Using Survival Prediction Techniques to Learn Consumer-Specific Reservation Price Distributions

by

Ping Jin

A thesis submitted in partial fulfillment of the requirements for the degree of

Master of Science
in
Statistical Machine Learning

Department of Computing Science

University of Alberta

Abstract

A consumer's "reservation price" (RP) is the highest price that s/he is willing to pay for one unit of a specified product or service. It is an essential concept in many applications, *e.g.*, personalized pricing, auction and negotiation. While consumers will not volunteer their RPs, we may be able to predict these values, based on each consumer's specific information, using a model learned from earlier consumer transactions.

This thesis proposes a novel framework of learning RP distributions that involves a model of formulating the relationship between consumers' RPs and their purchasing decisions, and a data collection method. Within this framework, we show a way to estimate the consumer-specific RP distribution using techniques from the survival prediction — here viewing the consumers' purchasing choices as the censored observations. To validate our new framework of RP, we run experiments on realistic data, with four survival methods. These models performed very well (under three different criteria) on the task of estimating consumer-specific RP distributions, which shows that our RP framework can be effective.

As we found that the multi-task logistic regression model (MTLR) dominated the other models under all three evaluation criteria, we explored ways to extend it, leading to extensions that are more general and more flexible. Moreover, we prove that it is the general regularizer, instead of the smoothness regularizer, that results in a smooth predicted distribution; this leads further simplification of the MTLR model.

Acknowledgements

First, I want to thank my supervisor Prof. Russell Greiner, for his patience, time, and great support during my last two-year master study. It is an honor for me to be his student. I have learned so much knowledge and research methodology on machine learning, which is a great and valuable experience for me.

The colleagues at University of Alberta have always been generous and helpful to me. I would like to thank Junfeng Wen, Dr. Chun-Nam Yu, and Dr. Karim Abou-Moustafa for their helpful suggestions and Muyu Wei and Prof. Gerald Hubl for their generous help and great idea about data collection for this thesis.

I would also like to express my gratitude to my committee members, Prof. Dale Schuurmans and Prof. Michael Bowling for their time, interest, and insightful comments on my thesis.

Finally, I want to thank my family, who have always been supporting me.

Table of Contents

List of Tables

List of Figures

Chapter 1

Introduction

1.1 Motivation

Reservation price (RP) is the highest price a consumer is willing to pay for one unit of a certain product or service [17], which is an important and widely used concept in both the economics and marketing literature. It is critical for designing various pricing strategies, such as personalized pricing [7, 1], one-to-one promotion [26], and optimal pricing [19]. Many other fields like auction [35, 22], ad exchange [2], negotiation, and the design and pricing of bundles [34, 28] also heavily rely on accurate estimations of consumers' RPs.

For example, suppose that we are interested in setting the price of a certain product ω to achieve maximum profit when selling it to a certain population. If we know the reservation price r_i of each subject i, then we can easily compute the overall purchasing probability function $PPF(v)$ over price v for this specific population as

$$PPF(v) = \frac{1}{n}\sum_{i}^{n}(1 - N_i(v)), \tag{1.1}$$

where n is the total number of subjects and $N_i(v) = \mathbf{I}\{v < r_i\}$ is a counting process for subject i (see Figure 1.1).

With the knowledge of $PPF(v)$, we can achieve the maximum expected profit by setting the price of ω to be

$$v^* = \arg\max_{v}((v - c) \cdot PPF(v)), \tag{1.2}$$

where c is the production cost of ω.

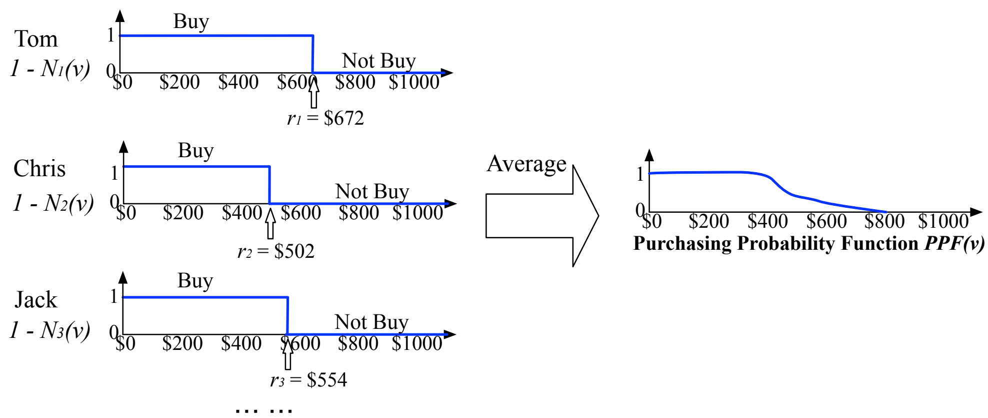

Figure 1.1: The overall purchasing probability function for a certain population.

Moreover, if it is allowed to sell ω at different prices to different subjects, *i.e.*, first degree price discrimination [32], then the seller's best strategy for maximum profit is to sell ω to the subjects at their individual reservation prices, *i.e.*, using Figure 1.1, sell to Tom at \$672, then Chris at \$502, etc. (Here, assume the production cost is under \$502.).

In the scenario of e-commerce, which has enjoyed booming development recently, the online retailers also have great interest in designing pricing strategies, understanding consumers' purchasing decisions, doing one-to-one promotion and so on, which rely on accurate estimation/elicitation of consumers' RPs. Additionally, online retailers usually have more information available than the traditional offline ones about their consumers, such as consumer-specific information (demographics, historical transactions and so on) and consumers' historical transactions, which may be related to consumers' RPs. This motivates us to find ways to better estimate consumers' RPs with this extra information.

1.2 Related Work

As revealing the true RPs will put consumers at a disadvantage in making deals with sellers, they would not volunteer to do it. This had led to a huge amount of research efforts in designing incentive compatible methods for eliciting fixed-point RPs [35, 4, 21, 16, 36]. Generally speaking, methods like the Becker-Degroot-Marschak (BDM) method strive to guarantee that consumers realize that revealing the true RP is their best strategy (see Figure 1.2), which is the key to accurate elicitation of consumers' RPs.

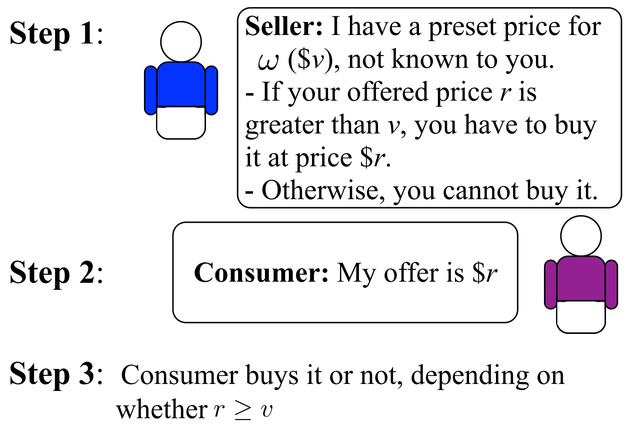

Figure 1.2: The Becker-Degroot-Marschak method.

However, it would be unrealistic to assume that a consumer's RP for a product always stays the same. Wang, et al [36] also notes that there is even uncertainty within an individual's RP, due to the consumer's uncertainty about his/her own preference [11] and the product performance [25].

Therefore, several different interpretations of the RPs have been proposed [13, 23, 33], which are associated with different probabilities of purchasing (see Figure 1.3):

(1) Floor RP: the maximum price at or below which the consumer will buy with 100% probability [13].

(2) Indifferent RP: the price at which a consumer is indifferent between the money and the product — *i.e.*, s/he will buy it with 50% probability [23].

(3) Ceiling RP: the minimum price at or above which the consumer will never buy it — *i.e.*, s/he has 0% probability of buying [33].

Furthermore, ICERANGE [36] embraced the inherent uncertainty of RP, by viewing a consumer's RP as a price range instead of a single price point,

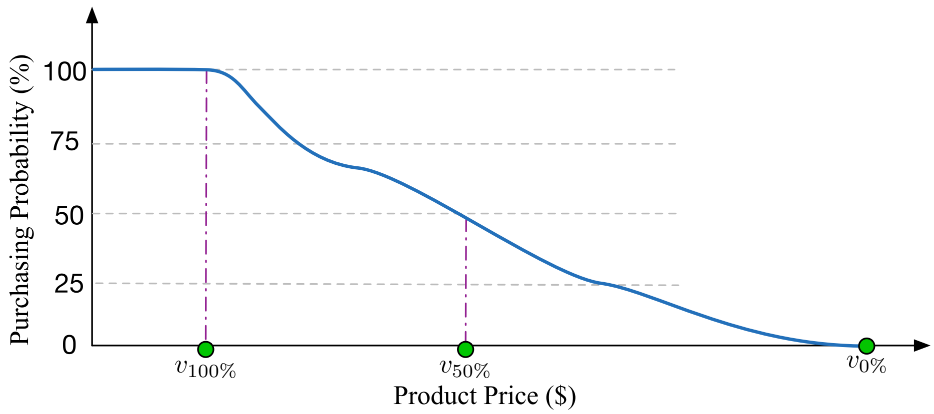

Figure 1.3: Three interpretations of reservation price: floor RP, indifferent RP, and ceiling RP.

which here means simultaneously eliciting several price points associated with different purchasing probabilities.

However, none of these methods deal with the challenge raised in the e-commerce scenario, *i.e.*, how to utilize the extra information in e-commerce to help the task of understanding consumers' RPs. What is worse, they also suffer from several drawbacks, which make them ineffective in the e-commerce scenario:

- First, consumers have little patience and no motivation to participate in the elicitation activity.
- Second, it is hard to validate if consumers realize that their best strategy is to tell their true RPs, which may lead to inaccurate elicitation of RPs.
- Third, these methods have no capability of making individual-level RP predictions, *i.e.*, each new consumer must go through the whole elicitation procedure to estimate his/her RP for the product of interest.

Therefore, we need a new model that can overcome these drawbacks and can effectively utilize the new information available in e-commerce setting to help the task of estimating consumers' RPs.

1.3 Contribution

In this thesis, we propose a novel framework of formulating the RP estimation problem, within which we explicitly define a consumer's RP as a random variable conditional on the consumer's features. This definition not only captures the inherent uncertainty of RP, but also allows us to use stochastic models to express the relationship between RP and consumer-specific features.

