

Three Essays on Pricing and Risk Management in Industrial Practice

by

Jie Ning

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Industrial & Operations Engineering)
in The University of Michigan
2013

Doctoral Committee:

Associate Professor Jussi S. Keppo, Co-Chair
Professor Romesh Saigal, Co-Chair
Professor Judy Jin
Assistant Professor Owen Wu

© Jie Ning 2013
All Rights Reserved

To my parents.

ACKNOWLEDGEMENTS

This thesis would not have been possible without the enormous help and encouragement from my advisors Professor Romesh Saigal and Professor Jussi Keppo. I am profoundly influenced by their adamant and unyielding strive for research of high caliber and rigor. Working with them makes the past five years the most rewarding and transformative period in my entire life, turning me from a casual thinker to an independent researcher, about to embark on an odyssey of my own.

I would also like to thank my committee members: Professor Judy Jin and Professor Owen Wu, for their great efforts and help in my prelim, dissertation and defense. My special thanks go to Professor Volodymyr Babich, the co-author of my second essay, for all the insightful suggestions and supports on the paper and my career.

I am especially grateful for all the help from our industrial and research collaborators: Dr. John Handley of Xerox, employees of Norsk Hydro Holland, Dr. Lalit Patil, Prof. Stein-Erik Fleten, and Dr. Geir Ringen. Without their support, none of the work in this thesis would happen.

I am greatly indebted to the love and help from my husband, who has always been there, applauding for my achievements and encouraging during my defeats. I also thank all my friends. With each other we share our stories—sad and merry, and from each other we derive lasting strengths that support us to forge ahead.

Finally, I would like to express my gratitude to all the faculty and staff members of the department, who have been extremely helpful whenever I reach out for suggestions. Among others, I especially thank Professor Robert Smith, Professor Xiuli

Chao and Professor Marina Epelman. It is their support and trust five years ago that opened up this new chapter of my life and put down the first touch of this dissertation.

TABLE OF CONTENTS

DEDICATION	ii
ACKNOWLEDGEMENTS	iii
LIST OF FIGURES	vii
LIST OF TABLES	viii
CHAPTER	
I. Introduction	1
II. Pricing Demand Risks in Just-in-Time Outsourcing	4
2.1 Introduction	4
2.2 Literature review	7
2.3 Model setup	9
2.4 Demand risk pricing	12
2.4.1 Demand risks over a subcontract	13
2.4.2 Pricing demand risks over the entire outsourcing horizon	23
2.5 Empirical study on an auto-parts CM	24
2.5.1 Model calibration	25
2.5.2 Demand risk of the outsourcing business	27
2.6 Concluding remarks	31
III. Empirical Study on Managed Print Services Pricing	32
3.1 Introduction	32
3.2 Literature review	37
3.3 Data	40
3.3.1 Pricing data set	41
3.3.2 Cost data set	42

3.3.3	Multi-price data set	43
3.4	Model setup: contracts and random processes	44
3.5	Price elasticity of printing demand	47
3.6	Customer’s pricing model	51
3.6.1	Model formulation	53
3.6.2	Empirical analysis	55
3.7	Optimal MPS pricing	62
3.7.1	Optimization model	63
3.7.2	Empirical analysis	74
3.8	Concluding remarks	81
IV. Price dynamics of used durable goods: A dynamic factor model approach		84
4.1	Introduction	84
4.2	Data	87
4.2.1	Data preparation	88
4.2.2	Descriptive statistics	91
4.3	Multi-level dynamic factors model	92
4.4	Industry and segment factors	97
4.4.1	Latent dynamic factors	98
4.4.2	Persistence properties of the factors	102
4.4.3	Variance decomposition	104
4.5	OEM analysis	106
4.5.1	Model on OEM and brand factors	108
4.5.2	OEM and brand factors	109
4.6	Concluding Remarks	112
V. Conclusion		115
BIBLIOGRAPHY		117

LIST OF FIGURES

Figure

4.1	Estimated industry and controlled segment factors. 1.64 standard deviation (i.e., 90% significance level) band. The red dot corresponds to Sept. 2008, the bankruptcy of Lehman Brothers.	101
4.2	Loaded industry factor (solid black), controlled segment factor (dashed black) and segment factor (dotted red).	103
4.3	Impacts of product recalls. (a-c) plot OEM 4's OEM and brand factors with 1.64 standard deviation (i.e., 90% significance level) band. In Jan. 2010 (red dot), several brand 1 product models were recalled.	111
4.4	Impacts of brand termination. (a-c) plot OEM 2's OEM and brand factors with 1.64 standard deviation (i.e., 90% significance level) band. Around the end of 2011 and 2010 (red dots), OEM 1 terminated brand 2 and OEM 2 terminated brand 3, respectively.	113

LIST OF TABLES

Table

2.1	Model parameters.	26
2.2	Risk ratios of outsourcing	30
3.1	Continuous variables and descriptive statistics	42
3.2	Definition of dummy variables and descriptive statistics	43
3.3	Descriptive statistics for BW and color volume series under multiple prices	44
3.4	Parameters estimates in model (3.11)	59
3.5	Model fitting statistics	61
3.6	Correlations between predicted and observed prices	77
4.1	Descriptive statistics on the models' growth rates within each segment	92
4.2	Loadings of the industry factor and variance of shocks to controlled segment factors	99
4.3	Persistence of latent factors	102
4.4	Variance decomposition	107
4.5	Persistence of OEM and brand factors	110

CHAPTER I

Introduction

Uncertainty affects both the demand and supply sides in industrial practice. Uncertainty in demand may increase the production costs. Uncertainty in earnings cash flows might result in high borrowing costs and budget problems. Variations in the external economic condition could greatly impact the product prices. This thesis presents three essays that price, manage, and understand the aforementioned uncertainties in different industrial settings.

In Chapter II, we study just-in-time outsourcing between an original equipment manufacturer (OEM) and a contract manufacturer (CM). Both the OEM and the CM have flexible production capacity and concave earnings in capacity usage. Under the just-in-time contract, the CM assumes demand risks from the OEM at the expense of her profit margins, because the CM is not explicitly compensated for the cost of the demand uncertainty. In this paper we price the demand risks from both agents' perspectives using marginal utility pricing theory, accounting for the demand pooling effect under flexible capacity. We show that when the outsourcing demand is positively correlated with the agent's existing business, the higher risk it carries, the more it benefits the OEM, and less the CM. Furthermore, we introduce risk ratio to measure the "riskiness" of the outsourcing business. Particularly, the risk ratio quantifies the extent to which the CM's gross profit margin is eroded. Finally, we

apply our model to a data set of a Norwegian auto-parts CM serving 11 OEMs. We report that while three OEMs' orders are more risky and costly than others, none of these outsourcing businesses' risk ratios exceeds the CM's gross profit margin. This agrees with our expectation given the CM's healthy business and prudence in selecting customers.

Managed services is a rapidly growing industry that offers IT infrastructure management for companies and institutions. In Chapter III, we conduct an in-depth study on managed services in the context of managed print services (MPS) by examining the contractual interactions between the MPS provider and his institutional customers using a proprietary data set. On the customer's side, we demonstrate that the customer's printing demand is insensitive to service prices over the observed contracts. Furthermore, we show that individual institutional customer's valuation of the service is dominated by the population valuation, despite of the heterogeneity in the customers' industries and negotiation processes. In particular, the population valuation of the service decreases with the printer fleet size, suggesting the market power of large companies. On the provider's side, we empirically show that given the customer's service valuation, the provider decides the optimal contracts accounting not only for the expected earnings but also for the variability of earnings; that is, the provider is risk-averse. Based on our model and data, all customers, irrespective of their industries and fleet sizes, bring in similar risk-adjusted earnings. This may indicate intense competition among the providers as different customers' risks are priced in a same way.

Chapter IV studies the comovements of resale prices of a particular type of used durable goods using US resale price data over 13 years. Starting from a multi-level dynamic factor model in state space presentation, we identify the comovements of products within each functionality segment, across the entire industry, and within each brand and OEM after the segment-level comovements are controlled for. We

observe that, despite the heterogeneity in product models and brands, there are strong comovements throughout the used goods market and within each segment—the latent industry and segment factors are able to capture up to 81.4% of the variation of a particular product’s price change. We also show that after the latent segment factors are controlled for, products from the same manufacturer still exhibit strong comovements, but not so much for products of the same brand. Finally, we observe no material impact on the OEM and brand factors when a brand is terminated. A big product recall, however, results in significant product price drops.

CHAPTER II

Pricing Demand Risks in Just-in-Time Outsourcing

2.1 Introduction

An increasing number of original equipment manufacturers (OEMs) opt to outsource production-related activities to contract manufacturers (CMs), resulting in rapid growth in contract manufacturing business. The international contract electronics industry, for example, has experienced a 33.4% increase in revenue, reaching \$347.3 billion in 2010 (Dinges 2011). Boom in contract manufacturing also gives rise to large CMs serving a number of OEMs around the globe. For instance, our collaborator, a Norwegian auto-parts CM, supplies bumper beams to 11 international OEMs.

This paper is motivated by our collaborator's just-in-time (JIT) contractual relationship with her OEM customers. The JIT contract is used quite often in contract manufacturing in various industries such as automotive and electronics (Huggins and Olsen 2003, Duenyas et al. 1997, Waters-Fuller 1996). It requires the CM to deliver exactly the OEM's orders at designated time and location. While inadequate delivery is severely penalized, excess production is at the CM's own cost. For the OEMs, the benefits of the JIT outsourcing include reduced inventory costs and shorter produc-

tion lead time (e.g., Ansari and Modarress 1990, Tracey et al. 1995). For the CMs, however, researchers found that they may bear the production and inventory costs, and demand risks transferred from the OEMs (Romero 1991, Fandel and Reese 1991, Dong et al. 2001). Indeed, one major concern of the CMs is the unstable outsourcing orders under the short response time allowed by the JIT contracts (Waters-Fuller 1996, González-Benito and Spring 2000). Some CMs would even quit the contractual relationship when the demand becomes too volatile (Matson and Matson 2007).

Our collaborator, the Norwegian auto-parts supplier, is among the CMs who are particularly concerned about the demand risks they have to bear in JIT outsourcing. With its earnings concave in production capacity usage, the Norwegian CM is reluctant to accept highly variable businesses as they could be more costly and less profitable. Although measures are taken to offer certain protection against demand fluctuations, for instance, setting thresholds on outsourcing orders such that, once reached, unit bumper beam price will be renegotiated, our collaborator's concerns are not fully addressed. As a result, the Norwegian CM only chooses OEMs with stable and predictable businesses.

High outsourcing demand variability itself, however, should not be a rejection rule if the CM is able to pool demands from different OEMs together. Demand pooling is made possible by the flexible production capacity, which is typical for CMs supplying standard product components such as automobile parts and personal laptop screens. In our collaborator's case, for example, while different car models require different bumper beams, they can all be produced using the same equipment with minor module changes incurring negligible additional costs. With demand pooling, the CM's overall capacity usage for all her customers could be relatively stable even when an individual OEM's demand is volatile. In this case, the cost caused by the OEM's random orders is also small. Thus, contracting with such an OEM still provides high profit margins even if the OEM places volatile orders and does not

compensate the CM for the risk cost explicitly.

On the other hand, from the OEM's perspective, JIT outsourcing is a perfect tool to transfer the unwanted random demand to the CM, thus reducing the capacity usage variability and consequently the production cost. However, not all demand randomness needs to be eliminated, especially when the OEM has flexible production capacity that allows demand pooling.

In this paper, we propose a model to price the CM's cost of bearing the demand risk and the OEM's benefit of transferring it under flexible production capacity. We show that, with demand pooling, the variability of the outsourcing business only partly determines the risk cost/benefit. We distinguish between the capacity risk and demand risk: the latter refers to the absolute variability of the outsourcing demand, while the former refers to the variability that cannot be offset by other demands utilizing the same production capacity. Under flexible capacity, only the capacity risk matters, as all idiosyncratic demand risks are diversified away. Particularly, the higher the capacity risk carried by the outsourcing demand, the more the outsourcing benefits the OEM and costs the CM. One interesting case is an outsourcing demand whose randomness can be completely offset by other sources. Such a demand carries zero capacity risk and, thus, creates zero cost for the CM and zero benefit for the OEM.

Based on the price of demand risks, we introduce a scale-free metric *risk ratio* to measure the relative cost of demand risks with respect to the expected revenue of the outsourcing business. For the CM, the higher the risk ratio, the more the CM's gross profit margin is eroded if the OEM does not compensate for the CM's bearing of the demand risk. For the OEM, the risk ratio measures the efficiency of outsourcing in terms of shifting the demand risk costs. Given two potential outsourcing demands, the one with the higher risk ratio creates more benefit for the OEM per unit outsourced.

In the empirical analysis, we apply our model to the Norwegian auto-parts CM

supplying bumper beams to 11 OEMs. We report that none of the OEMs' outsourcing businesses carries demand risk that significantly erodes the CM's profit margin. Our empirical result is consistent with this intuition given the long-standing contractual relationships between the CM and the OEMs in our data set.

The chapter proceeds as follows. Section 2 reviews related work in literature. Section 3 explains the model setup and Section 4 presents our pricing model. Section 5 applies the pricing model to a Norwegian auto-parts CM. Section 6 concludes.

2.2 Literature review

In this paper we propose a model to price the CM's future capacity usage, and then the demand risks under JIT contracts. Secomandi (2010) and Secomandi and Wang (2012) study pricing of natural gas pipeline capacity under firm (guaranteed) transport contracts. Under firm contracts, buyers of the pipeline capacity pay the demand rate (reservation fee) to reserve the maximum quantity of capacity to use and a commodity rate (execution fee) to use this capacity. Secomandi (2010) looks at point-to-point transport contracts while Secomandi and Wang (2012) focus on network contracts. In Secomandi (2010), Secomandi shows that under a linear earnings function, the pipeline capacity price, i.e. the demand rate, perceived by the buyer is equal to the price of a spread option on the natural gas futures prices at the delivery and receipt markets. Furthermore, he models the capacity value perceived by the pipeline company to be the equilibrium price that maximizes the welfare of the delivery and receipt markets. Proving that the equilibrium price coincides with the spread option price, Secomandi tests his model by comparing the realized demand rate with the prediction. In this paper, we study a problem where the randomness in demand has important cost implications for both the CM and the OEM. Hence, instead of taking an option perspective and pricing the right for future execution, we directly price the CM's cost of accepting and the OEM's benefit of removing a random

demand. This approach better matches the practitioners' interests. Furthermore, instead of focusing on the specific outsourcing business between the CM and the OEM, we look at the entire business portfolios of both agents. We adopt concave earnings function observed from the data set and price demand risks based on marginal utility pricing theory. In the end, we test our model by applying it to a Norwegian CM and report supportive result.

This paper studies the influence of correlated capacity usages on the price of demand risks. Demand correlation and risk pooling have been widely examined in the operations management literature. In supply chain coordination, Van Mieghem (1999) illustrates the influences of demand variability and correlation on the OEM's investment in a game-theoretic model. While Van Mieghem suggests that demand pooling may also affect the supplier, he does not analyze this issue further. In the area of capacity management, Van Mieghem (1998, 2007b, 2003) explore how to mitigate the mismatch between supply and demand via effective capacity management when the decision maker faces correlated demands. Chod and Rudi (2005) extend the model by allowing the decision maker to influence the demand via product pricing. Apart from reducing demand and supply mismatches, several studies have also considered the role of demand correlation and risk pooling in managing financial risks created by demand uncertainties. Gaur and Seshadri (2005) consider a risk-averse producer hedging financial risks with a market instrument whose price is correlated with the product demand. Zhou and Rudi (2010) examine the pricing of an over-the-counter financial contract designed to hedge risks in product markets. Chod et al. (2010) study the relationship between operation flexibility and financial hedging using an instrument whose payoff is negatively correlated with the demand. They report that operation and financial hedging can be complements or substitutes to each other depending on the demand correlation. Our work assumes that there is no mismatch between supply and demand: The OEM can always satisfy the demand, either by

producing in-house or outsourcing under a JIT contract. But via outsourcing, the OEM further benefits from the reduced overall capacity usage variability. We quantify this benefit and examine its dependence on capacity usage correlation.

Another main element of our model is demand risk reallocation in the supply chain. Under a JIT contract, the OEM effectively transfers his demand risks to the CM. Ülkü et al. (2007) examine whether the OEM or the CM should bear the financial risk imposed by demand variation to achieve maximum supply chain profit in a situation where the CM serves several OEMs with independent businesses. Cachon (2004) studies the inventory risk allocation between a retailer and a supplier using wholesale price contracts. In this paper, we evaluate the transferred demand risks from the OEM and the CM's own perspectives, rather than from the viewpoint of the entire supply chain.

Component outsourcing has also been studied in different situations. Grahovac and Parker (2003) show that OEMs outsource components of high degree of modularity to avoid over-investments even if the CM does not have any economy of scale. Mikkola (2003) reports that component outsourcing cultivates closer cooperation between suppliers and OEMs and this way raises innovation in future product design.

2.3 Model setup

Consider a component that is a standard part of a variety of products made by different OEMs, e.g., personal computer monitors, car bumper beams. This component can be customized to fit different products by simple changes of modules in a standard manufacturing procedure. Therefore, it uses flexible production resources.

An OEM (he) makes a range of products containing this component. Hence, the OEM can take advantage of the production flexibility and pool the market demand of different products together. With in-house production capacity installed, the OEM considers outsourcing part of the total demand to a CM via a JIT contract. The

CM (she) under consideration already supplies this component to a number of OEMs under JIT contracts when the new customer approaches. Thus, she also has flexible production capacity and can pool demand from different OEMs together.

We consider the case of JIT outsourcing where the OEM orders random number of product components at pre-determined time points over the contract horizon.¹ Whenever an order is released, the CM is required to fulfill it within a period shorter than the production lead time. Therefore, for the OEM, JIT outsourcing benefits him by replacing random in-house production capacity usage with risk-free deliveries. For the CM, while contracting with the new OEM brings in more business, it might also lead to more volatile usage of her production capacity. This is particularly true when the CM keeps little inventory of the component possibly for cost reduction purposes. Indeed, according to our collaborator who supplies bumper beams to a number of OEMs, “The bumper beams are directly loaded to the delivery trucks from the production line.”

Our goal is to quantify the OEM’s benefits and CM’s costs resulted from the transfer of the uncertain demand, or equivalently, stochastic capacity usage, under JIT outsourcing. To start, we make the following assumptions on the outsourcing process: (i) The OEM places stochastic orders Q_i at fixed time points $\{T_i\}$ ($i = 1, 2, \dots, N$) with $T_0 = 0$; (ii) the OEM’s orders are met immediately; and (iii) the CM keeps no inventory. Consequently, the deliveries/orders at different time points are not related. This allows us to view the JIT contract as a collection of independent subcontracts that become effective in a consecutive manner. In particular, subcontract i comes into effect at time T_{i-1} and expires at time T_i . At its expiration T_i , the OEM places the order Q_i and the CM immediately delivers it.

As the OEM’s orders Q_i ($i = 1, 2, \dots, N$) are contingent on the exogenous market demand, it can be modeled as a stochastic process that takes nonnegative values over

¹This is the JIT arrangement between our collaborator and her OEM customers.

the service horizon. At the expiration of subcontract i , both the OEM and the CM observe the latest demand Q_i and update their information on the stochastic demand process, as well as the prediction of the next outsourcing order Q_{i+1} . Therefore, at any point in time, the demand risk faced by the CM and shrugged off by the OEM is only the uncertainty of the next outsourcing demand. In other words, the benefit/cost of the transfer of stochastic demand in JIT outsourcing primarily depends on the its benefit/cost over each of the N independent subcontracts. Hence, we start by studying the demand risk over one subcontract.

Consider a subcontract and assume that it starts at time 0 and ends at time T without loss of generality. Denote the outsourcing demand at T by Q_T , with $Q_0 = q$. Based on the discussion above, we model Q_T as the terminal value of a nonnegative stochastic process Q_t ($t \in [0, T]$) dependent on the OEM's total market demand. Let M_T denote the rest of the OEM's market demand that is satisfied in-house at T . We model M_T as the terminal value of a nonnegative stochastic process M_t ($t \in [0, T]$) with $M_0 = m$. We assume that M_T is significantly greater than Q_T , indicating that the OEM usually outsources a relatively small portion of the total demand.

Before contracting with the new OEM, the CM already supplies the component to several other OEM customers under JIT contracts. Let S_T be the CM's total demand at T *without* the new outsourcing business, with $S_0 = s$. S_T is the terminal value of a nonnegative stochastic process S_t ($t \in [0, T]$). We assume S_T to be significantly greater than Q_T to capture the fact that the CM usually operates on a large business scale.

We use the quantities Q_T , S_T and M_T to measure the units of production capacity used by these demand, because satisfying each unit of the demand requires a fixed amount of production resources such as material, manufacturing time, and labor. For the same reason, we use the terms "demand" and "capacity usage" interchangeably, and call S_T and M_T as the original capacity usage in the following analysis. We

assume that Q_T , S_T and M_T all have continuous marginal distributions. Although the product is in discrete unit, the actual orders and capacity usage are usually in tens or hundreds of thousands. In addition, the contracts usually do not require the OEM to release his order in batches. Therefore, the continuous approximation is reasonable and convenient for analysis.

The outsourcing demand process Q_t ($t \in [0, T]$) is public information. This corresponds to the fact that the OEM usually grants the CM access to his demand database so as to make better production plans (Cachon and Fisher 2000, Li and Lin 2006); this is also the case with our collaborator in the auto-parts industry. M_t and S_t are private information to the OEM and the CM, respectively.

Assume that the CM and the OEM are profit maximizers. We learn from our data set that the CM's earnings are concave increasing in the outsourcing demand. The OEM is likely to use the same manufacturing equipments and follow the same procedure in his in-house production. Thus it is reasonable to postulate that his earnings are concave increasing in the market demand as well. Let a concave increasing function $U(\cdot)$ be their instantaneous earnings on the demand, or equivalently, capacity usage, at time T . The more concave $U(\cdot)$ is, the lower tolerance the OEM/CM has for capacity variability as less expected earnings can be obtained. Hence, we can interpret the concavity of the earnings as the OEM/CM's risk aversion parameter. $U(\cdot)$ is continuously differentiable on the nonnegative real line, since the capacity usage is never negative.

2.4 Demand risk pricing

In this section, we price the demand risks in JIT outsourcing from both the OEM and the CM's perspectives. We start by considering the demand risks over the subcontract on $[0, T]$. Then we price the demand risks over the entire service horizon using results from a single subcontract.

For ease of exposition, we use $Z \in \{S, M\}$ to denote either the CM ($Z = S$) or the OEM ($Z = M$) in the rest of the paper. When studying the subcontract on $[0, T]$, to make explicit the dependence of Q_T and Z_T on their initial values, q and z , we write them as Q_T^q and Z_T^z , respectively.

2.4.1 Demand risks over a subcontract

Consider the subcontract on $[0, T]$. The OEM and CM observe the latest demand q at time 0 and plan the production for the next demand that will be realized at T . Thus, they face inter-temporal demand risks over $[0, T]$. In the following, we first find the “present value” of the future outsourcing demand, or equivalently, production capacity usage, accounting for the inter-temporal demand risks. Then using this present value, we price the demand risks under the subcontract.

2.4.1.1 Present value of future production capacity

Our primary goal is to quantify the inter-temporal demand risks under flexible production capacity which enables demand pooling. This greatly resembles the standard pricing problem in finance, which prices the inter-temporal uncertainties of assets’ future payoffs in a variety of contexts (see Cochrane 2005 for an introduction to different asset pricing models). Therefore, we resort to standard financial models, particularly the marginal utility pricing theory (e.g., Davis 1998, 2001, Karatzas 1997, Staum 2008), and define the *present value* of the future outsourcing demand as follows.

Definition II.1. The present value of the outsourcing demand, Q_T^q , is the certainty equivalent quantity (CEQ) at which, at time 0, the OEM/CM is indifferent between remaining in the original business, Z_T^z ($z \in \{s, m\}$), and substituting a small amount of his/her original business for the outsourcing demand. More precisely, let $CEQ(Q_T^q)$

denote this present value. It satisfies

$$\left. \frac{\partial \mathbb{E} U \left(Z_T^{z-\delta} + \frac{\delta}{\text{CEQ}(Q_T^q)} Q_T^q \right)}{\partial \delta} \right|_{\delta=0} = 0. \quad (2.1)$$

We write the present values of Q_T^q perceived by the CM and OEM as $\text{CEQ}_s(Q_T^q)$ and $\text{CEQ}_m(Q_T^q)$, respectively. Z_T^z , Q_T^q and $\text{CEQ}_z(Q_T^q)$ are all in product units (e.g., number of bumper beams). Note that although the product unit is discrete, the actual orders and capacity usage are usually in large quantities, allowing us to approximate a few units as the infinitesimal amount δ in (2.1).

Theorem II.2. *Define $V(z) \equiv \mathbb{E}[U(Z_T^z)]$, where $z \in \{s, m\}$ is the initial capacity usage of the CM/OEM at time 0. If $V(z)$ is continuous and differentiable, the present value of the outsourcing demand, Q_T^q , is given by*

$$\text{CEQ}_z(Q_T^q) = \frac{\mathbb{E} \left[\frac{dU(Z_T^z)}{dZ_T^z} Q_T^q \right]}{\frac{dV(z)}{dz}}. \quad (2.2)$$

Proof. By Definition II.1, CEQ satisfies equation (2.1). Taylor expansion around $Z_T^{z-\delta}$ over small δ gives

$$U \left(Z_T^{z-\delta} + \frac{\delta}{\text{CEQ}_z(Q_T^q)} Q_T^q \right) = U(Z_T^{z-\delta}) + U'(Z_T^{z-\delta}) \frac{\delta}{\text{CEQ}_z(Q_T^q)} Q_T^q + o(\delta), \quad (2.3)$$

where $U'(x) \equiv \frac{dU(x)}{dx}$. Because $V(z) \equiv \mathbb{E}[U(Z_T^z)]$, above becomes

$$U \left(Z_T^{z-\delta} + \frac{\delta}{\text{CEQ}_z(Q_T^q)} Q_T^q \right) = V(z - \delta) + U'(Z_T^{z-\delta}) \frac{\delta}{\text{CEQ}_z(Q_T^q)} Q_T^q + o(\delta). \quad (2.4)$$

$V(z)$ is continuous and differentiable with respect to z . Therefore, by equation (2.1), we have

$$\left. -V'(z - \delta) + U'(Z_T^{z-\delta}) \frac{1}{\text{CEQ}_z(Q_T^q)} Q_T^q + o(\delta) \right|_{\delta=0} = 0. \quad (2.5)$$

This proves the theorem. □

Intuitively, when $T = 0$, $\text{CEQ}_z(Q_0^q) = q$. That is, the present value of the observed outsourcing demand q is still q , as no uncertainty is involved.

From Theorem II.2, the present value of Q_T^q can be written as

$$\text{CEQ}_z(Q_T^q) = \frac{\mathbb{E} U'(Z_T^z) \mathbb{E} Q_T^q + \text{cov}(U'(Z_T^z), Q_T^q)}{V'(z)}, \quad (2.6)$$

where $f'(x) = df(x)/dx$. We further rewrite $\text{cov}(U'(Z_T^z), Q_T^q)$ in (2.6) schematically as $\rho_u \sigma_q \sigma_u$, where ρ_u is the correlation between the outsourcing demand, Q_T^q , and the OEM/CM's marginal utility at the original capacity usage, $U'(Z_T^z)$. σ_q and σ_u are the standard deviations of Q_T^q and $U'(Z_T^z)$, respectively. Then (2.6) becomes

$$\text{CEQ}_z(Q_T^q) = \frac{\mathbb{E} U'(Z_T^z) \mathbb{E} Q_T^q + \rho_u \sigma_q \sigma_u}{V'(z)}. \quad (2.7)$$

Equation (2.7) suggests that the variability of the outsourcing demand influences the present value via the term $\rho_u \sigma_q \sigma_u$. As σ_u pertains to the OEM/CM's utility and original capacity usage, it is fixed. Hence, the dependence of $\text{CEQ}_z(Q_T^q)$ on the demand risk is primarily determined by the correlation coefficient ρ_u . This means that the common impression that, the more variable Q_T^q , the more risky the demand, and thus the lower its present value, is not true. Indeed, we will show that, under general modeling assumptions, the sign of ρ_u in (2.7) is the opposite of the correlation between the outsourcing demand and the OEM/CM's original capacity usage. This yields an intuitive dependence of $\text{CEQ}_z(Q_T^q)$ on σ_q that nicely resembles the standard results in risk diversification.

