.2. Heuristic Capacity Investment Policies

In this section, we compare the performance of the optimal capacity investment policy with that of heuristic capacity investment policies. Toward this end, we define the profit loss due to employing a heuristic policy (hp) as $\Delta_{hp} \doteq [(\Pi^* - \Pi(\mathbf{K}_{hp}))]/\Pi^*$, where Π^* is the optimal expected profit (as given by Proposition 2) and $\Pi(\mathbf{K}_{hp})$ is the expected total profit (as given by (8)) evaluated with the capacity portfolio $\mathbf{K}_{hp} = (K_{hp}^I, K_{hp}^O)$, which is chosen by the heuristic policy. Here we introduce the heuristic policies considered and provide analytical bounds on the profit loss Δ_{hp} with each heuristic policy. Later in Section 5.3 we use these analytical bounds to frame our numerical investigation of the heuristics in the context of the palm industry. For ease of notation, we define $\eta \doteq \beta^I/\beta^O$, the relative (processing-to-storage) capacity investment cost.

4.2.1. Heuristics Based on Ignoring the Production Yield Uncertainty.

We consider two heuristic policies in which the firm ignores production yield uncertainty and plans for capacity based on a single number representing the yield. In the *deterministic yield (maximum)*, or *DYM*, heuristic, capacity planning is based on the maximum possible yield. This policy is the one most often implemented by palm oil mills in practice. The optimal capacity investment with this policy, \mathbf{K}_{DYM} , can be obtained from Proposition 2 by replacing \bar{a} with a^h :

$$(K_{DYM}^I, K_{DYM}^O) = \begin{cases} \left(\frac{a^h(\bar{M}_1(a^h) - M_2)}{2\beta^I}, \frac{M_2}{2\beta^O} \right) & \text{if } \eta/(a^h)^2 > \bar{M}_1(a^h)/M_2 - 1, \\ \left(\frac{a^h\bar{M}_1(a^h)}{2\beta^I + 2(a^h)^2\beta^O}, \frac{(a^h)^2\bar{M}_1(a^h)}{2\beta^I + 2(a^h)^2\beta^O} \right) & \text{if } \eta/(a^h)^2 \leq \bar{M}_1(a^h)/M_2 - 1, \end{cases}$$

where $\bar{M}_1(a^h) \doteq (\delta/a^h)\mathbb{E}_0[-\tilde{p}_1^I - c + a^h\tilde{p}_1^O] + \mathbb{E}_0[\sum_{t=1}^{T-1} \delta^t \cdot \max(-\tilde{p}_t^O - h + \delta\mathbb{E}_t[\tilde{p}_{t+1}^O], (\delta/a^h)\mathbb{E}_t[-\tilde{p}_{t+1}^I - c + a^h\tilde{p}_{t+1}^O])]$.

In the *deterministic yield (average)*, or *DYA* heuristic, capacity planning is based on the average yield. The optimal capacity investment with this policy, \mathbf{K}_{DYA} , can be obtained from Proposition 2 by substituting \bar{a} for \bar{a} :

$$(K_{DYA}^I, K_{DYA}^O) = \begin{cases} \left(\frac{\bar{a}(\bar{M}_1(\bar{a}) - M_2)}{2\beta^I}, \frac{M_2}{2\beta^O} \right) & \text{if } \eta/(\bar{a})^2 > \bar{M}_1(\bar{a})/M_2 - 1, \\ \left(\frac{\bar{a}\bar{M}_1(\bar{a})}{2\beta^I + 2(\bar{a})^2\beta^O}, \frac{(\bar{a})^2\bar{M}_1(\bar{a})}{2\beta^I + 2(\bar{a})^2\beta^O} \right) & \text{if } \eta/(\bar{a})^2 \leq \bar{M}_1(\bar{a})/M_2 - 1, \end{cases}$$

where $\bar{M}_1(\bar{a}) \doteq (\delta/\bar{a})\mathbb{E}_0[-\tilde{p}_1^I - c + \bar{a}\tilde{p}_1^O] + \mathbb{E}_0[\sum_{t=1}^{T-1} \delta^t \cdot \max(-\tilde{p}_t^O - h + \delta\mathbb{E}_t[\tilde{p}_{t+1}^O], (\delta/\bar{a})\mathbb{E}_t[-\tilde{p}_{t+1}^I - c + \bar{a}\tilde{p}_{t+1}^O])]$.

If production yield uncertainty is ignored, then should capacity planning be based on the average yield or the maximum yield? Our next proposition shows

that the answer is the average yield when η is sufficiently high—that is, when the firm invests in a storage-dominating portfolio under the optimal policy as well as under both DYM and DYA heuristic policies.

Proposition 6. When $\eta > \max((a^h)^2(\bar{M}_1(a^h)/M_2 - 1), a^h\bar{M}_1(\bar{a})/M_2 - 1)$, $\Delta_{DYM} > \Delta_{DYA}$.

Recall from Proposition 2 that processing capacity in a storage-dominating portfolio is determined by $a^h(M_1 - M_2)$, which, in each period, depends on the difference between the processing margin $-p_t^I - c + \bar{a}p_t^O$ and the storage margin (per input) $a^h sm_t$, where $sm_t \doteq (-p_t^O - h + \delta E_t[\bar{p}_{t+1}^O])$. Because $-p_t^I - c + a^h p_t^O > -p_t^I - c + \bar{a}p_t^O$, a firm that uses the DYM heuristic will overestimate the processing margin and consequently overinvest in processing capacity (i.e., $K_{DYM}^I > K^{I*}$). There is no such effect on the processing margin for a firm that uses the DYA heuristic, although that firm will then underestimate the storage margin per input (i.e., $\bar{a}sm_t < a^h sm_t$) and consequently overinvest in processing capacity (i.e., $K_{DYA}^I > K^{I*}$). It turns out that overestimating the processing margin has a more significant impact than underestimating the storage margin, and so more significant processing capacity misspecification occurs when using the DYM heuristic (i.e., $K_{DYM}^I > K_{DYA}^I$). Therefore, the profit loss is higher than when using the DYA heuristic.