The common understanding of RP, *i.e.*, the highest price a consumer is willing to pay for one unit of a certain product or service, suggests that a consumer's purchasing decisions are closely related to his/her RP and thus we may indirectly infer consumers' RPs from their transaction data (the observation that a consumer decided to "buy" a product at some specified price) and non-transaction data (the observation that a consumer decided to "not buy" a product at some specified price). Motivated by this fact, we propose a consumer decision model, which formulates the relationship between consumers' RPs and their purchasing decisions. This idea of using purchasing decisions to infer RPs does not suffer from the first two drawbacks described in Section 1.2 any more, as we do not ask consumers to directly report their RPs.

Moreover, within this framework, the purchasing (resp., non-purchasing) observations in the RP setting are equivalent to right censored (resp., left censored) observations in the survival analysis setting, which indicates that we can utilize various survival techniques to model and learn the relationship between the RP and the features of a consumer from historical (non-)transaction data. This means that a seller, who knows the features of a consumer, can then make an individual-level prediction on the consumer's reservation price for a certain product, even if that consumer is new or has not bought the product of interest before.

Lastly, as MTLR (one of the survival models) achieves great performance in the experimental phase, we further develop a more general and succinct framework from MTLR, which provides more flexibility and reduces the training time tremendously.

1.4 Outline

Chapter 2 describes our framework of RP estimation. Section 2.1 introduces the formal definition of RP. Section 2.2 introduces the decision model that formulates the relationship between consumers' purchasing decisions and their RPs. Section 2.3 illustrates our way to collect (non-)transaction data, which can be used to learn the RP distributions.

Chapter 3 first describes the relationship between the RP estimation problem and survival analysis problem. Then Section 3.2 introduces four survival models that can be utilized to estimate RP, *e.g.*, Kaplan-Meier Estimator, Cox proportional hazard model, accelerated failure time model, and MTLR model.

In Chapter 4, we describe how we collect the data used in this thesis and some basic information about the four datasets. We also discuss several potential problems of data quality and the way to address them.

The empirical results of applying survival models to the RP estimation is presented in Chapter 5. Sections 5.1, 5.2 and 5.3 present the performance of survival models on the consumer-specific RP estimation under three different evaluation criteria: the mean absolute error of the RP predictions, the classification accuracy of purchasing choice, and the profit acquired with a simple pricing strategy. All results are based on ten-fold cross validation. The great performance of survival models on consumer-specific RP estimation task supports the effectiveness of our novel framework. Also we find that MTLR dominates the other approaches under all three evaluation criteria.

Chapter 6 presents the future work and the contributions of this thesis. Section 6.1 discusses three potential directions for future work. Then Section 6.2 summarizes the contributions of this thesis.

In Appendix A, we provide details about how we generalize MTLR and simplify the model.

Chapter 2

Framework of Reservation Price Estimation

The common understanding of RP — *i.e.*, the highest price a consumer is willing to pay for a certain unit of product or service — indicates that consumers' purchasing decisions on a certain product are closely related to their RPs of the product. In the BDM method, actual purchasing is also required for accurate measurements of RPs. This suggests that instead of directly asking consumers to report their RPs, we may be able to infer their RPs from their purchasing decisions, which are much easier to collect in practice. Therefore, in this chapter, we propose a consumer decision model that formulates the way consumers reach purchasing decisions and how it is related to their RPs. This decision model and a corresponding way of collecting data make up our framework of RP estimation. Within this framework, we can design new methods or utilize existing methods to learn the RP distributions from the observations of consumers' purchasing decisions, *i.e.*, (non-)transaction data.

2.1 Stochastic Setting of Reservation Price

First, we denote the random vector representing the features of consumers as $\vec{X}$ and a certain vector of feature values, corresponding to a single consumer, as $\vec{x}$. Now, we formally define two crucial random variables, *i.e.*, *consumer-specific reservation price* and *consumer-specific purchasing decision*; see below.

Definition 1 (Consumer-Specific Reservation Price) *For a certain prod-*

uct ω, the consumer-specific RP $R_\omega|\vec{x} \in \Re^{\geq 0}$ is a random variable conditioning on the features of a consumer $\vec{X} = \vec{x}$.

Definition 2 (Consumer-Specific Purchasing Decision) *If a product ω is offered at price v, the consumer-specific purchasing decision $A_{\omega,v}|\vec{x} \in \{buy, not\_buy\}$ is a binary random variable, conditioning on the features of a consumer $\vec{X} = \vec{x}$.*[1]

2.2 Consumer Decision Model

In this section, we propose a decision-making model that describes how the consumer's purchasing decision $A_{\omega,v}|\vec{x}$ is related to the RP $R_\omega|\vec{x}$.

When a consumer with features $\vec{X} = \vec{x}$ is faced with a certain offer, *i.e.*, a specific product ω is being sold at price v, s/he reaches her/his purchasing decision $a \sim A_{\omega,v}|\vec{x}$ in a two-step procedure

Step 1. Draw an "instant RP": an instant RP r is drawn from the distribution of $R_\omega|\vec{x}$

Step 2. Make a decision:

$$\text{Consumer will buy } \omega \text{ for price } v \quad \textbf{iff} \quad v \leq r$$

i.e.,

$$a = \mathbf{I}\{r \geq v\}$$

That is, we assume that after drawing an instant RP $r \sim R_\omega|\vec{x}$, the customer's decision is determined by the relationship between r and v (see Figure 2.1).

Then it is explicit that the relationship between the purchasing decision random variable $A_{\omega,v}|\vec{x}$ and the reservation price random variable $R_\omega|\vec{x}$ is

$$A_{\omega,v} \mid \vec{x} \;=\; \mathbf{I}\{v \;\leq\; R_\omega|\vec{x}\}. \tag{2.1}$$

[1]Note that $A_{\omega,v_1}|\vec{x}$ and $A_{\omega,v_2}|\vec{x}$ are two different random variables, for $v_1 \neq v_2$.

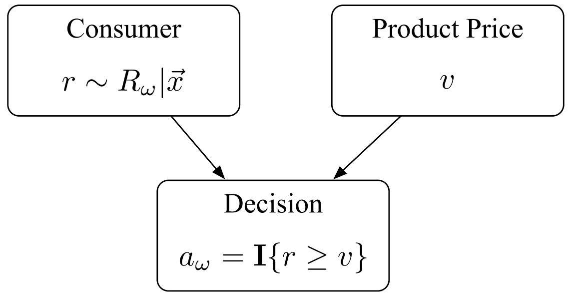

Figure 2.1: How consumers reach a purchasing decision when offered product ω at price v.

We can also derive the *purchasing probability function* $PPF_\omega(\vec{x}, v)$, *i.e.*, the probability that consumer $\vec{x}$ will buy product ω at price v, to be

$$PPF_\omega(\vec{x}, v) \triangleq Pr(\, A_{\omega,v} = 1 \,|\, \vec{x}\,) = Pr(\, R_\omega \geq v \,|\, \vec{x}\,) \;=\; 1 - F_{R_\omega|\vec{x}}(v) \tag{2.2}$$

where $F_{R_\omega|\vec{x}}(\cdot)$ is the cumulative distribution function (CDF) of consumer $\vec{x}$'s RP for product ω.

Note that this decision-making process strictly conforms with the common understanding of RP, *i.e.*, the highest price a consumer is willing to pay for a unit of a certain product or service.

2.3 Data Collection and Format

As shown in the previous section, RPs and purchasing decisions are closely related to each other, which suggests that we can indirectly infer consumers' RPs for a certain product from their purchasing decisions. Therefore, for a certain product ω, instead of directly asking consumers for their instant RPs r_i, we collect the (non-)transaction data, *i.e.*, consumers' decisions $a_{i,\omega}$ on whether they purchase ω at different prices v. Each observation in the dataset D_ω of product ω is a vector in the format $(x_i,\ a_{i,\omega},\ v_{i,\omega})$. An example dataset is shown in Table 2.1.

While the traditional RP models require consumers to be highly involved in a sophisticated elicitation procedure in order to make them understand that telling the true RP is their optimal choice, our data collection process is

features of consumers $\vec{x}_i$				price $v_{i,\omega}$	decision $a_{i,\omega}$
age	gender	monthly income	...		
18	male	\$200	...	\$3.50	0
26	female	\$3000	...	\$5.00	1
...	...	...	...	...	...
28	female	\$2000	...	\$4.5	0

Table 2.1: An example dataset of some product ω.

significantly simpler and does not make assumptions about consumers' understanding, as we do not ask consumers to report their RPs directly.

Chapter 3

Reservation Price Prediction Models

In this chapter, first we illustrate the relationship between the RP estimation problem and survival analysis problem and how we can learn the consumer-specific RP distribution from the (non-)transaction data via survival prediction techniques. We then introduce three popular survival models and one recent effort from the machine learning community in predicting patient-specific survival distributions.

3.1 Relation to Survival Analysis

The typical survival analysis focuses on time-to-event data, where the variable of interest is the death/event time T. In general, survival models strive to learn the survival function $\hat{S}_{T|\vec{x}}(t)$, which is defined as $1 - F_{T|\vec{x}}(t)$, from censored (left, right, or interval) data and event data. This task differs from ordinary regression as it must deal with censored observations, which are incomplete observations of the event time T.

There are many possible sources of censoring. For example, in a cancer study, where the variable of interest is the death time of patients with cancer, if an (alive) patient chooses to drop out of the study at time t, then we will never know his/her actual time of death T. All we know is that his/her death time is after his/her *censored time* t, which is only partial information about T. The phenomenon in this specific example is called *right censoring*, as the unknown event time is on the "right" side of censored time t. Similarly, we also have left censoring and interval censoring (see Figure 3.1).

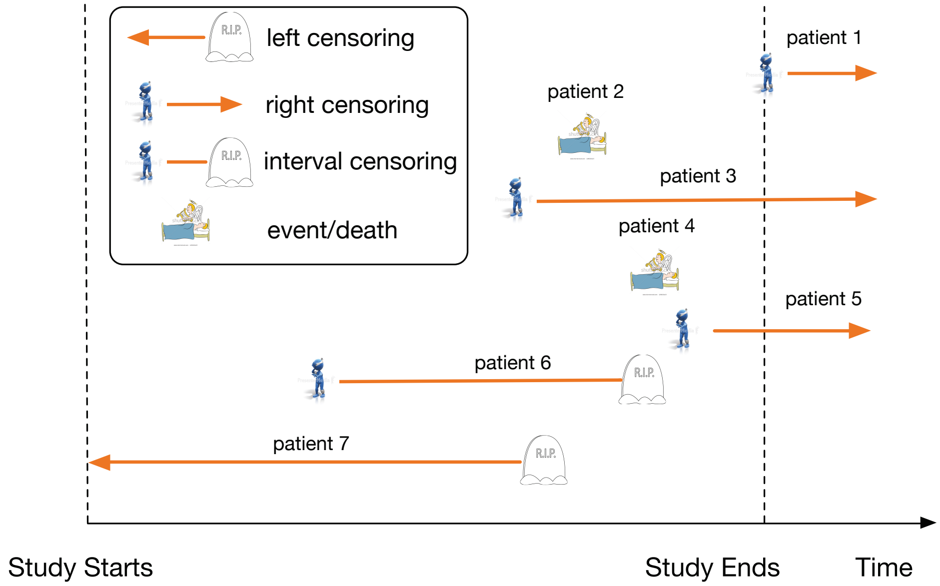

Figure 3.1: Left, right and interval censored observations.