Assumption II.3. *Assume that the outsourcing demand, Q_T^q , and the CM/OEM's original capacity usage, Z_T^z ($Z \in \{S, M\}$), are random variables with log-normal distributions.*

We note that Assumption II.3 is indeed the commonly-used modeling assumption in analysis of nonnegative stochastic processes, thus does not limit the applicability of our model and result in practice. Section 5 gives an example of modeling S_t and Q_t as processes with log-normal marginal distributions from real data.

Lemma II.4. *Let $\text{sgn}(\cdot)$ be the sign function that equals 1 if the argument is positive, -1 if negative, and 0 if zero. Under Assumption II.3 and a concave increasing earnings function that is continuously differentiable, we have*

$$\text{sgn}(\rho_u) = -\text{sgn}(\rho),$$

where $\rho \equiv \text{corr}(Q_T^q, Z_T^z)$.

Proof. Let $Q_T^q \sim LN(\mu_q, \delta_q)$, and $Z_T^z \sim LN(\mu_z, \delta_z)$. Define

$$Q^* \equiv \frac{\log(Q_T^q) - \mu_q}{\delta_q}, \quad Z^* \equiv \frac{\log(Z_T^z) - \mu_z}{\delta_z} \quad (Z \in \{S, M\}).$$

Then Q^* and Z^* are standard normal random variables. Let $\tilde{\rho} = \text{corr}(Q^*, Z^*)$. We have

$$Q^* = \tilde{\rho}Z^* + \sqrt{1 - \tilde{\rho}^2}\tilde{Z}^*,$$

where \tilde{Z}^* is a standard normal random variable independent of Z^* . Therefore,

$$Q_T^q = \exp \left\{ \mu_q + \delta_q \tilde{\rho} Z^* + \delta_q \sqrt{1 - \tilde{\rho}^2} \tilde{Z}^* \right\}.$$

Let $W \equiv e^{Z^*}$, $\tilde{W} \equiv e^{\tilde{Z}^*}$ for ease of exposition. W is independent of \tilde{W} . Then

$$\text{cov}(U'(Z_T^z), Q_T^q) = \text{cov} \left(U'(e^{\mu_z} W^{\delta_z}), e^{\mu_q} W^{\delta_q \tilde{\rho}} \tilde{W}^{\delta_q \sqrt{1 - \tilde{\rho}^2}} \right).$$

Because W and \tilde{W} are independent, the equality above reduces to

$$\text{cov}(U'(Z_T^z), Q_T^q) = \text{cov}(U'(e^{\mu z} W^{\delta z}), W^{\delta q \tilde{\rho}}) e^{\mu q} \mathbb{E} \tilde{W}^{\delta q} \sqrt{1 - \tilde{\rho}^2}.$$

Clearly $\text{sgn}(\rho_u) = \text{sgn}(\text{cov}(U'(Z_T^z), Q_T^q))$. Hence,

$$\text{sgn}(\rho_u) = \text{sgn}(\text{cov}(U'(e^{\mu z} W^{\delta z}), W^{\delta q \tilde{\rho}})).$$

Let $X \equiv e^{\mu z} W^{\delta z}$, then X is a log-normal random variable and $W^{\delta q \tilde{\rho}} = (e^{-\mu z} X)^{\tilde{\rho} \delta q / \delta z}$.

We have

$$\text{sgn}(\rho_u) = \text{sgn}(\text{cov}(U'(X), X^{\tilde{\rho} \delta q / \delta z})). \quad (2.8)$$

Let $F(x)$ be the cumulative distribution function of X . By Cuadras (2002),

$$\text{cov}(U'(X), X^{\tilde{\rho} \delta q / \delta z}) = \int_{\mathbb{R}^2} [F(x) - F^2(x)] dU'(x) d(x^{\tilde{\rho} \delta q / \delta z}).$$

Because $F(x) \in [0, 1]$, the integrand is nonnegative on the real line. Recall the utility function is concave increasing. Hence, $U'(X)$ is monotonically decreasing with X . Therefore,

$$\text{cov}(U'(X), X^{\tilde{\rho} \delta q / \delta z}) > 0 \quad \text{if } \tilde{\rho} < 0,$$

$$\text{cov}(U'(X), X^{\tilde{\rho} \delta q / \delta z}) < 0 \quad \text{if } \tilde{\rho} > 0.$$

Furthermore, when $\rho = 0$, Q_T and Z_T are independent and, thus, $\rho_u = 0$. Therefore,

$$\text{sgn}(\rho_u) = -\text{sgn}(\tilde{\rho}).$$

Next we relate the sign of $\tilde{\rho}$ to the sign of $\rho = \text{corr}(Q_T^q, Z_T^z)$. By definition,

$$\text{sgn}(\text{cov}(Q_T^q, Z_T^z)) = \text{sgn}(\text{cov}(e^{\delta_q Q^*}, e^{\delta_z Z^*})). \quad (2.9)$$

Furthermore, we know that for the two random variables Q^* and Z^* , their covariances can be computed by

$$\text{cov}(Q^*, Z^*) = \int_{R^2} [H(q^*, z^*) - F(q^*)G(z^*)] dq^* dz^*, \quad (2.10)$$

where $H(\cdot, \cdot)$ is the joint distribution, and $F(\cdot)$ and $G(\cdot)$ are marginal distributions of Q^* and Z^* .

For covariance of the functions of Q^* and Z^* , Cuadras (2002) shows that

$$\text{cov}(f(Q^*), g(Z^*)) = \int_{R^2} [H(q^*, z^*) - F(q^*)G(z^*)] df(q^*) dg(z^*). \quad (2.11)$$

Clearly the first-order derivatives of $e^{\delta_q Q^*}$ and $e^{\delta_z Z^*}$ are positive. Therefore, by (2.9–2.11), we have

$$\text{sgn}(\rho) = \text{sgn}(\tilde{\rho}) = -\text{sgn}(\rho_u).$$

□

Lemma II.4 and (2.7) indicate that, due to the risk pooling effect, only demand that is *correlated* with the contracting party's capacity usage is discounted in the present value. To distinguish this correlated demand variability from the absolute demand variability, we call the former the *capacity risk* of the outsourcing demand. Then Lemma II.4 states that, all else being equal, the present value of future outsourcing demand Q_T^q decreases in its *capacity risk*. Specifically we see that:

- When the outsourcing demand is positively (negatively) correlated with either contracting party's original capacity usage Z_T^z , the more variable it is, the lower

(higher) its present value. This agrees with the standard results from financial pricing models. Intuitively, a positively correlated Q_T^q magnifies the party's overall capacity usage fluctuations when added to the portfolio, and, thus, carries positive capacity risk. So it should have a low present value. Similarly, a negatively correlated Q_T^q offsets the original fluctuations (note that $Q_T^q \ll Z_T^z$ by construction), and, thus, carries negative capacity risk. So it should have a high present value.

- When the outsourcing demand is independent of either party's original capacity usage, we have

$$\text{CEQ}_z(Q_T^q) = \frac{\mathbb{E} U'(Z_T^z)}{V'(z)} \mathbb{E} Q_T^q = \text{CEQ}_z(\mathbb{E} Q_T^q). \quad (2.12)$$

That is, the present value of a random but independent Q_T^q is identical to that of a deterministic demand stable at $\mathbb{E} Q_T^q$, as both of them carry zero capacity risk. This demonstrates the perfect risk diversification effect. When the outsourcing demand is independent of the CM/OEM's original capacity usage, its uncertainties are completely idiosyncratic and, thus, can be diversified away by other demands using the same production resources.

2.4.1.2 Price of demand risks under the subcontract

We price the demand risk, or more explicitly, the capacity risk, over $[0, T]$ by comparing the present values of the following two outsourcing orders.

Order (i): Observed to be q at time 0. Observed to be Q_T^q at time 0_+ , immediately after the subcontract becomes effective at time 0. Remains at Q_T^q throughout $(0, T]$.

Order (ii): Observed to be q at time 0. Observed to be Q_T^q at time T .

Clearly orders (i) and (ii) differ only in the way the inter-temporal uncertainty is resolved. Thus, the difference in their present values is the price of the demand risk at time 0.

For order (i), the outsourcing demand first jumps from q to Q_T^q in an infinitesimal amount of time, then remains deterministic over $(0, T]$. During the demand jump from time 0 to 0_+ , the CM/OEM's original capacity usage remains constant at z by continuity. Thus, the outsourcing demand and the original capacity usage are independent at 0_+ . Over $(0, T]$, the outsourcing demand is fixed at the observed value Q_T^q and the capacity usage changes. Hence, the outsourcing demand is independent from the capacity usage at T as well. Therefore, order (i) is independent of the capacity usage on $[0, T]$. Let CEQ_z^c denote the present value of order (i). By Theorem II.2, we have

$$\text{CEQ}_z^c = \frac{\mathbb{E}[U'(Z_T^z)]}{V'(z)} \mathbb{E} Q_T^q = \text{CEQ}_z(Q_T^q). \quad (2.13)$$

Order (ii) corresponds to the actual outsourcing demand. By Theorem II.2, the present value of order (ii) is $\text{CEQ}_z(Q_T^q)$.

The difference between $\text{CEQ}_z(\mathbb{E} Q_T^q)$ and $\text{CEQ}_z(Q_T^q)$ is then the price of the demand risk at time 0. For simplicity we ignore the argument of CEQ in the rest of the paper. Formally we define the following:

Definition II.5. The price of the inter-temporal uncertainties carried by the outsourcing demand, Q_T^q , is

$$P_z = \text{CEQ}_z^c - \text{CEQ}_z, \quad (2.14)$$

where CEQ_z^c ($z \in \{s, m\}$) is defined in (2.13).

Recall that all else being equal, the more capacity risk the outsourcing demand carries, the lower its present value. Therefore,

- $P_z > 0$ if and only if Q_T^q carries positive capacity risk, i.e., is positively correlated with Z_T^s . In other words, adding Q_T^q to the production capacity deteriorates the

original risk exposure of either contracting party. For the OEM, this means that, if produced in-house, Q_T^q increases his original capacity usage variability and, thus, is better outsourced. P_m is the OEM's gain by replacing the stochastic in-house production with a risk-free delivery. Hence, P_m is the fair price for the risk from the OEM's perspective in the sense that, it is the *highest* amount the OEM is willing to pay for the benefits of the risk-free delivery. For the CM, $P_s > 0$ means that taking on the outsourcing business has an adverse impact on her original risk exposure, and P_s is the cost caused by such an increase in capacity usage variability. Consequently, P_s is the fair price for the risk from the CM's perspective in the sense that, it is the *lowest* compensation amount at which the CM is willing to enter the outsourcing business.

- $P_z < 0$ if and only if Q_T^q carries negative capacity risk. Put another way, incorporating Q_T^q to the portfolio mitigates the fluctuations in either contracting party's capacity usage. Specifically, producing Q_T^q in-house reduces the OEM's original capacity usage variability. As a result, the OEM will not outsource unless being offered *at least* $-P_m$ product units at time 0 for foregoing the risk-mitigation benefits of producing Q_T^q . On the other hand, the CM is willing to pay *at most* $-P_s$ to the OEM for the benefits of decreased capacity variability by adding Q_T^q to her production portfolio.
- $P_z = 0$ if and only if Q_T^q carries no capacity risk. Such an outsourcing demand is constant or is random but independent of Z_T^z . That is, it exerts neither positive nor negative influence on the original capacity usage variability of the OEM/CM.

Note that by Theorem II.2, when Q_T^q carries nonzero capacity risk, the CM and the OEM price the demand risks based on their original capacity usage, and perceive different cost P_s and benefit P_m of outsourcing. When Q_T^q carries no capacity risk,

however, both the OEM and the CM conclude $P_m = P_s = 0$, regardless of their original capacity usage, M_T^m and S_T^s .

While the risk price P_z is an intuitive measure for demand risks, its magnitude is not very informative for cost/benefit assessment when reported alone. This is because P_z 's magnitude depends on the scale of the outsourcing demand. Indeed, by (2.7) and (2.12), a large P_z may not be a concern at all if the outsourcing demand is constantly high. Similarly, a small P_z does not necessarily indicate low risk business, especially when demand is always low. In addition, when evaluating the cost/benefit of the outsourcing business in practice, it is often useful to compare the cost/benefit of demand risks with the outsourcing revenue. Hence, we define *risk ratio* to be the normalized price of demand risks as the following.

Definition II.6. The risk ratio, π_z ($z \in \{s, m\}$), of the outsourcing demand, Q_T^q , from either contracting party's perspective, is the relative risk price with respect to the expected outsourcing demand.

$$\pi_z \equiv \frac{P_z}{\mathbb{E} Q_T^q}. \quad (2.15)$$

In essence, the risk ratio measures the variability level of the outsourcing demand, and is thus similar to the commonly-used coefficient of variation, c_v . c_v is defined as $c_v \equiv \sigma/\mu$, where μ and σ are the mean and standard deviation of the random variable under consideration. In some sense, the risk ratio may be viewed as an extension of the coefficient of variation in a dynamic setting. However, we note that one important distinction between π_z and c_v is that the former accounts for the risk pooling effects.

The risk ratio, π_s , is a particularly relevant measure for the CM. It gives the cost of the demand risk relative to the expected outsourcing business. The higher the risk ratio, the more the CM's profit margin will be eroded if the OEM does not compensate for the demand risks she bears. Thus, the more the CM should worry

about the capacity risk.

For the OEM, the risk ratio, π_m , represents the benefit achieved via outsourcing relative to the scale of the outsourcing business. It measures the efficiency of outsourcing in terms of risk cost reduction in the sense that, given two potential outsourcing demands, the one with the higher risk ratio generates more benefits per unit outsourced. π_m can be used to select the part of the total market demand that will be outsourced when there exist factors that restrict the OEM from unlimited outsourcing. These factors include, for example, the OEM's concerns on the quality of the CM's products and political risks in a global supply chain.

2.4.2 Pricing demand risks over the entire outsourcing horizon

Having studied the demand risks transferred under a single subcontract, we are now ready to price the demand risks over the entire outsourcing horizon. Recall that the OEM places orders Q_i at fixed time points $\{T_i\}$ ($i = 1, 2, \dots, N$) with $T_0 = 0$. Then the JIT contract can be viewed as a collection of independent subcontracts. Subcontract i comes into effect at time T_{i-1} and expires at T_i , at which the outsourcing demand Q_i is observed and satisfied.

By Definitions II.5 and II.6, the price of demand risks over subcontract i is contingent on the outsourcing demand observed at the starting time T_{i-1} . To make explicit the dependence of Q_i (resp. Z_i) on its earlier realization Q_{i-1} (resp. Z_{i-1}), we write it as Q_i^{i-1} (resp. Z_i^{i-1}). Let CEQ_z^i be the present value of demand Q_i^{i-1} evaluated at T_{i-1} . Similarly, let P_z^i and π_z^i denote the price and risk ratio computed given the

demand up to time T_{i-1} . Then we have

$$\text{CEQ}_z^i = \frac{\mathbb{E}_{i-1} [U'(Z_i^{i-1})Q_i^{i-1}]}{V'(Z_{i-1})}, \quad (2.16)$$

$$\text{CEQ}_z^{ci} = \frac{\mathbb{E}_{i-1} [U'(Z_i^{i-1})]}{V'(Z_{i-1})} \mathbb{E}_{i-1} Q_i^{i-1}, \quad (2.17)$$

$$P_z^i = \text{CEQ}_z^{ci} - \text{CEQ}_z^i, \quad (2.18)$$

$$\pi_z^i = \frac{P_z^i}{\mathbb{E}_{i-1} Q_i^{i-1}}, \quad (2.19)$$

where $\mathbb{E}_{i-1}[\cdot]$ is the expectation conditional on the information up to time T_{i-1} .

At time 0, the risk ratios of all future subcontracts ($i > 0$) are random and contingent on future demands. An intuitive measure of the demand risks transferred over the entire service horizon is the average of the expected risk ratios of this subcontract portfolio. More explicitly, we define the following:

Definition II.7. The normalized price of demand risks in JIT outsourcing, Π_z ($z \in \{s, m\}$), is

$$\Pi_z \equiv \frac{1}{N} \sum_{i=1}^N \mathbb{E}_0[\pi_z^i]. \quad (2.20)$$

2.5 Empirical study on an auto-parts CM

In this section we apply our pricing model developed in Section 2.4 to a Norwegian auto-part CM who supplies bumper beams to 11 OEM customers under JIT contracts. Due to the large order quantities and short time window between order placement and delivery, the CM is granted access to the OEMs' demand databases for better production planning. The OEMs' orders are released on regular basis most of the time, usually in weeks. Hence, we can model the JIT contract as a collection of independent subcontracts, each of which spans over one week.

In the following, we first estimate the model parameters from the data set. Then

we study the cost of the demand risk from the CM's perspective. This way we examine whether the CM incurs significant costs due to uncertain demand when accepting the JIT contracts.

2.5.1 Model calibration

Our data set consists of the CM's daily production volume and monthly earnings from January 2003 to August 2008 for each of the 11 OEMs. Different OEMs started their outsourcing at different dates, but none of them stopped the outsourcing business over the observation period. The newest customer came in 2006, while 6 OEMs established the relationship prior to Jan. 2003. These observations imply that outsourcing is a profitable practice for the CM and OEMs alike.

According to our collaborator, it keeps almost zero inventory. Therefore, we can use the CM's weekly production volume to approximate each OEM's orders. Let Q_t^i represent OEM i 's outsourcing demand in week t ($i = 1, 2, \dots, 11; t = 1, 2, \dots$)². We assume that the OEM places the order Q_t^i at the end of week t and the CM produces and delivers Q_t^i instantaneously. Use S_t^i to denote the CM's original capacity usage *without* OEM i 's business Q_t^i . Because total orders from all 11 OEMs constitute about 98% of the CM's overall bumper beam manufacturing business, we approximate S_t^i by the sum of all the other OEMs' orders Q_t^j ($j \neq i$).

We find that all $\log(Q_t^i)$ and $\log(S_t^i)$ can be adequately modeled as time series with a constant mean and drift plus integrated MA(1,1) white noise processes ($t = 1, 2, \dots; i = 1, 2, \dots, 11$):

$$\begin{aligned} \log(Q_t^i) &= \log(q_i) + \mu_q^i t + \epsilon_t^i, & \Delta \epsilon_t^i &= \epsilon_t^i + \theta_i \epsilon_{t-1}^i, \\ \log(S_t^i) &= \log(s_i) + \mu_s^i t + \xi_t^i, & \Delta \xi_t^i &= a_t^i + \phi_i a_{t-1}^i, \\ \rho_i &= \text{corr}(\epsilon_t^i, a_t^i), \end{aligned} \tag{2.21}$$

²As we have multiple OEMs in the empirical study, our notation is slightly different from those in Section 2.4.

where q_i and s_i are the initial values of Q_t^i and S_t^i , respectively. $\Delta X_t \equiv X_t - X_{t-1}$, $e_t^i \sim N(0, \sigma_i^2)$, $a_t^i \sim N(0, \delta_i^2)$.

Table 2.1: Model parameters.

$(\log q_i, \mu_q^i)$, and $(\log s_i, \mu_s^i)$ are constant mean and trend coefficients for $\log Q_t^i$ and $\log S_t^i$, respectively. θ_i and ϕ_i are the MA coefficients of the integrated MA(1,1) white noise process driving $\log Q_t^i$ and $\log S_t^i$, respectively. σ_i and δ_i are the standard deviations of the normal innovations in the IMA(1,1) processes of $\log Q_t^i$ and $\log S_t^i$, respectively. ρ_i is the correlation coefficient between innovations for $\log Q_t^i$ and $\log S_t^i$.

OEM	Q_t^i				S_t^i				ρ^i
	$\log q_i$	μ_q^i	θ_i	σ_i	$\log s_i$	μ_s^i	ϕ_i	δ_i	
1	9.3	-0.004	-0.91	0.78	11.0	-0.02	-0.79	0.48	0.27
2	9.8	-0.007	-0.94	0.58	10.8	-0.0007	-0.85	0.56	0.48
3	9.1	0.03	-0.80	0.68	11.3	-0.01	-0.83	0.43	0.52
4	7.8	0.01	-0.81	1.00	11.2	-0.001	-0.88	0.34	0
5	8.0	0.01	-0.84	0.84	11.3	0.02	-0.55	0.31	0
6	9.7	0.007	-0.95	0.59	11.0	0.0004	-0.97	0.42	0.31
7	7.7	-0.22	-0.76	0.79	11.2	0.02	-0.56	0.33	0
8	8.5	-0.05	-0.70	0.96	11.4	-0.01	-0.4	0.22	0
9	7.8	0.06	-0.67	0.93	11.3	0.01	-0.70	0.28	0
10	8.7	-0.002	-0.996	0.69	11.2	-0.003	-0.86	0.38	0.17
11	8.6	-0.0006	-0.993	0.77	11.1	-0.006	-0.86	0.44	0

Table 2.1 shows parameter estimates for (2.21) from the data. We observe that the outsourcing quantities are highly volatile compared to CM's corresponding capacity usage processes. Furthermore, given the variances of innovations e_t^i and a_t^i ($i = 1, 2, \dots, 11$), the drift coefficients μ^i are not significantly different from zero. In fact, close-to-zero drift is expected considering that our time unit is in weeks. Otherwise, over the observation period of 296 weeks, Q_t^i and S_t^i processes would on average increase or decrease exponentially, which is not observed. Lastly, when estimating ρ_i , we set the estimates that are not significantly different from zero on a 5% confidence level to 0. The ρ_i values in Table 2.1 indicate that outsourcing demand processes have nonnegative levels of association with corresponding capacity usages. This is consistent with the fact that all OEMs are in the same industry and facing similar

market demand.

We assume that the CM has the following earnings function

$$U(S_t) = \begin{cases} \log(S_t) & \text{if } \alpha = 1 \\ \frac{A}{1-\alpha}(S_t)^{1-\alpha} & \text{if } \alpha \in [0, 1), \end{cases} \quad (2.22)$$

Parameter α indicates the concavity of the CM's earnings. The higher α is, the lower tolerance the CM has for capacity variability as less earnings can be obtained. Hence, we can view α as the risk-aversion indicator, with $\alpha = 0$ corresponding to risk-neutrality. The point estimate for α is 0.5, with estimation error 0.2.

2.5.2 Demand risk of the outsourcing business

We now apply the models developed in Section 2.4 to price the demand risks carried by the 11 OEMs' outsourcing businesses. For ease of exposition, we ignore the subscript s in e.g., CEQ_s , since all the data is from the CM and, consequently, the risk prices are from the CM's perspective.

As in Section 2.4, we first find the present value of a random demand Q_T^q over a subcontract on $[0, T]$.

Lemma II.8. *If the CM's earnings function takes the form (2.22), and her original capacity usage S_t^s and the outsourcing demand Q_t^q ($t \in [0, T]$) have the following dynamics*

$$d \log Y_t = \mu(t, G_t)dt + \sigma(t, G_t)dW_t, \quad Y_0 = y, \quad (2.23)$$

where G_t is any random or deterministic process that is functionally independent of Y_t and adapted to W_t . W_t can be a multi-dimensional Wiener process. Then the present value of Q_T^q is

$$CEQ = s \frac{\mathbb{E} [(S_T^s)^{-\alpha} Q_T^q]}{\mathbb{E} [(S_T^s)^{1-\alpha}]}. \quad (2.24)$$

Proof. By (2.23), we have

$$S_t^s = s \exp \left\{ \int_0^t \mu(s, G_s) ds + \sigma(s, G_s) dW_s \right\}.$$

Then by Theorem II.2

$$V(s) = \begin{cases} \log s + \mathbb{E} \left[\int_0^t \mu(s, G_s) ds \right] & \text{if } \alpha = 1 \\ \frac{(As)^{1-\alpha}}{1-\alpha} \mathbb{E} [h(G_t, W_t, t)] & \text{if } \alpha \in [0, 1) \end{cases}$$

where $h(G_t, W_t, t) \equiv \exp \left\{ (1-\alpha) \left(\int_0^t \mu(s, G_s) ds + \sigma(s, G_s) dW_s \right) \right\}$. Hence the denominator of CEQ in Theorem II.2 becomes

$$\frac{dV(s)}{ds} = \begin{cases} \frac{1}{s} & \text{if } \alpha = 1 \\ s^{-\alpha} A^{1-\alpha} \mathbb{E} [h(G_t, W_t, t)] = \frac{1}{s} \mathbb{E} [(AS_T^s)^{1-\alpha}] & \text{if } \alpha \in [0, 1) \end{cases}$$

Hence

$$\begin{aligned} \text{CEQ} &= \begin{cases} s \mathbb{E} \left[\frac{Q_T^q}{S_T^s} \right] & \text{if } \alpha = 1 \\ s \frac{\mathbb{E} [(S_T^s)^{-\alpha} Q_T^q]}{\mathbb{E} [(S_T^s)^{1-\alpha}]} & \text{if } \alpha \in [0, 1) \end{cases} \\ &= s \frac{\mathbb{E} [(S_T^s)^{-\alpha} Q_T^q]}{\mathbb{E} [(S_T^s)^{1-\alpha}]}. \end{aligned}$$

□

Clearly Q_t^i and S_t^i ($i = 1, 2, \dots, 11$; $t = 1, 2, \dots$) in (2.21) satisfy the dynamics in (2.23). Let π_{t-1}^i denote the risk ratio of the outsourcing demand Q_t^i evaluated at the end of week $t-1$ ($t = 1, 2, \dots$; $i = 1, 2, \dots, 11$). By (2.21) and Lemma II.8, straightforward algebraic manipulation yields

$$\pi_{t-1}^i = e^{\frac{2\alpha-1}{2}\delta_i^2} e^{-\phi_i a_{t-1}^i - \mu_s^i} \left(1 - e^{-\alpha\beta_i \delta_i^2} \right), \quad (2.25)$$

where $\beta_i \equiv \rho_i \sigma_i / \delta_i$ is the sensitivity of the outsourcing demand's driving white noise process to the white noise process of the CM's original capacity usage. a_{t-1}^i is the white noise of S_t^i at the end of week $t-1$. Recall that a_{t-1}^i is sampled from zero-mean normal distribution with variance δ_i^2 . Therefore, the risk ratios of the OEM i 's orders are iid log-normal random variables over the contract horizon when evaluated at time 0.

Let Π_i be the normalized price of the demand risks carried by OEM i 's outsourcing business ($i = 1, 2, \dots, 11$). By Definition II.7,

$$\Pi_i = \mathbb{E}_0 \pi_t^i = e^{\eta_i} \left(1 - e^{-\alpha \beta_i \delta_i^2} \right), \quad (2.26)$$

where $\eta_i = -\mu_s^i + \frac{1}{2}(2\alpha - 1 + \phi_i^2)\delta_i^2$. We note that when $\phi_i = 0$, this corresponds to the simple case where S_t^i is approximated as a geometric Brownian motion (GBM). In other words, the auto-regressive parts of $\Delta \log S_t^i$'s white noise process increases Π_i from the GBM value by a factor of $\exp\{\frac{1}{2}\phi_i^2\delta_i^2\}$.

From (2.25) and (2.26), the signs of both π_{t-1}^i and Π_i ($t = 1, 2, \dots, i = 1, 2, \dots, 11$) are determined by the factor $\left(1 - e^{-\alpha \beta_i \delta_i^2} \right)$. Particularly, we observe that:

- As expected, the risk ratio is positive if and only if the outsourcing demand is positively correlated with the CM's original capacity usage ($\beta_i > 0$), i.e., carries positive capacity risk. In this case, the CM's original risk is increased by the outsourcing business and, thus, incurs a positive risk cost. The risk ratio is zero if and only if the outsourcing demand is not correlated with the CM's capacity ($\beta_i = 0$).
- The price of demand risks increases with the correlation ρ_i between the innovations, e_t^i and a_t^i , driving Q_t^i and S_t^i . This result is intuitive because given the absolute demand variability σ_i , we expect the capacity risk to increase with the correlation coefficient, and consequently the risk price to increase as well.

- All else being equal, the more risk-averse the CM, the higher the risk ratio of the outsourcing demand. The risk-aversion parameter, α , reflects how much the CM weighs the capacity risk. For a risk-neutral CM ($\alpha = 0$), (2.25) and (2.26) give $\pi_{i-1}^i = 0$ and $\Pi_i = 0$, i.e., no risk price is expected, regardless of the capacity risk of the outsourcing demand.

Table 2.2 shows the CM’s average risk ratios with 11 OEMs. Compare the risk ratios with the absolute variability, σ_i , of OEM i ’s business ($i = 1, 2, \dots$) in Table 2.1, we see that OEMs with high risk ratios are actually the ones with relatively low σ . This indicates that making outsourcing decisions based on the demand risks might be misleading at times.

Table 2.2: Risk ratios of outsourcing

Π_i is the expected risk ratio of OEM i ’s outsourcing demand from the CM’s perspective, $i = 1, 2, \dots, 11$.