Given that it is better to base the capacity planning on the average yield, how significant, then, is the profit loss when using this heuristic policy? Proposition 7 provides bounds on the profit loss Δ_{DYA} when η is sufficiently high, that is, when the firm invests in a storage-dominating portfolio under the optimal and the DYA heuristic policies (i.e., $K^{O*} > a^h K^{I*}$ and $K_{DYA}^O > a^h K_{DYA}^I$), and when η is sufficiently low, that is, when the firm does not invest in a storage-dominating portfolio under the optimal and the DYA heuristic policies (i.e., $K^{O*} = a^h K^{I*}$ and $K_{DYA}^O = \bar{a} K_{DYA}^I$).

Proposition 7. Case (i). When $\eta > \max((a^h)^2(M_1/M_2 - 1), a^h\bar{a}(\bar{M}_1(\bar{a})/M_2 - 1))$, $\Delta_{DYA} \leq ((1 - \bar{a}/a^h)/(M_1/M_2 - 1))^2$.

Case (ii). When $\eta \leq \max((a^h)^2(M_1/M_2 - 1), \bar{a}^2\bar{M}_1(\bar{a})/M_2 - 1)$, if $\eta/(a^h)^2 + 1 \approx \eta/(a^h)^2$ then $1 - (\bar{a}/a^h)^2 \leq \Delta_{DYA} \leq 2(1 - \bar{a}/a^h)$.

In both cases, \bar{a}/a^h plays a critical role for the significance of the profit loss due to employing the DYA heuristic policy. For example, the proposed upper bound in each case—and thus the actual profit loss Δ_{DYA} —will be low if \bar{a}/a^h is close to 1 ($\bar{a} < a^h$ by definition). As \bar{a}/a^h decreases, ignoring yield uncertainty in capacity planning will have more severe consequences. This can be observed from Case (ii) where the proposed lower bound—and thus the actual profit loss—will be high if \bar{a}/a^h is low (which, in the next section, we will show to hold true in the context of the palm industry).

4.2.2. Deterministic Price Heuristic. Under the deterministic price (DP) heuristic, the firm ignores spot price uncertainty and plans for capacity based on expected spot prices. The optimal capacity investment with such a policy, K_{DP} , can be obtained from Proposition 2 by substituting the processing and storage margins in each period with their expected values:

$$(K_{DP}^I, K_{DP}^O) = \begin{cases} \left(\frac{a^h(\underline{M}_1 - \underline{M}_2)}{2\beta^I}, \frac{\underline{M}_2}{2\beta^O} \right) & \text{if } \eta/(a^h)^2 > \underline{M}_1/\underline{M}_2 - 1, \\ \left(\frac{a^h \underline{M}_1}{2\beta^I + 2(a^h)^2\beta^O}, \frac{(a^h)^2 \underline{M}_1}{2\beta^I + 2(a^h)^2\beta^O} \right) & \text{if } \eta/(a^h)^2 \leq \underline{M}_1/\underline{M}_2 - 1, \end{cases} \quad (11)$$

where

$$\begin{aligned} \underline{M}_1 &\doteq \frac{\delta}{a^h} E_0[-\bar{p}_1^I - c + \bar{a}\bar{p}_1^O] \\ &\quad + \left[\sum_{t=1}^{T-1} \delta^t \max \left(E_0[-\bar{p}_t^O - h + \delta E_t[\bar{p}_{t+1}^O]], \frac{\delta}{a^h} E_0[-\bar{p}_{t+1}^I - c + \bar{a}\bar{p}_{t+1}^O] \right) \right], \\ \underline{M}_2 &\doteq \sum_{t=1}^{T-1} \delta^t (E_0[-\bar{p}_t^O - h + \delta E_t[\bar{p}_{t+1}^O]])^+. \end{aligned}$$

Failing to account for spot price uncertainty leads the firm to underinvest in each capacity type, i.e., $K_{DP} < K^*$. This follows from $\underline{M}_1 < M_1$ and $\underline{M}_2 < M_2$, which can be easily established using Jensen's inequality. Recall that the firm's rationale for holding inventory is to benefit from output spot price fluctuations—instead of selling the output to the spot market, the firm stores the output for later sale to the spot market at what is expected to be a higher price. When the uncertainty in output spot price is ignored, such benefit will be less significant, and thus \underline{M}_2 , total expected storage profit over the planning horizon, will be considerably small. Therefore, the firm is likely to invest in the high yield-balanced portfolio when DP heuristic is employed (second case in (11)), a conjecture that we will show in the next section to hold true in the context of the palm industry. Proposition 8 provides a lower bound on the profit loss Δ_{DP} in the limiting case where $\underline{M}_2 = 0$.

Proposition 8. Assume $\underline{M}_2 = 0$.

Case (i). When $\eta \leq (a^h)^2(M_1/M_2 - 1)$, $\Delta_{DP} = (1 - \underline{M}_1/M_1)^2$.

Case (ii). When $\eta > (a^h)^2(M_1/M_2 - 1)$, if $\eta/(a^h)^2 + 1 \approx \eta/(a^h)^2$ and $M_1/M_2 - 1 \approx \underline{M}_1/M_2$, then $\Delta_{DP} > (1 - \underline{M}_1/M_1)^2$.

The lower bound on the profit loss equals the actual profit loss when η is low enough that the firm invests in a high yield-balanced portfolio under the optimal policy, i.e., Case (i). In both cases, \underline{M}_1/M_1 plays a critical

role for the significance of the profit loss due to employing DP heuristic policy. Because $\bar{M}_2 = 0$ by assumption, the expected storage margin in each period is practically zero, and thus \bar{M}_1 depends on the expected processing margin in each period. Under the price process specified in (10), this expected processing margin converges to a fixed quantity—where the input and output spot prices are at their long term means—after a certain number of periods. Therefore, this fixed quantity crucially determines \bar{M}_1 , and thus the loss of profit under the DP heuristic. We will show in the next section that this fixed quantity is sufficiently low, and the profit loss is small in the context of the palm industry.