Right censored observation: the unknown true event time is **after** a certain time t, *e.g.*, patients 1, 3 and 5 in Figure 3.1.

Left censored observation: the unknown true event time is **before** a certain time t. For example, in a breast cancer relapse study, a patient gets her first examination at the sixth month and is diagnosed as having already experienced a relapse. In this case, all we know is that the relapse happened in the first six months, *e.g.*, patient 7 in Figure 3.1. (Here, we only know that the subject was dead at the end, but not when she died.)

Interval censored observation: the unknown true event time is **in** a certain time range $[t_1, t_2]$. For example, in a breast cancer relapse study, where patients take monthly examinations (at different days), if a patient is diagnosed as having a relapse, then we only know that the relapse time is in the previous one month, *e.g.*, patient 6 in Figure 3.1.

In most survival studies, we typically have complete observations about some subjects — *i.e.*, we know when some actually died. In the RP estimation problem defined within our framework, however, we have no complete observations of consumers' RPs at all, since we do not ask consumers to report their true RPs. Instead, we only have purchasing transactions, where the consumers' true RPs are greater than or equal to the price of products v, and non-purchasing transactions, where the consumers' true RPs are less than the price of products v. That is, for $r \sim R_\omega \mid \vec{x}$:

Purchasing transaction: $a_\omega = 1 \;\Leftrightarrow\; r_\omega \geq v$;

Non-purchasing transaction: $a_\omega = 0 \;\Leftrightarrow\; r_\omega < v$.

If we take the RP R as the variable of interest instead of T, the purchasing (resp., non-purchasing) observations in the RP setting are equivalent to right censored (resp., left censored) observations in the survival analysis setting.

To be more clear, Table 3.1 shows the matching relationship between the terminologies in these two settings.

Survival Analysis	**Reservation Price**
Event time variable: T	Reservation price variable: R
Survival distribution: $f(t)$	RP distribution: $f(v)$
Survival function: $S_T(t)$	Purchasing probability function: $PPF(v)$
Left censored observation $c = 0$	Non-purchasing transaction $a = 0$
Right censored observation $c = 1$	Purchasing transaction $a = 1$

Table 3.1: Terminology/symbol matching table.

Then we can utilize survival models to learn the distribution of R using the (non-)purchasing transactions. An illustration of the whole learning system and how it works on new consumers is shown in Figure 3.2.

Section 3.2 below introduces four survival models: Kaplan-Meier estimator [18], Cox proportional hazard model [9], accelerated failure time model [37], and multi-task logistic regression model [38]. We will test their performance in Chapter 5, to evaluate the effectiveness of our framework of RP.

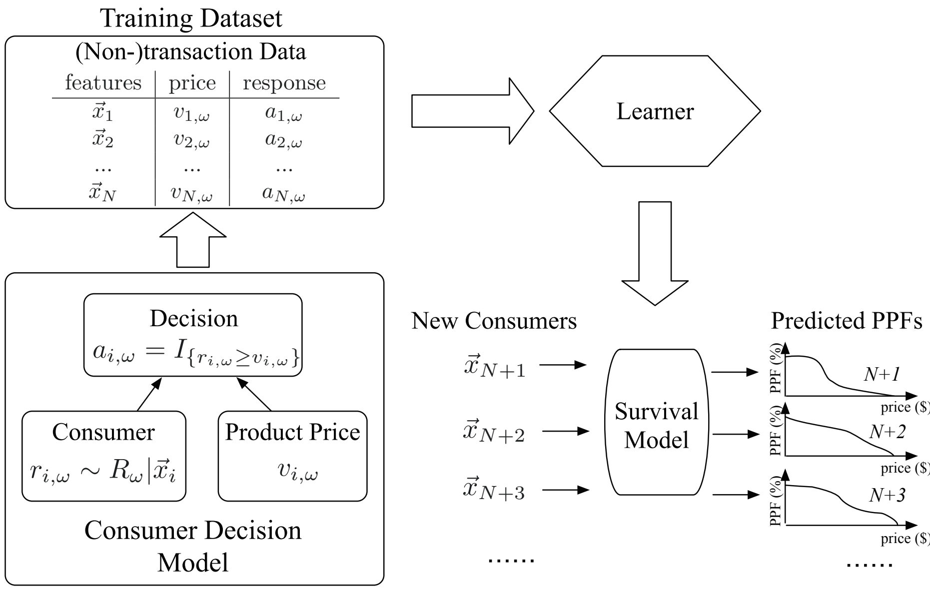

Figure 3.2: Illustration of the whole learning system and how it works on new instances.

3.2 Survival Models

3.2.1 Kaplan-Meier Estimator

The Kaplan-Meier (KM) estimator [18] is an empirical non-parametric model that estimates the survival function $S(t)$. This tool is designed for comparing the survival curves of two subpopulations in order to identify the risk factors, *i.e.*, the features important to survival, and is widely used in clinical study.

For a dataset consisting of only right censored and event data, the empirical estimate of $S(t)$ is

$$\hat{S}(t) = \prod_{j:\tau_j<t} \left(1 - \frac{d_j}{r_j}\right) \tag{3.1}$$

where $\tau_1, \tau_2, ..., \tau_K$ are the set of all K distinct death times in the dataset, d_j is the number of deaths at time τ_j and r_j is the number of subjects at risk right before τ_j (*i.e.*, number of subjects died or were censored at or after τ_j).

Since an RP dataset only consists of left and right censored data, but no event data, we have to resort to the Expectation-Maximization approach of Turnbull [31] to estimate the survival curve, where both left and right censored data are treated as interval censored data.

Notice that the KM estimator does not consider the features of the subjects, which means it predicts (summarizes) the same survival curve for all subjects and thus is not personalized. We still include it here for completeness, as it is one of the most widely accepted and used models in survival analysis.

3.2.2 Cox Proportional Hazards Model

The Cox proportional hazards (Cox) model is a semi-parametric model designed for comparing the survival time of two populations or to identify the risk factors critical to survival [9]. Unlike the KM model, the Cox model does use the subject features and works with the hazard function $\lambda(t)$ [1] instead of survival function.

$$\lambda(t) = \lim_{\Delta t \to 0} \frac{Pr(t \leq T < t + \Delta t | T \geq t)}{\Delta t} = \frac{f(t)}{S(t)} \tag{3.2}$$

The hazard function in the Cox model is modeled as

$$\lambda(t|x_i) \;\; = \;\; \lambda_0(t) \; \exp(\vec{x}_i^T \theta), \tag{3.3}$$

where $\lambda_0(t)$ is the baseline hazard function, and θ is learned from a data sample.

Here, the relative influence of each feature $x_{i,k}$ depends linearly ("proportionally") on the corresponding coefficient θ_k (albeit in the exponent). This is why it is called proportional hazard model.

One of the advantages of this model is that we can estimate θ by maximum partial likelihood estimation [9], which requires no knowledge of the baseline hazard $\lambda_0(t)$. This simplifies the task of identifying the risk factors. However when it comes to the prediction task, the proportional hazard assumption restricts the shapes of predicted survival curves of all patients to be the same, as shown in Figure 3.3. Its predictions on subjects' survival rates may thus be not calibrated [38].

Similar to Section 3.2.1, we treat both left and right censored data as interval censored data and utilize the Cox model designed for interval censored

[1] The hazard function is also called the "failure rate" since it reflects the subject's instantaneous rate of failure.

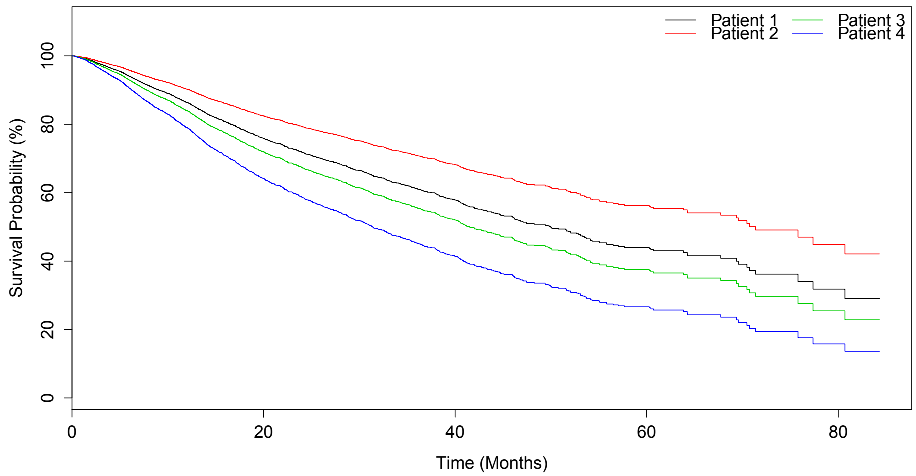

Figure 3.3: Survival curves for four patients from North Alberta Cancer Dataset generated by Cox model.

data to estimate our model [24, 14]. Specifically, we train the Cox model using the R package `intcox` [14].

3.2.3 Accelerated Failure Time Model (Tobit Model)

The accelerated failure time model (AFT) is a parametric model, which directly models the distribution of T with some parametric distribution [37], as shown below:

$$\log T_i \;=\; \theta^T \vec{x}_i \;+\; \delta\,\epsilon, \tag{3.4}$$

where δ is the scale parameter and ϵ is the error term. Different distributions of ϵ yield different forms of the AFT model. The most common distributions used for ϵ are the Weibull distribution, Gamma distribution and the standard normal distribution. With $\epsilon \sim \mathcal{N}(0,1)$ (Gaussian distribution with mean being 0 and standard deviation being 1), AFT model is actually equivalent to the well-known Tobit model in the economics literature [12].

In the AFT model, the effect of covariates is to accelerate/decelerate the scale of life time, while in the Cox model, the effect of covariates is to multiply the hazard by a constant. Figure 3.4 shows an example of the predicted survival curves of four patients in North Alberta Cancer Dataset, by the AFT model with ϵ following the log-normal distribution. In our experiments, we fit the AFT model with ϵ following the log-normal distribution using the function `survreg` in the R package `survival` [30].