OEM	1	2	3	4	5	6	7	8	9	10	11
Π (%)	5.4	8.4	7.9	0	0	4.1	0	0	0	2.3	0

Because the Norwegian CM produces many other products besides bumper beams, its overall gross profit margin is not an informative metric to assess the cost of demand risks. We use Delphi’s gross profit margin, 18.14%, to approximate that of our collaborator’s JIT outsourcing business. We test the risk ratios computed from our model by comparing them to 18.14%. We observe that the risk ratios from all 11 OEMs are moderate, indicating that the CM is not assuming significant demand risks without being compensated. This is consistent with our expectation given the CM’s long-standing healthy contractual relationship with the OEMs.

2.6 Concluding remarks

At the best of our knowledge, this paper is the first attempt to price demand risks in component outsourcing in JIT literature. Both the OEM and the CM have flexible production capacities and concave earnings on their capacity usage, and are profit maximizers. We show that when the outsourcing demand is positively correlated with the agent's existing business, the higher risk it carries, the more the outsourcing costs the CM and benefits the OEM. Furthermore, we introduce a scale-free quantity *risk ratio* to measure the "riskiness" of the outsourcing business, accounting for the demand pooling effect under flexible capacity. For the CM, the risk ratio gives the cost of the demand risk relative to the expected outsourcing business. The higher the risk ratio, the more the CM's profit margin will be eroded if the OEM's payment is primarily based on the realized demand. Finally, we apply our model to a data set of a Norwegian auto-parts CM serving 11 OEMs. We report that the CM's risk ratios of these outsourcing businesses are small. This agrees with our expectation given the CM's healthy business and prudence in selecting customers.

This study can be extended in many ways. One interesting direction to pursue is incorporating the outsourcing demand as OEM's decision variable depending on its in-house production capacity, inventory level, etc. This will lead to an optimization problem with the decision being the optimal capacity/inventory from the risk perspective. Another interesting direction is extending the model to scenarios where the OEM's order arrives randomly over the contract horizon. Intuition has that price of demand risks becomes higher as there are now two sources of randomness: the uncertainty in demand quantities and the uncertainty of the order placement.

CHAPTER III

Empirical Study on Managed Print Services

Pricing

3.1 Introduction

Increasingly, companies and institutions are outsourcing the management of their information technology (IT) infrastructure to reduce cost and improve operations, giving rise to the so-called *managed services* industry. Broadly defined, *managed services* is the management provided by a third-party contractor, that is, the managed services provider, over the institutional customer's infrastructural IT operations such as printing, network security, and data center monitoring. The trend of adopting managed services is sweeping all sectors of the economy from manufacturing, banking, to education and health care. In a recent survey on 400 US professionals (CompTIA 2011), 62% responded that their businesses plan to increase the spending on managed services over the next two years. Answering to this demand, numerous managed services providers emerge, and more and more traditional manufacturers expand their existing operations to this new arena (Jacob and Ulaga 2008, Oliva and Kallenberg 2003, Salonen 2011, Cohen et al. 2006). One of the biggest challenges confronting the practitioners is how to price these services. Indeed, according to a recent research report (Walsh 2011), managed services providers are watching their profit margins

“erode precipitously” due to “poor pricing practices” and intense competition.

Pricing in managed services is complicated by the unique features that distinguish it from services to individual consumers. First, managed services are infrastructural in nature. That is, the operations covered under managed services are mainly supportive and might be viewed as utilities by the customer. Therefore, the customer’s service demand is less sensitive to the service prices than service demand from individual consumers. Second, for a majority types of managed services, service contracts are customized rather than being a standard menu as in individual consumer services. Contract customization arises because the managed services customers are enterprises with different IT environment and service level requirements. Indeed, the fact that the contracts are customized and flexible to meet different needs is emphasized by managed services providers.

In this paper we study in-depth a particular type of managed services—the managed print services (MPS). The goal of our study is to understand the contractual interactions between the MPS provider and the institutional customers in practice. This understanding, in turn, provides insights into pricing of general managed services.

MPS provides comprehensive management of the institutional customer’s printer fleet and covers supplies and maintenance for a monthly fee. Most MPS customers (e.g., universities, banks, manufacturers) are not in the professional printing industry and, thus, MPS is an infrastructural service unrelated to their core operations. Just like other managed services, MPS has been growing explosively over the recent years. Photizo (2011) estimates MPS global market to be \$30 billion in 2011 and forecasts a double-digit annual growth rate for the following 5 years. There are a number of providers competing in the MPS market, including Xerox, HP, Ricoh and Lexmark.

We consider an MPS provider (he) serving a number of institutional customers. The provider faces the fundamental problem of managing revenue and costs from his

entire MPS portfolio (from multiple customers), when bidding for the business of a new customer (she) among other service providers. Because a customer usually owns hundreds or even thousands of printers with heterogeneous models, work loads and costs, each provider negotiates service contracts with the customer. In this case, the provider needs to design service contracts that achieve two goals: to win the bid, and to optimize his own objective function. The MPS contracts under consideration are a more complex variant of the three-part contracts commonly studied in marketing and operations management literature (Wilson 1997, Iyengar et al. 2008, Sundararajan 2004). A standard three-part contract consists of a fixed price, a variable price and a free-usage allowance. MPS contracts feature five parameters for each printer: the fixed price, the variable prices per black & white (BW) and color print, and the free-usage allowances for BW and color prints. In practice, the contract terms are decided before the service starts and remain unchanged over the service horizon or until a renegotiation takes place. Different printers may share the same contract terms. We observe in practice that while the fixed price and allowances could be zero if the customer so prefers, the variable prices are always positive.

To understand the contractual interactions between the provider and the institutional customers, we develop two models. First, we propose a model of how much a customer would pay for the service; or equivalently, the customer's service valuation (hereafter, the "customer's pricing model"). Predictions from this model yield the maximum amount the provider, whose data we use, can charge, and can be interpreted as the provider's winning bids that are acceptable to the customer. We discover the following. (i) The customer's service demand, that is, the print volume, is inelastic in the service prices over the observed range. We believe this result arises from the infrastructural nature of printing and anticipate it to apply to other types of managed services as well. Under inelastic demand, we do not observe moral hazard type behavior where lower prices result in higher print volumes (Laffont and Martimort 2001).

(ii) An arbitrary customer’s payment for the service of a particular printer is the sum of two parts: the payment of a “representative” customer, that is, the population service valuation, and a random amount attributed to unobserved customer-specific characteristics. This customer’s pricing model adequately characterizes (in-sample $R^2 = 92\%$) and reliably forecasts (out-of-sample $R^2 = 65\%$) the observed customers’ payments, yielding the set of provider’s bids that are acceptable to the customer. Specifically, the high explanatory power of the out-of-sample predictions indicates that the unobserved customer-specific characteristics (e.g., negotiation peculiarities, geographic locations) have weak influences on the customer’s payment. That is, the population valuation dominates an individual institutional customer’s service valuation. (iii) Both the population service valuation and the customer-specific random deviation are affine in the device’s print volume. In other words, the population perceives a fixed value of owning the service, and variable values of using the service for each BW and color page printed. A particular customer deviates from the population perceived fixed and variable values randomly. We observe that the customer’s perceived fixed value has no implication on her perceived variable values, and vice versa, while the customer’s perceived variable values per BW and color page are highly correlated: if her perceived value per color page is 10% higher than the population value, then there is a 75% probability that she attaches an above population value per BW page as well. (iv) The population perceived fixed value of MPS decreases with the customer’s fleet size. For example, for a laser Xerox printer with median work load¹, the population valuation of the service of a 200-printer fleet is 6% lower than that of a 50-printer fleet. This might indicate the existence of the customer’s market power: the larger the customer’s device fleet, the more powerful the customer is in negotiation, hence the less she is likely to pay.

Second, within the set of winning bids yielded by the customer’s pricing model, we

¹1,700 BW pages and 1,900 color pages per month (see Table 3.1).

turn attention to the supply side and postulate that risk-aversion plays an important role in the provider's decision on the optimal contract to offer. We formulate an optimization problem that captures the risk-aversion of the provider, considers portfolio effects from the existing contracts, and uses the customer's pricing model as an input. Under a mean-variance objective function, we show that the optimal contracts do not depend on the exact value of the provider's risk-aversion parameter, as long as he is risk-averse. That is, the optimal contracts are the ones that minimize the provider's total profit variability and charge the predicted customer's service valuation. In the empirical analysis, we confirm that our predictions match the observed contracts significantly better than the benchmark linear regression model. Furthermore, we observe that the risk-adjusted earnings (i.e., monthly average earnings normalized by the earnings variability) do not depend on the customer's industry or fleet size. Put another way, all customers bring in similar risk-adjusted monthly earnings for the provider. This might indicate intense competition among the providers as all customers' earnings risks are priced in a same way.

One important finding of our empirical analysis is that the MPS provider is better modeled to be risk-averse than risk-neutral. Our risk-aversion assumption is motivated by the fact that, for the providers, the MPS earnings constitute a significant part of his overall profit. Indeed, in a recent survey on 435 managed services firms in 19 countries (Kaseya 2011), 82% of the firms report that they derive more than 25% of annual revenue from managed services, with 23% deriving more than 75%. Therefore, besides the preference for higher profits, the provider is likely to favor stable service earnings due to financial frictions and constraints such as short-term funding costs and budget limitations. On the other hand, however, because the provider serves many clients with independent service demands, would he still care about the earnings uncertainties that could be diversified away? Therefore, at the portfolio level, the risk-aversion of the provider is not obvious and it is important to verify whether

earnings variability is a factor in the provider’s decision. Furthermore, many previous theoretical studies have assumed that the provider is risk-neutral. By empirically testing risk-aversion against risk-neutrality, we find that risk-neutral provider might not be a good assumption in practice.

The customer’s pricing model and the provider’s optimization model constitute our pricing framework for MPS providers. Their outputs can be used in practice when providers bid for the service of a new customer. We note that due to confidentiality issues, it is difficult to obtain data from multiple providers, and our empirical results are derived using the data set from Xerox. However, given the wide variety of customers in the MPS portfolio of Xerox, who is a leading provider in the US and Europe with a significant market share, we expect this modeling approach and the qualitative insights to apply to other MPS providers and possibly to other types of managed services as well.

The rest of the paper proceeds as follows. In Section 2 we review related literature. In Section 3 we describe our data set in detail. In Section 4 we discuss managed print services in more detail and present the model setup. In Section 5 we study the price elasticity of the printing demand. In Section 6 we formulate the customer’s pricing model and present corresponding empirical analysis. Section 7 develops and solves the provider’s optimization model, and discusses the empirical results. Section 8 concludes.

3.2 Literature review

We develop a pricing model that yields reliable estimates of the customer’s service valuation. This objective shares commonalities with marketing and economics papers that study individual consumer utilities derived from product or service purchasing. One widely adopted approach in these studies is the discrete choice model (Berry 1994, McFadden and Train 2000, Chintagunta et al. 2003). In managed services,

however, the provider proposes customized contracts rather than standard menus and does *not* know the contract offerings made by his competitors. As a result, the provider has no information on the customer's choice set, unlike the individual consumer service providers such as grocery stores and phone service carriers. Therefore, it is particularly difficult for the provider to estimate the customer's service valuation using standard discrete choice models. In this paper, we propose an alternative model of the customer's service valuation and show that it has high explanatory power in case of managed print services.

MPS pricing falls into the area of pricing with multi-part tariff contracts. In this context, one standard assumption in the marketing literature is that the consumer's utility, and consequently her service usage, depends on service prices (Lambrecht et al. 2007, Iyengar et al. 2008, 2011, Goettler and Clay 2010). In the MPS case, however, we show that the institutional customers' service demand is inelastic in service prices using data set provided by Xerox. As a result, we model the customer's service usage as an exogenous stochastic process and show that the customer's service valuation is adequately captured with a model independent of the service prices.

A prevalent assumption in the literature on optimal pricing is that the service provider is risk-neutral. In the context of multi-part contract, examples include Png and Wang (2010) who study optimal two-part tariff contract for a profit maximizer with risk-averse buyers, Essegai et al. (2002) that compare different tariff contracts for a risk-neutral provider offering access services with capacity constraint, and Sundararajan (2004) on the optimality of offering both fixed payment and usage-based contracts on information goods. More broadly, Chu et al. (2008) focus on the optimal bundling pricing strategies of a multiproduct firm, and Huang and Sundararajan (2005) examine the form of optimal pricing plan for utility computing. In this paper, however, we find the optimal multi-part contracts for a *risk-averse* provider, and provide empirical evidence that the risk-aversion matters.

The existing literature on risk management in services is mainly on IT and computing services. Under a linear pricing scheme, Paleologo (2004) examines the profit-maximizing unit price for a utility computing service seller who imposes distributional constraint on his gross profit margin. Using similar constraints, Kauffman and Sougstad (2008) consider the service level that maximizes an IT service provider's expected profit under uncertain cost, assuming that the customer's willingness to pay can be completely extracted by the provider. These works are similar to ours in that they assume risk-averse providers and take service demand as external variables. However, they do not consider the interaction between the customer and service provider, while we compute the provider's optimal contracts based on the predicted customer's service valuation.

There is a strong interest in the contracts used in after-sales services, amidst the wave of servicization that swept many manufacturing industries (Cohen et al. 2006). Kim et al. (2007a), Kim et al. (2007b), and Kim et al. (2010) study performance-based contracting in military and aviation industries using a principal-agent framework where the customer finds the optimal contract structure to offer incentives to the provider for equipment reliability improvement. In a recent empirical study, Guajardo et al. (2012) compare the product reliability under performance-based contracts and time and material contracts in aircraft engine services. In MPS, however, there is only one contract structure adopted in practice—the multi-part tariff contract. This paper focuses on identifying the optimal contract parameters under the multi-part tariff structure. In particular, we show that there is no moral hazard as the customer's service demand is inelastic in service prices. This is different from the principal-agent framework and the empirical analysis in Guajardo et al. (2012) where the customer's choice between performance-based and cost and material contracts is endogenized.

3.3 Data

In this section, we describe the data provided by Xerox on its MPS to a number of institutional customers. We explain the pre-processing steps and characterize the resulting *pricing data set*, *cost data set*, and *multi-price data set*. We carry out empirical analysis of the customer’s pricing model and the provider’s optimization model using the *pricing data set* and *cost data set*, respectively. The *multi-price data set* is used to study the sensitivity of the printing demand to service prices.

Our raw data consists of three different files: one containing print volume on devices in our sample during September 2006–August 2012 (hereafter, the “volume file”); one containing each device’s contract in use during September 2007–August 2012 (hereafter, the “contract file”); and one on the service cost (hereafter, the “cost file”). The list of variables in the volume file includes: printer ID, meter reads of cumulative print volume, meter type (BW print or color print), and date of the record. The contract file provides more information on printers and customers. The list of variables is: printer ID, printer model, customer name, manufacturer of the printer, contract in use, and date of the record. The cost file includes supplies cost on a subset of devices. The list of variables is: printer ID, printer model, name of the consumables (e.g., toner, catridges, fuser), cost of the consumables, and number of pages printed while the consumables last.

We carry out the following steps to clean the data. For the volume file, we compute the monthly BW and color print volume of each device and remove inconsistent volume entries (e.g., negative monthly BW/color volume due to meter resets). For the contract file, we remove abnormal entries with zero contract prices or with multiple different contract records on the same date. For the cost file, we compute each printer’s consumables cost per BW and color print, respectively.

3.3.1 Pricing data set

We take the latest contract of each printer in the contract file and find the monthly volume records under these contracts in the volume file. When merging them together, we discard printers with fewer than 10 monthly BW or color volume records to ensure adequate sample size for estimation purposes. The resulting *pricing data set* consists of 3,075 printers from 26 institutional customers (private and public companies and institutions). Each printer has a two-dimensional print volume time series (BW, color). Unit root tests (augmented Dickey-Fuller test and Phillips-Perron test) reject the null and accept the stationarity alternative on the BW and color volume series at 95% significance level. Furthermore, the zero autocorrelation assumption for BW and color volume series are not rejected in the data set at 95% significance level. Therefore, we use the volume time series to estimate the marginal distribution of each device’s monthly print volume. Specifically, we estimate the mean, standard deviation and correlation of the BW and color print volume for all 3,075 devices, and summarize the results in Table 3.1. Print volume is the only continuous variable at the printer level. At the customer level, we compute each customer’s total number of printers, that is, the fleet size, and the average aggregate monthly BW and color print volume. Table 3.1 summarizes the customer-level variables.

Besides the print volume, the pricing data set also contains information of each device’s contract, customer, printer model and manufacturer. Based on this, we define a number of dummy variables at the printer and customer levels. We first categorize the 26 customers into 5 industrial segments based on the North America Industry Classification System (NAICS): Manufacturing; Professional, Scientific, and Technical Services; Finance and Insurance; Information Technology; and Other Services. Next we categorize the 73 printer models observed in the data set into 3 classes based on their technology and functionality: laser printer; solid-ink printer; and laser multi-functional printer (MFP). This classification is motivated by the discussion with the

Table 3.1: Continuous variables and descriptive statistics

(μ_b, μ_c) , (σ_b, σ_c) , and ρ are the mean, standard deviation and correlation of the monthly BW and color print volume of each device. This table summarizes these marginal distribution statistics over all 3,075 printers. N is the number of printers each customer owns in the original contract file, before cleaning the data. $(\mu_b^{agg}, \mu_c^{agg})$ represent the average aggregate BW and color monthly volume of each customer.

Unit	No. of obs.	Variable	Mean	Median	Std. dev.	Min	Max
Printers	3,075	μ_b	3,324	1,701	4,231	0.2	41,600
		μ_c	3,252	1,872	4,324	0.9	59,540
		σ_b	1,242	683	1,657	0.6	15,569
		σ_c	1,354	799.8	1,784	2.2	22,454
		ρ	0.53	0.60	0.32	-0.92	0.99
		N	393.3	284.3	294.0	19	839
Customers	26	μ_b^{agg}	393,100	235,800	427,977	2,053	1,858,000
		μ_c^{agg}	384,600	238,900	449,247	5,811	1,688,000

MPS provider. In practice, different printer models within the same class incur very similar service costs for the provider and function approximately in the same way for the customers. Therefore, it is more informative to consider 3 rather than 73 classes of printers in the analysis. Similar reasoning applies to classifying the customers into industrial segments. Lastly, we categorize the contracts into 3 types: a pay-per-use contract with zero fixed price and allowances; a two-part tariff with zero allowances; and a three-part tariff with both positive prices and allowances. Table 3.2 summarizes the dummy variables.

3.3.2 Cost data set

After merging the volume, contract and cost files, we obtain the *cost data set* consisting of 1,035 printers owned by 8 customers. Printers in the cost data set are a subset of those in the pricing data set. We do not provide descriptive statistics for the service costs for confidentiality reasons.

Table 3.2: Definition of dummy variables and descriptive statistics

Shorthand for industrial segments: “Prof. Serv.”: Professional, Scientific, and Technical Services; “Fin. & Ins.”: Finance and Insurance; “IT”: Information Technology; “Others”: Other Services. Shorthand for printer class: “MFP”: multi-functional printer.

Dummy variable	Levels	No. of printers	No. of customers
Industrial segment	Manufacturing	1,941	12
	Prof. Serv.	296	5
	Fin. & Ins.	687	5
	IT	40	2
	Others	111	2
Printer class	Laser	343	-
	Solid-ink	907	-
	Laser MFP	1,825	-
Contract type	Pay-per-use	1,682	-
	Two-part tariff	1,275	-
	Three-part tariff	118	-
Printer manufacturer	Xerox	3,045	-
	non-Xerox	30	-

3.3.3 Multi-price data set

In the contract file, we observe that contracts change only in the variable prices per BW and color print. The fixed price and allowances remain the same throughout the observation horizon (September 2007–August 2012). We are interested in the contract’s influence on the print volume. Therefore, we consider only contracts with multiple variable prices in the cleaned contract file.

Out of a total of 3,435 printers in the cleaned contract file, 650 devices from 8 customers have at least two different BW/color prices. After merging with the volume file, we obtain the *multi-price data set* that consists of the monthly print volume and corresponding contracts of 326 printers owned by 7 customers. Because the BW and color variable prices do not always change simultaneously, the multi-price data set includes 279 BW volume series and 292 color volume series with two or three different BW and color variable prices, respectively. Specifically, 312 series are under

two different variable prices while 259 are under three variable prices. We partition each volume series into two or three segments indexed by the corresponding variable prices. Then within each segment, the BW or color variable price remains constant. In time order, we call the segments of a particular printer’s BW or color volume series the *early segment*, the *intermediate segment* (if the total number of segments is three), and the *last segment*. Table 3.3 summarizes the total and segment lengths of these BW and color volume series.

Table 3.3: Descriptive statistics for BW and color volume series under multiple prices

Time series on each printer’s monthly volume is partitioned into 2 or 3 segments based on the corresponding BW or color variable price. Within each segment, the corresponding variable price remains constant. In time order, these segments are the *early segment*, the *intermediate segment* (if the total number of segment is three), and the *last segment*, respectively.

Variable	No. of obs.	Mean	Median	Std. dev.	Min	Max
Length of volume series	571	35.43	37	10.50	8	55
Length of early segment	571	5.9	5	4.0	1	20
Length of inter. segment	259	13.0	12	5.6	3	20
Length of last segment	571	23.6	24	11.5	1	50

3.4 Model setup: contracts and random processes

In the following, we first discuss the MPS contracts in detail. Then we model the print volume and service cost of each device.

Consider a service provider managing printers for M institutional customers. Customer i owns N_i printers ($i = 1, 2, \dots, M$), each covered by one contract. We use the customer-printer index (i, j) to denote printer j in customer i ’s fleet ($j = 1, 2, \dots, N_i$; $i = 1, 2, \dots, M$). All printers can print in both BW and color. Subscripts “ ijk ” represent the BW ($k = b$) and color ($k = c$) aspect of printer (i, j) . We sometimes suppress these indices for the ease of exposition.

Let $[0, T_{ij}]$ be the observation horizon of printer (i, j) . The contract in month t ($t = 1, 2, \dots, T_{ij}$) is $x_{ij}(t) \equiv (F_{ij}, C_{ij}(t)^T, A_{ij}^T)$, where F_{ij} is the fixed monthly

payment, $A_{ij}^T \equiv (A_{ijb}, A_{ijc})$ is the free-usage allowances for BW and color prints, and $C_{ij}(t)^T \equiv (C_{ijb}(t), C_{ijc}(t))$ is the variable price per BW and color print once the allowances are exceeded at month t . The fixed price and allowances are static because in our data (see Section 3.3.3), only variable prices change over time. Naturally, we require $x_{ij}(t) \geq 0$ and write this as $x_{ij}(t) \in \mathbb{R}_+^5$, where \mathbb{R}_+ is the nonnegative real line. These contracts are more complicated than the commonly seen three-part pricing contracts, because they cover both the BW and color printing. Based on the number of nonzero parts in such a contract, we define the following contract types (we ignore the time argument for simplicity): (i) Pay-per-use: $x_{ij} = (0, C_{ijb}, C_{ijc}, 0, 0)$; (ii) Two-part tariff: $x_{ij} = (F_{ij}, C_{ijb}, C_{ijc}, 0, 0)$; (iii) Three-part tariff: $x_{ij} = (F_{ij}, C_{ijb}, C_{ijc}, A_{ijb}, A_{ijc})$.

Now we characterize each device’s monthly print volume. While service payments reportedly affect service demands of individual consumers (Png and Wang 2010, Lambrecht et al. 2007, Goettler and Clay 2010), this need not be the case for institutional customers whose core businesses are unrelated to printing. Intuitively, printing acts like a utility for these companies. Furthermore, the workload of a particular printer depends on the employees directly operating it, who are not price-sensitive because they are not being charged for the pages printed. Finally, according to the MPS provider, the intense competition among MPS providers makes it impossible to charge high prices in practice. Therefore, the customer’s internal printing demand might be inelastic in service prices, at least within a reasonable price range. We confirm this intuition by analyzing the sensitivity of the printing demand to service prices using the multi-price data set in Section 3.5. Consequently, we model printer (i, j) ’s monthly print volume as an exogenous stochastic process $D_{ij}(t)$, where $D_{ij}(t) \equiv (D_{ijb}(t), D_{ijc}(t))^T$ and $t = 1, 2, \dots, T_{ij}$.

We assume that the volume time series $D_{ij}(t)$ of a particular printer is stationary with zero autocorrelation—hypothesis that is not rejected at 95% significance level by the data. Under this assumption, a particular volume series has identical marginal

distributions, that is, $D_{ij}(t) \sim f_{D_{ij}}$ for all $t = 1, 2, \dots, T_{ij}$, where T_{ij} is the observation horizon.

Next, we assume that the print volume of different printers have zero covariance, that is, $\text{corr}(D_{ijk}, D_{i'j'k'}) = 0$ for $(i, j) \neq (i', j')$ and for all $k, k' \in \{b, c\}$. Formally, for each printer, we define the mean and variance of D as the following

$$\mu_{ij} = \mathbb{E} D_{ij}, \quad \Sigma_{ij} \equiv \text{var}(D_{ij}) = \begin{pmatrix} \sigma_{ijb}^2 & \rho_{ij}\sigma_{ijb}\sigma_{ijc} \\ \rho_{ij}\sigma_{ijb}\sigma_{ijc} & \sigma_{ijc}^2 \end{pmatrix}, \quad (3.1)$$

where $\sigma_{ijk}^2 = \text{var}(D_{ijk})$ ($k \in \{b, c\}$), ρ_{ij} is the Pearson's correlation between the BW and color monthly prints of device (i, j) .

The customer's payment for printer (i, j) in month t is

$$P_{ij}(t; x_{ij}) = F_{ij} + C_{ij}^T [D_{ij}(t) - A_{ij}]^+, \quad (3.2)$$

where $z^+ \equiv \max\{z, 0\}$. From properties of $D_{ij}(t)$, the printer's monthly payment, $P_{ij}(t; x_{ij})$, has identical marginal distribution $P_{ij}(t; x_{ij}) \sim f_{P_{ij}(x)}$ for all t and no autocorrelation.

Now we consider the MPS provider's service cost. In MPS, the total cost of serving a customer mainly consists of two parts: labor cost, and supplies and parts cost. For customers with many Xerox printers, like the majority of Xerox MPS businesses, the total labor cost is essentially the salaries of technicians who conduct on-site repairs and preventative maintenance. Therefore, it is fixed with respect to the customer's printing demand. The total supplies cost (e.g., toner) is roughly linear in the total print volume. The total parts cost (e.g., color tubes and cartridges) depends on each *individual* printer's volume and printer-specific factors like the average print job length and usage frequency. Therefore, the total supplies and parts cost of serving a customer is the sum of the supplies and parts cost of each printer. The overall cost of

serving multiple customers, consequently, is the sum of labor and supplies and parts costs on each individual customer.

Consider the supplies and parts cost of printer (i, j) . Let $S_{ij}(t) \in \mathbb{R}_+$ be the random supplies and parts cost in month t incurred by the provider $(t = 1, 2, \dots, T_{ij})$. We assume that, similar to the print volume, $S_{ij}(t)$ has identical marginal distribution $S_{ij}(t) \sim f_{S_{ij}}$ for all $t = 1, 2, \dots, T_{ij}$ with zero autocorrelation. Let $\Gamma_{ij} \in \mathbb{R}^2$ represent the covariance between the monthly cost and demand. Let $\gamma_{ij} \in \mathbb{R}^2$ denote the sensitivity of the cost to demand. We have

$$\Gamma_{ij} \equiv \begin{pmatrix} \text{cov}(S_{ij}, D_{ijb}) \\ \text{cov}(S_{ij}, D_{ijc}) \end{pmatrix}, \quad \gamma_{ij} \equiv \begin{pmatrix} \Gamma_{ijb}/\sigma_{ijb}^2 \\ \Gamma_{ijc}/\sigma_{ijc}^2 \end{pmatrix}. \quad (3.3)$$

3.5 Price elasticity of printing demand

In this section we examine the price elasticity of a customer’s printing demand. Because we observe no change in the fixed price and allowances of any printer over the observation horizon, the problem reduces to examining the dependence of each printer’s volume series on the variable prices. We choose this approach over the cross-sectional analysis across different printers because it allows us to control better for unobserved factors at the printer level, thus obtaining more reliable results. We study the BW and color printing demand separately. We consider the influence of BW (color) variable price on the corresponding BW (color) volume, because this is the first-order effect compared to the color (BW) variable price from the same printer².

We carry out the analysis using the multi-price data set, which consists of 279 BW volume series and 292 color volume series under two or three different variable prices. For the ease of exposition, we focus on BW printing demand in the following.