4.2.3. No By-product Heuristic. Under the no by-product (NB) heuristic, the firm does not account for by-product revenue when planning for capacity. The optimal capacity investment with such a policy, \mathbf{K}_{NB} , can be obtained from Proposition 2 by substituting the effective processing cost $c = \underline{c} - a^B p^B$ for \underline{c} , i.e., substituting M_1 for

$$\begin{aligned} \bar{M}_1^{NB} \doteq & \frac{\delta}{a^h} \mathbb{E}_0[-\tilde{p}_1^I - \underline{c} + \bar{a}\tilde{p}_1^O] \\ & + \mathbb{E}_0 \left[\sum_{t=1}^{T-1} \delta^t \max \left(-\tilde{p}_t^O - h + \delta \mathbb{E}_t[\tilde{p}_{t+1}^O], \right. \right. \\ & \left. \left. \frac{\delta}{a^h} \mathbb{E}_t[-\tilde{p}_{t+1}^I - \underline{c} + \bar{a}\tilde{p}_{t+1}^O] \right) \right]. \end{aligned}$$

Intuitively, failing to account for by-product revenue leads the firm to underestimate the processing margin in each period and, thus, to underinvest in capacity, i.e., $\mathbf{K}_{NB} \leq \mathbf{K}^*$. This follows from $\bar{M}_1^{NB} < M_1$. Proposition 9 gives an upper bound on the profit loss when η is high enough that the firm invests in a storage-dominating portfolio under both the optimal and the heuristic policies.

Proposition 9. Let $\chi \doteq M_1/M_2$ and $\underline{\chi} \doteq \bar{M}_1^{NB}/M_2$. When $\eta > (a^h)^2(\chi - 1)$, $\Delta_{NB} \leq (1 - \underline{\chi}/\chi)(1 - (\underline{\chi} - 1)/(\chi - 1))$.

In this case, failing to account for by-product revenue leads to underinvestment in processing capacity and no change in storage capacity. Because the processing capacity is determined by the maximum of the processing margin and the storage margin in each period, the proposed upper bound crucially depends on the relative magnitude of the two margins. When both margins are of the same order of magnitude, underestimating the processing margin—that is, by not accounting for the by-product revenue—is less consequential. Suppose, for example, that underestimation of the processing margin leads to a 10% reduction in χ —that is, $\chi = 0.9\bar{\chi}$. Then, the upper bound on $\Delta_{NB} \times 100$ is only 1% (i.e., the profit loss is of extremely low magnitude). We will show in the next section that this example is not relevant in the context of the palm industry.

4.2.4. High Yield–Balanced Portfolio Heuristic. Under this heuristic, the firm chooses its capacity investment portfolio by assuming $K^O = a^h K^I$. The optimal capacity investment under this policy is characterized by $K_{HYBP}^I = a^h M_1 / (2\beta^I + 2(a^h)^2 \beta^O)$ and $K_{HYBP}^O = (a^h)^2 M_1 / (2\beta^I + 2(a^h)^2 \beta^O)$. The following proposition provides an upper bound for the profit loss experienced under this heuristic policy.

Proposition 10. Let $\eta > (a^h)^2(M_1/M_2 - 1)$ such that $\Delta_{HYBP} > 0$. If $\eta/(a^h)^2 + 1 \approx \eta/(a^h)^2$, then $\Delta_{HYBP} \leq (\eta/(a^h)^2) / (\eta/(a^h)^2 + (M_1/M_2)^2)$.

The proposed upper bound—and thus the actual profit loss—will be low if $\eta/(a^h)^2$ is sufficiently large, so $\eta/(a^h)^2 + 1 \approx \eta/(a^h)^2$ holds, and M_1/M_2 is very large (in other words, the processing margin is substantially higher than the storage margin in each period) and so $(M_1/M_2)^2$ significantly outweighs $\eta/(a^h)^2$. We will show in the next section that these conditions hold in the context of the palm industry.

5. Numerical Analysis: Application to the Palm Industry

In this section, we discuss an application of our model in the context of the palm industry. In this industry, a palm oil mill processes palm fresh fruit bunches to produce crude palm oil and palm kernel. The fresh fruit bunches first pass through receiving and sterilization stations, where high-pressure steam is applied. The palm fruits are then separated from the bunches at the threshing station before being crushed at the pressing station to produce palm kernel and crude palm oil, from which water and waste are then removed via centrifuge. The crude palm oil is transferred to storage tanks prior to dispatch from the mill. In the context of our model, the palm fresh fruit bunch is the input, the crude palm oil is the output, and the palm kernel is the by-product. The joint capacity of the receiving, sterilization, threshing, pressing, and centrifuge stations corresponds to K^I ; the CPO storage tank capacity corresponds to K^O .

The rest of this section is organized as follows. In Section 5.1 we describe the data and calibration on which our numerical experiments will be based. Section 5.2 investigates the effect of spot price uncertainty on the firm's optimal capacity investment policy and profitability. Finally, in Section 5.3 we compare the performance of optimal and heuristic capacity investment policies.

5.1. Data, Model Calibration, and Computation for Numerical Experiments

Our focal unit of analysis is a palm oil mill located in Southeast Asia. Within this region, Malaysia and Indonesia share many characteristics; they are the two largest players in the palm oil industry, accounting

for 86% of world palm oil production for the 2013–2014 period (U.S. Department of Agriculture 2015, Table 11). Our numerical experiments use publicly available data from the Malaysian Palm Oil Board complemented by proprietary and publicly available data from palm oil mills located in Malaysia. Hereafter, we shall often use “RM” to denote the Malaysian ringgit (currency) and “mt” to denote metric ton (equal to 1,000 kg, or about 1.1 U.S. tons). Throughout this section, we use \hat{x} to denote the calibrated value for parameter x .