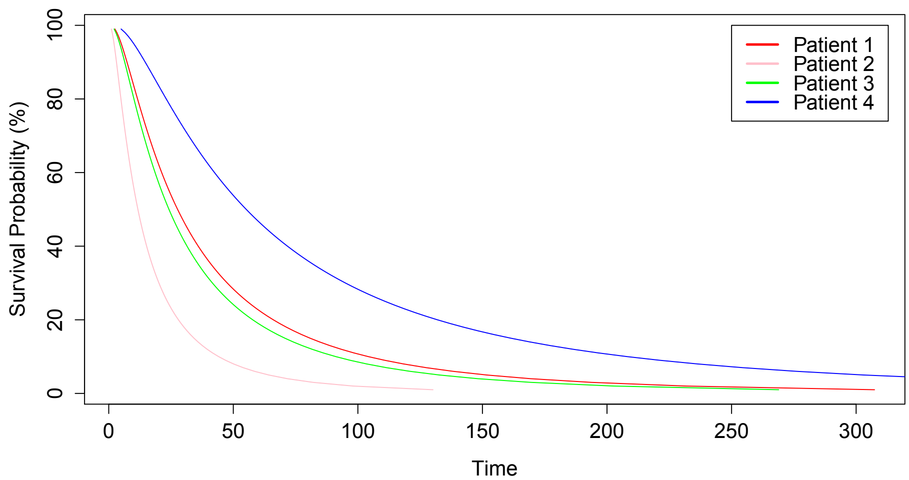

Figure 3.4: Survival curves for four patients from the North Alberta Cancer Dataset [38] generated by the AFT model.

3.2.4 Multi-Task Logistic Regression

Multi-task logistic regression is a recent effort from the machine learning community to produce a patient-specific survival function, which works well, according to several criteria [38].

Unlike the Cox model, MTLR does not have an explicit assumption or restriction about the hazard function or the shape of survival curves. Survival curves of different individuals can be very different and can intersect with each other. This offers greater prediction capacity and flexibility. Figure 3.5 shows the predicted survival functions from MTLR for fifteen patients from the North Alberta Cancer Dataset [38].

MTLR first discretizes the continuous time axis into $K + 1$ time points $\{\tau_0, \tau_1, \tau_2, \ldots, \tau_K\}$, with $\tau_0 = 0$ and $\tau_K = \infty$, and then transforms the survival function prediction task into a sequence of binary classification tasks, by constructing initially a logistic regression model for each time point τ_j, $j =$

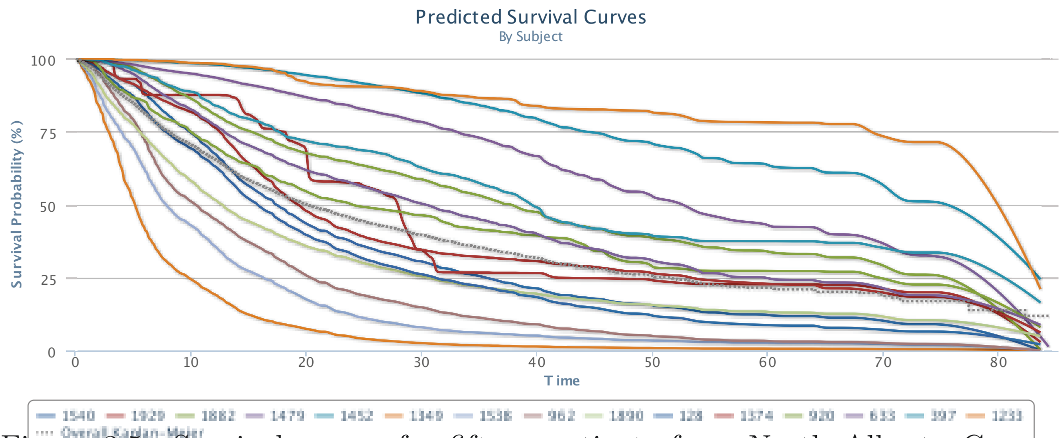

Figure 3.5: Survival curves for fifteen patients from North Alberta Cancer Dataset generated by MTLR model.

$1, ..., K-1$:

$$Pr(\, y_j = 0 \,|\, \vec{x}\,) \;=\; \Big(1 + \exp(\vec{x}^T \cdot \vec{\theta}_j + b_j)\Big)^{-1} \tag{3.5}$$

where $\vec{\theta}_j$ and b_j are the parameters associated with the j^{th} time point and $y_j = \mathbf{I}\{T < \tau_j\}$ indicates if the subject $\vec{x}$ has incurred an event before τ_j.

Then if we (for now) treat the classifiers as independent, we have the probability mass function (PMF) of $\vec{y}$ as

$$\tilde{Pr}(\,\vec{y} \mid \vec{x}) \;=\; \frac{\exp\Big(\sum_{k=1}^{K-1}(\vec{x}^T \cdot \vec{\theta}_k + b_k)y_k\Big)}{\prod_{k=1}^{K-1}(1 + \exp(\vec{x}^T \cdot \vec{\theta}_k + b_k))} \tag{3.6}$$

However, as we must prevent the case that $y_j = 1$ and $y_{j+1} = 0$ from holding (that is, after someone dies, that person cannot come back alive), the normalization term is the summarization of the unnormalized “probability” of these K legal $\vec{y}$s, which are $(1, 1, \ldots, 1, 1)$, $(0, 1, \ldots, 1, 1)$, $\ldots$, $(0, 0, \ldots, 0, 1)$, and $(0, 0, \ldots, 0, 0)$. The final form of the PMF of T is

$$\begin{aligned}
& Pr(\,\tau_{j-1} \leq T < \tau_j \,|\, \vec{x}\,) \\
&= \tilde{Pr}(\,\vec{y} = (y_1 = 0, \ldots, y_{j-1} = 0, y_j = 1, ..., y_{K-1} = 1) \mid \vec{x}) \\
&= \frac{\exp\Big(\sum_{k=1}^{K-1}(\vec{x}^T \cdot \vec{\theta}_k + b_k)y_k\Big)}{Z(\Theta, B, \vec{x})} \\
&= \frac{\exp\Big(\sum_{k=j}^{K-1}(\vec{x}^T \cdot \vec{\theta}_k + b_k)\Big)}{Z(\Theta, B, \vec{x})},
\end{aligned} \tag{3.7}$$

where $\Theta = (\vec{\theta}_1, \vec{\theta}_2, ..., \vec{\theta}_{K-1}), B = (b_1, b_2 ... b_{K-1})$ and

$$Z(\Theta, B, \vec{x}) = \sum_{j=1}^{K} \exp\left(\sum_{k=j}^{K-1} (\vec{x}_i^T \cdot \vec{\theta}_k + b_k)\right)$$

is the normalization term.

Then it is trivial to derive the log-likelihood function of a dataset $D = \{[\vec{x}_i, t_i]\}$, where the first N_e instances are uncensored and the remaining N_c are censored.

$$\begin{aligned} LL(D; [\Theta, B]) = & \sum_{i=1}^{N_e} \sum_{k=1}^{K-1} (\vec{x}_i^T \cdot \vec{\theta}_k + b_k) y_k(t_i)) \; + \\ & \sum_{i=N_e+1}^{N_e+N_c} \log\left[\sum_{j=1}^{K} c_j(t_i) \exp\left(\sum_{k=j}^{K-1} (\vec{x}_i^T \cdot \vec{\theta}_k + b_k)\right)\right] - \\ & \sum_{i=1}^{N_e+N_c} \log\left(Z(\Theta, B, \vec{x}_i)\right), \end{aligned} \tag{3.8}$$

where $y_k(t_i) = \mathbf{I}\{t_i < \tau_k\}$, $c_j(t_i) = \mathbf{I}\{t_i < \tau_j\}$ for right censored observations and $c_j(t_i) = \mathbf{I}\{t_i \geq \tau_{j-1}\}$ for left censored observations.

Yu *et al.* [38] proposed finding the parameters $[\Theta, B]$ that optimize the log-likelihood together with two regularizers: $\sum_{k=1}^{K-1} \|\theta_k\|$ and $\sum_{k=1}^{K-2} \|\theta_{k+1} - \theta_k\|$. The first regularizer is designed to reduce the chance of overfitting. The second one is designed to smooth the prediction and control the model capacity, which is commonly seen in multi-task learning [6]. However, we prove that the second regularizer does not do its job (of smoothing the predicted distribution), and in fact may even harm the prediction performance. This proof can be found in Appendix A.

Moreover, we have also reformulated MTLR into a multi-class softmax classifier and developed a more general form of it, which allows for more flexibility. Unfortunately, though we benefit greatly in training speed by saving about 90% of training time, this new progress does not lead to a better (or worse) performance under the three evaluation criteria used in this thesis. Therefore, this further analysis appears in Appendix A.

Chapter 4

Dataset

4.1 Data Collection

While there are many datasets of financial transactions, they all report only the actual purchases, but not the "non-purchases" — *i.e.*, not situations where a consumer has declined an offer. For our stochastic RP setting, we need a dataset that contains both purchases and non-purchases. While the donation dataset used in KDD Cup 1998 [15] does provide "non-donate" transactions, these non-donations only happen when a donor's "reservation donation" is zero, which means that this dataset provides no left-censored observations.

We therefore designed and executed our own online survey on Qualtrics, asking subjects to provide information about themselves, and about their interest in purchasing each of four different specific-types of chocolate bar.

Here, we acquire one dataset for each type of chocolate, leading to four datasets in total. For each consumer, we collected 41 features, *e.g.*, the consumer's demographics information, and preference towards the chocolate brand and flavor, the time when s/he ate her/his last meal and so on.[1] Note that the subjects did not purchase any product in the survey; they just provided information, for which they were paid.[2]

In our survey, instead of recording the responses of consumers to offers of products at certain prices, as described in Section 2.3, we directly ask for their

[1] For more details about the survey, please visit `https://qtrial2014.az1.qualtrics.com/SE/?SID=SV_0kycgJjTgOj5Z8p`.

[2] We obtained the appropriate ethics permission for this study with human participants.

instant RPs $r_{i,\omega}$ and used this to determine their responses $a_{i,\omega}$ to certain offers, *i.e.*, (non-)transaction data, following the decision-making process proposed in Section 2.1. N.b., our *learning* algorithms do NOT use those $r_{i,\omega}$'s — instead, they just use the (non-)transaction data; see Figure 3.2. We only collected the $r_{i,\omega}$ values in the study as a way of evaluating our learners.