For a particular printer, we can view its BW variable price as a “treatment” that

²Indeed, we found that the color (BW) variable price to be insignificant for BW (color) volume when added as explanatory variables to our model.

takes continuous values. Within the period that a particular “treatment” is in effect, we record the responses, that is, the monthly BW print volume. Our task is to test if the price “treatment” affects the response. Because different printers change their BW variable price at different months, this is a longitudinal study with unbalanced design. In addition, we face a multilevel clustered data set with three levels: customer — printer — monthly volume records. Therefore, it is natural to use linear mixed-effects models (e.g., Fitzmaurice et al. 2004) which allow randomization at each level.

Formally, we model the BW volume of printer (i, j) at month t as follows:

$$\log D_{ijb}(t) = (\theta_0^b + w_{0ij}^b) + (\theta_1^b + w_{1i}^b) \log C_{ijb}(t) + v_{ij}^b(t), \quad (3.4)$$

where C_{ijb} is the variable price per BW print of printer (i, j) , $t = 1, 2, \dots, T_{ij}^b$, $v_{ij}^b(t) \sim N(0, \delta_m^2)$ is the month-level printer-specific error, $w_{0ij}^b \sim N(0, \delta_p^2)$ is the printer-level random effect that remains unchanged over $[0, T_{ij}^b]$ given a particular printer (i, j) , $w_{1i}^b \sim N(0, \delta_c^2)$ is the customer-level random effect that remains unchanged for all printers (i, j) ($j = 1, 2, \dots, N_i$) owned by customer i . Random variables $v_{ij}^b(t)$, w_{0ij}^b , and w_{1i}^b are mutually independent.

By definition, the price elasticity of the demand is $(dD/D)/(dC/C)$. Therefore, θ_1^b gives the demand elasticity of an arbitrary printer (see Hughes et al. 2006). A well-known problem when estimating demand equations of the form (3.4) is the endogeneity of the price, which leads to biased and inconsistent estimates of the price elasticity of the demand θ_1^b . Instrumentation, or equivalently, two stage least squares, is the common approach to account for the endogeneity (see Wooldridge 2010). The ideal instrumental variable is the one that has high correlation with the variable price but is independent of the external shocks affecting the print volume. In the case of MPS, it is particularly challenging to find a good instrumental variable at the printer level due to the lack of data. We use the public revenue data reported by

each customer as the instrumental variable because on the one hand, revenue affects a company’s budget for infrastructural costs and thus the price they would pay for MPS; on the other hand, revenue is most likely unrelated to internal printing. Similar choice of instrument can be found in for example Schultz and Tansel (1997), Keng and Huffman (2010) and Stone et al. (2007). We report the results under both the basic model (3.4) and instrumentation.

The printer- and customer-level random effects w_{0ij} and w_{1i} are a variant of the consumer-specific random effects commonly assumed in microeconometrics and marketing literature (McFadden and Train 2000, Chintagunta et al. 2003, Iyengar et al. 2008). In these works, the consumer-specific random effects are used as a parsimonious way of accounting for consumer heterogeneity, compared with introducing customer-level dummy variables. The independence among the random effects from different consumers indicates the conventional assumption of mutually independent consumers. In our model (3.4), the printer- and customer-level random effects have the similar interpretation. More explicitly:

1. By introducing the printer-level random effect w_{0ij}^b , we assume that for an arbitrary printer, its average monthly BW volume under zero BW price is a normal random variable oscillating around the population mean θ_0^b with variance δ_p^2 . This captures the fact that printers used by different companies and employees are likely to have different average work loads. Besides incorporating the printer-specific deviations in a parsimonious way, w_{0ij}^b also accounts for the printer-level clustering among the monthly errors $v_{ij}^b(t)$.
2. Intuitively, different customers’ BW printing demands have different price elasticities. By introducing the customer-level random effect w_{1i}^b , we assume that for an arbitrary customer, her price elasticity randomly varies around the population mean θ_1^b with variance δ_c^2 . Similar to w_{0ij}^b , this random effect accounts for the error clustering at the customer level.

3. *Conditional independence assumption:* We make two conditional independence assumptions at the printer level and customer level, which are explained in detail as follows.

- (a) w_{0ij}^b : We assume that, given customer i , the unobserved variables affecting her printer j 's BW volume ($j = 1, 2, \dots, N_i$) are exogenous and thus can be modeled as independent and identically distributed (i.i.d.) random shocks. This gives the printer-level independence within customer i 's fleet. This is also intuitive because in practice, within the same company, employees using different printers might be from distinct departments and for various reasons.
- (b) w_{1i}^b : We assume that the unobserved variables affecting the price elasticity of customer i 's volume ($i = 1, 2, \dots, M$) are exogenous and thus can be modeled as i.i.d. random shocks. This gives the customer-level independence.

Lastly, the random variable $v_{ij}^b(t)$ is the ordinary i.i.d. error term in regression. The i.i.d. assumption of $v_{ij}^b(t)$ states that the volume series of a particular printer has no serial correlation, a hypothesis that is not rejected at 95% confidence level in the preliminary analysis.

Under model (3.4), the question of whether the BW print volume is elastic in BW variable price reduces to testing if the population mean of the demand elasticity is zero. That is, if $\theta_1^b = 0$. Running the regression on the 279 BW volume series gives a p-value of 0.14 for θ_1^b under the basic model (3.4). Using the instrumental variable, the p-value for θ_1^b is 0.94. Therefore, our data implies that the BW printing demand is inelastic in the BW variable price.

We fit a similar model of color printing demand using the 292 color volume series. We report a p-value of 0.28 (0.18) for the population mean θ_1^c under the basic model

(instrumentation). Therefore, the multi-price data set indicates that the color printing demand is not elastic in the color variable price either. That is, printing demand is inelastic in the service price over the observed price range.

The result of inelastic demand is quite different from findings in Hughes et al. (2006), where the price elasticity of gasoline demand is significantly negative both under the basic model and with crude oil production disruptions as instruments. We would like to point out that this result should not be surprising given that printing acts as a utility for these companies, and that the intense competition among service providers significantly constrains the service prices offered to/accepted by the customer. We note that although this finding is based on the winning contracts of Xerox, given the significant market share of Xerox in MPS, we believe that this result is representative of the entire MPS industry. Furthermore, it might apply to other infrastructural services with intense market competition as well.

Remarks on the robustness and generality of results: We have run (3.4) by incorporating additional printer- and customer-level random effects on the slope and intercept. We have also added other printer- and customer-level regressors including customer's industry and fleet size, and assumed linear functional dependence of the print volume on variable price in model (3.4). We note that under the linear form of (3.4), $\theta_1 = 0$ indicates that both the mean and variability of a printer's monthly print volume do not depend on the variable price. We find p-values for θ_1^b to be between 0.08–0.94 and p-values for θ_1^c between 0.12–0.95 under the basic model and instrumentation. Due to space constraints, detailed models and results are not shown, but are available on request.

3.6 Customer's pricing model

The inelastic printing demand allows us to separate the supply and demand sides of the problem. In particular, the service demand, that is, the customer's print

volume, now becomes an exogenous stochastic process within a reasonable range of service prices. This range is defined by the observed contracts in the data set, or, equivalently, the contracts accepted by Xerox customers. Consequently, when constrained within the set of acceptable contracts, the customer's print volume is inelastic in service prices and the specific service prices are completely determined by the supply side, that is, the service provider.

In this section, we characterize the set of contracts acceptable to the customer, or the winning contracts in bidding from the MPS provider's perspective. Particularly, we assume that the customer is risk-neutral, because infrastructural services are usually small parts of their cash flows and are unrelated to their core businesses. An acceptable contract then satisfies the condition that the present value of payment from the customer does not exceed the maximum amount she would pay. We call this amount the customer's service valuation. We note that the customer's service valuation could depend on provider characteristics, for example reputation, sales force, relationship with the customer. In this paper, we focus on different customers' valuation of the service offered by the *same* provider. While constraining to one MPS provider controls for the unobserved provider effects in the empirical analysis, it might also limit the generality of the results, because similar data from multiple MPS providers is unavailable due to confidentiality issues. However, given the wide variety of customers in the MPS portfolio of Xerox, who is a leading provider in the US and Europe with a significant market share, we expect our modeling approach and the qualitative insights to apply to other MPS providers and possibly to other types of managed services.

In the following, we first model the customer's service valuation using the print volume, customer and printer characteristics, and the customer's preferences between zero- or positive-allowance contracts. This is our customer's pricing model that yields the maximum amount the customer would pay. The predicted payments can be

achieved by combining the five parameters of the MPS contracts in a variety of ways. Thus, the customer’s pricing model yields a set of acceptable contracts to the customer, or winning bids for the provider, over which the demand inelasticity holds. Applying the customer’s pricing model to the pricing data set (see Section 3.3.1), we find that it adequately characterizes and reliably predicts the observed payments. We then derive insights in the customer’s general perception of the value of MPS provided by Xerox.

3.6.1 Model formulation

Customer i ($i = 1, 2, \dots, M$) pays $P_{ij}(t; x_{ij})$ for the service of her printer j ($j = 1, 2, \dots, N_i$) at month t ($t = 1, 2, \dots, T_{ij}$) under contract x_{ij} . Let $V_{ij}(t)$ represent customer i ’s corresponding valuation of the service offered by *a particular provider*. $V_{ij}(t)$ is the maximum amount that customer i would pay for the provider’s MPS to printer j .

Rather than modeling customer’s decisions explicitly, which is difficult, we capture customer’s observed actions in a simple statistical model. We assume that $V_{ij}(t)$ can be described by the following customer’s pricing model:

$$V_{ij}(t) = (\alpha_0 + U_i^T \kappa + G_i^T \phi_0 + W_{ij}^T \psi_0 + u_{0i}) + (\alpha_1 + \phi_1 G_i + \psi_1 W_{ij} + u_{1i})^T D_{ij}(t), \quad (3.5)$$

where U_i and G_i are customer i ’s characteristics while W_{ij} contains printer (i, j) ’s characteristics, $u_{1i} \equiv (u_{1ib}, u_{1ic})^T$, $(u_{0i}, u_{1i}^T) \stackrel{\text{iid}}{\sim} N(0, \Lambda)$, α , κ , ϕ , ψ and Λ are model parameters with appropriate dimensions ($i = 1, 2, \dots, 26$). In the following, we first provide the intuition behind (3.5), then discuss the variables and model parameters in detail.

The pricing model (3.5) assumes that the customer’s valuation of the MPS to a particular printer is an affine function on the printer’s volume, with the intercept

and slope varying with the customer and printer characteristics. Specifically, the intercept, $(\alpha_0 + U_i^T \kappa + G_i^T \phi_0 + W_{ij}^T \psi_0 + u_{0i})$, can be interpreted as the monthly fixed value of owning the service perceived by customer i . The fixed value of MPS can arise from the benefits obtained by streamlining the customer’s printing fleet, freeing up resources that are occupied when there are multiple service vendors, etc. The 2-dimensional slope $(\alpha_1 + \phi_1 G_i + \psi_1 W_{ij} + u_{1i})$ denotes the variable value of using the service for each BW and color print perceived by customer i . This linear dependence assumes that each print produced by printers with the same characteristics owned by the same customer has the same value. We make this assumption based on the MPS’s infrastructural nature: MPS derives its value from satisfying the customer’s internal document needs, rather than the very behavior of printing itself. Consequently, similar internal document has approximately the same value for the customer. To test for model misspecification and robustness, we also extend (3.5) to include higher order terms of the print volume. The result, however, is not as good as (3.5).

U_i , G_i and W_{ij} are public information on the customer and printer. U_i is a 3-by-1 vector consisting of customer i ’s fleet size, and the total BW and color monthly printing volume (see Tables 3.1 and 3.2). G_i is a 5-by-1 dummy vector indicating the industrial segment customer i is in. That is, whether customer i is in manufacturing, professional services, finance and insurance, IT or other services industries. W_{ij} is a 8-by-1 dummy vector consisting information on printer (i, j) ’s manufacturer, technology and functionality, and contract type (see Table 3.2). We include the contract type into our pricing model because in practice, customers usually have preferences among pay-per-use, two-part tariff and three-part tariff contracts (recall the definition of contract types in Section 3.4). Incorporating it into model (3.5) allows us to gain insights on how customers with different preferences price the same “product”, that is, the service.

u_{0i} and u_{1i} are customer-level random variables that account for customer het-

erogeneity (Goettler and Clay 2010, Chintagunta et al. 2003, Iyengar et al. 2008) on the perception of the service. u_{0i} and u_{1i} are time-invariant and orthogonal to the randomness in customer i 's printing demand. Similar to models in Section 3.5, we make the conditional independence assumption among printers owned by the same customer and account for clustering at the customer level by introducing u_{0i} and u_{1i} . We also examined models with industry-level random effects. However, we find that the variance of the industry-level effect is not significantly different from zero (see Section 3.6.2.2 for details).

We note that the customer herself need not think of her value process in terms of (3.5). In fact, she is more likely to evaluate the service of her entire device fleet, rather than to each individual printer. However, pricing model (3.5) leads to a statistical model with printers as observation units, whereas a similar model of the overall service lead to a model with customers as observation units. In practice, the number of printers is several magnitudes greater than the number of institutional customers. Consequently, the pricing model (3.5) on printers provides more reliable estimates on the model parameters and, thus, more reliable predictions for the perceived value from a new customer. In Section 3.6.2, we estimate the model parameters of (3.5) based on the information on the provider's existing customers and their accepted contracts, and show how the resulting model has high explanatory power.

3.6.2 Empirical analysis

In the following, we present our empirical study on the customer's pricing model. We first establish a statistical model from (3.5). After that we draw managerial insights in the customer's service valuation using the estimated model parameters and report model fitting results.

3.6.2.1 Linear mixed-effects model

We first relate the pricing model (3.5) to the realized monthly payments from customers to the MPS provider. This yields a linear mixed-effects model with the realized average monthly payments as response variables, average monthly print volumes, customer and printer's characteristics as explanatory variables.

Let $\hat{x}_{ij} \equiv (\hat{F}_{ij}, \hat{C}_{ij}^T, \hat{A}_{ij}^T)$ be the observed contract of printer (i, j) . By (3.2), the realized payment of printer (i, j) in month t , $P_{ij}(t; \hat{x}_{ij})$, is given by $P_{ij}(t; \hat{x}_{ij}) = \hat{F}_{ij} + \hat{C}_{ij}^T (D_{ij}(t) - \hat{A}_{ij}^T)^+$. Following the property of $D_{ij}(t)$, $P_{ij}(t; \hat{x}_{ij})$ has identical marginal distribution for all t . The average monthly payment over the service horizon \bar{P}_{ij} is

$$\bar{P}_{ij} = \hat{F}_{ij} + \frac{1}{T_{ij}} \sum_{t=1}^{T_{ij}} \hat{C}_{ij}^T \left[(D_{ij}(t) - \hat{A}_{ij}^T)^+ \right]. \quad (3.6)$$

Let x_{ij}^* be the provider's theoretically optimal contract of printer (i, j) . Intuitively, the theoretically optimal expected monthly payment, $\mathbb{E} P_{ij}(x_{ij}^*)$, is the maximum amount the risk-neutral customer would pay (see Section 7.1.1). By the customer's pricing model (3.5), this maximum amount is

$$\mathbb{E}_D V_{ij} = (\alpha_0 + U_i^T \kappa + G_i^T \phi_0 + W_{ij}^T \psi_0 + u_{0i}) + (\alpha_1 + \phi_1 G_i + \psi_1 W_{ij} + u_{1i})^T \mu_{ij}, \quad (3.7)$$

where $\mu_{ij} \equiv \mathbb{E} D_{ij}$, as defined in Section 3.4. We use $\mathbb{E}_D[\cdot]$ to make explicit that the expectation is over the print volume, D .

We assume that the realized average monthly payment of printer (i, j) , \bar{P}_{ij} , is shifted from the theoretically optimal value, $\mathbb{E} P_{ij}(x_{ij}^*)$, by a random amount due to the unmodeled exogenous shocks that affect the finalized contracts on printer (i, j) . Let ϵ_{ij} represent the random error on printer (i, j) , with ϵ_{ij} being independent zero-mean random variables. Formally we have

$$\bar{P}_{ij} = \mathbb{E} P_{ij}(x_{ij}^*) + \epsilon_{ij}. \quad (3.8)$$

Combining (3.7) and (3.8), the customer's expected monthly valuation, $\mathbb{E}_D V_{ij}$, is related to the realized average monthly payment as follows:

$$\bar{P}_{ij} = (\alpha_0 + U_i^T \kappa + G_i^T \phi_0 + W_{ij}^T \psi_0 + u_{0i}) + (\alpha_1 + \phi_1 G_i + \psi_1 W_{ij} + u_{1i})^T \mu_{ij} + \epsilon_{ij}. \quad (3.9)$$

We note that the unobserved factors, captured by ϵ and the random effects u , do not cause omitted-variable bias in our parameter estimates because they are exogenous to the system. Specifically, by Section 3.5, the printing demand is inelastic in service prices. In addition, the customer- and printer-level characteristics, U and W , are fixed physical properties. Therefore, there exists no variable that can affect the response and explanatory variables simultaneously. This means that unobserved variables will not introduce bias to model parameter estimates. Incorporating more variables helps reduce the unexplained response variability, but does *not* affect inferences made on the current model (3.9) (Wooldridge 2010).

In the pricing data set, we observe each device's print volume and payment over at least 10 months. Given the stationarity of print volume series, we estimate μ_{ij} by the arithmetic mean of the observations on printer (i, j) , and denote it by \bar{D}_{ij} . Let the errors due to this estimation be ξ_{ij} . We assume that ξ_{ij} are independent zero-mean random variables. By (3.9), we have

$$\bar{P}_{ij} = (\alpha_0 + U_i^T \kappa + G_i^T \phi_0 + W_{ij}^T \psi_0 + u_{0i}) + (\alpha_1 + \phi_1 G_i + \psi_1 W_{ij} + u_{1i})^T \bar{D}_{ij} + \epsilon_{ij} + \xi_{ij}.$$

Let $e_{ij} \equiv \epsilon_{ij} + \xi_{ij}$. Then e_{ij} are independent random variables. We assume $e_{ij} \stackrel{\text{iid}}{\sim} N(0, \sigma_e^2)$. We now derive the following regression model:

$$\bar{P}_{ij} = (\alpha_0 + U_i^T \kappa + G_i^T \phi_0 + W_{ij}^T \psi_0 + u_{0i}) + (\alpha_1 + \phi_1 G_i + \psi_1 W_{ij} + u_{1i})^T \bar{D}_{ij} + e_{ij}. \quad (3.10)$$

Equation (3.10) is a linear mixed-effects model with the observed average monthly

payment as the response variable, the average monthly print volume, customer and printer’s characteristics as explanatory variables.

3.6.2.2 Parameter estimates and model fitting

We start with the model (3.10) and do model selection based on the Akaike Information Criterion (AIC) and the out-of-sample fitting results. In (3.10), candidate customer characteristics U are industrial segment, fleet size, and the total BW and color monthly printing volume (see Tables 3.1 and 3.2). Candidate printer characteristics W are its manufacturer, technology and functionality, and contract type (see Table 3.2). We emphasize that model selection based on the out-of-sample fitting result is critical for model validation purposes and guarantees a relatively robust and parsimonious model. More importantly, it determines if the customer’s pricing model (3.5) is successful in terms of uncovering a new customer’s service valuation.

Our final pricing model selected based on AIC and out-of-sample fitting results is the following:

$$\bar{P}_{ij} = (\alpha_0 + \kappa \log N_i + W_{ij}^T \psi_0 + u_{0i}) + (\alpha_1 + \psi_1 W_{ij} + u_{1i})^T \bar{D}_{ij} + e_{ij}, \quad (3.11)$$

where $e_{ij} \stackrel{\text{iid}}{\sim} N(0, \sigma_e^2)$, $(u_{0i}, u_{1i}^T) \stackrel{\text{iid}}{\sim} N(0, \Lambda)$ is customer-specific random shocks, W_{ij} consists of information on printer (i, j) ’s manufacturer, technology and functionality, and contract type, N_i is the fleet size of customer i . Table 3.4 shows the parameter estimates for model (3.11). The baseline case is the non-Xerox laser printer with zero-allowance contract.

By the fixed-effects estimates in Table 3.4, the fixed and variable values of MPS to an arbitrary printer are strictly positive. Furthermore, the variable values per BW print are significantly lower than that per color print. Lastly, the intercept estimate for the positive-allowance contracts is positive, while the slope estimates

Table 3.4: Parameters estimates in model (3.11)

The upper panel gives the estimates for fixed-effect coefficients, that is, κ , ϕ and ψ . The lower panel gives the estimates for random-effects, that is, the standard deviation and correlation of u_0 and u_1 . Baseline estimates are for non-Xerox laser printer with zero-allowance contracts. Entries in the “Intercept” column are the coefficient estimates for the corresponding characteristic in the intercept term of (3.10), that is, κ and ψ_0 . Similar interpretation applies to BW and color slope columns. N is the fleet size of each customer. All fixed-effects regressors are categorical except the fleet size, which is continuous. *Significance codes:* * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Fixed effects	Intercept		BW slope		Color slope	
	Estimate	SE	Estimate	SE	Estimate	SE
Baseline, α	-25	101	0.040***	0.010	0.260***	0.020
Multi-functional, ψ	108***	14	-0.004	0.006	-0.016***	0.006
Solid-ink, ψ	6	16	0.006	0.008	-0.016**	0.008
Xerox, ψ	102	66	-0.022**	0.010	-0.160***	0.020
Positive-allowance, ψ	182***	69	-0.030***	0.010	-0.020**	0.010
$\log N$, κ	-10	19				
Random effects	u_0		u_{1b}		u_{1c}	
Std. dev., σ_u	106		0.01		0.04	
	$\text{corr}(u_0, u_{1b}) = -0.26$,		$\text{corr}(u_0, u_{1c}) = 0.045$,		$\text{corr}(u_{1b}, u_{1c}) = 0.76$	

are negative. Recall that the baseline corresponds to the customer’s perceived value under zero-allowance contracts. This result shows that when choosing a positive-allowance contract, customers attach a higher fixed value but lower variable values to the service, consistent with our intuition.

By the estimates for the random effects u_0 and u_1 , we see that the standard deviation of the intercept random effect, u_0 , is much greater than that of the slope random effect, u_1 , after normalized by their respective fixed-effect estimates at baseline. This indicates that there is stronger heterogeneity in the perceived fixed values of the service than the perceived variable values among institutional customers. One important factor affecting the customer’s perceived fixed value is the savings obtained in her financial and human resources when purchasing the MPS. Therefore, this higher variation in the perceived fixed value may arise from the wide variety of ways by which different companies manage their printers before adopting MPS. Furthermore,

the weak correlations between u_0 and u_{1b} , u_0 and u_{1c} indicate that a particular customer’s perceived fixed value, u_0 , has no implication on her perceived variable values, u_1 , for the service and vice versa. That is, customers evaluate the MPS in terms of two *separate* parts: the fixed value of owning it, and the variable value when using it. In contrast, her perceived variable value per BW and color page are highly correlated; if the customer’s perceived value per color page is 10% higher than the population average, there is a 75% probability that she attaches an above average value per BW page.

We note that although the in-sample estimate for the coefficient of $\log N$ is not significantly different from zero, incorporating $\log N$ does improve the predictive power of the model for a new customer’s service valuation, both in and out of sample. This might be due to the limited number of customers (26) in our pricing data set. In fact, we observe that in our cross-validation at the customer level, that is, if we leave one customer out each time, the mean of the 26 point estimates for the coefficient of $\log N$ is -9.92 with p-value < 0.001 . The negative $\log N$ coefficient suggests the market power of large institutional customers: the larger the customer’s device fleet, the more powerful the customer is in negotiation, hence the less she is likely to pay³.

Table 3.5 shows both the in-sample and out-of-sample fitting results of the final pricing model. The in-sample statistics are computed when using all the 3,075 printers in the pricing data set to fit the model. The out-of-sample statistics are computed by doing cross-validation at the customer level. Specifically, we leave one customer out as the “new customer,” estimate the pricing model parameters by running regressions on the rest 25 customers, and then predict the expected service valuation, \bar{V}_{ij} , of the “new customer.” These are the out-of-sample estimates for the service valuation of

³While customers with large fleet might have dense device distribution, thus lowering the fixed service cost per printer for the provider and leading to lower payment, we note that it is not likely to be the reason for the negative $\log N$ coefficient because, according to the provider, big institutional customers usually have branch offices at different places. As a result, the large fleet is possibly distributed over geographically dispersed locations, and, thus, does not necessarily decrease the fixed service cost.

the “new customer.” We rotate this procedure over all 26 customers, pool these out-of-sample estimates together (payments for 3,075 printers in total), and compare them with the 3,075 observed average monthly payments \bar{P}_{ij} . We note that this out-of-sample estimation procedure mimics what happens in practice if the pricing model is used, because the provider then has no information on the new customer’s earlier MPS contracts. Under the final pricing model (3.11), the out-of-sample estimates for customer n ($n \in \{1, 2, \dots, 26\}$) are given by:

$$\bar{V}_{nj} = (\hat{\alpha}_0 + \hat{\kappa} \log N_n + W_{nj}^T \hat{\psi}_0) + (\hat{\alpha}_1 + \hat{\psi}_1 W_{nj})^T \bar{D}_{nj}, \quad (3.12)$$

where $j = 1, 2, \dots, N_n$, $\hat{\alpha}$, $\hat{\kappa}$ and $\hat{\psi}$ denote the point estimates for these model parameters from the rest 25 customers. In the cross-validation, the prediction for the new customer’s expected service valuation, \bar{V}_{ij} , is at the population level. That is, all random effects of the new customer are set to 0 as they cannot be estimated due to lack of information. Hence, \bar{V}_{ij} can be viewed as the population valuation of the service. We note that here we take the contract types on the new customer’s fleet to be known. This information is indeed conveyed to the provider in practice when he designs contracts for the customer.

Table 3.5: Model fitting statistics

	Mean absolute error	Mean absolute percentage error	R-square
In-sample	55.5	32%	92%
Out-of-sample	151	53%	65%

By Table 3.5, the pricing model that takes the form of (3.11) has high explanatory power, both in- and out-of-sample. In particular, recall that by the pricing model (3.5), an arbitrary customer’s service valuation is the sum of the population valuation and a customer-specific random amount. The high out-of-sample explanatory power indicates that the population valuation dominates the customer’s service valuation,

because we use the population valuation, \bar{V} , as the out-of-sample estimates. This is a remarkable result considering customer heterogeneity and the fact that Xerox faces different competitors when bidding for the business of each customer, indicating that these peculiarities have weak influences on the customer’s payments/winning bids. With the help of statistical modeling, we can reliably predict a new customer’s service valuation, which is the key input to the provider’s optimization problem.

Notes on parameter estimates and model specification: Besides (3.10), we also consider other model constructs, such as adding higher-order moments of print volume and industry-level random effects. In particular, we test the within-industry clustering by testing if the variance of the industry-level random effect is zero. We cannot reject the null hypothesis of no industry-level clustering. The fact that the industrial segment does not matter could be due to the infrastructural nature of MPS.

3.7 Optimal MPS pricing

In Section 3.6, we characterized the set of acceptable contracts to the customer using the customer’s pricing model. Constrained within this set, the print volume is inelastic in service prices and the pricing problem is completely determined by the supply side, that is, the provider. In this section, we study how the provider views these acceptable contracts and investigate whether the provider’s risk preference affects his contract decisions in practice. In the following, we first formulate an optimization problem that assumes a risk-averse provider, considers portfolio effects from the existing contracts, and uses the customer’s pricing model as an input. This model yields the optimal contracts for the provider. Next we test empirically the structural assumptions we made on the provider’s decision making process. In particular, we validate the optimization model by examining its explanatory power of the observed contracts, and test our risk-aversion assumption against the risk-neutrality alternative.

3.7.1 Optimization model

We first propose the provider's portfolio optimization problem. Then we find the optimal contracts given the customer's service valuation.