5.1.1. Calibration for Price Process Parameters. In our computational experiments, each period corresponds to a weekday in practice. We use the daily prices of FFB and CPO reported by the MPOB from January 1, 2006, to December 31, 2013. This period encompasses 1,940 weekdays. The daily FFB price varies as a function of the palm fruit’s origin (i.e., the north, south, west, or east subregion of the Malaysian Peninsula) and quality (i.e., Grade A, B, or C), so we use the average of FFB prices across subregions and grades. The daily CPO prices vary as a function of the delivery month (i.e., to be delivered in the same month, next month, etc.). Consistent with our model, we use the CPO prices that correspond to immediate delivery (i.e., within in the same month). The daily prices used in our calibration (in ringgit per metric ton) are plotted in Figure 1(a). According to the price process specified in (10), the daily spot prices evolve as follows:

$$\begin{aligned}\tilde{p}_t^I &= e^{-\theta^I} p_{t-1}^I + (1 - e^{-\theta^I}) \bar{p}^I + \sigma^I \sqrt{\frac{1 - e^{-2\theta^I}}{2\theta^I}} \tilde{z}^I, \\ \tilde{p}_t^O &= e^{-\theta^O} p_{t-1}^O + (1 - e^{-\theta^O}) \bar{p}^O + \sigma^O \sqrt{\frac{1 - e^{-2\theta^O}}{2\theta^O}} \tilde{z}^O,\end{aligned}\quad (12)$$

where $(\tilde{z}^I, \tilde{z}^O)$ follows a standard bivariate normal distribution with correlation ρ . The expressions in (12) can be viewed as a system of simultaneous equations of $(\tilde{p}_t^I, \tilde{p}_t^O)$ on (p_{t-1}^I, p_{t-1}^O) ; that is, $\tilde{p}_t^j = \alpha^j p_{t-1}^j + \varphi^j + \tilde{\epsilon}^j$ for $j \in \{I, O\}$. Because the error terms $(\tilde{\epsilon}^I, \tilde{\epsilon}^O)$ are correlated, we use the “seemingly unrelated” regression (SUR; see Zellner 1962) to estimate α^j , φ^j , and the covariance matrix of $(\tilde{\epsilon}^I, \tilde{\epsilon}^O)$. We can then use these estimates together with (12) to obtain $\hat{\theta}^I = 0.00345$, $\hat{p}^I = 532.75$, $\hat{\sigma}^I = 8.60$, $\hat{\theta}^O = 0.00437$, $\hat{p}^O = 2689.87$, $\hat{\sigma}^O = 39.08$, and $\hat{\rho} = 0.734$. According to the McElroy’s R^2 , the SUR equations can explain 99.36% of the variation in the spot prices observed.

5.1.2. Calibration for Production Yield Parameters. The most granular data from MPOB are the monthly average production yields (extraction rates) in the Malaysian Peninsula. As plotted in Figure 1(b), the CPO yield from January 2006 to December 2013 ranges from 18.51% to 20.37%, with a mean of 19.72% and a standard deviation of 0.43%. Proposition 2 suggests that the average production yield \hat{a} and the high yield

realization a^h are sufficient for numerical computation. Accordingly, we set $\hat{a} = 19.72\%$, which is the average yield in our data set, and $\hat{a}^h = 20.37\%$, which is the highest yield recorded in our data set. As discussed in Section 2, we assume that the production yield and spot price distributions are statistically independent. To verify the reasonableness of this assumption, we examine the correlation between CPO yield and the CPO price change lagged by k months for $k \in [1, \dots, 5]$ and find that (results not reported here) this correlation lies in the range $[-0.08, -0.02]$.

5.1.3. Calibration for Other Operational Parameters.

For processing cost, we set $\hat{c} = 40$ RM per metric ton of FFB, which is representative of the palm industry. For instance, in the 2013 annual report of Sime Darby (a major palm producer in Malaysia), the average “mill cost” between 2008 and 2012 is given as 199.75 RM/mt of CPO, which corresponds to 39.39 RM/mt of FFB at the average production yield of 19.72%.⁵ As in Section 4, we continue to assume a nonnegative processing margin for all price realizations. To verify that this is a reasonable assumption, we examine the observed processing margin $ap^O + a^B p^B - c - p^I$ (with $c = 40$) using the daily FFB (p^I) and CPO (p^O) prices together with the monthly CPO (a) and palm kernel (a^B) production yield reported by the MPOB. We find that the observed margin is strictly positive for all 1,940 weekdays excepting only 2 (on which the margins are -0.83 and -0.17). For the inventory holding cost, we use $\hat{h} = 1$ RM/day per metric ton of CPO, which is approximately 10% of the CPO value if the inventory is held in storage for an entire year (based on the long-term CPO price level $\hat{p}^O = 2,689.87$ and counting 250 weekdays annually). For the palm kernel by-product, we use the overall average of the data reported by the MPOB within our time frame for both the price ($\hat{p}^B = 1,510.70$ RM/mt) and the production yield ($\hat{a}^B = 5.53\%$). These values entail an effective processing cost of $\hat{c} = \hat{c} - \hat{a}^B \hat{p}^B = -39.47$ RM/mt of FFB. Capacity cost parameters β^I and β^O are calibrated based on the following capacity cost information obtained from a palm oil mill located in Malaysia. The cost of processing facilities (fruit receiving, sterilization, threshing, pressing, and centrifuge stations) with a capacity of 30 mt of FFB per hour (or 300 mt of FFB daily in our model if we assume there are 10 production hours per day) is 6,723,940 RM; the cost of storage tank that can hold 2,000 mt of CPO is 969,570 RM. Given this information, we estimate $\hat{\beta}^I = 75$ and $\hat{\beta}^O = 0.25$. We do not consider fixed costs (e.g., land) in our numerical experiments. For the discount factor δ , we take an annual compound interest rate r and set $\delta = (1 + r)^{-1/250}$ (i.e., based again on 250 weekdays per year). In the baseline scenario we assume that $\hat{r} = 10\%$.