4.2 Data Quality

To ensure that our data quality is good and the reported RPs are accurate, our online survey included five attention-check questions, one RP understanding question and a two-step RP elicitation procedure [29]. We eliminated any subject who failed any attention-check or RP understanding question or who showed any inconsistency in his/her answers about RP. We also eliminated blatantly ridiculous responses — *e.g.*, a subject willing to pay $10000 for a 100g chocolate bar.

At last, 722 responses (out of 1080 submissions) qualified for each of the four chocolates. The basic statistics appear in Table 4.1.

	Lindt	Godiva	Valrhona	Hersheys
Mean of RP	3.88	4.84	2.94	1.48
Std of RP	1.89	2.92	2.08	1.05
Retail Price	6.00	10.00	7.50	2.00

Table 4.1: Summary of four chocolate datasets.

While we tried to produce a dataset with good quality, the hypothetical response bias [36] cannot be completely avoided, as there were no real purchases.

The good news is, as our goal is to evaluate the performance of survival models within our novel framework of RP, this systematic bias will not be a serious issue. When online retailers later collect (non-)transaction data in practice, the consumers will be making purchases, which will mitigate this hypothetical response bias.

4.3 Generation of (Non-)transaction Data

After acquiring the true RPs and features of the consumers, we simulated a (non-)transaction data collection session by first sampling one query price $v_{i,\omega}$ for each consumer from a stretched Chi-Square distribution[3] and then determining the consumer's response $a_{i,\omega}$ following the decision-making process defined in Section 2.1. That is, the consumer's purchasing decision is simply

$$a_{i,\omega} = \begin{cases} \text{yes} & \text{if } v_{i,\omega} \leq r_{i,\omega} \\ \text{no} & \text{otherwise} \end{cases} \tag{4.1}$$

This led to a database, for each product, that had 722 instances, whose i^{th} row described the i^{th} consumer using 41 features. It also had an offer price $v_{i,\omega}$ for each i^{th} instance of product ω, and the response bit $a_{i,\omega}$. That is, the format of this dataset strictly conforms with the example dataset shown in Table 2.1 (see also Figure 3.2).

Note in particular that the dataset does *not* include the consumer's RP $r_{i,\omega}$ nor did we use the true RP data in training nor in the hyper-parameter selection via cross validation. The true RP data is only available in the final testing phase for evaluating the RP prediction performance.

[3]First we set the parameter k in χ_k to be the mean of the RPs in the whole datasets and then used a linear mapping to match the variance of the distribution χ_k^2 with the variance of the RPs in the whole dataset.

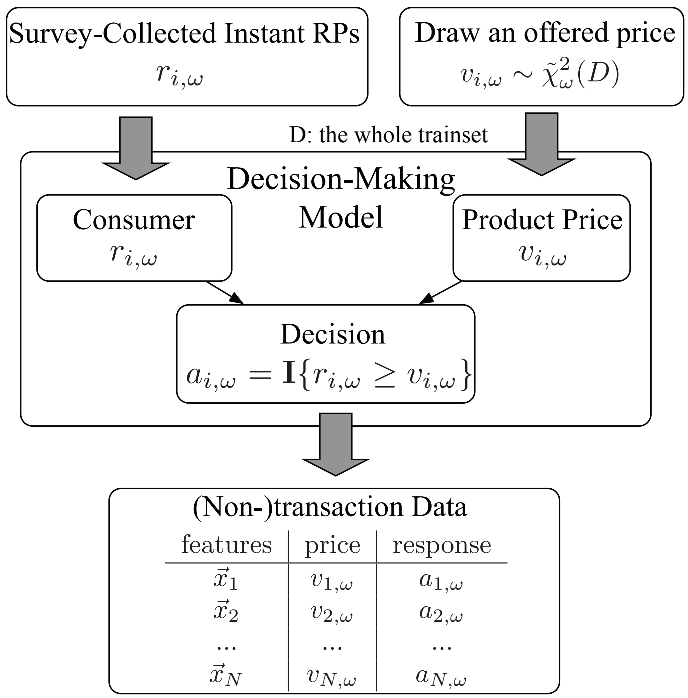

Figure 4.1: Procedure of generating (non-)transaction data.

Chapter 5

Experimental Results

In this section, we present empirical results for three reasonable evaluation criteria. All results are based on ten-fold cross-validation: for each type of chocolate and for each of our four survival models (KM, Cox, AFT, and MTLR), we train a model on 9/10 of the subjects; we then use that learned model to produce a "RP distribution" for each of the remaining 1/10 of the subjects. For MTLR, within each fold we execute an internal five-fold cross validation to select the best hyper-parameter, *e.g.*, regularization constant.

5.1 Mean Absolute Error

Given the learned CDF $\hat{F}_{R_\omega|\vec{x}_i}(v)$ of consumer $\vec{x}_i$'s RP for product ω, we use the median RP as the prediction for consumer $\vec{x}_i$'s RP value[1]:

$$\text{Median}(\vec{x}_i) = \hat{r}_{i,\omega} \quad s.t.\ \hat{F}_{R_\omega|\vec{x}_i}(\hat{r}_{i,\omega}) \ = \ 0.5 \tag{5.1}$$

As we have collected the consumers' true instant RP $r_{i,\omega}$, we can compute the mean absolute error (MAE) of our predicted RPs.

$$MAE_\omega \ = \ \frac{1}{N}\sum_{i=1}^{N} |\hat{r}_{i,\omega} - r_{i,\omega}| \tag{5.2}$$

where N is the number of consumers.

Note that we cannot use this criterion in cross validation to select hyper-parameters, because the learners do not have access to the true RP.

[1] We use the *median* price point as the RP prediction, as it is more robust than *mean*.

Brand	Lindt	Godiva	Valrhona	Hersheys
KM	2.03(0.29)	2.34(0.44)	1.63(0.22)	0.88(0.12)
Cox	1.71(0.22)	2.15(0.27)	1.45(0.18)	0.73(0.08)
AFT	1.36(0.13)	2.03(0.17)	1.37(0.18)	0.68(0.06)
MTLR	**1.23**(0.12)	**1.79**(0.15)	**1.25**(0.16)	**0.60**(0.08)
Baseline	1.40(0.21)	2.14(0.27)	1.54(0.18)	0.72(0.10)

Table 5.1: Ten-fold cross validation MAE. Each cell gives the "mean(std dev)" over the ten-fold cross validation, of running a particular learner on a dataset. **Bold** values are the best performance across the four models for that dataset. Table 5.2 uses the same format.

Table 5.1 shows the the ten-fold cross validation MAE for all four survival models. The bottom row is a strong "cheating" baseline, which computes the median value of RP in the training set, then uses that as a prediction over the test set. This "cheating" baseline utilizes the consumers' true RPs, which are not available to the learners, and so should not be used in training. The results are visualized in Figure 5.1.

We see that the MTLR and AFT model achieve lower MAE than the cheating baseline, on all four chocolate datasets. Moreover, their standard deviations of the MAE for all four datasets are also very small, which indicates that the performance of MTLR and AFT is stable. Figure 5.2 shows a scatter plot of the true reservation prices and the predicted reservation prices of MTLR over the Godiva dataset. The predictions on the whole dataset come from the ten-fold cross validation procedure. The correlation coefficient of the predictions and true reservation prices is 0.5321.

It is worth noting that among all four models, MTLR performs the best in each of the four chocolate datasets. (One-side t-test shows that MTLR was significantly better than others here, each at $p < 0.05$.)

5.2 Binary Classification Accuracy

This evaluation criterion tests if the learned models can accurately predict the consumer's response to our offer of ω at price v. This too is very important in

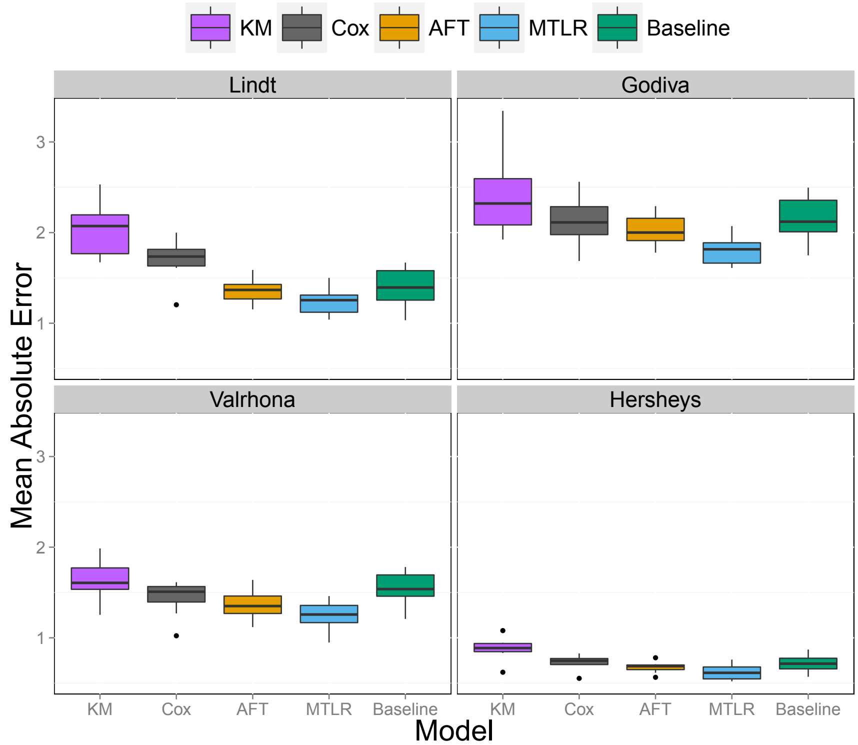

Figure 5.1: Mean absolute error, MAE (ten-fold CV).

real applications. Here, each predictor predicts the response using

$$\hat{a}_{i,\omega} \quad = \quad \begin{cases} \text{yes} & \text{if } PPF_{\omega}(\vec{x}_i, v) = 1 - \hat{F}_{R_{\omega}|\vec{x}_i}(v) \;\; \geq \;\; 0.5 \\ \text{no} & \text{otherwise} \end{cases} \tag{5.3}$$

We then compute the classification accuracy as

$$ACC \quad = \quad \frac{1}{N} \sum_{i=1}^{N} \mathbf{I}\{\hat{a}_{i,\omega} = a_{i,\omega}\} \tag{5.4}$$

Brand	Lindt	Godiva	Valrhona	Hersheys
KM	0.57(0.04)	0.68(0.07)	0.70(0.06)	0.62(0.08)
Cox	0.64(0.08)	0.67(0.11)	0.76(0.09)	0.75(0.04)
AFT	0.57(0.04)	0.69(0.06)	0.77(0.04)	0.76(0.03)
MTLR	**0.81**(0.04)	**0.80**(0.06)	**0.81**(0.04)	**0.81**(0.03)
Baseline	0.53	0.50	0.53	0.54

Table 5.2: Classification accuracy (ten-fold cross validation).

