Pricing Virtual Goods: Using Intervention Analysis and Products' Usage Data

by

Lin Yang

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Master of Applied Science
in
Management Sciences

Waterloo, Ontario, Canada, 2014

AUTHOR'S DECLARATION

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

The rapid growth of online games enables firms to charge players for virtual goods they sell for use within their online game environments. Determining prices for such virtual goods is inherently challenging due to the absence of explicit supply curve as the marginal cost of producing additional virtual goods is negligible. Utilizing sales data, we study daily revenue of a firm operating a virtual world and selling cards. Explicitly, we analyze the impact of new product releases on revenue using ARIMA with intervention model. We show that during initial days after a new product release, the firm's daily revenue significantly increases. Using a quality measure, based on the Elo rating method, we can determine the relative good prices based on good usage. Applying this method we show that the rating of a product can be a good proxy for units sold. We conclude that our quality-based measure can be adopted for pricing other virtual goods.

Acknowledgements

I am deeply indebted to Dr. Benny Mantin and Dr. Stanko Dimitrov for their guidance and constant supervision. Their positive attitude and encouragement helped me immensely to achieve my goal. I will always remember the time Dr. Stanko Dimitrov spent to help me develop an academic and professional thinking in research. I am equally grateful to Dr. Benny Mantin for his meticulous attention to detail and his critical thinking when doing research. This dissertation would not be complete without the support of my two supervisors.

I would like to thank Dmytro Korol who dedicated invaluable inputs to our project during his co-op term. We had really intensive and meaningful discussions during the progress. I benefited greatly from his passion for research. My gratitude towards our industry partner who was providing necessary data information, and all of the individuals at the organization that supported and helped us throughout the process.

It is my pleasure to thank Amy E. Greene who offered her time and shared experiences to help me with improving writing skills almost everyday when I started writing my thesis.

I owe all my achievements to my family. I am grateful to my parents, Chuanfeng Yang and Yuzhi Wang, and my brother, Mu Yang. Their unswerving support has helped me overcome difficult times. I have drawn strength and inspiration from them and am fortunate to find my role models so close to me in life.

Finally, I would like to thank Dr. Brian Cozzarin and Dr. Selcuk Onay for their careful reviewing and constructive criticism of this thesis.

Table of Contents

List of Figures

List of Tables

Chapter 1
Introduction

The inception of the Internet allowed entrepreneurs to create new economies based on goods that have no physical properties besides some bits on a hard drive or network, commonly referred to as *digital goods*. In this work we focus on a subset of digital goods called *virtual goods*, goods that are exchanged and used only within online virtual environments. Pricing virtual goods is of great interest to companies operating virtual environments; the standard supply curve is no longer applicable as it can be set by the manufacturer at little to no additional cost. Working with an industry partner who sells virtual goods, we investigate how virtual goods can be priced in order to increase revenue. We characterize the impact of marketing externalities on revenue and show quality of virtual goods needs to be considered when setting prices. Accordingly, we propose a measure of quality using virtual goods' *usage data*, how players use the goods throughout gameplay.

Virtual goods can vary dramatically from being simple goods, such as a card in a game, to rather complex goods, such as a complete virtual 3D environment like a home. Companies that design and manage virtual environments have different business models. Generally, these are categorized as either *pay-to-play*, where payment is required to access the virtual world, or *free-to-play*, where access to the game is free but virtual goods are purchased during gameplay. Our interest is in determining pricing of virtual goods in order to increase revenue under the free-to-play model. In various free-to-play games, players need a collection of items such as cards or weapons. During the game, players can purchase additional items either by buying individual items or through a bundle of unknown items referred to as *packs*, as is the case with our industry partner. In the latter case, the identity of the items is revealed after the purchase (although the probability of receiving any one item may be known). Additionally, as a virtual game evolves, new items are released and become available for purchase. We focus on pricing virtual goods when two purchase channels are available (individual and packs) and new items are released to the markets.

One approach to determine prices is to the find the market equilibrium using the appropriate supply and demand curves. However, with many virtual goods, the marginal cost can often be assumed to be zero, raising the need for alternate approaches. Further, release of new goods can stimulate demand and alter choices made by players. Accordingly, we show how to approximate the quality of newly released virtual goods and use that information to set their initial prices.

Although this work is intended to price cards that are played in an online game and enable the firm to increase its revenue from selling cards to players, the methods used may be applicable to different settings. Namely, our method can be widely applied to pricing digital goods with limited historical usage data, especially when it is difficult to define the supply and the demand functions and price may be a function of past usage.

The remainder of the thesis is organized as follows: Chapter 2 reviews background information on virtual worlds while Chapter 3 provides basic information on the game we study. Chapter 4 then explores the daily revenue of the firm using time-series analysis. The impact of new product releases on the daily revenue is determined by applying ARIMA modeling with intervention analysis. Chapter 5 presents our Set-based Elo rating method to evaluate product quality derived from product usage. Finally, Chapter 6 concludes and further highlights the practical benefits from the study and methods used in this study.