5.1.4. Numerical Computation. It follows from Proposition 2 that \mathbf{K}^* and Π^* depend on $\mathbb{E}_0[(-\tilde{p}_t^O - h + \delta \mathbb{E}_t[\tilde{p}_{t+1}^O])^+]$ and $\mathbb{E}_0[\max(-\tilde{p}_t^O - h + \delta \mathbb{E}_t[\tilde{p}_{t+1}^O], (\delta/a^h) \cdot \mathbb{E}_t[-\tilde{p}_{t+1}^I - c + \bar{a}\tilde{p}_{t+1}^O])]$ for $t \in [1, T-1]$. It can be proven that at period 0, $(-\tilde{p}_t^O - h + \delta \mathbb{E}_t[\tilde{p}_{t+1}^O], \mathbb{E}_t[-\tilde{p}_{t+1}^I - c + \bar{a}\tilde{p}_{t+1}^O])$ for $t \in [1, T-1]$ follow a bivariate normal distribution, so these expressions can be written in closed form (using the moments as well as the probability density function and cumulative distribution function of the standard normal distribution). Numerical computation can therefore be carried out in an efficient manner. We initialize the FFB and CPO prices at the beginning of the planning horizon to their last available values in the data set: $p_0^I = 528.5$ RM/mt and $p_0^O = 2,570.5$ RM/mt. We consider a five-year planning horizon, which is equivalent to 1,250 weekdays (i.e., $\hat{T} = 1,250$).

5.1.5. Baseline Scenario. In our baseline scenario the optimal capacity investment is given by the storage-dominating portfolio with $K^{I*} = 858.91$ mt/day and $K^{O*} = 1,653.66$ mt (where $M_1 = 633,308.421$ and $M_2 = 826.83$), and the optimal expected profit is 56,012,483.86 RM over the five-year planning horizon.

5.2. Effects of Spot Price Uncertainty

Here we illustrate our analytical sensitivity results, as discussed in Section 4.1, using numerical studies in the context of the palm industry. We analyze (but refrain from plotting here) the effect of changing price correlation in our baseline scenario for $\rho \in [0.5, 0.975]$ in increments of 0.025. We observe in all instances that \mathbf{K}^* is given by the storage-dominating portfolio, and so, in line with Proposition 3, K^{O*} is not affected by ρ , while K^{I*} and Π^* are both decreasing in ρ .

Figure 3 plots the effects of changing input price volatility in our baseline scenario for $\sigma^I \in [-50\%, 50\%]$ of the baseline value $\hat{\sigma}^I = 8.60$ in 5% increments. We can


see that \mathbf{K}^* is always given by the storage-dominating portfolio ($K^{O*} > a^h K^{I*}$ as observed in the first panel), and so, in line with Proposition 4, K^{O*} is unaffected by σ^I . In Figure 3, as σ^I increases in its specified range, we observe a unique σ^I threshold where K^{I*} decreases in σ^I when σ^I is below this threshold and increases in σ^I otherwise. (The decreasing behavior is less visible than the increasing behavior in the figure because the decreasing behavior is less significant in magnitude.) The same pattern also holds for the effect of σ^I on Π^* . These observations are consistent with Proposition 4, which proves that both K^{I*} and Π^* are decreasing (increasing) in σ^I when σ^I is sufficiently low (high).

Figure 4 plots the effects of changing output price volatility in our baseline scenario for $\sigma^O \in [-50\%, 50\%]$ of the baseline value $\hat{\sigma}^O = 39.08$ in 5% increments. If σ^O is low, then \mathbf{K}^* is given by the high yield-balanced portfolio ($K^{O*} = a^h K^{I*}$ as observed in the first panel). Otherwise, \mathbf{K}^* is given by the storage-dominating portfolio, and, in line with Proposition 5, K^{O*} increases with σ^O . We again observe a unique σ^O threshold where K^{I*} decreases in σ^O when σ^O is below this threshold and increases in σ^O otherwise. The same pattern also holds for the effect of σ^O on Π^* . These observations are also consistent with Proposition 5.

Two remarks are in order. First, we observe from Figure 4 that when \mathbf{K}^* is given by the storage-dominating portfolio, as σ^O increases, K^{O*} changes to a larger extent than Π^* —that is, the optimal profit is more robust to changes in σ^O than the optimal storage capacity level. Second, comparing the second panels in Figures 3 and 4 reveals that the optimal profit is less sensitive to changes in input price volatility than to changes in output price volatility.


To summarize, we find that the palm oil mill benefits from a lower spot price correlation and also from a lower (higher) FFB or CPO price volatility when

Figure 3. Effects of Changing Input Spot Price Volatility (σ^I) on Optimal Levels of Processing Capacity (K^{I*}) and Storage Capacity (K^{O*}) and on Optimal Expected Profits, Where $\sigma^I \in [-50\%, 50\%]$ of the Baseline Value $\hat{\sigma}^I = 8.60$ in 5% Increments



Note. In the two panels, baseline scenario is indicated by the circle (●) aligned (horizontally) with 0.

Figure 4. Effects of Changing Output Spot Price Volatility (σ^O) on Optimal Levels of Processing Capacity (K^{I*}) and Storage Capacity (K^{O*}) and on Optimal Expected Profits, Where $\sigma^O \in [-50\%, 50\%]$ of the Baseline Value $\hat{\sigma}^O = 39.08$ in 5% Increments



Note. Circles again indicate the baseline scenario.

this volatility is low (high). How do the changes in spot price uncertainty affect the processing and storage capacities of a typical palm oil mill (that invests in a storage-dominating portfolio)? Our results highlight the significant differences in how spot price uncertainty affects each capacity type. In particular, we find that the optimal processing capacity decreases with an increase in price correlation or an increase (a decrease) in CPO or FFB volatility when this volatility is low (high). In contrast, the optimal storage capacity increases with an increase in CPO volatility, but it is not affected otherwise.