3.7.1.1 Service provider's portfolio optimization model

At time 0, the provider's existing portfolio consists of M customers. The provider bids for the business of a new customer, the $(M+1)$ -th customer, by offering contracts that optimize the provider's expected objective value while still being acceptable to the customer. The contract acceptance condition is the customer's individual rationality constraint. Under the risk-neutrality assumption, a contract of printer j is accepted by the $(M+1)$ -th customer if the present value of the payments does not exceed the maximum amount she would pay:

$$\sum_{t=1}^{T_{M+1,j}} e^{-rt} \mathbb{E} P_{M+1,j}(t; x_{ij}) \leq \sum_{t=1}^{T_{M+1,j}} e^{-rt} \mathbb{E} V_{M+1,j}(t) \quad \forall j \in \{1, 2, \dots, N_{M+1}\}, \quad (3.13)$$

where $V_{M+1,j}(t)$ is the customer's service valuation under the pricing model, r is the discount rate and $T_{M+1,j}$ is the service horizon. The time series $\{P_{M+1,j}(t; x_{M+1,j})\}$ follows identical marginal distribution $f_{P_{M+1,j}}$ at each time point. Similarly, the process $V_{M+1,j}(t)$ has marginal distribution $f_{V_{M+1,j}}$. Therefore, $\mathbb{E} P_{M+1,j}(t; x_{M+1,j}) = \mathbb{E} P_{M+1,j}(x_{M+1,j})$ and $\mathbb{E} V_{M+1,j}(t) = \mathbb{E} V_{M+1,j}$ for all t . The individual rationality constraint (3.13) reduces to

$$E P_{M+1,j}(x_{M+1,j}) \leq \mathbb{E} V_{M+1,j} \quad \forall j \in \{1, 2, \dots, N_{M+1}\}. \quad (3.14)$$

Per our discussion with Xerox, who is a leading provider in the US and Europe with a significant market share, MPS contributes a significant part of the earnings in its entire business portfolio. Due to this large weight and the size of each in-

stitutional customer’s business, Xerox has an incentive to control the variability of the cash inflows from MPS for financial reasons, such as short-term funding costs and budget constraints. Therefore, we model the MPS provider with a risk-averse objective function. Specifically, we assume the following:

Assumption III.1. *The provider has mean-variance objective function $U_s(z; \lambda)$ on the discounted profits and losses from his overall portfolio:*

$$\mathbb{E}U(z; \lambda) = -\lambda \text{var}(z) + \mathbb{E}[z], \quad (3.15)$$

where z is the discounted profits and the constant $\lambda \in \mathbb{R}_+$ is the risk-aversion parameter.

Parameter λ reflects the service provider’s attitude towards profit uncertainty—the greater is λ , the more risk-averse is the provider. When $\lambda = 0$, the provider is risk-neutral. In Section 3.7.2.2, we provide evidence for the risk-aversion assumption by testing it against risk-neutrality.

There are various ways of modeling risk-aversion. In particular, there are two parallel approaches: utility theory and risk measures. The two strive to achieve similar things (i.e., capture decision maker’s preferences towards uncertainty in outcomes), but are based on different assumptions and one does not subsume the other. The utility theory, based on the seminal work by von Neumann and Morgenstern (1947) is usually favored by economists. The theory of risk measures is favored by risk managers, actuaries, and finance practitioners. The foundation of risk measures dates back to Markowitz (1952) and is further developed by, for example, Artzner et al. (1999). In operations management, both approaches are represented. One of the simplest risk measures is variance in (3.15). Because of its simplicity it has been widely applied (see Kim et al. 2007a, Van Mieghem 2007a). The groundbreaking work by Markowitz (1952) that ignited the development of the portfolio theory is

based on this metric. Mean-variance objective in (3.15) has a connection with the utility theory. Specifically, if the utility function is quadratic and uncertainty is Gaussian, then the expected utility maximization is equivalent to the maximization of mean-variance objective (3.15). Nonetheless, one should be cautious that the mean-variance objective function also suffers limitations such as negative marginal valuation and satiation for sufficiently large wealth (Caldentey and Haugh 2006).

Let $O(t)$ denote the profit process of the provider's original business from M customers at time t . Recall that the total service cost for a particular customer consists of labor cost, which is independent of the print volume, and supplies and parts cost. The supplies and parts cost depends on the volume and is the sum of the cost on each printer. Under the mean-variance utility (3.15), we can normalize the fixed labor cost to zero without loss of generality. Recall that in Section 3.4, we use $S_{ij}(t)$ to represent the supplies and parts cost of printer (i, j) . Then the total cash flow from the existing M customers in month t is

$$O(t) \equiv \sum_{i=1}^M \sum_{j=1}^{N_i} (P_{ij}(t; x_{ij}) - S_{ij}(t)). \quad (3.16)$$

Note that $O(t)$ has identical marginal distribution f_O for all t , and no autocorrelation, because $P_{ij}(t; x_{ij})$ and $S_{ij}(t)$ are unautocorrelated stationary processes. Let $X \equiv (x_{M+1,1}, \dots, x_{M+1,N_{M+1}})$ be the collection of all contracts for the $(M+1)$ -th customer's printers. Define $I(t; X)$ as the provider's profit process from the incoming $(M+1)$ -th customer,

$$I(t; X) \equiv \sum_{j=1}^{N_{M+1}} P_{M+1,j}(t, x_{M+1,j}) - S_{M+1,j}(t). \quad (3.17)$$

$I(t; X) \sim f_{I(X)}$ for all t and has no serial correlation for the same reasons as $O(t)$ has. Total cash flow to the provider is $O(t) + I(t; X)$.

Let τ be the provider's planning horizon. By (3.14) and (3.15), the provider solves

the following optimization problem:

$$\max_X \quad -\lambda \text{var} \left(\sum_{t=1}^{\tau} e^{-rt} [I(t; X) + O(t)] \right) + \mathbb{E} \left[\sum_{t=1}^{\tau} e^{-rt} (I(t; X) + O(t)) \right], \quad (3.18a)$$

$$\text{s.t.} \quad \mathbb{E} P_{M+1,j}(x_{M+1,j}) \leq \mathbb{E} V_{M+1,j} \quad \forall j \in \{1, 2, \dots, N_{M+1}\}, \quad (3.18b)$$

$$X \equiv (x_{M+1,1}, \dots, x_{M+1,N_{M+1}}) \quad \forall j \in \{1, 2, \dots, N_{M+1}\}. \quad (3.18c)$$

That is, the provider selects contracts for the new customer so as to maximize the expected discounted cumulative profits from all $M + 1$ customers and minimize the profit variance.

3.7.1.2 Printer-specific optimization

We can write the provider's optimization problem (3.18) as follows.

Proposition III.2. The service provider's optimization problem (3.18) for the profits and losses from all the customers can be written as the following variance minimization problem for the profits and losses from the new customer with binding individual rationality constraints (ignoring customer index $M + 1$):

$$\min_X \quad \sum_{j=1}^N \text{var} (P_j(x_j) - S_j) \quad (3.19a)$$

$$\text{s.t.} \quad \mathbb{E} P_j(x_j) = \mathbb{E} V_j \quad \forall j \in \{1, 2, \dots, N\}, \quad (3.19b)$$

$$X \equiv (x_1, x_2, \dots, x_N) \quad x_j \in \mathbb{R}_+^5 \quad \forall j \in \{1, 2, \dots, N\}, \quad (3.19c)$$

where X is the set of contracts offered to the new customer.

Proof. Recall that $I(t; X)$ and $O(t)$ are stochastic processes with zero serial correlation. Furthermore, $I(t; X) \sim f_{I(X)}$, $O(t) \sim f_O$ for all $t \in \{1, 2, \dots, \tau\}$. Therefore, the

objective function in (3.18) reduces to

$$\max_X \quad -\lambda \sum_{t=1}^{\tau} e^{-2rt} \text{var}(I(X) + O) + \sum_{t=1}^{\tau} e^{-rt} \mathbb{E}[I(X) + O]. \quad (3.20)$$

In addition, $\text{corr}(D_{ij}, D_{i'j'}) = 0$ for all $(i, j) \neq (i', j')$, indicating $\text{cov}(I(X), O) = 0$ and $\text{cov}(P_j(x_j) - S_j, P_{j'}(x_{j'}) - S_{j'}) = 0$. Therefore, (3.20) becomes

$$\max_X \quad -\lambda e^{-2r} \frac{1 - e^{-2r\tau}}{1 - e^{-2r}} \sum_{j=1}^N \text{var}(P_j(x_j) - S_j) + e^{-r} \frac{1 - e^{-r\tau}}{1 - e^{-r}} \sum_{j=1}^N \mathbb{E}[P_j(x_j) - S_j]. \quad (3.21)$$

By (3.2), the Karush-Kuhn-Tucker (KKT) condition on fixed monthly price, F_j requires

$$-1 + u - u_F = 0,$$

where u and u_F are the KKT multipliers for the individual rationality constraint on printer j and $F_j \geq 0$, respectively. Because $u_F \geq 0$, we have $u > 0$. By the complementarity condition, the individual rationality constraint on printer j is binding at the optimum. Then in the objective (3.21), $\sum_{j=1}^N \mathbb{E}[P_j(x_j) - S_j] = \sum_{j=1}^N \mathbb{E}[V_j - S_j]$ is independent of contract x . Therefore, the optimization problem (3.18) reduces to variance minimization with binding constraints (3.19). \square

Optimization problem (3.19) is difficult to solve analytically. However, because the objective function is the sum of the objectives on individual printers and the constraints are per printer, in (3.19) one cannot do better than optimize each printer separately. This way we decompose (3.19) to N simpler problems, one for each printer:

$$\min_{x \in \mathbb{R}_+^5} \quad \text{var}(P_j(x) - S_j), \quad (3.22a)$$

$$\text{s.t.} \quad \mathbb{E} P_j(x) = \mathbb{E} V_j. \quad (3.22b)$$

Problem (3.22) states that under an exogenous upper bound $\mathbb{E}V_j$ on printer j 's monthly payment, the provider looks for the best contract, x_j , so that his monthly profit on printer j equals the upper bound and has the lowest variance. Thus, an ideal contract offsets all the fluctuations in the provider's service cost, if possible.

3.7.1.3 Optimal contracts

We now solve the printer-specific optimization problem (3.22). In the following analysis, we suppress the customer and printer indices for simplicity.

Under given print allowances, A , the service payment is a function of the fixed and variable prices. To make this dependence explicit we let $P_A(F, C)$ represent the service payment. Then the optimization problem (3.22) for an arbitrary printer

$$\min_{F \geq 0, C \geq 0} \text{var}(P_A(F, C) - S), \quad (3.23a)$$

$$\text{s.t. } \mathbb{E} P_A(F, C) = \mathbb{E} V, \quad (3.23b)$$

has a concave objective value function in the fixed and variable prices, (F, C^T) . Therefore, KKT conditions are both sufficient and necessary. We write the optimal fixed and variable prices under given allowances A as $F^*(A)$ and $C^*(A)$, respectively.

Consider first the extreme case where $A = (\infty, \infty)^T$. Intuitively, when the customer has infinite allowances, she only makes a constant monthly payment equal to the fixed price, F . Thus, to extract the customer's service valuation, the provider sets $F^*(\infty, \infty) = \mathbb{E} V$, $C^*(\infty, \infty) = (0, 0)^T$. The optimal contract under infinite allowances is a fixed payment contract, $x^* = (\mathbb{E} V, 0, 0, \infty, \infty)$. In terms of risk mitigation, such a contract does not affect the provider's profit fluctuation. We use the fixed payment contract as the benchmark in our analysis. Let \mathcal{U}_0 be the expected objective value of the provide under the optimal fixed payment contract: $\mathcal{U}_0 \equiv \text{var}(S)$.

By the payment definition (3.2), we write the optimization (3.23) as follows

$$\min_{F \geq 0, C \geq 0} \mathcal{U}_0 - [2\text{cov}(S, C^T(D - A)^+) - \text{var}(C^T(D - A)^+)], \quad (3.24a)$$

$$\text{s.t. } F + C^T \mathbb{E}[(D - A)^+] = \mathbb{E}V. \quad (3.24b)$$

The service cost does not depend on the contract term, x . Therefore, the provider selects the optimal fixed and variable prices, $F^*(A)$ and $C^*(A)$, so that his profit variability is minimized. In the following, we solve for the optimal variable and fixed prices, $C^*(A)$ and $F^*(A)$ and derive the improvement in the expected objective value of a finite-allowance contract over \mathcal{U}_0 . We first define some notation. When $A < (\infty, \infty)^T$, let

$$\tilde{D} \equiv (D - A)^+, \quad \tilde{\mu} \equiv \mathbb{E} \tilde{D}, \quad \text{and} \quad \tilde{\Sigma} \equiv \text{var}(\tilde{D}) = \begin{pmatrix} \tilde{\sigma}_b^2 & \tilde{\rho}\tilde{\sigma}_b\tilde{\sigma}_c \\ \tilde{\rho}\tilde{\sigma}_b\tilde{\sigma}_c & \tilde{\sigma}_c^2 \end{pmatrix}, \quad (3.25)$$

where $\tilde{\sigma}_k^2 = \text{var}(\tilde{D}_k)$ ($k \in \{b, c\}$) and $\tilde{\rho} = \text{corr}(\tilde{D}_b, \tilde{D}_c)$. Similarly, let

$$\tilde{\Gamma} \equiv \begin{pmatrix} \text{cov}(S, \tilde{D}_b) \\ \text{cov}(S, \tilde{D}_c) \end{pmatrix}, \quad \tilde{\gamma} \equiv \begin{pmatrix} \tilde{\Gamma}_b/\tilde{\sigma}_b^2 \\ \tilde{\Gamma}_c/\tilde{\sigma}_c^2 \end{pmatrix}. \quad (3.26)$$

Using notation (3.25) and (3.26), optimization problem (3.24) becomes

$$\max_{F \geq 0, C \geq 0} 2C^T \tilde{\Gamma} - C^T \tilde{\Sigma} C, \quad (3.27a)$$

$$\text{s.t. } F + C^T \tilde{\mu} = \mathbb{E}V. \quad (3.27b)$$

Proposition III.3. Given the print allowances A , the following is the solution of (3.27).

(i) *Positive fixed price:* If the fixed price is positive under the variable prices $C^+(A)$

defined below, that is, $\mathbb{E}V - \tilde{\mu}^T C^+(A) > 0$, then $C^*(A) = C^+(A)$ and $F^*(A) = \mathbb{E}V - \tilde{\mu}^T C^+(A)$, and the corresponding optimal expected objective values $\mathcal{U}^*(A) = \mathcal{U}^+(A)$.

Condition	$C^+(A)$	$\mathcal{U}^+(A)$
$\tilde{\rho} < \min \left\{ \frac{\tilde{\gamma}_b \tilde{\sigma}_b}{\tilde{\gamma}_c \tilde{\sigma}_c}, \frac{\tilde{\gamma}_c \tilde{\sigma}_c}{\tilde{\gamma}_b \tilde{\sigma}_b} \right\}$	$\tilde{\Sigma}^{-1} \tilde{\Gamma}$	$\mathcal{U}_0 + \lambda \tilde{\Gamma}^T \tilde{\Sigma}^{-1} \tilde{\Gamma}$
$\tilde{\rho} \geq \frac{\tilde{\gamma}_c \tilde{\sigma}_c}{\tilde{\gamma}_b \tilde{\sigma}_b}$	$(\tilde{\gamma}_b, 0)^T$	$\mathcal{U}_0 + \lambda \tilde{\gamma}_b^2 \tilde{\sigma}_b^2$
$\tilde{\rho} \geq \frac{\tilde{\gamma}_b \tilde{\sigma}_b}{\tilde{\gamma}_c \tilde{\sigma}_c}$	$(0, \tilde{\gamma}_c)^T$	$\mathcal{U}_0 + \lambda \tilde{\gamma}_c^2 \tilde{\sigma}_c^2$

(ii) *Zero fixed price:* If the fixed price is negative or zero under the variable prices $C^+(A)$, that is, $\mathbb{E}V - \tilde{\mu}^T C^+(A) \leq 0$, then $C^*(A) = C^0(A)$ as defined below, $F^*(A) = 0$, and the corresponding optimal expected objective values $\mathcal{U}^*(A) = \mathcal{U}^0(A)$.

Condition	$C^0(A)$	$\mathcal{U}^0(A)$
$\tilde{\Sigma}^{-1} \left(\tilde{\Gamma} - \frac{\tilde{\mu}^T \tilde{\Sigma}^{-1} \tilde{\Gamma} - \mathbb{E}V}{\tilde{\mu}^T \tilde{\Sigma}^{-1} \tilde{\mu}} \tilde{\mu} \right) > 0$	$\tilde{\Sigma}^{-1} \left(\tilde{\Gamma} - \frac{\tilde{\mu}^T \tilde{\Sigma}^{-1} \tilde{\Gamma} - \mathbb{E}V}{\tilde{\mu}^T \tilde{\Sigma}^{-1} \tilde{\mu}} \tilde{\mu} \right)$	$\mathcal{U}_0 + \lambda \left[\tilde{\Gamma}^T \tilde{\Sigma}^{-1} \tilde{\Gamma} - \frac{(\tilde{\mu}^T \tilde{\Sigma}^{-1} \tilde{\Gamma} - \mathbb{E}V)^2}{\tilde{\mu}^T \tilde{\Sigma}^{-1} \tilde{\mu}} \right]$
$\tilde{\rho} > \frac{1}{\tilde{\sigma}_b \tilde{\sigma}_c \mathbb{E}V} \left(\tilde{\mu}_b \tilde{\Gamma}_c - \tilde{\mu}_c \tilde{\Gamma}_b + \frac{\tilde{\mu}_c}{\tilde{\mu}_c} \tilde{\sigma}_b^2 \mathbb{E}V \right)$	$\left(\frac{\mathbb{E}V}{\tilde{\mu}_b}, 0 \right)^T$	$\mathcal{U}_0 + \lambda \tilde{\sigma}_b^2 \left[\tilde{\gamma}_b^2 - \left(\frac{\mathbb{E}V}{\tilde{\mu}_b} - \tilde{\gamma}_b \right)^2 \right]$
$\tilde{\rho} > \frac{1}{\tilde{\sigma}_c \tilde{\sigma}_b \mathbb{E}V} \left(\tilde{\mu}_c \tilde{\Gamma}_b - \tilde{\mu}_b \tilde{\Gamma}_c + \frac{\tilde{\mu}_b}{\tilde{\mu}_c} \tilde{\sigma}_c^2 \mathbb{E}V \right)$	$\left(0, \frac{\mathbb{E}V}{\tilde{\mu}_c} \right)^T$	$\mathcal{U}_0 + \lambda \tilde{\sigma}_c^2 \left[\tilde{\gamma}_c^2 - \left(\frac{\mathbb{E}V}{\tilde{\mu}_c} - \tilde{\gamma}_c \right)^2 \right]$

Proof. We prove Proposition III.3 as the following.

Let $u_F \geq 0$, $u_b \geq 0$ and $u_c \geq 0$ be the KKT multipliers of the inequality constraints $F \geq 0$, $C_b \geq 0$ and $C_c \geq 0$, respectively. Let v be the KKT multiplier to the equality constraint $F + C^T \tilde{\mu} = \mathbb{E}V$. The KKT condition on F yields $u_F = v$. Then the KKT condition on C is

$$-2\tilde{\Gamma} + 2\tilde{\Sigma}C + u_F \tilde{\mu} - \begin{pmatrix} u_b \\ u_c \end{pmatrix} = 0.$$

We proceed by discussing whether $u_F = 0$ or $u_F > 0$. That is, we solve the

following two sets of equations on the variable price C :

$$-2\tilde{\Gamma} + 2\tilde{\Sigma}C - \begin{pmatrix} u_b \\ u_c \end{pmatrix} = 0, \quad (3.28a)$$

$$C \geq 0, \quad u_b \geq 0, \quad u_c \geq 0, \quad (3.28b)$$

and

$$-2\tilde{\Gamma} + 2\tilde{\Sigma}C + u_F\tilde{\mu} - \begin{pmatrix} u_b \\ u_c \end{pmatrix} = 0, \quad (3.29a)$$

$$C^T\tilde{\mu} = \mathbb{E}V, \quad (3.29b)$$

$$C \geq 0, \quad u_b \geq 0, \quad u_c \geq 0. \quad (3.29c)$$

Let $C^{*i}(A)$ and $C^{*ii}(A)$ be the solutions to (3.28) and (3.29), respectively. Both (3.28) and (3.29) have unique solutions because they are the KKT conditions for the following convex optimization

$$\max_{C \geq 0} 2C^T\tilde{\Gamma} - C^T\tilde{\Sigma}C, \quad (3.30)$$

$$\max_{C \geq 0, C^T\tilde{\mu} = \mathbb{E}V} 2C^T\tilde{\Gamma} - C^T\tilde{\Sigma}C. \quad (3.31)$$

For $C^{*i}(A)$ to be feasible in the original problem (3.27), we further require $F \geq 0$.

That is,

$$\mathbb{E}V - \tilde{\mu}^T C^{*i}(A) \geq 0. \quad (3.32)$$

If (3.32) holds, then $(\mathbb{E}V - \tilde{\mu}^T C^{*i}(A), C^{*i}(A))$ satisfy all the KKT conditions for the original optimization problem (3.27) and, thus, is optimal. Similarly, for $C^{*ii}(A)$ to

be feasible in (3.27), we further require

$$u_F > 0. \quad (3.33)$$

If (3.33) holds, $(0, C^{*ii}(A))$ satisfy all the KKT conditions and, thus, is optimal.

With a concave objective function, (3.23) has a unique solution on $(F \geq 0, C \geq 0)$. Indeed, the feasibility conditions (3.32) and (3.33) complement each other. Furthermore, we will see that if the equality holds in (3.32), then $C^{*i}(A) = C^{*ii}(A)$. In the following, we first characterize the solutions to optimization problems (3.30) and (3.31), and after that we give the solution to the original problem (3.27).

Lemma III.4. *Given the print allowances A , $C^{*i}(A)$ solves the optimization problem (3.30) and it is given by*

$$C^{*i}(A) = \begin{cases} \tilde{\Sigma}^{-1}\tilde{\Gamma} & \text{if } \tilde{\rho} < \min \left\{ \frac{\tilde{\gamma}_b\tilde{\sigma}_b}{\tilde{\gamma}_c\tilde{\sigma}_c}, \frac{\tilde{\gamma}_c\tilde{\sigma}_c}{\tilde{\gamma}_b\tilde{\sigma}_b} \right\}, \\ (\tilde{\gamma}_b, 0)^T & \text{if } \tilde{\rho} \geq \frac{\tilde{\gamma}_c\tilde{\sigma}_c}{\tilde{\gamma}_b\tilde{\sigma}_b}, \\ (0, \tilde{\gamma}_c)^T & \text{if } \tilde{\rho} \geq \frac{\tilde{\gamma}_b\tilde{\sigma}_b}{\tilde{\gamma}_c\tilde{\sigma}_c}. \end{cases} \quad (3.34)$$

Lemma III.4 describes the solution to the optimization problem (3.30) when $C^{*i}(A)$ falls in one of the mutually exclusive sets: $(C_b > 0, C_c > 0)$, $(C_b > 0, C_c = 0)$ and $(C_b = 0, C_c > 0)$, respectively. These three sets and $(C_b = 0, C_c = 0)$ are partitions of the feasible region $C \geq 0$ of the optimization problem (3.30). However, under finite allowances, a contract with no variable payment, i.e., $C = 0$, is never optimal as long as the sensitivity of the service cost to print volume is nonzero. Therefore, the solution to the convex optimization problem (3.30) must be in one and only one of the three sets. Inequalities on the correlation coefficient $\tilde{\rho}$ give the conditions under which the solution falls in the corresponding set, or equivalently, takes the corresponding $C^{*i}(A)$ value.

Lemma III.5. *Given the print allowances A , $C^{*ii}(A)$ solves optimization problem*

(3.31) and it is given by

$$C^{*ii}(A) = \begin{cases} \tilde{\Sigma}^{-1} \left(\tilde{\Gamma} - \frac{\tilde{\mu}^T \tilde{\Sigma}^{-1} \tilde{\Gamma} - \mathbb{E}V}{\tilde{\mu}^T \tilde{\Sigma}^{-1} \tilde{\mu}} \tilde{\mu} \right) & \text{if } \tilde{\Sigma}^{-1} \left(\tilde{\Gamma} - \frac{\tilde{\mu}^T \tilde{\Sigma}^{-1} \tilde{\Gamma} - \mathbb{E}V}{\tilde{\mu}^T \tilde{\Sigma}^{-1} \tilde{\mu}} \tilde{\mu} \right) > 0, \\ \left(\frac{\mathbb{E}V}{\tilde{\mu}_b}, 0 \right)^T & \text{if } \tilde{\rho} > \frac{1}{\tilde{\sigma}_b \tilde{\sigma}_c \mathbb{E}V} \left(\tilde{\mu}_b \tilde{\Gamma}_c - \tilde{\mu}_c \tilde{\Gamma}_b + \frac{\tilde{\mu}_c}{\tilde{\mu}_b} \tilde{\sigma}_b^2 \mathbb{E}V \right), \\ \left(0, \frac{\mathbb{E}V}{\tilde{\mu}_c} \right)^T & \text{if } \tilde{\rho} > \frac{1}{\tilde{\sigma}_c \tilde{\sigma}_b \mathbb{E}V} \left(\tilde{\mu}_c \tilde{\Gamma}_b - \tilde{\mu}_b \tilde{\Gamma}_c + \frac{\tilde{\mu}_b}{\tilde{\mu}_c} \tilde{\sigma}_c^2 \mathbb{E}V \right). \end{cases} \quad (3.35)$$

Lemma III.5 describes the solution to the optimization problem (3.31) when $C^{*ii}(A)$ falls in ones of the mutually exclusive sets $(C_b > 0, C_c > 0)$, $(C_b > 0, C_c = 0)$ and $(C_b = 0, C_c > 0)$, respectively. These three sets are partitions of the region $S = \{C : C \geq 0, C^T \tilde{\mu} = \mathbb{E}V\}$ in (3.31). Therefore, the solution to the convex optimization problem (3.31) must be in one and only one of the three sets. Inequalities on the correlation coefficient $\tilde{\rho}$ give the conditions under which the solution falls in the corresponding set, or equivalently, takes the corresponding $C^{*ii}(A)$ value.

Proposition III.3 is proved by combining Lemma 3.34 and Lemma III.5. \square

In practice, the customer decides whether to use a zero- or positive-allowance contract of a particular printer. By Proposition III.3, if zero-allowance contract is preferred, then the optimal contract is $x = (F^*(0, 0), C^*(0, 0), 0, 0)$. If positive-allowance contract is preferred, then the optimal contract has to be solved numerically. Specifically, we need to compute and compare the optimal expected objective value $\mathcal{U}^*(A)$ under all positive allowances using Proposition III.3. Fortunately, as observed in practice, the set of positive allowances is rather limited, possibly due to administrative concerns. The fixed and variable prices, on the other hand, takes continuous values. Therefore, finding the optimal contract requires solving continuous optimization problems on (F, C) under a discrete set of allowances, and comparing the corresponding objective values.

3.7.2 Empirical analysis

Because service cost is an input to the optimization model, we carry out the empirical study using the cost data set (see Section 3.3.2), which consists of 1,305 printers from eight customers with consumables costs. The cost data set consists of the latest contract, monthly print volume, and consumables costs per page of each printer. We assume that the monthly service cost is affine in print volume D : $S = (\gamma^o)^T D + Z$, where Z is independent of both D_b and D_c , $\gamma^o \equiv (\gamma_b^o, \gamma_c^o)^T$ are the provider's service costs per BW and color print.

We first compare the provider's risk-adjusted monthly earnings from these eight customers. This provides insights into how the MPS provider prices earnings risks of different customers. Under the mean-variance objective, the provider measures the earnings variability using the variance. Therefore, we mean the provider's risk-adjusted earnings of printer (i, j) by normalizing the average monthly earnings by the earnings' standard deviation. Let $I_{ij}(t) \equiv P_{ij}(t; x_{ij}) - S_{ij}(t)$ represent printer (i, j) 's monthly earnings in month t under observed contract x_{ij} . Then $I_{ij}(t)$ has identical marginal distribution $f_{I_{ij}}$, where $I_{ij} \equiv P_{ij}(x_{ij}) - S_{ij}$. Let I_{ij}^{RA} be the risk-adjusted earnings of printer (i, j) . By definition, $I_{ij}^{RA} \equiv \frac{\bar{I}_{ij}}{\sqrt{\text{var}(I_{ij})}}$. To test if the risk-adjusted profit is related to customer-specific characteristics such as industrial segment and fleet size, we establish the following mixed-effects model

$$I_{ij}^{RA} = G_i^T \beta_0 + \beta_1 \log N_i + u_i + e_{ij}, \quad (3.36)$$

where G_i is the 5-by-1 dummy vector indicating the industrial segment of customer i ($i = 1, 2, \dots, 8$), N_i is the fleet size of customer i , u_i is the customer-level random effect accounting for customer-level clustering, and e_{ij} is the random error. u_i and e_{ij} are mutually independent zero-mean normal random variables.

Under the observed contracts, we find that neither β_0 or β_1 is statistically signifi-

cant for the 1,035 printers. This result means that all the customers bring in similar risk-adjusted monthly earnings for the provider. Thus, all customers' earnings risks are priced in a same way possibly due to intense competition among the providers.

In the following, we turn to examine our modeling assumptions. We first validate the optimization model of the provider's decision making process. Next we compare the risk-aversion assumption to the risk-neutrality alternative. We apply our optimization model in a way that mimics the practice. Specifically, when computing the optimal contracts for a particular customer, we take her as a new customer and predict her service valuation using all the other customers' data as in the out-of-sample test procedure described in Section 3.6.2.2. This predicted maximum average monthly payment, \bar{V} , along with the customer's contract type preference, the print volume and service cost, are taken as inputs of the optimization model (3.19), which yields the optimal contracts.