Table 5.2 and Figure 5.3 show that the classification accuracies of all four survival models (even the non-personalized KM) are significantly better than the “random guess” baseline, *i.e.*,

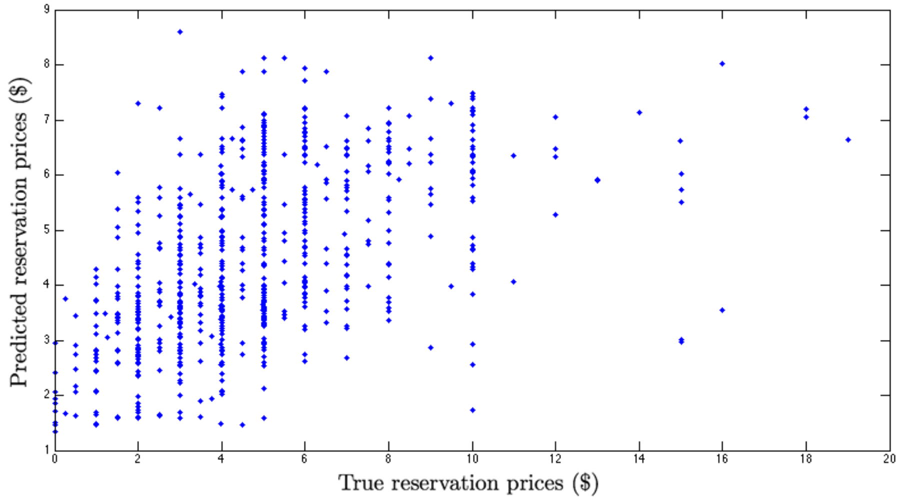

Figure 5.2: Scatterplot of true reservation prices and predicted median reservation prices on the Godiva dataset.

$$baseline = \frac{max\{\#purchasing\ transactions, \#non\text{-}purchasing\ transactions\}}{\#all\ transactions}$$

Among those models, MTLR again performs the best on all four datasets and this result is statistically significant ($p < 0.05$ in t-test).

5.3 Profit Using a Simple Pricing Strategy

RP is a very important concept in designing pricing strategies for various purposes, *e.g.*, high customer loyalty or more profit, as discussed in Section 1.1.

In this section, we want to evaluate whether survival models within our RP framework can lead to real profit in practice. The pricing strategy we use here is simple and intuitive, which aims to maximize the expected profit and relies on good estimates of the PPF.

As we have a predicted purchasing probability function $P\hat{P}F_\omega(\vec{x}_i, v) = 1 - \hat{F}_{R_\omega|\vec{x}_i}(v)$ for each consumer $\vec{x}_i$, we know the predicted expected profit by offering ω at price v to $\vec{x}_i$ would be $(v - c) \cdot P\hat{P}F(\vec{x}_i, v)$ where c is the seller's cost to produce ω. Here, the seller should therefore offer the product ω to $\vec{x}_i$ at the price $\hat{v}_i(c)$ with maximum expected profit:

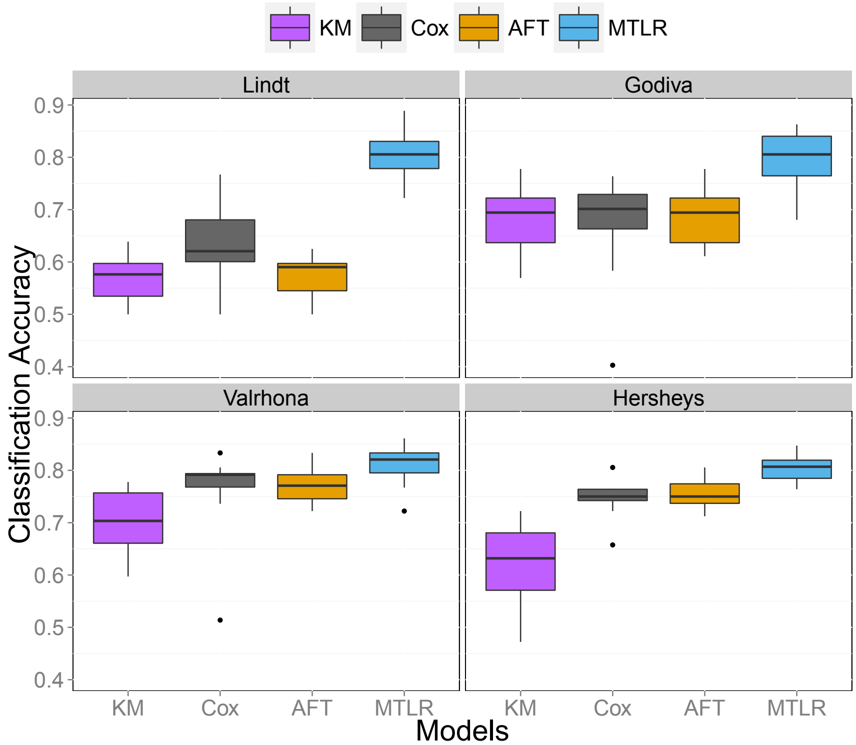

Figure 5.3: Classification accuracy (ten-fold cross validation).

$$\hat{v}_i(c) \quad = \quad \arg\max_v \{(v - c) \cdot P\hat{P}F_\omega(\vec{x}_i, v)\} \tag{5.5}$$

The true mean profit $PFT(c)$ with cost being c for product ω is computed as

$$PFT_\omega(c) \quad = \quad \frac{1}{N} \sum_{i=1}^{N} \mathbf{I}\{\hat{v}_i(c) \leq r_{i,\omega}\} \cdot (\hat{v}_i(c) - c) \tag{5.6}$$

However, since we are not sure about the true manufacturing cost c of chocolates, we compute (ten-fold CV) mean profit, for each cost $c \in \{\$0,\ \$0.1, \$0.2, ..., \$4.9,\ \$5.0\}$; see Figure 5.4.

For comparison, we also evaluate the performance of three probabilistic classifiers, *i.e.*, naïve Bayes (NB), logistic regression (LR) and random forest (RF) on the profit criterion, as these models can also be used to estimate PPF, though in a different manner. To do this, we define the purchasing decision variable as $A_c \in \{0, 1\}$ and use $(\vec{X}, V)$ as the input variables, where V is the product price variable. That is, the PPF can be represented as

$$PPF'_\omega(\vec{x}_i, v) \;=\; Pr(A_c = 1 | \vec{X} = \vec{x}_i, V = v)$$

However, note that for RF and NB, the $PPF'_\omega(\vec{x}_i, v)$ is not guaranteed to be monotonically non-increasing over price v.

The "Retail" lines in Figure 5.4 show the profit associated with selling the product at the standard price (see Table 4.1). We see in Figure 5.4 that this naïve approach, *which is the standard approach*, returns very low profit — indeed, for many costs, essentially all of the predictors do much better.

To further evaluate whether our idea, of estimating the consumer-specific RP instead of uniform RP for all consumers, helps in practice, we computed the profit of selling the product at six different fixed prices ($1.0, $2.5, $4.0, $5.5, $7.0, $8.5) and compare them with the profit earned by MTLR. As shown in Figure 5.5, at almost all cost points, the profit of MTLR is higher than selling the product at any fixed price. This result strongly support our idea of estimating RP on individual level.

5.4 Discussion

First, the fact that two survival models (MTLR, AFT), with no knowledge of the true RP, can beat the cheating baseline on the MAE evaluation criterion, shows that even without direct measurement of consumers' true RPs, but only censored observations, we can still have pretty good estimations of consumers' RPs. It suggests that our way of collecting data may work in practice for the RP estimation task.

Moreover, it is worth noting that the performance on all three evaluation criteria of the personalized models — *i.e.*, MTLR, AFT, and Cox — are generally much better than the non-personalized one, *i.e.*, KM. This can serve as strong empirical support for our idea of modeling consumer-specific RPs instead of a uniform RP. Moreover, the fact that the personalized MTLR obtained more profit than selling the product at any fixed price further bolsters our personalized setting.

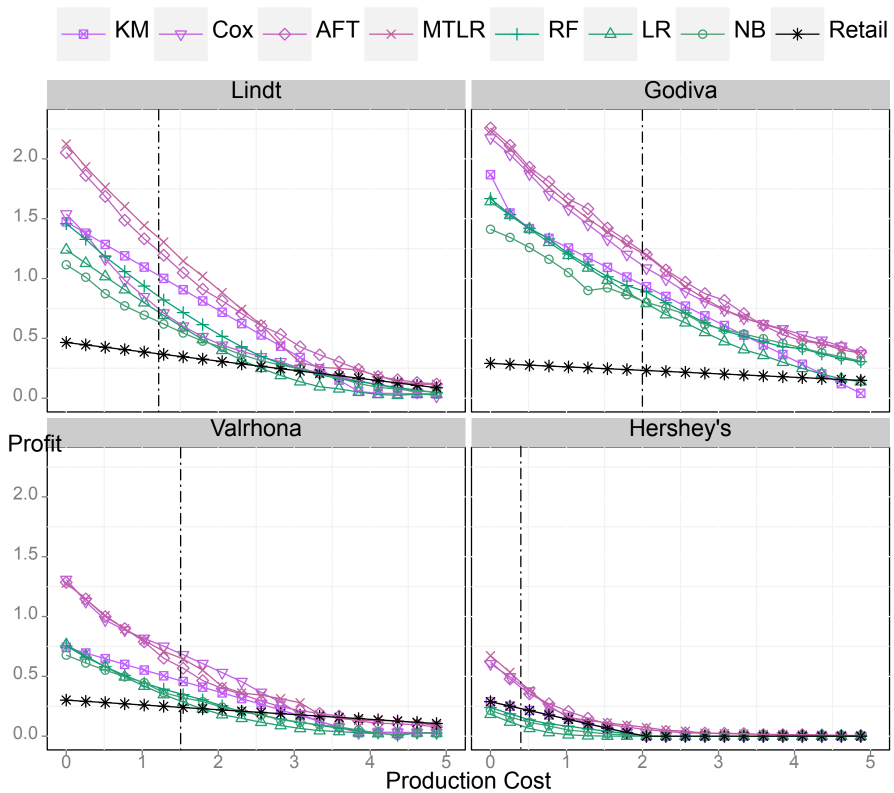

Figure 5.4: Profits across production costs from \$0 to \$5 (ten-fold cross validation). The dashed line indicates the likely production cost of the chocolate, which is %20 of its retail price [10].