5.3. Performance of Heuristic Capacity Investment Policies

Here we numerically compare the performance of the optimal capacity investment policy with that of heuristic capacity investment policies discussed in Section 4.2 in the context of the palm industry. Recall that the profit loss due to employing a heuristic policy (hp) is defined as $\Delta_{hp} \doteq [(\Pi^* - \Pi(\mathbf{K}_{hp})) / \Pi^*]$, where Π^* is the optimal expected profit, and $\Pi(\mathbf{K}_{hp})$ is the expected profit evaluated with the capacity portfolio $\mathbf{K}_{hp} = (K_{hp}^I, K_{hp}^O)$, which is chosen by the heuristic policy. We relate our numerical results to the analytical results (bounds on the profit loss with each heuristic policy) presented in Section 4.2.

We extend our numerical instances to assess the sensitivity of our results to several key parameters. In particular, for $\eta = \beta^I / \beta^O$, we consider $\eta \in [-30\%, 30\%]$ of the baseline value $\hat{\eta} = 300$ in 10% increments. We also consider the maximum yield $a^h \in [20.37\%, 22.37\%]$ in 0.05% increments as well as holding costs $h \in \{0.5, 1, 2\}$ and interest rates $r \in \{0\%, 10\%, 20\%\}$. Altogether, we evaluate 315 numerical instances.

Before discussing our key findings, we make two crucial observations based on our numerical results.

First, $\eta / (a^h)^2$ is sufficiently greater than 1 (with an average value of 6,590, while ranging between 4,196 and 9,399), and so the condition $\eta / (a^h)^2 + 1 \approx \eta / (a^h)^2$ in Propositions 7 and 8 is satisfied. Second, M_1 / M_2 is very high in our numerical instances (with an average value of 2,838, while ranging from 439 to 10,676), and so the condition $M_1 / M_2 - 1 \approx M_1 / M_2$ in Proposition 8 is also satisfied. Another implication of this observation is that the processing margin is significantly higher than the storage margin in each period. Recall that, in each period, M_1 is characterized by the maximum of the processing and storage margins, while M_2 is characterized simply by the storage margin. These observations will be critical in delineating the intuition behind the results that follow.

Table 1 summarizes the percentage profit loss $\Delta_{hp} \times 100$ incurred under each heuristic policy using a classification of the numerical instances based on the optimal capacity investment policy (storage dominating or high yield balanced). We now present our key findings.

1. *If the production yield uncertainty is ignored, then capacity planning should be based on the average yield and not on the maximum yield.* In all numerical instances, the profit loss under the DYM heuristic is greater than the corresponding loss under the DYA heuristic. This observation is consistent with Proposition 6. As discussed in Section 4.2, a firm that uses the DYM heuristic will overestimate the processing margin, while a firm that uses the DYA heuristic will underestimate the storage margin per input. Yet, because the processing margin dominates the storage margin in our numerical instances, this latter underestimation is less consequential—that is, a firm that uses the DYA heuristic does not significantly deviate from the optimal processing capacity level. Consistent with these arguments, we observe that the *maximum* absolute percentage misspecification of processing capacity with

Table 1. Performance of Heuristic Capacity Investment Policies in the Palm Industry

Optimal policy (% instances)	Percentage loss (%) = $\Delta_{hp} \times 100$				
	DYM	DYA	DP	NB	HYBP
K^* is storage dominating (87.9%)	67.68	0	5.95	65.12	0.57
$\eta/(a^h)^2 > (M_1/M_2 - 1)$	(6.99, 161.97)	(0, 0)	(5.35, 8.78)	(61.86, 66.96)	(0, 3.50)
K^* is high yield balanced (12.1%)	73.98	14.53	5.31	67.32	0
$\eta/(a^h)^2 \leq (M_1/M_2 - 1)$	(7.83, 162.65)	(1.78, 23.70)	(5.30, 5.36)	(66.14, 67.51)	(0, 0)

Note. For each of these heuristics, the boldface values report the average percentage loss observed in the relevant numerical instances, while the other values report the minimum and the maximum percentage loss observed.

the DYA heuristic, i.e., $|(K^{I*} - K_{DYA}^I)/K^{I*}| \times 100$, is only 0.06% in our numerical instances, while the *minimum* absolute percentage misspecification of processing capacity with the DYM heuristic is 26.97%. Therefore, the profit loss is lower when using the DYA heuristic than the profit loss when using the DYM heuristic.

2. *Ignoring the production yield uncertainty while using the average yield when planning for capacity does not affect profitability when the relative (processing-to-storage) capacity investment cost is high; otherwise, it leads to a significant loss of profit.* Table 1 shows that in all numerical instances with sufficiently high η (i.e., such that the firm invests in a storage-dominating portfolio with the optimal policy), the loss under the DYA heuristic is negligible (0%), whereas in the numerical instances with sufficiently low η (i.e., such that the firm invests in a high yield-balanced portfolio with the optimal policy), the average profit loss is 14.53%. Because processing capacity under the DYA policy differs little from the optimal level, these observations are driven by the magnitude of storage capacity misspecification. In a storage-dominating portfolio, per Proposition 2, storage capacity is determined by the storage margin per output, i.e., $-p_t^O - h + \delta E_t[\tilde{p}_{t+1}^O]$. Because this margin is not affected by the DYA heuristic, the firm chooses the same storage capacity level as under the optimal policy. Therefore, there is no loss of profit. In a high yield-balanced portfolio, the firm underinvests in storage capacity under the DYA policy because that capacity is determined by $K^O = \bar{a}K^I$ rather than by $K^O = a^h K^I$. In the relevant numerical instances, we observe an average percentage misspecification of storage capacity with the DYA heuristic, i.e., $|(K^{O*} - K_{DYA}^O)/K^{O*}| \times 100$, of 7.23%. That storage capacity misspecification leads to a sizable profit loss. These numerical observations are consistent with our analytical results in Proposition 7. In particular, the proposed upper bound with sufficiently high η (Case (i))—and thus, the actual profit loss Δ_{DYA} —is practically zero because M_1/M_2 is very large (as discussed before), and it significantly outweighs \bar{a}/a^h (which ranges between 0.88 and 0.97 in our numerical instances). Moreover, the proposed lower bound with sufficiently low η (Case (ii)) on $\Delta_{DYA} \times 100$

is larger than 5.91%—that is, the actual loss is significant. In this case, the proposed upper bound on $\Delta_{DYA} \times 100$ ranges from 6% to 24%. We observe that this upper bound is tight in our numerical instances where the maximum deviation from the actual $\Delta_{DYA} \times 100$ is only 0.01%.