Chapter 2
Virtual Worlds: Background

In this chapter we first review related work on the online game industry, and we proceed to discuss two revenue models of online games. Instead of studying the players' purchasing motivations, researchers tend to investigate the attributes of virtual goods that stimulate demand. Interesting findings from the related work inspire us to explore a pricing method using players' usage data collected under a certain game environment.

One of the most popular segments in virtual worlds is online game. It is estimated that the worldwide revenue from online games was $21 and $24 billion in 2012 and 2013, respectively, and is expected to reach $35 billion by 2018 (DFC Intelligence, 2013). This revenue comes from a variety of sources, such as digital delivery, subscriptions, and social networking services. As mentioned in Chapter 1, online games broadly have two revenue models: pay-to-play and free-to-play. Before 2005, most online games used different variants of pay-to-play. In general, companies used a combination of a one-time access fee and a periodic subscription fee (Meagher and Teo, 2005). This approach allowed players to select a pricing method that works best for them, and allowed companies to extract a higher surplus from consumers.

In a free-to-play setting, as the name suggests, accessing the game is free of charge. However, oftentimes players may be exposed to advertising or be incentivized to make in-game purchases of goods and services. As such, game companies and researchers examine games and product attributes that lead to players spending money during a game. For example, Lehdonvirta (2005) determines that it is players' motivation for playing a game that leads players to purchase various in-game virtual good. In a follow-up study, Lehdonvirta (2009) discusses the functional, hedonic and social attributes of virtual goods. The author finds that it is virtual goods' qualities, as measured by those three attributes, that lead to players spending money within a game. In this thesis we do not consider the sociological or physiological aspects of virtual-goods, instead we assume that the goods are given and fixed. Our objective is to set goods' initial prices as they are released in order to increase revenue over the method currently used by our industry partner.

Nojima (2007) reveals that game players that spend real money to acquire virtual goods tend to play the game longer and assess the value of virtual goods during gameplay. Motivated by Nojima's finding, we use product usage in the virtual environment as a proxy for the product quality. Hamari and Lehdonvirta (2010) propose that game design can be integrated with marketing, exposing players to products throughout gameplay, such as in-game promotions (Hamari, 2009). Hamari and Lehdonvirta (2010) further suggest setting expiration dates or degradation rates for virtual goods to mimic real-world goods. Oh and Ryu (2007) propose that some virtual goods be sold for real money while others be sold for virtual currency (in-game currency). The reason is that actual major performance upgrades need to be obtained through gameplay rather than being purchased, however, purchases can help reduce the time needed to reach gameplay goals.

Lehdonvirta (2005) considers operational issues when selling virtual goods. The author proposes using a dynamic price updating mechanism to improve revenue from virtual goods as demand changes over time. There are other digital goods that have similar attributes, i.e., insignificant marginal cost, as virtual goods, such as software service. Bala and Carr (2010) investigate both fixed and usage-based pricing schemes of software service in a competitive setting. In recent times, instead of selling goods, firms charge a subscription service fee and customers decide whether to renew the service after each time period (Bala, 2012). Researchers continue to study subscription-based digital goods that are dynamically priced at each stage of their life cycle in order to maximize total revenue. The general consensus is that optimal prices follow a decreasing trend (Yao *et al*, 2012).

Price updating mechanisms can be placed into one of two categories: *post-price* mechanisms and *price-discovery* mechanisms (Elmaghraby and Keskinocak, 2003)[1]. Post-price mechanisms have sellers set goods' prices, for example price tags are one such mechanism. Price-discovery mechanisms involve customers and sellers discovering the price together, for example an auction is one such mechanism. In this thesis we cover a post-price mechanism for selling a collection of virtual goods. We propose incorporating product usage in a post-price mechanism.

[1] Elmaghraby and Keskinocak (2003) define the distinction between these two pricing mechanisms. Yet, we recognize that the core of their paper is dedicated to reviewing the literature on dynamic pricing mechanism of perishable goods.

Chapter 3
Gameplay Background

In this chapter we provide some additional details on the game we consider for this study. In the online game, players battle with each other head-to-head using a subset of cards acquired during gameplay. The game puts hard restrictions on the makeup of the subset of cards selected, called a *deck*. Each player is represented by a hero in a battle, while the cards—representing weapons, potions, or spells—work on the targeted hero. The ultimate goal in this game is to damage the opponent's hero by taking away the hero's health points. The first player to take all of his opponent's hero's health points wins the battle. Players play in a sequence of games either with each other, against other teams of players or against a computer opponent. The core functionality of the game is the same, but the global objectives change depending on the type of gameplay selected by a player. For example, in team play, a team with the most head-to-head wins is said to win the team battle.

Transactions in the game are executed either using premium or common in-game currencies. The premium in-game currency is usually acquired through conversion of real money into this currency. Players receive some premium in-game currency upon registration and also through regular gameplay; however, the total amount of premium in-game currency players can obtain in non-purchase transactions is very limited. Common in-game currency is earned during regular gameplay and cannot be purchased. Accordingly, since premium in-game currency is actually purchased by players using real money, our analysis only considers the premium currency. We limit our attention to the total revenue obtained on a daily basis.[2]

[2] The interplay between sales obtained from packs versus individual cards is explored in the appendix.