3.7.2.1 Test of the optimization model

We compare the predicted contracts from the optimization model (3.19) with the corresponding actual contracts. Let x_{ij}^* and \hat{x}_{ij} be the theoretical and observed contracts on printer (i, j) , respectively. We observe only zero-allowance contracts in the cost data set. That is, all the eight customers prefer zero-allowance contracts on these printers. Therefore, we set the allowances in x^* and \hat{x} to zero for all the printers under consideration. We compute correlations between the corresponding contract parameters in x^* and \hat{x} . These correlations quantify the explanatory power of our optimization model (3.19). To benchmark the performance of model (3.19), we also predict the contract parameters using a simple linear regression model with no structural assumptions on the provider's decision making process. We call this the null model and use $\tilde{x} \equiv (\tilde{F}, \tilde{C}_b, \tilde{C}_c, 0, 0)$ to denote the predicted contract from the null model. Specifically, we formulate the null model for contract of printer (i, j) as

follows.

$$\tilde{C}_{ijb} = Y_{ij}^T \eta_b + \omega_{ijb}, \quad (3.37a)$$

$$\tilde{C}_{ijc} = Y_{ij}^T \eta_c + \omega_{ijc}, \quad (3.37b)$$

$$\tilde{F}_{ij} = \bar{V}_{ij} - \tilde{C}_{ij}^T \mu_{ij}, \quad (3.37c)$$

where $\omega_{ijb} \stackrel{\text{iid}}{\sim} N(0, \delta_{nb}^2)$, $\omega_{ijc} \stackrel{\text{iid}}{\sim} N(0, \delta_{nc}^2)$, and $Y_{ij}^T \equiv (\mu_{ij}^T, \sigma_{ijb}, \sigma_{ijc}, \rho, \gamma_{ij}^T, \bar{V}_{ij})$ is the vector of inputs of the optimization model (3.19). Because the expected monthly payment under the optimal contract equals the maximum amount acceptable to the customer, \bar{V} , there are only two degrees of freedom in \tilde{x} . This is reflected in the equality on the fixed price \tilde{F} in (29c). In essence, the null model assumes linear dependence of the contract parameters on the inputs of the optimization model. Besides the basic linear form of (3.37), we also tested models with higher-order terms. However, we do not observe significant changes from the result when using (3.37). Comparing the explanatory power of the null model with the optimization model highlights the improvements gained from structural modeling. More importantly, the null model is a benchmark that, in some sense, “quantifies” the complexity of the actual contracting process, thus facilitating the evaluation of the optimization model.

We obtain optimal contracts x^* and null predictions \tilde{x} following the procedure below.

Step 1: Select customer i ($i = 1, 2, \dots, 8$) as the new customer.

Step 2: Predict service valuation of all her printers, \bar{V}_{ij} ($j = 1, 2, \dots, N_i$), by (3.12) as described in the out-of-sample test of Section 3.6.2.2.

Step 3.1: Find x_{ij}^* ($j = 1, 2, \dots, N_i$) by solving the optimal problem (3.38) below for printer j under $A = (0, 0)^T$. The estimated maximum average monthly payment, \bar{V}_{ij} , correlated cost, γ_{ij} , and observed print volume, D_{ij} , are model

inputs.

$$\min_x \text{var}(P_{ij}(x) - S_{ij}), \quad (3.38a)$$

$$\text{s.t. } \mathbb{E} P_{ij}(x) = \bar{V}_{ij}, \quad (3.38b)$$

$$x \in \{(F, C^T, 0, 0) : F \in \mathbb{R}_+, C \in \mathbb{R}_+^2\}. \quad (3.38c)$$

Step 3.2: Find \tilde{x}_{ij} ($j = 1, 2, \dots, N_i$) by first estimating null model parameters η_b and η_c using data from the other 7 customers, then predicting the variable and fixed prices based on the estimated parameters and model (3.37).

Following this procedure, we obtain 1,035 optimal contracts x^* and 1,035 null predictions \tilde{x} . Because x^* , \tilde{x} and \hat{x} all have zero allowances, we compare the fixed and variable prices of the predicted x^* and \tilde{x} , with the observed \hat{x} , respectively. Table 3.6 presents the correlations between the corresponding terms in the predicted and observed contracts. The poor explanatory power of the null model and the significant improvement of the optimization model validate our structural formulation on the provider's decision making process. It indicates that our optimization model captures major factors in the actual contracting process.

Table 3.6: Correlations between predicted and observed prices

	Fixed price (95% C.I.)	BW variable price (95% C.I.)	Color variable price (95% C.I.)
Optimal contract: x^*	0.55 (0.51, 0.60)	0.16 (0.10, 0.22)	0.11 (0.05, 0.17)
Contract from regression: \tilde{x}	0.16 (0.10, 0.22)	-0.24 (-0.18, -0.30)	-0.67 (-0.63, -0.70)

3.7.2.2 Test of risk-aversion assumption

In our optimization model, one critical assumption is that the provider is risk-averse with mean-variance objective: $U(z) = \mathbb{E}[z] - \lambda \text{var}(z)$, where λ is the risk-

aversion parameter and z is the cumulative discounted earnings. While on the one hand, Xerox has an incentive to control the variability of the MPS cash inflows for financial reasons given the significant weight of MPS in its business portfolio; on the other hand, since he serves many clients with independent printing demands, he might not concern about the earnings uncertainties when deciding the contracts as they could be diversified away. Therefore, the risk-aversion of the provider is not immediate and it is important to verify whether risk-aversion is a factor in the provider's decisions or not. In the following analysis, we test the risk-aversion assumption in the context of MPS by testing the null hypothesis below:

Hypothesis 1. The optimal contracts obtained under the risk-neutrality (RN) assumption match the observed contracts as well as those under the risk-aversion (RA) assumption.

To test Hypothesis 1, we first characterize the optimal contracts under the RN assumption. Next we introduce the *normalized squared distance* as a measure for the “match” between the predicted and observed contracts in Hypothesis 1. Finally, we formulate Hypothesis 1 mathematically using the normalized squared distance and test it by a Student's t -test.

A risk-neutral provider has $\lambda = 0$. Under our constrained optimization model, this means that any contract that charges the customer at the maximum acceptable amount is optimal for the provider. Formally, for printer (i, j) , any contract (F, C^T, A^T) that satisfies

$$F + C^T \mathbb{E} [(D_{ij} - A)^+] = \bar{V}_{ij} \quad (3.39)$$

is equally likely to be selected by the provider to be the optimal contract. This is because under the RN assumption, the average monthly earning is the only criterion by which the provider selects the contracts. Furthermore, recall that all printers in

the cost data set are preferred to have zero-allowance contracts. We set $A = 0$ in (3.39) and obtain

$$F + C^T \mu_{ij} = \bar{V}_{ij}, \quad (3.40)$$

where F and C are nonnegative. Recall that the expected monthly print volume, μ_{ij} , and the predicted customer's payment, \bar{V}_{ij} , are both printer-specific and fixed. Therefore, (3.40) defines a unique bounded plane $\mathcal{S} = \{(F, C_b, C_c)\}$ for each printer. Any point (F, C_b, C_c) in plane \mathcal{S} is equally likely to be selected by the provider and offered to the customer. In other words, because both the provider and the customer are RN, the winning bids are also the optimal bids for the provider. In contrast, under the RA assumption the provider selects one unique point in the RN plane \mathcal{S} to be the optimal contract.

Now we specify the meaning of “match” between the predicted and observed contracts. We define it as the normalized squared distance between the predicted and observed contracts as follows.

Definition III.6. The normalized squared distance d between the predicted contract (F, C_b, C_c) and observed contract $(\hat{F}, \hat{C}_b, \hat{C}_c)$ is

$$d \equiv \left(\frac{F - \hat{F}}{L_f} \right)^2 + \left(\frac{C_b - \hat{C}_b}{L_b} \right)^2 + \left(\frac{C_c - \hat{C}_c}{L_c} \right)^2, \quad (3.41)$$

where $L_f = \bar{V}$, $L_b = \bar{V}/\mu_b$ and $L_c = \bar{V}/\mu_c$ are the maximum values of F , C_b and C_c under equality (3.40), respectively.

Under the RA assumption, we have one unique prediction for each printer, thus one unique distance. Let the distance between the predicted and observed contracts of printer l be d_l^A ($l = 1, 2, \dots, 1035$). d_l^A measures the “match” of the RA optimal contract to the observed contract of printer l . Recall that this predicted contract under the RA assumption is one unique point in the RN plane \mathcal{S} .

Under the RN assumption, all points in the RN plane \mathcal{S} are optimal. Therefore, there are infinitely many optimal contracts for each printer, thus infinitely many distances. Let $d_l^N(F, C_b, C_c)$ denote the distance between the observed contract of printer l and one particular predicted contract (F, C_b, C_c) in the optimal plane \mathcal{S}_l defined by (3.40) ($l = 1, 2, \dots, 1035$).

To formulate Hypothesis 1 mathematically, we further characterize the “match” of the RN predictions using the average distance from the RN plane \mathcal{S} to the true contract:

$$\overline{d_l^N} \equiv \frac{\iiint_{\mathcal{S}_l} d_l^N(F, C_b, C_c) dF dC_b dC_c}{\iiint_{\mathcal{S}_l} dF dC_b dC_c}. \quad (3.42)$$

Intuitively, this is the sum of distances from points in the RN plane weighted by their corresponding probabilities of being chosen as the RN contract. Therefore, $\overline{d_l^N}$ represents the average “match” of the RN plane \mathcal{S}_l to the observed contract of printer l . We compute the integral in (3.42) numerically by dividing the plane \mathcal{S}_l into evenly spaced small areas.

Now we are ready to compare the “match” of the RN predictions with the “match” of the RA predictions. In particular, we define Δ_l to be the difference between the average distances of the RN and RA contracts for printer l : $\Delta_l \equiv \overline{d_l^N} - d_l^A$. Δ_l is a random variable, contingent on the printer under consideration. Under Hypothesis 1, the mean of Δ should be 0. That is, for an arbitrary printer in the population, the randomly picked contract under the RN assumption is expected to be as close to the true contract as the RA contract. Our 1,035 printers provide 1,035 samples of Δ . Therefore, Hypothesis 1 can be written as the following:

Hypothesis 1a.

$$(H_0) \quad \mathbb{E}[\Delta] = 0,$$

$$(H_1) \quad \mathbb{E}[\Delta] < 0.$$

Hypothesis 1a is a one-sided test with the alternative hypothesis being that the RA assumption matches the data better than the RN assumption over the population.

We are interested in the testing if the mean of Δ is 0. Assume that Δ_l of different printers are i.i.d. normal random variables with unknown mean and variance (plots of $\{\Delta_l\}$ exhibit features of a normal random variable). Then we can test Hypothesis 1a using a Student's t -test. The t -statistic is -15.2 , thus rejecting H_0 and accepting the alternative H_1 . We also tested the hypothesis using normalized absolute distance and the t -statistic is -13.7 . These results⁴ mean that for an arbitrary printer, the predictions under the RA assumption match the observations better than that under the RN assumption. That is, the service provider is better modeled to be risk-averse than risk-neutral.

3.8 Concluding remarks

We examine contractual interactions between the managed print service provider and customers using proprietary data sets provided by Xerox. We develop two models: one on the customer's valuation of the service, and the other on the provider's selection among the set of contracts acceptable to the customer obtained from the first model.

In the first model on the customer's side, we show that the customer's printing demand is inelastic in service prices over the observed price range. Furthermore, we propose a customer's pricing model that generates reliable forecasts for the customers' service payments. We find that customers in different industries evaluate the MPS in the same way and pay less for each printer when the fleet size is large. The customer's service payments also depend on her preferences between zero- and positive-allowance contracts, and printer-specific characteristics. We note that due to

⁴We also tested Hypothesis 1 by comparing the *chances* that an RN contract has smaller distance than the RA contract. In particular, we computed the probability p that a randomly-picked RN contract is better than the RA contract, and tested $(H_0) \mathbb{E}[p] = 0.5$ against $(H_1) \mathbb{E}[p] < 0.5$. We carried out a percentile bootstrap procedure (Cameron and Trivedi 2005) and rejected (H_0) under the normalized squared and absolute distances at 95% significance level.

confidentiality issues, it is difficult to obtain data from multiple providers, and our empirical results are derived using the data set from Xerox. However, given the wide variety of customers in the MPS portfolio of Xerox, who is a leading provider in the US and Europe with a significant market share, we expect this modeling approach and the qualitative insights to apply to other MPS providers and possibly to other types of managed services as well.

In the second model on the provider’s selection of the optimal contracts, we formulate a contract optimization problem that captures the risk-aversion of the provider, considers portfolio effects from the existing contracts, and uses customer’s pricing model as an input. We validate our modeling assumptions empirically. Furthermore, we observe that the risk-adjusted earnings (i.e., monthly average earnings normalized by the earnings variability) do not depend on industry or fleet sizes. Put another way, all customers bring in similar risk-adjusted monthly earnings for the provider. This may indicate intense competition among the providers as all customers’ earnings risks are priced in a same way.

To the best of our knowledge, this paper is the first in-depth study of pricing in managed services. Our approach provides a new way of tackling information asymmetry between service providers and customers. It offers a viable construct for studies on managed services for institutional customers. One interesting extension for future research is to consider the problem in the dynamic setting and incorporate the option of renegotiation. In practice, the sales representatives tend to “seal the deal” first and adjust the terms later. This is analogous to a real option whose value is related to the profit foregone before renegotiation. Another possible direction to pursue in the future research is to account for administrative cost associated with the number of unique contracts in a service portfolio. In the paper, both the number of unique contracts and the set of printers each contract covers are external variables. Introduction of administrative costs will endogenize them as decision variables and explore

the trade-off between the improvement on portfolio risk and reduction on expected profit.

CHAPTER IV

Price dynamics of used durable goods: A dynamic factor model approach

4.1 Introduction

The price dynamics of used durable goods provides valuable information on the movement patterns of the market, and how a specific product will be affected by an operational decision from the higher management or a shock in the external economic condition. In this paper, we study the price dynamics of a particular type of durable goods (exact name nondisclosable due to confidentiality reasons) using resale price data throughout the country. Our econometric exercise serves two goals and, thus, naturally divides into two parts: (i) Identify the comovements among the resale prices across the entire market as well as from products with similar functionalities. Examine the effects of external shocks on the resale prices. (ii) Identify the comovements among products of the same brand and original equipment manufacturer (OEM) after controlling for the fluctuations due to external effects. Study the impacts of brand-/OEM-level shocks on the resale prices.

Our research problem falls within the broad realm of making inferences and forecasting using a panel data set with a large number of predictors in macroeconomics. The standard and widely used approach is to extract a few number of *dynamic latent*

factors from the panel data set that explain a significant fraction of the variability of each time series, i.e., the dynamic factor models (Breitung and Eickmeier 2006, Stock and Watson 2011, 2006). In the literature, there are mainly two approaches to dynamic factor models: a non-parametric one that primarily relies on the method of principal component (Eickmeier 2009, Sala 2003, Eickmeier and Breitung 2006, Forni and Reichlin 1996) and a parametric approach that casts the model into the state space framework¹ (Gregory et al. 1997, Kose et al. 2003, Diebold et al. 2006, Moench et al. 2011). In the former, principal component analysis (PCA) first extracts a few *static* factors (Bai and Ng 2002), which are then subject to vector autoregression study (VAR) and identification schemes to further identify factor dynamics and shocks. In the latter, researchers impose structural assumptions on the data and factor dynamics first and then estimate the model parameters and latent factors. In comparison, the PCA method is less computationally intensive and shown to have nice properties in forecasting (Stock and Watson 2002). The state space method, on the other hand, is better for inference purposes and provides a more holistic way for model estimation. In this paper, as we are primarily interested in making inferences on the latent factors, the state space approach is adopted.

Our work has the following features. First, compared with similar studies in macroeconomics, this research problem is complicated by its operational details. There are three essential metrics associated with each used product: its product model, age and usage level. While all *new* goods of the same product model are identical, special care has to be taken to explicitly control for the depreciation effects due to aging and usage in our study, so as to obtain well-defined factors and derive meaningful insights. Second, we group the products into a number of segments based on their functionality and distinguish between the market-wide comovements (captured by the *industry factor*) and segment-specific comovements. This multi-level dynamic

¹For a comprehensive introduction on state space models, see Durbin and Koopman (2001).

factor model captures the fact that due to strong heterogeneity, products within different segments might have their own dynamics in addition to the industry factor. It shares some commonality with Kose et al. (2003), and particularly Moench and Ng (2011), Moench et al. (2011), and Diebold et al. (2008). Third, we further examine the OEM- and brand-level comovements after controlling for the external segment factors. This study is particularly relevant to industrial practitioners because, while the industry and segment factors are helpful in understanding how the market as a whole responds to exogenous shocks such as a financial downturn, they provide little guidance on how a particular business decision or OEM/brand-specific incident affects the associated products. The study on the OEM and brand comovements, on the other hand, sheds some light in this direction.

In this paper, we show that, despite the heterogeneity in product models and brands, there are strong comovements throughout the used goods market and within each segment. The latent industry and segment factors are able to capture up to 81.4% of the variation of a particular product's price changes. This dependence on the latent factors, or equivalently, the level of within-segment comovement, however, varies with segment. In particular, the segment with fancy and pricey products exhibits weaker comovement than others. This is an intuitive result because ordinary goods targets the majority of the population and is thus more sensitive to changes in the external economic condition, which in turn renders the brand-level heterogeneity secondary effects in the products' price dynamics. The fancy and pricey products, on the other hand, target a much smaller crowd generally with higher income. As a result, their price dynamics is not as contingent on the general economic condition and the heterogeneity of the product models comes into play.

We report that the industry factor is very persistent. The persistence of segment-specific comovements after controlled for the industry factor (i.e., the *controlled* segment factors), however, vary. The persistence property of a dynamic factor character-

izes how long into the future a one-time shock's effect lasts. Therefore, the industry factor and some controlled segment factors have long memories. As an example, we study the lasting impacts of the 2008 financial crisis on the entire market and each segment.

In the OEM analysis, we report that the OEM- and brand-level comovements are much less persistent than the industry and segment factors. Furthermore, we find that: (i) Termination of a brand does not have a material impact on the resale prices of the associated products. (ii) Big recalls due to product malfunction induce significant price drops in the associated brand. These negative impacts, however, do not spill-over to other brands owned by the same OEM.

The rest of the paper is organized as follows. Section 2 discusses the data and measures to account for the aging and usage depreciation in detail. Section 3 presents our model. Section 4 explains the estimation results for the industry and segment factors. Section 5 discusses the OEM- and brand-level analysis. Section 6 concludes.

4.2 Data

Our data set contains nationwide resale prices of a large class of used durable goods during 2000:1–2012:12 (name of the goods nondisclosable due to confidentiality reasons). Each product sold has one basic attribute—*model*, which specifies the product's functionality and design, and signifies the product's brand and OEM. Take the automobile industry for example. Focus, Corolla and Camero are all vehicle models, whose brands are Ford, Toyota and GMC, and whose OEMs are Ford, Toyota and GM, respectively. One OEM could produce multiple brands. For example, besides the Ford brand, OEM Ford also owns Lincoln. Besides the Toyota brand, OEM Toyota also produces Lexus. Besides the GMC brand, OEM GM also manufactures Chevrolet, Cadillac, etc. Our data set covers a variety of brands and models within the class of durable goods under consideration.

From the data set, we know the identification number of each product sold, the sale month and price, as well as the product’s usage level, brand, model, and the production year of the model. Here the usage level is a continuous metric. It could be the operation hours in cases of used capital equipment, or the mileage in cases of used vehicles. The production year of the model, on the other hand, is related to the model generation. This metric is most relevant in industries with regular new product launches such as electronics and automobile. The data set also contains assessment on the overall condition of each product—whether it is fair, thus sold with insurance; or junk, thus sold without insurance. We only consider products under normal condition which are sold with insurance in our study. We define the *age* of a product as the difference between its resale time and the production year of the model, as we do not have information on when a particular product is first purchased/leased. To avoid significant depreciation effects and extremely old models, we only consider products that are 0 to 6 years old. These products constitute more than 90% sales volume of the original data set. This way we obtain a data set with 39,406,520 used products sold over 13 years.

4.2.1 Data preparation

Our goal is to study the price dynamics of this used goods market on a macro level. We hope to identify and understand latent dynamic factors that drive the observed resale price changes. While similar problems have been widely studied in macroeconomics and finance literature (Eickmeier 2009, Kose et al. 2003, Diebold et al. 2008), the operations nature of our research question requires special attention in data preparation in order to derive meaningful insights. In the following, we discuss the complications of our data set and outline the procedure we follow to prepare our data for further analysis.

Instead of time series of well-defined financial indexes and economic variables, we

observe resale prices of individual products. Therefore, before employing any econometric technique, we need to aggregate the individual product prices into a number of time series. A natural way to do this is at each month to aggregate the resale prices of all products of the same model together. This way we get one resale price time series per model. The complication, however, is that the used products of a particular model sold at different months could have very different age and usage mixes. If not accounted for, this change in the products' age and usage composition could lead to fluctuations that are purely due to depreciation effects, thus contaminating the true latent factors. Therefore, for a particular product model i , we control the age and usage level as the following:

1. We estimate the model's usage depreciation coefficient by running the following simple linear regression at each month t ($t = 1, 2, \dots, 156$):

$$p_{lt}^i = \alpha_{0t}^i + \alpha_{1t}^i \text{Age}_{lt}^i + \beta_t^i U_{lt}^i + e_{lt}^i. \quad (4.1)$$

Here p_{lt}^i is the price of model i 's l -th product sold at month t , U_{lt}^i is its usage level, Age_{lt}^i is its age, and $e_{lt}^i \sim N(0, \sigma_i^2)$ is the model error. Because the age of product l is defined as the difference (in months) between the resale time t and the production year of the model, Age_{lt}^i in (4.1) only varies by the model year of the product (e.g., whether it is a 2010 Ford Focus or a 2011 one, if model i under consideration is Ford Focus), and thus only takes a few discrete values. Therefore, we treat it as a categorical variable in (4.1).

2. We normalize the products' resale prices at month t to a pre-determined reference usage level using the corresponding estimates for β_t^i .²
3. We compute the average normalized prices for model i , P_{it} , at month t . This

²We use the median of all reliably estimated $\{\beta_t^i\}$ values for the normalization if the fitting of (4.1) at the month under consideration is poor or the reference usage level lies significantly higher or lower than the observed range of usage levels.

way we obtain a time series of the average price of this particular model's used products under the reference usage level.

4. We define the *age index*, \tilde{A}_{it} , at month t to be the average age of all model i products sold in month t .

We carry out this procedure for all models within our data set and obtain two time series for each product model: an average resale price P_{it} , and an age index \tilde{A}_{it} . Here the price time series P_{it} should be interpreted as the average resale price of model i at the reference usage level with age \tilde{A}_{it} . We will later use the age index \tilde{A}_{it} to control for the aging effects in our dynamic factor model.

We note that in practice, the usage depreciation (β_t) is usually affected by external factors such as the economic condition and the total product supply at month t . Therefore, running regression (4.1) at each month automatically controls for these external factors and yields depreciation coefficients that are conditional on all external factors at month t ($t = 1, 2, \dots, 156$). These depreciation coefficients hence allow us to preserve the influences of external factors on the normalized resale prices for further investigation. One potential danger with this month-by-month regression approach is that the resale prices within each month could be too noisy. However, as our data set contains the nationwide resale prices, it provides adequate sample size (hundreds or even thousands of items per month) for reliable monthly estimation and normalization.

We observe that the time series of average resale price P_{it} are unit root processes. Therefore, we follow the convention in macroeconometrics (Kose et al. 2003, Eickmeier 2009) and study the growth rates of the resale price of each model. In particular, we let Y_{it} represent the percentage growth of model i 's resale price in month $t + 1$ from month t :

$$Y_{it} = (\log P_{it+1} - \log P_{it}) \times 100. \quad (4.2)$$

We observe that the age indexes \tilde{A}_{it} of two consecutive months have close values. This is consistent with the intuition that the age mixes of model i 's used products changes gradually in time, as old models age and new models enter the used market. Therefore, we define the age index associated with the growth rate Y_{it} as the following

$$A_{it} = \frac{1}{2}(\tilde{A}_{it+1} + \tilde{A}_{it}). \quad (4.3)$$

Then Y_{it} should be interpreted as the percentage growth of model i 's resale price in month $t + 1$ from month t at age A_{it} and the reference usage level.

Finally, we de-mean both the growth rates and age indexes and denote the corresponding zero-mean processes as y_{it} and a_{it} , respectively.

After data preparation, we obtain a total of 235 models' monthly growth rates over 13 years (155 months) from 9 OEMs' 24 brands. Preliminary analysis shows that all growth rates processes are stationary. Our following analysis are based on y_{it} and a_{it} .

4.2.2 Descriptive statistics

Because different models have different launch and termination times, our panel data is unbalanced. We group the models into five segments based on the product functionality and an international classification system. For example, in the case of used vehicles, we can categorize the models into passenger cars, SUVs, pick-ups etc.; in the case of personal computers, we can categorize them into laptops, desktops, netbooks etc. Let N_s be the number of product models, or equivalently, number of time series within segment s ($s = 1, 2, \dots, 5$). Let O_s and B_s be the number of OEMs and brands within segment s , respectively. Table 4.1 gives the values of N , O , B and the descriptive statistics on the time series y_{it} within each segment.

Table 4.1: Descriptive statistics on the models' growth rates within each segment

This table presents the number of models, brands and OEMs within each segment, as well as the summary statistics of the length of growth time series. N_s denotes the total number of models/time series within segment s ($s = 1, 2, \dots, 5$). O_s and B_s represent the number of OEMs and brands within segment s . TS is the shorthand for time series.

Segment	N_s	O_s	B_s	Length of TS				
				Mean	Median	Std. dev.	Min	Max
1	100	9	23	113.7	129.5	36.7	39	155
2	20	7	8	138.0	148.0	25.4	72	155
3	22	9	15	86.1	77.0	36.5	39	153
4	81	9	24	100.3	99.0	35.8	29	155
5	12	7	10	120.2	132.0	34.5	59	155

4.3 Multi-level dynamic factors model

Our research interests lie in studying the comovements among different product models' price growth rates. We carry out the study employing a multi-level dynamic factor model (Moench et al. 2011, Moench and Ng 2011). Particularly, we establish a state space model within each of the five segments to study their respective segment-wide factors. Then we pool the five segment factors together and identify one industry-wide factor. There are two commonly used approaches in dynamic factor model literature: the state space method (Diebold et al. 2006, Kose et al. 2003, Koopman et al. 2011) and the principal component analysis (PCA) approach (Bai and Ng 2002, Eickmeier 2009, Forni and Reichlin 1996). We adopt the former because, first of all, it can easily handle unbalanced panel data (Durbin and Koopman 2001). Furthermore, it allows for structural modeling of the system and factor dynamics, and thus is particularly appealing when our goal is to make inferences on the used goods market. Lastly, the state space presentation provides an integrated framework that estimates the factor dynamics and other model parameters together. In the following, we first propose our model, and then discuss it in detail.

1. On the segment level, we assume that within segment s ($s = 1, 2, \dots, 5$), model

i 's price growth rate y_{it}^s is driven by a corresponding dynamic segment factor f_{st} ($i = 1, 2, \dots, N_s$):

$$y_{it}^s = \gamma_i^s(f_{st} + c_{st}) + \lambda_i^s a_{it} + e_{it}^s, \quad (4.4)$$

$$f_{st} = \tilde{\phi}_s f_{st-1} + \tilde{u}_{st}, \quad (4.5)$$

$$c_{st} = - \sum_{i=1}^{11} c_{s,t-i}, \quad (4.6)$$

$$e_{it}^s \sim NID(0, \delta_{si}^2), \quad \tilde{u}_{st} \sim NID(0, 1),$$

$$\text{cor}(\tilde{u}_{st}, e_{it}^s) = 0, \quad \text{cor}(e_{it}^s, e_{jt}^s) = 0 (i \neq j).$$

Here in (4.4) c_{st} is the seasonality pattern of segment s models' price growth rates, making f_{st} the de-seasonalized segment factor. We assume c_{st} to be deterministic³ throughout the observation horizon in (4.6). γ_i^s is product model i 's time-invariant loading on the corresponding segment factor $f_{st} + c_{st}$. Recall that a_{it} is the de-measured age index of used model i . It is used in (4.4) to control for the varying age index of the growth rates. We note that in (4.4)–(4.6), the only observed data is the time series of each model's price growth rate y_{it} and its de-measured age index a_{it} . All latent segment factors, the associated seasonality, measurement errors and state innovations are unobservable.