Additionally, the great performance of survival models (all much better than the baselines) on the *ACC* and *PFT* evaluation criteria suggests that our way of estimating RP can be helpful in reality for predicting if consumers would accept an offer or not and achieve better profit. This is extremely useful for online retailers who want to conduct private promotions or general first degree price discrimination [27].

Among all survival models, MTLR did extremely well — better than the others — on all three evaluation criteria, on all four chocolate datasets. This suggests that MTLR is probably a good choice for predicting RP distribution predictions, in general. Besides right censored and event data, MTLR can also handle left censored and interval censored data without modifications, while most packages of KM and Cox only deal with event and right censored data.

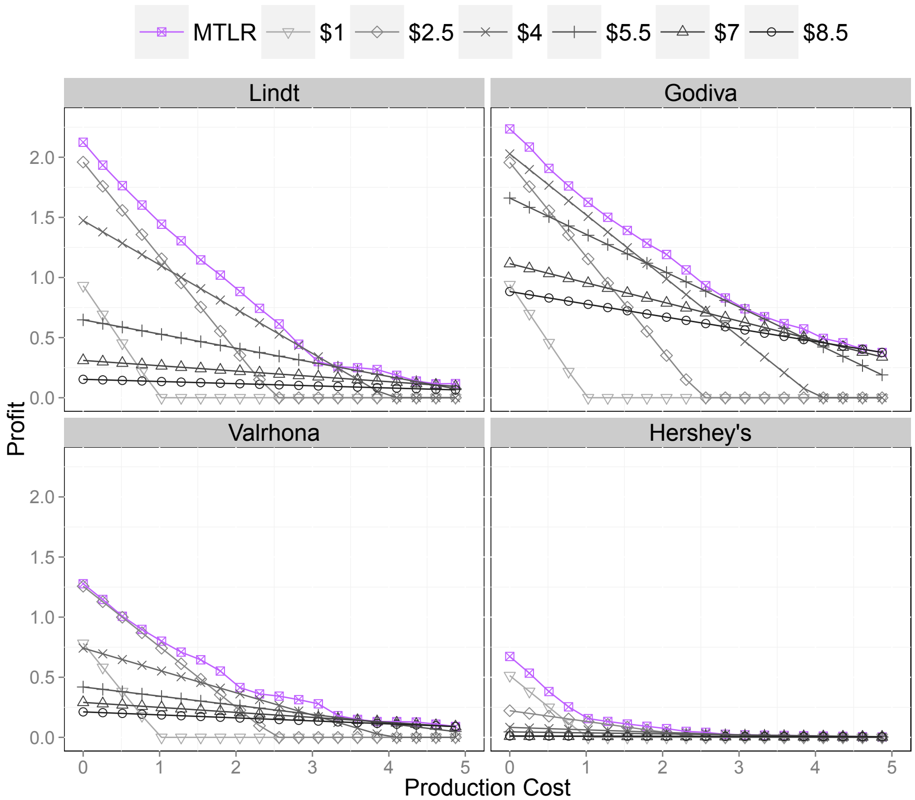

Figure 5.5: Profit comparison between fixed-prices and MTLR across production costs from \$0 to \$5 (ten-fold cross validation).

Note that these strong results are based on four similar and relatively small datasets, of only 722 subjects. We anticipate getting significantly better results if we uses larger datasets, with more consumers, and perhaps more information about each.

Chapter 6

Conclusion

6.1 Future Work

6.1.1 Transaction-Specific RP Estimation

One easy but meaningful extension is to integrate the product features $\vec{Y}$ in the dataset, with which we could achieve *(consumer, product)-specific* reservation price estimation. This is an easy extension to our framework and involves just estimating $R|\vec{x}, \vec{y}$ instead of $R|\vec{x}$. By doing this, we can

- learn the model for different products at the same time, instead of treating each as an independent dataset.
- transfer the knowledge of RP between similar products, *e.g.*, we can even estimate a consumer's RP for a new product, as long as we know the product features.

However, as we only have data about four similar products right now, it is not realistic for us to experiment on this task.

A more ambitious goal is to include other information, such as transaction time and transaction location. Again, our framework can easily handle this case too.

6.1.2 Unbalanced Data

In real settings, most consumers will turn down most products — that is, most consumers will not accept most offers from online retailers. This means that most datasets will be (seriously) unbalanced [20], where the degree of unbalance will depend on several factors, such as the promotion strategy, distribution of offer price and the product itself. We plan to further study this direction, to see if survival models can handle such very unbalanced datasets.

6.1.3 Online Predictor

Suppose we have two consumers $\vec{x}_i$ and $\vec{x}_j$ where $\vec{x}_i$ and $\vec{x}_j$ are very similar, then we find that $\vec{x}_i$ declines our offer for ω at $v_i = \$5$. Should we then offer ω to $\vec{x}_j$ at a price higher than \$5? Probably not, as $\vec{x}_i$ and $\vec{x}_j$ are similar. This example argues that we should generate the offers sequentially, utilizing the previous responses, as this may be better than generating the offers v_i in a batch mode. This leads to many interesting contextual bandit issues, and associated analyses [3]. We plan to extend our system to this on-line context.

6.2 Contributions

Motivated by the new demands of e-commerce, my thesis has proposed a novel framework of estimating consumer-specific reservation price, which consists of a consumer decision-making model, and a corresponding data collection method.

This framework has three major advantages over the traditional elicitation methods in the marketing literature, which help it meet the new demands of the e-commerce scenario:

- It captures the inherent uncertainty of reservation price.
- It connects the RP estimation task to survival prediction, which allows us to use survival models (standard and novel) to perform an individual-level RP prediction based on consumer-specific information.
- It is much easier and more practical for online retailers to implement our framework than the traditional elicitation method. As our data collection method does not ask consumers to report their true RPs, but

indirectly infers consumers' RPs based on historical (non-)transaction data.

The experimental results show that survival prediction models, especially MTLR, perform well on this task under three different criteria. This empirically shows that our framework of RP is meaningful and useful in practice. Given this success based on a relatively small dataset, we anticipate that others may try this approach on larger datasets.

Last, motivated by the great performance of MTLR on the RP estimation task, we further explored this model, and succeeded in developing a more general framework, which provides more flexibility. We also prove that the smoothness regularizer $\|\theta_j - \theta_{j-1}\|$ in the original MTLR model is not useful, as the smoothness of predicted distribution is actually controlled by $\|\theta_j\|$ instead; see Appendix A.

Bibliography

[1] Neeraj Arora, Xavier Dreze, Anindya Ghose, James D Hess, Raghuram Iyengar, Bing Jing, Yogesh Joshi, V Kumar, Nicholas Lurie, Scott Neslin, et al. Putting one-to-one marketing to work: Personalization, customization, and choice. *Marketing Letters*, 19(3-4):305–321, 2008.

[2] Santiago Balseiro, Jon Feldman, Vahab Mirrokni, and S Muthukrishnan. Yield optimization of display advertising with ad exchange. In *Proceedings of the 12th ACM conference on Electronic commerce*, pages 27–28. ACM, 2011.

[3] Gábor Bartók and Csaba Szepesvári. Partial monitoring with side information. In *Algorithmic Learning Theory*, pages 305–319. Springer, 2012.

[4] Gordon M Becker, Morris H DeGroot, and Jacob Marschak. Measuring utility by a single-response sequential method. *Behavioral science*, 9(3):226–232, 1964.

[5] Christopher M Bishop et al. *Pattern recognition and machine learning*, volume 1. springer New York, 2006.

[6] Rich Caruana. *Multitask learning*. Springer, 1998.

[7] Vidyanand Choudhary, Anindya Ghose, Tridas Mukhopadhyay, and Uday Rajan. Personalized pricing and quality differentiation. *Management Science*, 51(7):1120–1130, 2005.

[8] Corinna Cortes and Vladimir Vapnik. Support-vector networks. *Machine learning*, 20(3):273–297, 1995.

[9] David R Cox et al. Regression models and life tables. *JR stat soc B*, 34(2):187–220, 1972.

[10] Donna Ferguson. The real cost of... a chocolate bar. https://uk.finance.yahoo.com/news/the-real-cost-of—-a-chocolate-bar-115659540.html, 2013. [Online; accessed 21-Jan-2015].

[11] Gregory W Fischer, Mary Frances Luce, and Jianmin Jia. Attribute conflict and preference uncertainty: Effects on judgment time and error. *Management Science*, 46(1):88–103, 2000.

[12] William H Greene. *Econometric analysis*. Pearson Education India, 2003.

[13] John R Hauser and Glen L Urban. The value priority hypotheses for consumer budget plans. *Journal of Consumer Research*, 12(4):446, 1986.

[14] Volkmar Henschel, Christiane Heiß, and Ulrich Mansmann. intcox: Compendium to apply the iterative convex minorant algorithm to interval censored event data, 2009.

[15] S. Hettich and S. D. Bay. The UCI KDD Archive [http://kdd.ics.uci.edu]. Irvine, CA: University of California, Department of Information and Computer Science, 1999. [Online; accessed 30-Sept-2014].

[16] Kamel Jedidi and Z John Zhang. Augmenting conjoint analysis to estimate consumer reservation price. *Management Science*, 48(10):1350–1368, 2002.

[17] Shlomo Kalish and Paul Nelson. A comparison of ranking, rating and reservation price measurement in conjoint analysis. *Marketing Letters*, 2(4):327–335, 1991.

[18] Edward L Kaplan and Paul Meier. Nonparametric estimation from incomplete observations. *Journal of the American statistical association*, 53(282):457–481, 1958.

[19] Rajeev Kohli and Vijay Mahajan. A reservation-price model for optimal pricing of multiattribute products in conjoint analysis. *Journal of Marketing Research*, pages 347–354, 1991.

[20] Sotiris Kotsiantis, Dimitris Kanellopoulos, Panayiotis Pintelas, et al. Handling imbalanced datasets: A review. *GESTS International Transactions on Computer Science and Engineering*, 30(1):25–36, 2006.

[21] Klaus M Miller, Reto Hofstetter, Harley Krohmer, and Z John Zhang. How should consumers' willingness to pay be measured? an empirical comparison of state-of-the-art approaches. *Journal of Marketing Research*, 48(1):172–184, 2011.

[22] Mehryar Mohri and Andres Munoz Medina. Learning theory and algorithms for revenue optimization in second-price auctions with reserve. *arXiv preprint arXiv:1310.5665*, 2013.

[23] Sridhar Moorthy, Brian T Ratchford, and Debabrata Talukdar. Consumer information search revisited: Theory and empirical analysis. *Journal of Consumer Research*, pages 263–277, 1997.

[24] Wei Pan. Extending the iterative convex minorant algorithm to the Cox model for interval-censored data. *Journal of Computational and Graphical Statistics*, 8(1):109–120, 1999.