3. *Ignoring the spot price uncertainty when planning for capacity leads to a relatively small loss of profit.* Table 1 shows that the average profit loss under the DP heuristic is 5.87% (which is obtained from the two average profit losses 5.95% and 5.31% reported in the table) in our numerical instances. As discussed in Section 4.2, when the output spot price uncertainty is ignored, the incentive for holding inventory will be lower, and thus total expected storage profit over the planning horizon (as captured by \bar{M}_2) will be considerably small. Consistent with this observation, in all numerical instances we observe $\bar{M}_2 \approx 0$ and the firm invests in high yield-balanced portfolio when DP heuristic is employed. Because $\bar{M}_2 = 0$ condition is satisfied, the analytical results in Proposition 8 are relevant. In particular, the proposed lower bound on $\Delta_{DP} \times 100$ ranges from 5.34% to 5.47% in numerical instances with sufficiently high η (Case (ii)). One may argue that the profit loss under the DP heuristic is not as high as what is expected. This result crucially depends on two key observations. First, the processing margin when input and output spot prices are at their long-term means is sufficiently high in the palm industry (37.16 RM/mt in our baseline scenario). Recall that this margin is the main determinant of \bar{M}_1 because the expected storage margin in each period is practically zero. Second, FFB (input) and CPO (output) prices are highly positively correlated ($\hat{\rho} = 0.734$ in our baseline scenario) and so the variability of processing margin in each period is low. Recall that higher processing margin variability is beneficial for the firm under the optimal policy. Therefore, ignoring the price uncertainty does not lead to a sizable loss of profit.

4. *Although by-product revenue constitutes a small portion of a palm oil mill's total revenues, ignoring it during capacity planning substantially reduces the firm's profit.* In our baseline scenario, the processing cost is $\hat{c} = 40$ RM/mt and the effective processing cost is $\hat{c} =$

$\hat{c} - \hat{a}^B \hat{p}^B = -39.47$ RM/mt, and so the by-product revenue is 79.47 RM/mt. When the production yield is assumed to be at its average and the output spot price is assumed to be at its long-term mean, by-product revenue constitutes to only 13% of the total revenues. Yet, as confirmed by the values reported in Table 1, ignoring by-product revenue when planning for capacity leads to substantial profit loss—an average of 65.39% in our numerical instances. As discussed in Section 4.2, the relative magnitude of the processing and storage margins plays a key role in the profit loss experienced under this heuristic policy. In our numerical instances, because the processing margin significantly outweighs the storage margin, not accounting for the by-product revenue in the processing margin leads (on average) to a 81% reduction in M_1/M_2 , i.e., $\bar{M}_1^{NB}/M_2 = 0.19(M_1/M_2)$. Therefore, the proposed upper bound in Proposition 9 is not small (the average lower bound on $\Delta_{NB} \times 100$ is 65.61%).

5. *Using a high yield-balanced portfolio in capacity planning is a near-optimal heuristic policy.* Table 1 shows that for those numerical instances in which η is high enough to result in positive profit loss (because the optimal policy is storage dominating), the average profit loss is only 0.57%. This observation can be explained by our proposed upper bound on the actual profit loss Δ_{HYBP} in Proposition 10. This upper bound is low because, as established previously, M_1/M_2 is very large (since the processing margin is substantially higher than the storage margin) and so $(M_1/M_2)^2$ significantly outweighs $\eta/(a^h)^2$. In the relevant numerical instances, we observe that the average upper bound on $\Delta_{HYBP} \times 100$ is only 0.73% within a range from 0.02% to 3.88%.

6. Conclusion

This paper contributes to the operations management literature by studying the joint processing and storage capacity investment decisions of a commodity processor in the context of agricultural industries. Previous work on commodity processors has focused on operating decisions (e.g., processing and inventory) and usually assumed exogenously given capacity levels for processing and storage resources. We study how these capacity levels are chosen. Toward this end, we develop a stylized multiperiod model to help devise a characterization of capacity investment policy that has useful ramifications in terms of numerical computation and sensitivity analysis. We provide insights on how spot price uncertainty shapes the firm's capacity investment policy and profitability, and insights on the benefits of using the optimal capacity investment policy rather than heuristic policies.

Our work has several limitations due to our specific modeling assumptions. Further research is needed to validate the relevance of our insights when those

assumptions are relaxed. First, we assume that the capacity levels, once chosen, remain fixed during the planning horizon. In practice, however, it is not uncommon for firms to own multiple processing facilities (e.g., Wilmar in the palm industry). In that case, the firm can temporarily increase one facility's processing capacity by shifting processing to other facilities. Incorporating flexible capacity into our framework would be a promising avenue for future research. Second, our model assumes that the firm does not face frictions in transportation. In practice, however, there may be constraints (on input procurement or output sales) that arise from limits to transportation capacity (Devalkar et al. 2011). There could also be marginal costs associated with transferring output from the storage facility to the market, resulting in a spread between the output's marginal spot procurement cost and marginal spot sales revenue (Kazaz and Webster 2011). When these frictions in transportation are of a significant nature, they will provide another rationale for the firm to hold inventory.⁶ Incorporating these frictions into our model should prove to be an interesting avenue for future research. Third, our model calibration was based on the palm industry. Because other oilseeds and grain industries share common characteristics with the palm industry—for instance, a significantly higher cost of processing capacity than of storage capacity and the expected dominance of the processing margin over the storage margin—we expect that the majority of our findings are valid for those industries as well.⁷ That being said, future research is still needed to verify this conjecture by using our paper's methodology to calibrate the model based on a different agricultural industry.

Acknowledgments

The authors are grateful to the review team for offering many constructive suggestions for improvement.

Endnotes

¹ In a general processing environment, the convexity of the capacity investment costs can also be attributed to limits on production technology and increasing managerial complexity or maintenance cost with additional investment.