2. On the industry level, we assume that each segment factor f_{st} are driven by one

³We also fitted the model with stochastic seasonality: $c_{st} = - \sum_{i=1}^{11} c_{s,t-i} + \tilde{u}_{ct}^s$, where $\tilde{u}_{ct}^s \sim NID(0, \delta_c^2)$. We observe that δ_c^2 is small ($\sim 10^{-5}$) with large estimation error. Furthermore, our model parameter estimates and smoothed states are robust to the seasonality specification.

industry-wide factor:

$$f_{st} = \theta_s f_t + g_{st}, \quad (4.7)$$

$$f_t = \psi f_{t-1} + u_t. \quad (4.8)$$

$$g_{st} = \phi_s g_{st-1} + u_{st}, \quad (4.9)$$

$$u_t \sim NID(0, 1), \quad u_{st} \sim NID(0, \sigma_s^2),$$

$$\text{cor}(u_{st}, u_t) = 0, \quad \text{cor}(u_{st}, u_{lt}) = 0 (s \neq l).$$

Here u_t and u_{st} ($s = 1, 2, \dots, 5$) represent the industry-wide and segment-specific shocks, respectively.

We note that in the model above, both f_{st} in (4.4) and g_{st} in (4.7) pertain to a particular segment s ($s = 1, 2, \dots, 5$). The key difference between the two is that while the former accounts for *all* comovements among the product models' growth within segment s , the latter only captures the comovements among these products *after* the industry-wide factor f_t is controlled. As we will see later, both f_{st} and g_{st} play important roles in our study. To distinguish between the two, we call f_{st} the segment factor and g_{st} the controlled segment factor in the rest of the paper.

The multi-level dynamic factor model above is a linear Gaussian system. We now discuss the model specification in detail.

1. Both the industry and segment factors are autoregressive processes of order 1 (AR(1)). Processes with higher AR orders can also be considered. But we find that the estimated factors using the Kalman filter and smoother are robust to the model specification of the factor dynamics.
2. The idiosyncratic errors of each product model, e_{it}^s in (4.4), are strict white noise processes. We also fitted model with AR(1) idiosyncratic errors e_{it}^s , but the estimated AR coefficients are not significant and the results are similar to

those from the current model.

3. *Factor orthogonality:* The segment factors f_{st} are orthogonal to the model-specific idiosyncratic errors e_{it}^s . Similarly, the industry factor f_t is orthogonal to the controlled segment factors g_{st} . In addition, the idiosyncratic errors of different product models are assumed to be orthogonal to each other. So are the controlled segment factors. These orthogonality assumptions are conventional in state space models (Kose et al. 2003, Otrok and Whiteman 1998, Gregory et al. 1997), partly for identification purposes.
4. *Identification requirements:* Both the segment factors f_{st} and the industry factor f_t are assumed to have unit unconditional variances so that the magnitude of the loadings γ_i^s in (4.4) and θ_s in (4.7) are identifiable. To ensure the identifiability of their signs, we further impose the constraint that the first elements of the loading vectors γ^s and θ are positive.

If we combine (4.4) and (4.7) on the segment and industry levels together, we have ($i = 1, 2, \dots, N_s$, $s = 1, 2, \dots, 5$)

$$\begin{aligned}
 y_{it}^s &= \gamma_i^s(\theta_s f_t + g_{st} + c_{st}) + \lambda_i^s a_{it} + e_{it}^s \\
 &= \chi_i^s f_t + \gamma_i^s(g_{st} + c_{st}) + \lambda_i^s a_{it} + e_{it}^s.
 \end{aligned}
 \tag{4.10}$$

This is the form of the multi-level state space model used in Kose et al. (2003), Stock and Watson (2009), Gregory et al. (1997) with the counterparts of χ_i^s assumed to be free parameters. Compared with the earlier “bottom-up” model that works from the segment level up, (4.10) provides a “top-down” model that directly extracts the industry factor, the controlled segment factors that are orthogonal to the industry factor, and the idiosyncratic errors that are orthogonal to all factors. One advantage of the formulation in (4.4) – (4.9) over (4.10), as noted by Moench et al. (2011),

is the parsimony of the model by imposing the constraint of $\chi_i^s = \theta_s \gamma_i^s$. Given the dimension of our data set (155×235), the computational task is formidable should formulation (4.10) be adopted. A more important reason of our employing the formulation in (4.4)–(4.9), however, is that besides the controlled segment factors g_{st} , we are interested in obtaining the segment factors f_{st} as well. Specifically, factors f_{st} provide a convenient way to control for the segment-level comovements when we carry out the OEM and brand analysis in Section 4.5. We note that in (4.10), the controlled segment factor and its associated seasonality always appear simultaneously and it is thus tempting to combine them into one latent factor. However, because the segment-specific comovements are highly seasonal, by extracting the seasonality part c_{st} , we obtain much more accurate estimates for the AR(1) coefficients $\tilde{\phi}_s$ in (4.5). Furthermore, incorporating the seasonality as a latent state instead of de-seasonalizing the original observations by differencing allows us to explicitly estimate each segment's c_{st} and obtain insights.

For ease of exposition in the rest of the paper, we write (4.4) and (4.7) in matrix form. On the segment level, we have

$$Y_t^s = \Gamma_s(f_{st} + c_{st}) + \text{diag}(\Lambda_s)A_t^s + e_t^s, \quad (4.11a)$$

$$(1 - \tilde{\phi}_s L)f_{st} = \tilde{u}_{st}, \quad (4.11b)$$

$$c_{st} = - \sum_{i=1}^{11} c_{s,t-i}, \quad (4.11c)$$

$$\tilde{u}_t = I_{N_s}, \quad e_t^s = \Delta_s, \quad s = 1, 2, \dots, 5, \quad (4.11d)$$

where $Y_t^s = (y_{1t}^s, y_{2t}^s, \dots, y_{N_s t}^s)$, $A_t^s = (a_{1t}^s, \dots, a_{N_s t}^s)$, $\Gamma_s = (\gamma_1^s, \dots, \gamma_{N_s}^s)$, $\Lambda_s = (\lambda_1^s, \dots, \lambda_{N_s}^s)$, $e_t^s = (e_{1t}^s, \dots, e_{N_s t}^s)$, $\Delta = \text{diag}(\delta_{s1}^2, \dots, \delta_{sN_s}^2)$, I_n is the n -by- n identity matrix, L is the lag operator ($Lf_t = f_{t-1}$).

On the industry level:

$$f_t^S = \Theta f_t^I + G_t, \quad (4.12a)$$

$$(1-\psi L)f_t^I = u_t^I, \quad (4.12b)$$

$$(I_5 - \text{diag}(\Phi)L)G_t = u_t^S, \quad \text{var}(u_t^S) = \Sigma, \quad (4.12c)$$

where $f_t^S = (f_{1t}, \dots, f_{5t})$, $f_t^I = f_t$, $G_t = (g_{1t}, \dots, g_{5t})$, $\Theta = (\theta_1, \dots, \theta_5)$, $\Phi = (\phi_1, \dots, \phi_5)$, $u_t^S = (u_{1t}, \dots, u_{5t})$, $\Sigma = \text{diag}(\sigma_1^2, \dots, \sigma_5^2)$.

4.4 Industry and segment factors

We estimate our two-level state space model using a multi-step approach. We start with the state space models (4.11) within each segment and obtain the estimated segment factors \hat{f}_{st} using Kalman smoother. Then we treat the estimated \hat{f}_{st} as “data” in the industry state space model (4.12) and estimate the industry and controlled segment factors. Given the dimension of our data set (235 times series in total), this sequential estimation method greatly reduces the computational burden than if estimated the entire model at once. Similar approach is also used in a series of papers on yield curve dynamics such as Diebold et al. (2008) and Diebold and Li (2006). In particular, Diebold et al. (2008) establish a multi-level factors model on the bond prices issued by different governments. In the first step, they estimate the latent factors for each country using Kalman smoother. In the second step, they use the set of estimated country factors as the “data” to compute the global factor.

Within segment s ($s = 1, 2, \dots, 5$), the parameters to be estimated are: the loading vectors Γ_s and Λ_s , the N_s variances of marginal distribution of the idiosyncratic errors e_t^s , and the AR coefficient $\tilde{\phi}_s$. These are a total of $3N_s + 1$ parameters. By Table 4.1, this number ranges from 37 to 301 in the five segments we have. On the industry level, we need to estimate the segment loadings Θ , the AR coefficients Φ and ψ ,

and the variance matrix Σ_S of u_t^S . That is, 26 free parameters in the second step. We estimate all model parameters under the maximum loglikelihood principle. More explicitly, we maximize the logarithmic of the joint probability of all observed growth rates y_{it} , which has an analytical expression based on our models (4.11) and (4.12) (see e.g. Durbin and Koopman (2001)). Solving this optimization problem presents a challenging computational task. We use the package by Ghalanos and Theussl (2012) written in the R environment. To ensure the robustness of our results to the initial conditions of our optimization, we estimate the every model using 30 times with randomized initial conditions and use the estimates with the maximum loglikelihood as the reference. We found that the point estimates of the model parameters from the rest 29 runs are not significantly different from the reference estimates at 95% significance level using z-tests. We also found that the segment and industry factors from the rest 29 runs are highly correlated with the factors under the reference run, with correlation over 0.9. These results confirm that our results are robust to the initial conditions of the algorithm and thus likely to be the global optimum over the feasible region. In the following, we report the results of the reference run. All the factors are estimated using Kalman smoother under these parameters.

4.4.1 Latent dynamic factors

This section studies the industry factor f_t^I and controlled segment factors G_t in (4.12). We consider the controlled segment factors rather than the segment factors because the former is more informative on segment-level comovements when the industry factor is also examined.

Table 4.2 presents the parameter estimation results for (4.12). We see that all segments have significant loadings on the industry factor. Recall that the shocks to industry factor u_t has unit variance. Thus, these large loadings indicate that there are remarkable comovements among different segments. In addition, comparing with

the estimated variance of the controlled segment factor shocks u_{st} ($s = 1, 2, \dots, 5$), we see that industry-wide shocks account for a significant portion of the variability in segment factors f_{st} .

Table 4.2: Loadings of the industry factor and variance of shocks to controlled segment factors

	Seg. 1	Seg. 2	Seg. 3	Seg. 4	Seg. 5
Industry loading θ	0.76 (0.07)	0.69 (0.08)	0.58 (0.06)	0.81 (0.07)	0.79 (0.06)
Controlled segment factor shocks σ^2	0.36 (0.06)	0.43 (0.07)	0.24 (0.03)	0.25 (0.05)	0.16 (0.04)

Figure 4.1 presents the estimated industry and controlled segment factors with the significance bands at 1.64 standard deviations. The red point corresponds to Sept. 2008, when Lehman Brothers filed bankruptcy. By figure (4.1a), we see that at the height of the 2008 financial crisis, the estimated industry factor plunged over 4.5%, the most substantial drop over the entire 13-year observation horizon. This indicates that the entire class of the used goods under consideration experienced a remarkable price decline during financial crisis. In particular, from the segment loadings in Table 4.2, segments 1, 4 and 5 have the strongest comovements with the industry factors and are thus affected the most among all five segments.

Figures (4.1b – 4.1f) present the controlled segment factors g_{st} ($s = 1, \dots, 5$) in (4.7) and (4.12). By construction, g_{st} is orthogonal to the industry factor f_t^I . Therefore, it captures the latent factors peculiar to each segment, thus allowing further investigation and comparison of the segment-specific price dynamics. While we cannot disclose the segment names due to confidentiality reasons, we note that segments 1 and 3 are mostly for individual uses with segment 3 products being fancier and more pricey; segments 2 is mainly for industrial/commercial uses, segment 4 is primarily for family uses, and segment 5 for both family and commercial uses. In Figure 4.1, we observe significant controlled factors for all segments except the hybrid segment

(segment 5). Along with the low variance value of segment 5 shocks reported in Table 4.2, this indicates that the latent factor for segment 5, f_{5t} , is not significantly different from the industry factor. This observation is likely due to the hybrid nature of segment 5 products. Controlled factors of the rest four segments all have distinctive dynamics, especially in the way they respond to the financial crisis. The controlled factor of ordinary individual goods (segment 1) suffer a steep and prolonged drop during the financial crisis. Compared with other segments and particularly the pricey individual goods (segment 3), segment 1 controlled factor takes the longest time to recover to the original level before crisis. The controlled factor of commercial and family goods (segments 2 and 4), on the other hand, are hit by the crisis five months earlier than all other segments.

Figure 4.1 sheds some light on the pattern of industry-wide and segment-specific comovements. In practice, however, we are also interested in the comovements among different product models within the same segment, that is, the segment factor f_{st} itself ($s = 1, \dots, 5$). In addition, we'd like to see how the industry and controlled segment factors compare with each other within each segment. Therefore, we plot the *loaded* industry factors $\theta_s f_t^I$, the controlled segment factors g_{st} , and the segment factor f_{st} in Figure 4.2. By equation (4.7), $f_{st} = \theta_s f_t^I + g_{st}$. Hence, these plots visualize the respective contributions of the industry and controlled factors to the price growth fluctuation within each segment. We observe that the segment-specific factors of ordinary individual goods and commercial goods (segment 1 and 2) are of the same magnitude as the corresponding loaded industry factors, whereas the industry factor dominates the controlled segment factor within other segments. Furthermore, combining the two factors together, we see that the price of individual goods (segments 1 and 3) experiences a sudden decline in Sept. 2008, while the price of commercial and family goods (segment 2 and 4) drops five months earlier. This difference in the response times might be due to the fact that the 2008 financial crisis was developed

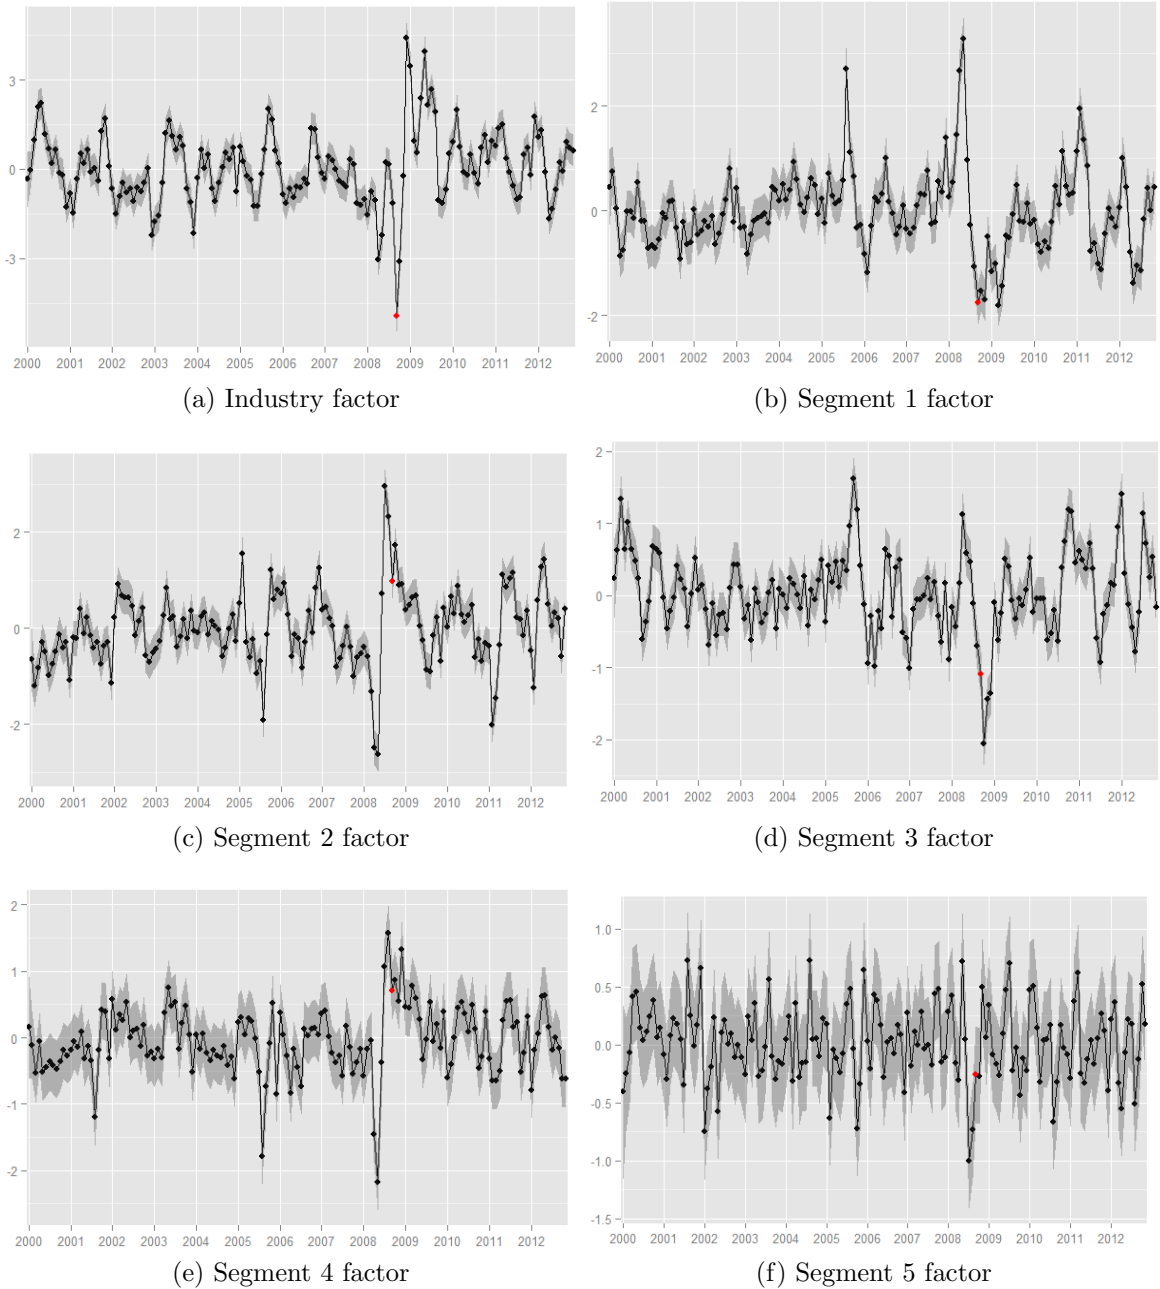


Figure 4.1: Estimated industry and controlled segment factors. 1.64 standard deviation (i.e., 90% significance level) band. The red dot corresponds to Sept. 2008, the bankruptcy of Lehman Brothers.

from subprime mortgage crisis and, thus, influenced companies financing and families first before the critical point when Lehman Brothers' bankruptcy shocked the entire economy. In addition, among all segments, the ordinary individual goods sees the biggest and longest decline following the financial crisis. This is also consistent with

macroeconomic reports on the lagging consumer confidence and high unemployment rates in the long recession period afterwards. In particular, because the ordinary individual goods are the least expensive segment of the entire industry, the majority of their buyers earn low/medium wages and are thus more likely to be affected by the crisis.


4.4.2 Persistence properties of the factors

The persistence of a time series characterizes how autocorrelated the series is. A crude way of describing a univariate time series' persistence is to use the autocorrelation plot. The slower the autocorrelation function decays, the greater the persistence. By Figure 4.1, we can already see strong persistence in the industry and segment factors. In this section, we measure the persistence explicitly by the AR(1) coefficients $\tilde{\phi}_s$, ψ and ϕ_s ($s = 1, \dots, 5$) in (4.5), (4.8) and (4.9). Similar approaches have also been used in e.g. Kose et al. (2003), Eickmeier (2009).


Table 4.3: Persistence of latent factors

<i>Industry factor</i>					
AR(1) coefficient ψ	0.62 (0.07)				
<i>Controlled segment factors</i>					
	Seg. 1	Seg. 2	Seg. 3	Seg. 4	Seg. 5
AR(1) coefficient ϕ_s	0.66 (0.07)	0.58 (0.07)	0.57 (0.08)	0.49 (0.09)	0.2 (0.1)
<i>Segment factors</i>					
	Seg. 1	Seg. 2	Seg. 3	Seg. 4	Seg. 5
AR(1) coefficient $\tilde{\phi}_s$	0.56 (0.07)	0.52 (0.07)	0.54 (0.10)	0.53 (0.07)	0.53 (0.08)


By the upper and middle panels of Table 4.3, both the industry and controlled segment factors are quite persistent. This result indicates that whenever a shock occurs on the industry or segment level, its effect on the corresponding factor will last several months before fading out. Specifically, for individual, commercial and




(a) Segment 1




(b) Segment 2



(c) Segment 3



(d) Segment 4



(e) Segment 5

Figure 4.2: Loaded industry factor (solid black), controlled segment factor (dashed black) and segment factor (dotted red).

family goods (segments 1 to 4), industry-wide shocks and segment-specific shocks are equally persistent in terms of their effects on the resale prices. For goods in the hybrid segment (segment 5), however, industry-wide shocks are much more persistent.

The lower panel of Table 4.3 presents the persistence of the segment factors f_{st}

directly. As expected from Figure 4.2 and the upper panels of Table 4.3, the segment factors are highly persistent.

The ultimate goal of our persistence study is to evaluate the impact of a random shock on the latent factors, and consequently the resale prices, over time. There are roughly two types of shocks based on their scope of influence: A “local” shock that only pertains to a subset of segments (e.g., supply/demand shocks of commercial goods); and a “global” shock (e.g., financial crisis, natural disaster) that affects the entire used goods market. Apparently, in cases of a “local” shock, we can evaluate its impacts into the future using the persistence properties of the corresponding controlled segment factors. In cases of a “global” shock, however, it can be particularly challenging to decouple the shock into the orthogonal industry- and segment-level shocks, u_t and u_{st} ($s = 1, \dots, 5$), defined in (4.8) and (4.9). Indeed, as an *ex post* analysis, Figure 4.1 illustrates how distinctive u_t and u_{st} can be in case of a financial crisis. As a result, evaluating the future impacts of a “global” shock via the industry and controlled segment factors becomes difficult. This complication, however, could be easily avoided by using persistence properties of the segment factors directly. Therefore, Table 4.3 provides the necessary information to estimate the future impacts of a random shock on the resale prices.

4.4.3 Variance decomposition

In our multi-level dynamic factors model, we assume that (i) there is strong comovement among the product models within the same segment, which is captured by the segment factor; and (ii) there is strong comovement among different segments, which is captured by the industry factor. Section 4.4.1 illustrates that there is indeed strong comovement among the segment factors. In this section we further quantify the extent to which each segment factor is driven by the common industry factor, as well as the comovement level among the product models within the same segment.

To do this, we need to account for the way different segments (product models) loads the corresponding latent factors and their own idiosyncratic errors. Variance decomposition is a standard approach in factor analysis that serves this purpose (Kose et al. 2003, Gregory et al. 1997, Otrok and Whiteman 1998).

We assume that the controlled segment factors are orthogonal to the industry factor, and that the product-specific errors are orthogonal to the segment factors in (4.12) and (4.11). The estimated factors from the finite sample, however, could be correlated even if they are not in the population. Indeed, we observe weak correlation ($|\rho| \in (10^{-5}, 0.3)$) between (g_t^S, f_t^I) and (f_{st}, e_{it}^s) . Therefore, we orthogonalize g_t^S and f_t^I by regressing the segment factors on the industry factor f_t^I , and replacing the loadings with the regression coefficient and the controlled segment factors with the regression error (Diebold et al. 2008). Similarly, we orthogonalize f_{st} and e_{it}^s by regressing product i 's growth rate on the seasonal factor $(f_{st} + c_{st})$ and the age index a_{it} , replacing the loadings and the idiosyncratic errors with the regression coefficients and residuals. Then on the industry level, by (4.7) we have

$$\text{var}(f_{st}) = \theta_s^2 \text{var}(f_t) + \text{var}(g_{st}). \quad (4.13)$$

On the segment level, by (4.4) we have

$$\text{var}(y_{it}^s) = (\gamma_i^s)^2 \text{var}(f_{st}) + \text{var}(e_{it}^s). \quad (4.14)$$

We note that in the measurement equation (4.4), the seasonality c_{st} and the age index a_{it} are both deterministic, thus not entering (4.14).

By the state equations (4.5), (4.8) and (4.9), the unconditional variances of the industry and segment factors are

$$\text{var}(f_t) = \frac{1}{1 - \psi^2}, \quad \text{var}(f_{st}) = \frac{1}{1 - \tilde{\phi}_s^2}. \quad (4.15)$$

Then the proportion of the segment s variability that can be explained by the industry factor is ($s = 1, \dots, 5$)

$$w_s = \frac{\theta_s^2/(1 - \psi^2)}{\theta_s^2/(1 - \psi^2) + \text{var}(g_{st})}. \quad (4.16)$$

The proportion of the product model i 's variability that can be explained by the segment factor f_{st} is ($i = 1, \dots, N_s$)

$$w_i^s = \frac{(\gamma_i^s)^2/(1 - \tilde{\phi}_s^2)}{(\gamma_i^s)^2/(1 - \tilde{\phi}_s^2) + \text{var}(e_{it}^s)}. \quad (4.17)$$

Table 4.4 presents the variance decomposition results. The upper panel shows that the industry factor accounts for a significant portion of each segment factor's variation, confirming our earlier observation from Figure 4.2. The lower panel reports the range of the segment factor's contributions to each product's price growth variability. We see that, despite the large number of product models within segment 1 and 4 (100 and 81, respectively, see Table 4.1), the latent factors still explain remarkable fractions of individual product's variation. This result indicates that there are strong comovements within these two segments, i.e., ordinary individual goods and family goods. In comparison, although segment 3 of pricey individual goods includes only 22 models, the explanatory power of the corresponding latent factor is much lower than other factors, implying stronger heterogeneity among the pricey individual products. This observation is consistent with our intuition that fancy and up-scale products tend to distinguish from each other more than ordinary ones, thus attracting different consumers and have different resale price dynamics.

4.5 OEM analysis

Section 4.3 examines the comovements among products with similar functionalities by grouping them into a few segments. In this section, we study the comovements of products with the same brand/OEM. In practice, one brand/OEM can cover sev-

Table 4.4: Variance decomposition

For each segment factor/product price, we decompose its variation into two parts: one from the industry/segment-level comovement, the other from the segment/product-specific errors. The upper panel presents the proportion of each segment factor's variation that can be explained by the industry factor. The lower panel presents the proportion of each product's price growth variation that can be explained by the corresponding segment factor. ("perc." is a shorthand for percentile.)

<i>Industry level</i>					
	Seg. 1	Seg. 2	Seg. 3	Seg. 4	Seg. 5
w_s (%)	60.6	55.7	71.8	82.3	92.3
<i>Segment level</i>					
	5th perc.	25th perc.	Median	75th perc.	95th perc.
Seg. 1: w^1 (%)	6.2	13.7	20.8	37.1	52.2
Seg. 2: w^2 (%)	15.2	18.6	31.6	48.7	81.4
Seg. 3: w^3 (%)	7.2	9.7	14.3	23.5	28.4
Seg. 4: w^4 (%)	5.8	13.7	23.6	34.2	45.5
Seg. 5: w^5 (%)	26.3	29.2	32.6	42.1	58.9

eral segments. Take the automotive industry for example. OEM Toyota has two brands, Toyota and Lexus, each of which produces a range of models in passenger car, SUV and pick-up segments. From earlier analysis, we already know that there is strong comovement among products with similar functionalities. Therefore, to avoid confounding effects from the segment dimension, we examine the OEM- and brand-level comovements *after* the earlier identified segment factors have been controlled. The meaning of these OEM- and brand-level comovements then is two-fold. First, the price dynamics of all products within a particular segment break into smaller blocks in brands and/or OEMs. Second, after controlling for the segment factors, products of the same brand and/or OEM but different segments still exhibit strong comovements. In other words, we aim to answer the following questions: Are there OEM- and brand-specific factors driving the product's price dynamics on top of the segment factors? And if so, more importantly, how would an incident such as product recall and brand termination affect the OEM and brand factors on the resale prices?

In the following, we will first propose our model, and then discuss the results.

4.5.1 Model on OEM and brand factors

As we are now interested in the OEM- and brand-level latent factors, we index the growth rate of product model i 's resale price by its brand b and OEM o in addition to the segment s . We consider 6 large OEMs observed in our data set, each of which produces 2 to 5 brands with 14 to 52 product models. Notation-wise, we write y_{it}^s as y_{it}^{sob} to reflect the brand and OEM information. Furthermore, let \hat{f}_{st} and \hat{c}_{st} denote the estimated segment factor and its seasonality from the multi-level model ($s = 1, \dots, 5$). Our measurement equation for the latent factors is then

$$y_{it}^{sob} = \eta_i^o h_{ot} + \eta_i^{ob} h_{bt}^o + \tilde{\lambda}_i a_{it}^{sob} + \chi_i (\hat{f}_{st} + \hat{c}_{st}) + \xi_{it}^{sob}, \quad (4.18)$$

$$\xi_{it}^{sob} \sim NID(0, \tilde{\sigma}_i^2).$$

where $o = 1, \dots, 6$, $s = 1, \dots, 5$, $b = 1, \dots, B_o$, $i = 1, \dots, N_o$. B_o and N_o denote the number of brands and product models manufactured by OEM o .