[25] Roland T Rust, J Jeffrey Inman, Jianmin Jia, and Anthony Zahorik. What you don't know about customer-perceived quality: the role of customer expectation distributions. *Marketing Science*, 18(1):77–92, 1999.

[26] Greg Shaffer and Z John Zhang. Competitive one-to-one promotions. *Management Science*, 48(9):1143–1160, 2002.

[27] Benjamin Reed Shiller. First degree price discrimination using big data. Technical report, 2013.

[28] Bernard L Simonin and Julie A Ruth. Bundling as a strategy for new product introduction: Effects on consumers' reservation prices for the bundle, the new product, and its tie-in. *Journal of Business Research*, 33(3):219–230, 1995.

[29] S Smith. 4 ways to ensure valid responses for your online survey. http://www.qualtrics.com/blog/online-survey-valid-responses/, 2013. [Online; accessed 30-Sept-2014].

[30] Terry M Therneau. *A Package for Survival Analysis in S*, 2014. R package version 2.37-7.

[31] Bruce W Turnbull. Nonparametric estimation of a survivorship function with doubly censored data. *Journal of the American Statistical Association*, 69(345):169–173, 1974.

[32] David Ulph and Nir Vulkan. *Electronic commerce and competitive first-degree price discrimination*. University of Bristol, Department of Economics, 2000.

[33] H.R. Varian. *Microeconomic Analysis*. Norton International edition. W.W. Norton, 1992.

[34] R Venkatesh and Vijay Mahajan. 11 the design and pricing of bundles: a review of normative guidelines and practical approaches. *Handbook of pricing research in marketing*, page 232, 2009.

[35] William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of finance*, 16(1):8–37, 1961.

[36] Tuo Wang, R Venkatesh, and Rabikar Chatterjee. Reservation price as a range: An incentive-compatible measurement approach. *Journal of Marketing Research*, 44(2):200–213, 2007.

[37] LJ Wei. The accelerated failure time model: a useful alternative to the cox regression model in survival analysis. *Statistics in medicine*, 11(14-15):1871–1879, 1992.

[38] Chun-Nam Yu, Russell Greiner, Hsiu-Chin Lin, and Vickie Baracos. Learning patient-specific cancer survival distributions as a sequence of dependent regressors. In *NIPS*, pages 1845–1853, 2011.

[39] Ji Zhu and Trevor Hastie. Kernel logistic regression and the import vector machine. *Journal of Computational and Graphical Statistics*, 14(1), 2005.

Appendix A

Multi-task Logistic Regression

A.1 Reformulation of MTLR

First, instead of viewing MTLR as a sequence of logistic regressors as discussed in Section 3.2.4, we reformulate the survival problem into a multi-class classification problem and motivate MTLR as a multi-class softmax classifier. This simplifies the next two sections and gives some insight about how this model works.

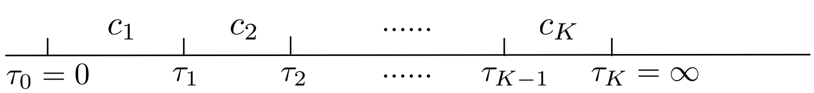

Figure A.1: K intervals and classes.

Similar to [38], we discretize the time axis into K intervals as shown in Figure A.1, but here we treat each interval as a class. We define a new random variable C out of the event time variable T as

$$C = c_j \quad \text{iff} \quad \tau_{j-1} \leq T < \tau_j, \tag{A.1}$$

then each uncensored patient is associated with a class label, while for censored patients, we know that their true class label $C|\vec{x} \ \in \ \{c_j \mid T_c < \tau_j\}$, where T_c denotes the censored time.

It is intuitive that we would like the model to have the smoothness property that if $\tau_j - \tau_k$ is small,

$$Pr(C = c_j \mid \vec{x}) \quad \approx \quad Pr(C = c_k \mid \vec{x}). \tag{A.2}$$

Suppose that we have a unique probability expression for each class, then one intuitive way to achieve the smoothness property is to guarantee the probability expressions of classes for close intervals are similar to each other. This observation suggests the candidate way of modeling the probability for different class c_j

$$Pr(C = c_j|\vec{x}) = \frac{\exp(\vec{x}^T\theta_j + ... + \vec{x}^T\theta_{K-1})}{Z(\Theta, \vec{x})}, \tag{A.3}$$

where $Z(\Theta, \vec{x}) = \sum_{j=1}^{K} \exp(\vec{x}^T\theta_j + ... + \vec{x}^T\theta_{K-1})$ is the normalization term. Here, for simplification, we omit the intercept terms $B = \{b_1, b_2, ..., b_{K-1}\}$, which are usually not important in a softmax model.

The difference between $Pr(C = c_j|\vec{x})$ and $Pr(C = c_k|\vec{x})$ in the exponential term is

$$\sum_{t=j}^{k-1} \vec{x}^T\theta_t$$

which means that the closer j is to k, the more similar the probability expression of $Pr(C = c_j|\vec{x})$ is to that of $Pr(C = c_k|\vec{x})$. Therefore, the smoothness property is achieved.

It is not hard to see that this way of coding described above is exactly the same with the MTLR model described in Section 3.2.4.

A.2 Smoothness

In the last section, we motivate MTLR from a different perspective and transform the whole problem into a multi-class classification problem. In this section, we provide more detailed analysis about the smoothness property of the predicted probability distribution.

Supposed that τ_j and τ_{j+1} are very close to each other, then a good model has to predict

$$\hat{Pr}(C = c_j \mid \vec{x}) \quad \approx \quad \hat{Pr}(C = c_k \mid \vec{x}).$$

as justified in the previous section.

Here we compute the ratio of $Pr(C = c_{j-1}|\vec{x})$ and $Pr(C = c_j|\vec{x})$, which is

$$r_j = \frac{Pr(C = c_{j-1}|\vec{x})}{Pr(C = c_j|\vec{x})} = \exp(\vec{x}^T \theta_j) \tag{A.4}$$

Thus a sufficient condition for r_j to be close to one, which means that $Pr(C = c_{j-1}|\vec{x}) \approx Pr(C = c_j|\vec{x})$, is that $\|\theta_j\|$ is small.

Surprisingly, this result indicates that the general regularizer $\sum_{k=1}^{K-1} \|\theta_j\|$ not only prevents overfitting, but also controls the smoothness of the predicted probability distribution.

Now we examine how the original smoothness regularizer term $\|\theta_j - \theta_{j+1}\|$ works. As $r_j = \exp(\vec{x}^T \theta_j)$ and $r_{j+1} = \exp(\vec{x}^T \theta_{j+1})$, if θ_j is close to θ_{j+1}, then we have

$$r_j = \exp(\vec{x}^T \theta_j) \approx \exp(\vec{x}^T \theta_{j+1}) = r_{j+1}$$

which leads to

$$\frac{Pr(C = c_{j-1}|\vec{x})}{Pr(C = c_j|\vec{x})} \approx \frac{Pr(C = c_j|\vec{x})}{Pr(C = c_{j+1}|\vec{x})} \tag{A.5}$$

It is unclear if this property will help, as it suggests a constant ratio of the probabilities for adjacent times.

We therefore tried omitting this regularizer, and instead just sought the parameters that optimized the log-likelihood with the general regularizer $\|\Theta\|$ only. Our experimental results, over several survival datasets show that the performance of MTLR under several evaluation criteria, *e.g.*, concordance index, mean absolute error, mean square error and log-likelihood, does not degenerate without the smoothness regularizer.

One major advantage of eliminating this regularizer is that we could save lots of training time, as we now have one less regularizer constant to tune in the cross validation phase.

A.3 Generalization of MTLR

This section illustrates how we can generalize MTLR into a more general form and discusses several special cases.

Here with some simple linear algebra operations, we can rewrite the $Pr(C = c_j|\vec{x})$ as

$$
\begin{aligned}
\Pr(C = c_j|\vec{x}) &= \frac{\exp(\vec{x}^T \sum_{k=j}^{K-1} \theta_k)}{Z(\Theta, \vec{x})} \\
&= \frac{\exp(\vec{x}^T \Theta \mathbf{G} \vec{y}_j)}{Z(\Theta, \vec{x})}
\end{aligned}
\tag{A.6}
$$

where $\mathbf{G} = (\vec{g}_1, \vec{g}_2, ..., \vec{g}_K)$ and $\vec{g}_j$ is a $(K-1)$-dim vector with first $j-1$ elements being zero. We call $\mathbf{G}$ the coding matrix. See Figure A.2 (left) for an example with $K = 4$. The definition of $\vec{y}$ can be found in Section 3.2.4.

1	0	0	0
1	1	0	0
1	1	1	0

G of MTLR
$K = 4$
(left)

1	0	0	0
0	1	0	0
0	0	1	0

G of MNLR
$K = 4$
(mid)

1	0.5	0	0
1	1	0.5	0

G of $m = 2, K = 4$

(right)

Figure A.2: Examples of possible coding matrices.

By using different $\mathbf{G}$s (not restricted to square matrices), we can derive different multi-class classifiers out of this general formula.

For example, if $\mathbf{G}$ in Figure A.2(mid) is adopted, then this model turns into the *multinomial logistic regression* (MNLR) [5]. In this case, the smoothness regularizer would be quite helpful to smooth the distribution prediction.

We can even use an $m \times K$ dimensional $\mathbf{G}$, where $m \neq K - 1$ (*e.g.*, Figure A.2 (right)). In this case, the dimension of Θ would also be altered to be $p \times m$ instead of $p \times (K-1)$. If we have $m < K - 1$, then we will have less parameters in Θ to tune, which speeds up the training process.

With this extension, one can design his/her own $\mathbf{G}$ with appropriate regularizers to fulfill his/her own need, which comes in handy. Note that the two basic principles for choosing $\mathbf{G}$ are that

- No two classes share the same probability expression (which means the columns of $\mathbf{G}$ must be different)

- Classes representing time intervals close to each other should have similar probability (smoothness property) — which means G's adjacent columns should be similar.

As similar probability expressions would be a sufficient condition for smoothness, generally we want the probability expressions of close time intervals to be similar. Meanwhile, an obvious exception is the MNLR model. The probability expressions for all classes are different from each other, but the smoothness regularizer $\|\theta_j - \theta_{j-1}\|$ can still serve the purpose of restricting the probabilities at close time intervals to be close. (So for MNLR, one might want to continue using smoothness regularizer.)

This general formulation of MTLR also allows us to easily kernelize the model. Also this model is related to import vector machine (IVM) [39] — an extension of support vector machine [8] — which gives probabilistic output for multi-class classification problem. Exploring the relationship between IVM and MTLR is also of future interest.