² In our model, decisions are made under the true pricing measure that reflects the firm's actual expectations of the spot market prices. A stream of papers (e.g., Devalkar et al. 2011) considers models in which decisions are made under the risk-neutral pricing measure that reflects the spot price expectation in a competitive equilibrium.

³ In the unabridged version of this paper (which is available from http://ink.library.smu.edu.sg/cgi/viewcontent.cgi?filename=0&article=4187&context=lkcsb_research&type=additional), we relax this assumption and show that it is not a critical assumption in our model.

⁴ In the literature, correlated mean-reverting processes are also used to model the evolution of the natural logarithm of commodity spot prices; see, for example, Secomandi (2010b) and Secomandi and Wang (2012).

⁵ See http://www.simedarby.com/clients/simedarby_group/assets/contentMS/img/template/editor/Sime_Darby_AR2013_Lowres.pdf (accessed July 7, 2017).

⁶In our model, the firm's rationale for holding inventory is to benefit from fluctuations in the output spot price, which is one of the most common reasons for holding inventory in agricultural industries (Westlake 2005).

⁷As a first attempt to address the relevance of our insights outside of the palm industry, we conduct additional numerical experiments in which we alter several key parameters that were kept constant at their calibrated values in Section 5. The details of this analysis are relegated to the unabridged version of this paper. We find that our main insights about the heuristic capacity investment policies continue to hold except for one: in contrast with the palm setting, ignoring the spot price uncertainty when planning for capacity leads to a substantial loss of profit.

References

- Allen SJ, Schuster EW (2004) Controlling the risk for an agricultural harvest. *Manufacturing Service Oper. Management* 6(3):225–236.
- Boyabatlı O (2015) Supply management in multiproduct firms with fixed proportions technology. *Management Sci.* 61(12):3013–3031.
- Boyabatlı O, Kleindorfer PR, Koontz SR (2011) Integrating long-term and short-term contracting in beef supply chains. *Management Sci.* 57(10):1771–1787.
- Boyabatlı O, Wee KE, Avci B (2016) Farm-yield management when production rate is yield dependent. Working paper, Singapore Management University, Singapore.
- Chang L, Sani ARA, Basran Z (2003) An economic perspective of oil extraction rate in the oil palm industry of Malaysia. *Oil Palm Indust. Econom. J.* 3(1):25–31.
- Devalkar S, Anupindi R, Sinha A (2011) Integrated optimization of procurement, processing, and trade of commodities. *Oper. Res.* 59(6):1369–1381.
- Devalkar S, Anupindi R, Sinha A (2017) Dynamic risk management of commodity operations: Model and analysis. Working paper, University of Michigan, Ann Arbor.
- 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 Sci.* 60(12):2908–2926.
- Fackler PL, Livingston MJ (2002) Optimal storage by crop producers. *Amer. J. Agricultural Econom.* 84(3):645–659.
- Goel A, Tanrisever F (2017) Financial hedging and optimal procurement policies under correlated price and demand. *Production Oper. Management*, ePub ahead of print July 19, <http://dx.doi.org/10.1111/poms.12723>.
- Kazaz B (2004) Production planning under yield and demand uncertainty with yield-dependent cost and price. *Manufacturing Service Oper. Management* 6(3):209–224.
- Kazaz B, Webster S (2011) The impact of yield-dependent trading costs on pricing and production planning under supply uncertainty. *Manufacturing Service Oper. Management* 13(3):404–417.
- Kleindorfer PR, Wu DJ (2003) Integrating long-and short-term contracting via business-to-business exchanges for capital-intensive industries. *Management Sci.* 49(11):1597–1615.
- Kouvelis P, Li R, Ding Q (2013) Managing storable commodity risks: Role of inventories and financial hedges. *Manufacturing Service Oper. Management* 15(3):507–521.
- Lai G, Wang MX, Kekre S, Scheller-Wolf A, Secomandi N (2011) Valuation of storage at a liquefied natural gas terminal. *Oper. Res.* 59(3):602–616.
- Mehrotra M, Dawande M, Gavirneni S, Demirci M, Tayur S (2011) Production planning with patterns: A problem from processed food manufacturing. *Oper. Res.* 59(2):267–282.
- Pei PP-E, Simchi-Levi D, Tunca TI (2011) Sourcing flexibility, spot trading, and procurement contract structure. *Oper. Res.* 59(3):578–601.
- Plambeck EL, Taylor TA (2013) On the value of input efficiency, capacity efficiency, and the flexibility to rebalance them. *Manufacturing Service Oper. Management* 15(4):630–639.
- Secomandi N (2010a) Optimal commodity trading with a capacitated storage asset. *Management Sci.* 56(3):449–467.
- Secomandi N (2010b) On the pricing of natural gas pipeline capacity. *Manufacturing Service Oper. Management* 12(3):393–408.
- Secomandi N, Wang MX (2012) A computation approach to the real option management of network contracts for natural gas pipeline transport capacity. *Manufacturing Service Oper. Management* 14(3):441–454.
- U.S. Department of Agriculture (2015) Oilseeds: World markets and trade—October 2015. Accessed July 7, 2017, <http://usda.mannlib.cornell.edu/usda/fas/oilseed-trade//s/2015/oilseed-trade-10-09-2015.pdf>.
- Van Mieghem JA (2003) Capacity management, investment, and hedging: Review and recent developments. *Manufacturing Service Oper. Management* 5(4):269–302.
- Westlake MJ (2005) Addressing marketing and processing constraints that inhibit agrifood exports: A guide for policy analysts and planners. Report, Food and Agriculture Organization of the United Nations, Rome.
- Zellner A (1962) An efficient method of estimating seemingly unrelated regressions and tests for aggregation bias. *J. Amer. Statist. Assoc.* 57(298):348–368.
- Zhou Y, Scheller-Wolf AA, Secomandi N, Smith S (2014) Managing wind-based electricity generation in the presence of storage and transmission capacity. Working paper, Carnegie Mellon University, Pittsburgh, PA.