In the measurement equation (4.18), we used the factors $\hat{f}_{st} + \hat{c}_{st}$ estimated from (4.11) as exogenous variables. This approach is similar to the widely used factor-augmented regression and forecasting (e.g. Stock and Watson 2001, Ludvigson and Ng 2007, Bai and Ng 2008, Moench and Ng 2011). In particular, Bai and Ng (2006) prove that the least squares estimates from these factor-augmented regressions are consistent and asymptotically normal, and that the first-stage estimation of the factors does not affect the consistency of the second-stage parameter estimates. Our OEM model above extends some of the existing work in factor-augmented regression in that we further impose a two-level dynamic factor model on the regression errors.

Specifically, we assume AR(1) dynamics for the OEM and brand factors h_{ot} and h_{bt}^o :

$$h_{ot} = \kappa^o h_{ot-1} + u_{ot}, \quad (4.19a)$$

$$h_{bt}^o = \kappa_b^o h_{bt-1}^o + u_{bt}^o, \quad (4.19b)$$

where both u_{ot} and u_{bt}^o are orthogonal strict white noise processes with unit variance.

4.5.2 OEM and brand factors

We estimate the model parameters and factors as before using the MLE approach with 30 randomized initial parameter values. We find that the parameter estimates for the loadings, η_1 , η_2 and χ , the coefficient of the age index, $\tilde{\lambda}$, and the idiosyncratic error variance $\tilde{\sigma}_i^2$ are robust to the initial parameter values of the optimization. The OEM factors estimated from different runs are also highly correlated. The brand factors under each OEM, however, do not always exhibit high correlation. This result suggests that after controlling for the segment factors, while there is still strong comovement among products of the same OEM, the brand-level comovement is not always present. In other words, these brands are not significantly different from each other after controlling for the segment factors, thus could be captured by a common OEM factor. In the following, we report the result with the highest log-likelihood of all runs.

Table 4.5 shows the estimates for AR(1) coefficients κ^o and κ_b^o in (4.19) ($o = 1, \dots, 6$, $b = 1, \dots, B_o$). Only results of factors that are robust to the initial values are reported, otherwise indicated with a “–” sign. Comparing with the persistence results in Table 4.3, we see that most factors are much less persistent than the segment factors, except for OEMs 1 and 4. This result suggests that for most products, the segment factors capture their most persistent, or low frequency comovements. Consequently, when an OEM- or brand-level shock occurs, most OEMs and brands

recover quickly to the prior-shock state while OEMs 1 and 4 endure a lasting impact. This high persistence of OEM 1 and 4's factors might indicate that unlike other brands and OEMs, these two OEMs have developed their own particular images among the consumers.

Table 4.5: Persistence of OEM and brand factors

We report the AR(1) coefficients of the OEM and brand factors. κ^o is the AR(1) coefficient for OEM o . κ_i^o ($i = 1, 2, \dots, 5$) is the AR(1) coefficient for OEM o 's i -th brand. Entries “–” mean that the corresponding brand factors are not robust to the optimization initial conditions, thus their results not presented. Blanks are used when the corresponding brands do not exist. For example, entries for OEM 1's brands 4 and 5 are left blank as OEM 1 only produce three brands.

	κ^o	κ_1^o	κ_2^o	κ_3^o	κ_4^o	κ_5^o
OEM 1	0.52(0.08)	–	–	0.4(0.1)		
OEM 2	-0.1(0.1)	–	–	0.4(0.1)		
OEM 3	0.3(0.1)	–	–	-0.1(0.1)	–	-0.82(0.09)
OEM 4	0.6(0.1)	0.3(0.1)	0.4(0.1)			
OEM 5	-0.2(0.1)	0.5(0.1)	–			
OEM 6	0.3(0.1)	-0.1(0.1)	-0.65(0.1)			

The state space model (4.18-4.19) on different OEMs provides a convenient way of studying how OEM- and brand-level shocks affect the products' resale price dynamics. This is a particularly relevant question in business operations. In the following, we examine the following two incidents: product recalls and brand termination.

Over our observation horizon, OEM 4 announced a large-scale product recall on almost 50% of its brand 1 product models. Given the persistence results in Table 4.5, we expect the impact of this recall to last a relatively long period. However, to what extent is the brand 1 factor influenced? Are there significant spill-over effects so that the brand 2 factor and even the OEM factor also suffered? These are the questions of particular interest and we answer them next. By the estimated OEM factor in Figure 4.3a, this recall has not caused a wide-spread resale price decline for all products from OEM 4. The impact on the brand being recalled, however, is significant and lasting as shown in Figure 4.3b. The brand 1 factor dropped more than 4% at the

recall announcement, indicating a sizable decline in the resale prices of all brand 1 products. The brand 2 products, on the other hand, were only slightly affected, if at all (Figure 4.3c). This result implies that there is little spill-over effects of the large-scale recall on brand 1. This might be linked to our earlier observation that brand 1 and 2 products are driven by separate brand factors (Table 4.5), suggesting that they exhibit significantly different price dynamics and probably perceived distinctively by the consumers.

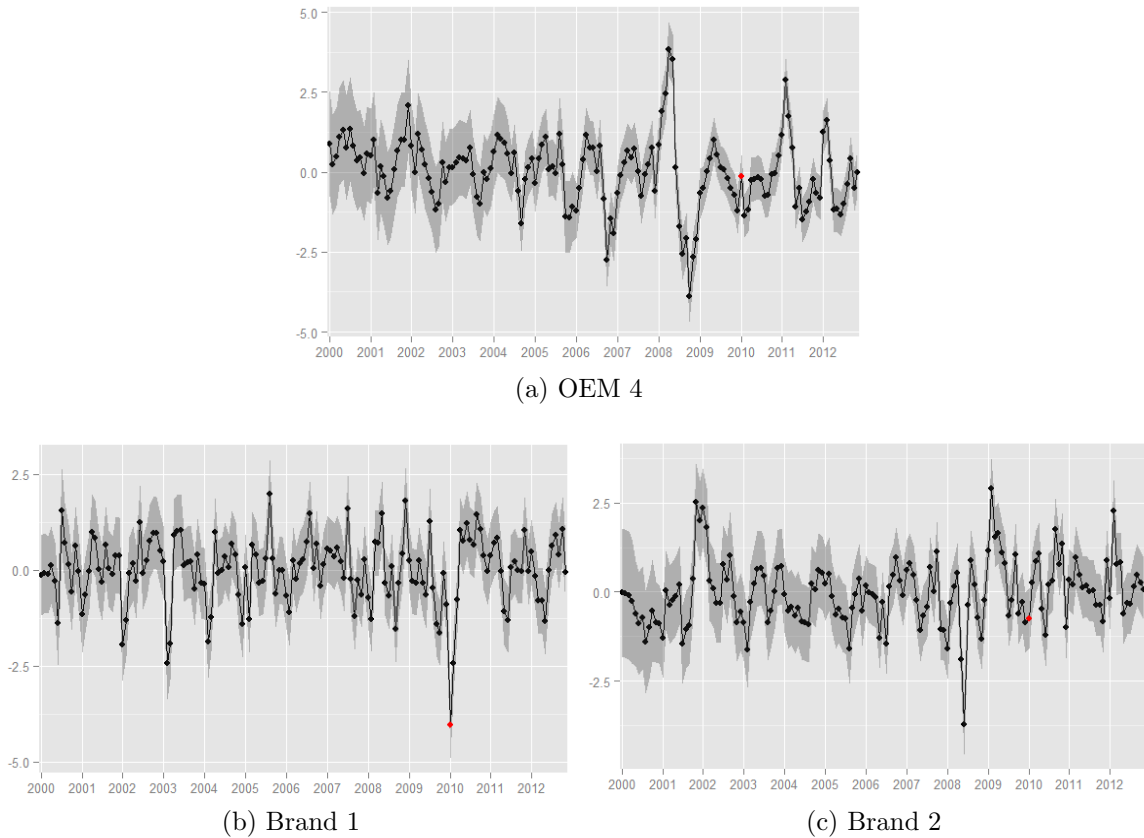


Figure 4.3: Impacts of product recalls. (a-c) plot OEM 4's OEM and brand factors with 1.64 standard deviation (i.e., 90% significance level) band. In Jan. 2010 (red dot), several brand 1 product models were recalled.

During 2000–2012, two OEMs terminated their brands. OEM 1 terminated production for its brand 2 around the end of 2011, while OEM 2 terminated its brand 3 around the end of 2010. We are interested in examining how these brand termination affects the products' price dynamics. On one hand, the termination might result

in decreasing resale prices as when all operations related to a particular brand are stopped, repairing and servicing its products becomes harder for the buyers. This additional maintenance cost could then be priced in when these products are sold. On the other hand, although the brand is terminated, the OEM still operates normally. In fact, when OEMs terminate a brand in practice, they always accompany the termination announcement with detailed product repair and maintenance support to alleviate such concerns from the consumers. If effective, this policy might lead to small additional costs in product maintenance, if at all. Therefore, whether the brand termination has significant impacts on the OEM and corresponding brand factors is not clear, and thus worthwhile to understand.

Figure 4.4 shows the estimated OEM factors of OEM 1 and 2, and the estimated brand 3 factor of OEM 2. Because we do not observe significant comovement among OEM 1's brand 2 products, we do not show the corresponding brand factor. We note that the large uncertainty of the estimated brand factor at the early period (2000–2004) is due to a lack of data within that time frame. From Figure 4.4, we cannot detect any unusual movements around the termination time in the OEM and brand factors. This indicates that when the OEM operates normally, brand termination does not have a material impact on the resale prices of related products or products of other brands.

4.6 Concluding Remarks

In this paper we study the comovements of resale prices of a particular type of used durable goods using the US data during 2010:1 – 2012:12. In particular, we identify the comovements of products within each functionality segment, across the entire industry, and within each brand and OEM after the segment-level comovements are controlled for. At the best of our knowledge, this is the first attempt of applying dynamic factor models developed in macroeconomics to a problem with an operational




Figure 4.4: Impacts of brand termination. (a-c) plot OEM 2's OEM and brand factors with 1.64 standard deviation (i.e., 90% significance level) band. Around the end of 2011 and 2010 (red dots), OEM 1 terminated brand 2 and OEM 2 terminated brand 3, respectively.

nature.

In the industry and segment-level analysis, we find that despite the heterogeneity in brand, OEM, etc., the resale prices of the products are strongly affected by the corresponding segment factor. This dependence on the segment factor, or equivalently, the level of within-segment comovement, however, varies with segment. In particular, the segment with fancy and pricey products exhibits weaker comovement than others. Furthermore, we illustrate that all segment factors are quite persistent, indicating lasting impacts from random shocks.

In the OEM analysis, we report that there are strong comovements among products of the same OEM after their corresponding segment factors are controlled for.

However, given a particular OEM, its products under different brands do not always distinguish from each other. We show that the OEM and brand factors are much less persistent than the segment factors, consistent with our intuition. Finally, we find that brand termination does not have a material impact on the corresponding products' resale prices, while a big recall could induce a significant price drop.

CHAPTER V

Conclusion

In this dissertation, I study three types of uncertainties in industrial practice: the demand uncertainty, the earnings uncertainty and the external market uncertainty. In particular, Chapter II prices the demand uncertainties in the just-in-time outsourcing between an OEM and a CM; Chapter III propose a model for the provider to manage his earnings uncertainties using contracts in managed print services; Chapter IV aims to understand the uncertainties of the external market trend and market responses using resale prices of used durable goods.

Our study of demand risk pricing in just-in-time outsourcing is one of the first attempts to price demand risks in component outsourcing in JIT literature. We account for the flexible production capacity feature of component outsourcing and show that when the outsourcing demand is positively correlated with the agent's existing business, the higher risk it carries, the more the outsourcing costs the CM and benefits the OEM. One possible extension of this study is incorporating the outsourcing demand as OEM's decision variable depending on its in-house production capacity, inventory level, etc. This will lead to an optimization problem with the decision being the optimal capacity/inventory from the risk perspective. Another interesting direction to pursue is extending the model to scenarios where the OEM's order arrives randomly over the contract horizon. Intuition has that price of demand

risks becomes higher as there are now two sources of randomness: the uncertainty in demand quantities and the uncertainty of the order placement.

In the empirical study on the managed print services, we examine contractual interactions between the managed print service provider and customers using proprietary data sets provided by Xerox. We develop two models: one on the customer's service payments, and the other on the provider's selection among the set of contracts acceptable to the customer obtained from the first model. To the best of our knowledge, this paper is the first in-depth study of pricing in managed services. Our approach provides a new way of tackling information asymmetry between service providers and customers under inelastic service usage. It offers a viable construct for studies on managed services for institutional customers.

Chapter IV examines the price dynamics of resale prices of a type of used durable goods using dynamic factor models. We identify the underlying market trends of the industry of interest and quantify the impact of external economic shocks such as the financial crisis. Furthermore, we study the impact of how certain operational decisions of a company's senior management could impact the resale prices of the products. Specifically, we find that, while a large-scale product recall leads to a significant drop in product prices, this price-drop effect is rather "local" and does not spill over to other brands produced by the same manufacturer. We also find that the termination of a brand does not have a material impact on the corresponding products' resale prices, possibly due to the fact that the OEM behind the brand still exists and, thus, the maintenance services to those used products are not affected by the brand termination.

BIBLIOGRAPHY

BIBLIOGRAPHY

- Ansari, A., B. Modarress. 1990. *Just-In-Time Purchasing*. The Free Press, New York.
- Artzner, P., F. Delbaen, J.M. Eber, D. Heath. 1999. Coherent measures of risk. *Mathematical Finance* **9(3)** 203–228.
- Bai, Jushan, Serena Ng. 2002. Determining the number of factors in approximate dynamic factor model. *Econometrica* **70** 191–221.
- Bai, Jushan, Serena Ng. 2006. Confidence intervals for diffusion index forecasts and inference with factor-augmented regressions. *Econometrica* **74** 1133–1150.
- Bai, Jushan, Serena Ng. 2008. Forecasting economic time series using targeted predictors. *Journal of Econometrics* **146** 304–317.
- Berry, Steven T. 1994. Estimating discrete-choice models of product differentiation. *RAND Journal of Economics* **25(2)** 242–262.
- Breitung, J., S. Eickmeier. 2006. Dynamic factor models. *Modern Econometric Analysis*. Springer Berlin Heidelberg.
- Cachon, G.P. 2004. The allocation of inventory risk in a supply chain: push, pull, and advance-purchase discount contracts. *Management Science* **50(2)** 222–238.
- Cachon, G.P., M. Fisher. 2000. Supply chain inventory management and the value of shared information. *Management Science* **46(8)** 1032–1048.
- Caldentey, Rene, Martin Haugh. 2006. Optimal control and hedging of operations in the presence of financial markets. *Mathematics of Operations Research* **31** 285–304.
- Cameron, A. Colin, Pravin K. Trivedi. 2005. *Microeconometrics: Methods and applications*. Cambridge University Press.
- Chintagunta, Pradeep K., Jean-Pierre Dubé, Vishal Singh. 2003. Balancing profitability and customer welfare in a supermarket chain. *Quantitative Marketing and Economics* **1** 111–147.
- Chod, J., N. Rudi. 2005. Resource flexibility with responsive pricing. *Operations Research* **53(3)** 532–548.
- Chod, J., N. Rudi, J.A. Van Mieghem. 2010. Operational flexibility and financial hedging: Complements or substitutes. *Management Science* **56** 1030–1045.
- Chu, Chenghuan Sean, Phillip Leslie, Alan Sorensen. 2008. Nearly optimal pricing for multiproduct firms. Working Paper 13916, National Bureau of Economic Research.
- Cochrane, John. H. 2005. *Asset pricing*. Princeton University Press.
- Cohen, Morris A., Narendra Agrawal, Vipul Agrawal. 2006. Winning in the aftermarket. *Harvard Business Review* **May** 129–138.
- CompTIA. 2011. Trends in managed services. URL <http://www.comptia.org>.

- Cuadras, C.M. 2002. On the covariance between functions. *Journal of Multivariate Analysis* **81** 19–27.
- Davis, M. H. A. 1998. *Option pricing in incomplete markets*, chap. 14. Cambridge University Press, 216–226.
- Davis, M. H. A. 2001. Pricing weather derivatives by marginal value. *Quantitative Finance* **1** 305–308.
- Diebold, Francis X., Canlin Li. 2006. Forecasting the term structure of government bond yields. *Journal of Econometrics* **130** 337–364.
- Diebold, Francis X., Canlin Li, Vivian Z. Yue. 2008. Global yield curve dynamics and interactions: A dynamic nelson siegel approach. *Journal of Econometrics* **146** 351–363.
- Diebold, Francis X., Glenn D. Rudebusch, S. Borağan Aruoba. 2006. The macroeconomy and the yield curve: A dynamic latent factor approach. *Journal of Econometrics* **131** 309–338.
- Dinges, T. 2011. Electronics contract manufacturing market growth slows in 2011.
- Dong, Y., C. R. Carter, M. E. Dresner. 2001. Jit purchasing and performance: an exploratory analysis of buyer and supplier perspectives. *Journal of Operations Management* **19** 471–483.
- Duenyas, I., W. J. Hopp, Y. Bassok. 1997. Production quotas as bounds on interplant JIT contracts. *Management Science* **43**(10) 1372–1386.
- Durbin, James, Siem Jan Koopman. 2001. *Time series analysis by state space methods*. Oxford Univ Press.
- Eickmeier, S., J. Breitung. 2006. How synchronized are new eu member states with the euro area? evidence from a structural factor model. *Journal of Comparative Economics* **34** 538–563.
- Eickmeier, Sandra. 2009. Comovements and heterogeneity in the euro area analyzed in a non-stationary dynamic factor model. *Journal of Applied Econometrics* **24** 933–959.
- Essegaier, Skander, Sunil Gupta, John Z. Zhang. 2002. Pricing access services. *Marketing Science* **21** 139–159.
- Fandel, G., J. Reese. 1991. Just-in-time logistics of a supplier in the car manufacturing industry. *International Journal Production Economics* **24** 55–64.
- Fitzmaurice, Garrett, Nan Laird, James Ware. 2004. *Applied Longitudinal Analysis*. John Wiley & Sons.
- Forni, Mario, Lucrezia Reichlin. 1996. Dynamic common factors in large cross-sections. *Empirical Economics* **21** 27–42.
- Gaur, V., S. Seshadri. 2005. Hedging inventory risk through market instruments. *Manufacturing and Service Operations Management* **7**(2) 103–120.
- Ghalanos, Alexios, Stefan Theussl. 2012. *Rsolnp: General Non-linear Optimization Using Augmented Lagrange Multiplier Method*. R package.
- Goettler, R.L., K.B. Clay. 2010. Tariff choice with consumer learning and switching costs. URL <http://ssrn.com/abstract=885573>.
- González-Benito, J., M. Spring. 2000. Jit purchasing in the spanish auto components industry - implementation patterns and perceived benefits. *International Journal of Operations & Production Management* **20**(9) 1038–1061.

- Grahovac, J., G. Parker. 2003. Component modularity, external economies of scale, and outsourcing decisions in the supply chain. Working paper, Freeman School of Business, Tulane University.
- Gregory, Allen W., Allen C. Head, Jacques Raynauld. 1997. Measuring world business cycles. *International Economic Review* **38** 677–701.
- Guajardo, Jose A., Morris A. Cohen, Sang-Hyun Kim, Serguei Netessine. 2012. Impact of performance-based contracting on product reliability: An empirical analysis. *Management Science* **58(5)** 961–979.
- Huang, Ke Wei, Arun Sundararajan. 2005. Pricing models for on-demand computing. Working paper, CeDER Working Paper No. 05-26.
- Huggins, E.L., T.L. Olsen. 2003. Supply chain management with guaranteed delivery. *Management Science* **49(9)** 1156–1167.
- Hughes, Jonathan E., Christopher R. Knittel, Daniel Sperling. 2006. Evidence of a shift in the short-run price elasticity of gasoline demand. Tech. rep., NBER.
- Iyengar, Raghuram, Kamel Jedidi, Skander Essegaier, Peter J. Danaher. 2011. The impact of tariff structure on customer retention, usage, and profitability of access services. *Marketing Science* **30** 820–836.
- Iyengar, Raghuram, Kamel Jedidi, Rajeev Kohli. 2008. A conjoint approach to multipart pricing. *Journal of Marketing Research* **45(2)** 195–210.
- Jacob, Frank, Wolfgang Ulaga. 2008. The transition from product to service in business markets: An agenda for academic inquiry. *Industrial Marketing Management* **37(3)** 247–253.
- Karatzas, I. 1997. *Lectures on the Mathematics of Finance*. American Mathematical Society, Providence, RI.
- Kaseya. 2011. MSP global pricing survey. URL <http://www.kaseya.com>.
- Kauffman, Robert J., Ryan Sougstad. 2008. Risk management of contract portfolios in it services: The profit-at-risk approach. *Journal of Management Information Systems* **25** 17–48.
- Keng, Shao-Hsun, Wallace E. Huffman. 2010. Binge drinking and labor market success: A longitudinal study on young people. *Journal of Population Economics* **23** 303–322.
- Kim, Sang-Hyun, Morris A. Cohen, Serguei Netessine. 2007a. Performance contracting in after-sales service supply chains. *Management Science* **53(12)** 1843–1858.
- Kim, Sang-Hyun, Morris A. Cohen, Serguei Netessine. 2007b. Reliability or inventory? Contracting strategies for after-sales product support. *Proceedings of 2007 International Conference on Manufacturing & Service Operations Management*.
- Kim, Sang-Hyun, Morris A. Cohen, Serguei Netessine. 2010. Contracting for infrequent restoration and recovery of mission-critical systems. *Management Science* **56(9)** 1551–1567.
- Koopman, Siem Jan, André Lucas, Bernd Schwaab. 2011. Modeling frailty-correlated defaults using many macroeconomic covariates. *Journal of Econometrics* **162** 312–325.
- Kose, M. Ayhan, Christopher Otrok, Charles H. Whiteman. 2003. International business cycles: World, region and country-specific factors. *The American Economic Review* **93** 1216–1239.
- Laffont, Jean-Jacques, David Martimort. 2001. *The Theory of Incentives: The Principal-Agent Model*. Princeton University Press.

- Lambrecht, A., K. Seim, B. Skiera. 2007. Does uncertainty matter? Consumer behaviour under three-part tariffs. *Marketing Science* **26(5)** 698–710.
- Li, S., B. Lin. 2006. Accessing information sharing and information quality in supply chain management. *Decision Support Systems* **42(3)** 1641–1656.
- Ludvigson, Sydney C., Serena Ng. 2007. The empirical riskreturn relation: A factor analysis approach. *Journal of Financial Economics* **83** 171–222.
- Markowitz, Harry M. 1952. Portfolio selection. *The Journal of Finance* **7(1)** 77–91.
- Matson, J. E., J. E. Matson. 2007. Just-in-time implementation issues among automotive suppliers in the southern usa. *Supply Chain Management: An International Journal* **12(6)** 432–443.
- McFadden, Daniel, Kenneth Train. 2000. Mixed mnl models for discrete response. *Journal of Applied Econometrics* **15** 447–470.
- Mikkola, J.H. 2003. Modularity, component outsourcing, and inter-firm learning. *R&D Management* **33(4)** 439–454.
- Moench, Emanuel, Serena Ng. 2011. A hierarchical factor analysis of u.s. housing market dynamics. *Econometrics Journal* **14** 1–24.
- Moench, Emanuel, Serena Ng, Simon Potter. 2011. Dynamic hierarchical factor models. Tech. rep., Federal Reserve Bank of New York Staff Reports.
- Oliva, Rogelio, Robert Kallenberg. 2003. Managing the transition from products to services. *International Journal of Service Industry Management* **14(2)** 160–172.
- Otrok, Christopher, Charles H. Whiteman. 1998. Bayesian leading indicators: Measuring and predicting economic conditions in iowa. *International Economic Review* **39** 997–1014.
- Paleologo, G.A. 2004. Price-at-risk: A methodology for pricing utility computing services. *IBM Systems Journal* **43** 20–31.
- Photizo. 2011. Managed print services market continues to demonstrate strong growth worldwide. URL <http://photizogroup.com/>.
- Png, IPL, Hao Wang. 2010. Buyer uncertainty and two-part pricing: Theory and applications. *Management Science* **56(2)** 334–342.
- Romero, B. P. 1991. The other side of jit in supply management. *Production and Inventory Management Journal* **32(4)** 1–2.
- Sala, Luca. 2003. Monetary transmission in the euro area: A factor model approach. Manuscript.
- Salonen, Anna. 2011. Service transition strategies of industrial manufacturers. *Industrial Marketing Mangement* **40(5)** 683–690.
- Schultz, T. Paul, Aysit Tansel. 1997. Wage and labor supply effects of illness in cote d’ivoire and ghana: Instrumental variable estimates for days disabled. *Journal of development economics* **53** 251–286.
- Secomandi, Nicola. 2010. On the pricing of natural gas pipeline capacity. *Manufacturing and Service Operations Management* **12(3)** 393–408.
- Secomandi, Nicola, Mulan X. Wang. 2012. A computational approach to the real option management of network contracts for natural gas pipeline transport capacity. *Manufacturing and Service Operations Management* **14(3)** 441–454.

- Staum, J. 2008. *Incomplete markets*, chap. 12. Handbooks in Operations Research and Management Science, Elsevier, 511–562.
- Stock, James H., Mark W. Watson. 2001. Forecasting output and inflation: The role of asset prices. *Journal of Economic Literature* **47** 1–48.
- Stock, James H., Mark W. Watson. 2002. Forecasting using principal components from a large number of predictors. *Journal of the American Statistical Association* **97** 1167–1179.
- Stock, James H., Mark W. Watson. 2006. Forecasting with many predictors. *Handbook of Economic Forecasting*, vol. 1. North Holland.
- Stock, James H., Mark W. Watson. 2009. The evolution of national and regional factors in u.s. housing construction. *Volatility and Time Series Econometrics: Essays in Honour of Robert F. Engle*. Oxford University Press.
- Stock, James H., Mark W. Watson. 2011. Dynamic factor models. *Oxford Handbook of Economic Forecasting*. Oxford University Press.
- Stone, Patricia W., Cathy Mooney-Kane, Larson Elaine L., Diane K. Pastor, Jakc Zwanziger, Andrew W. Dick. 2007. Nurse working conditions, organizational climate, and intent to leave in icus: An instrumental variable approach. *Health Services Research* **42** 1085–1104.
- Sundararajan, Arun. 2004. Nonlinear pricing of information goods. *Management Science* **50** 1660–1673.
- Tracey, M., C. L. Tan, M. Vonderembse, E. J. Bardi. 1995. A reexamination of the effects of just-in-time on inbound logistics. *International Journal of Logistics Management* **6(2)** 25–37.
- Ülkü, S., L.B. Toktay, E. Yücesan. 2007. Risk ownership in contract manufacturing. *Manufacturing and Service Operations Management* **9(3)** 225–241.
- Van Mieghem, J. A. 2007a. Risk mitigation in newsvendor networks: Resource diversification, flexibility, sharing, and hedging. *Management Science* **53(8)** 1269–1288.
- Van Mieghem, J.A. 1998. Investment strategies for flexible resources. *Management Science* **44(8)** 1071–1078.
- Van Mieghem, J.A. 1999. Coordinating investment, production and subcontracting. *Management Science* **45(7)(7)** 954–971.
- Van Mieghem, J.A. 2003. Capacity management, investment, and hedging: Review and recent developments. *Manufacturing and Service Operations Management* **5(4)** 269–302.
- Van Mieghem, J.A. 2007b. Risk mitigation in newsvendor networks: resource diversification, flexibility, sharing and hedging. *Management Science* **53** 1269–1288.
- von Neumann, John, Oskar Morgenstern. 1947. *Theory of games and economic behavior*. Princeton University Press.
- Walsh, Lawrence M. 2011. Countering artificial commoditization and poor pricing practices in managed services. Tech. rep., The 2112 Strategy Group.
- Waters-Fuller, N. 1996. The benefits and costs of jit sourcing: A study of scottish suppliers. *International Journal of Physical Distribution & Logistics Management* **26(4)** 35–50.
- Wilson, Robert B. 1997. *Nonlinear pricing*. Oxford University Press.

- Wooldridge, M. Jeffrey. 2010. *Econometric Analysis of Cross Section and Panel Data*. The MIT Press.
- Zhou, J., N. Rudi. 2010. Pricing of over-the-counter hedging contracts for real investments. Working paper.