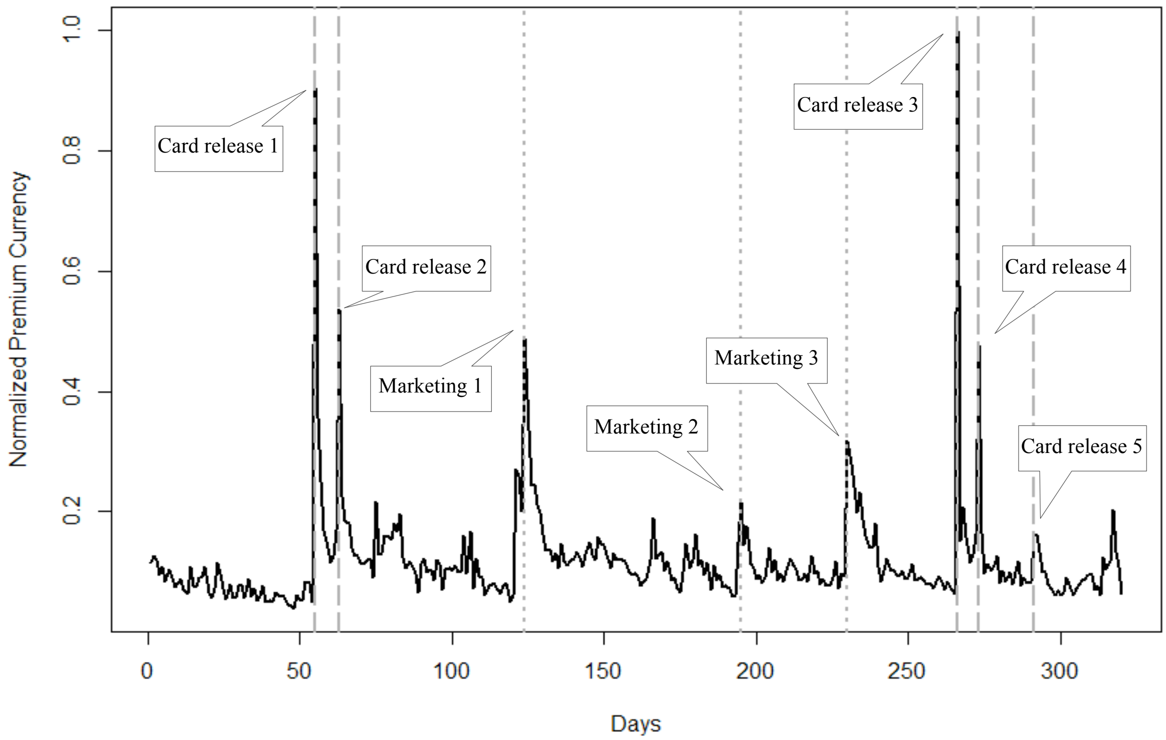

Figure 1: Daily total revenue and timing of new card releases and other external events

In Figure 1, we present the normalized daily revenue over 320 days for our industry partner.[3] From this figure, one can observe eight spikes in the daily revenue each corresponding to an external intervention. These spikes are individually labeled. Generally, we distinguish between two types of spikes: releases of new cards and other event types. Five of the spikes are of the former type referred to as Card release in this study. The latter type of events (three instances) includes advertising and card attribute changes referred to as Marketing. In the figure, long dashed gray lines are new card releases and dot grey lines are other marketing events.

[3] Note: the revenue has been normalized to the request of our industry partner. Additionally, the data has been cleaned to remove any anomalies such as removal of employees' accounts and automated purchases made by computer programs.

Chapter 4
Exploring Daily Revenue Spikes Through Time-series Analysis

Releases of new cards have a substantial impact on players' purchasing decisions, and consequently on the firm's revenue. In this chapter we carry out a time-series analysis to determine the impact of new releases on revenue. As mentioned in the previous chapter, we observe multiple spikes in daily revenue which can be categorized into two types: new card releases and other events (such as advertising campaigns), both of which can be classified as external interventions. Since we have a sequence of daily observations and we seek to extract insights regarding the effects of the external events on the daily revenue, we proceed by carrying out a time-series analysis. The time-series analysis accounts for the fact that a data sequence may have a possible internal structure, such as autocorrelation, trend or seasonal pattern. As the data (Figure 1) reveals no obvious trend, but exhibits a certain degree of intertemporal relationship, we make use of the *autoregressive integrated moving average* (ARIMA) model to determine the type of intertemporal relationship that exists in the data. Additionally, since we are interested in the effects induced by the external events, we embed intervention analysis in our ARIMA estimations. Specifically, our intervention analysis will help estimate the values of the spikes in generating incremental revenues. Incidentally, although our interest lies in determining the duration of the spikes, time-series analysis also provides a method for predicting impacts of similar spikes. These spikes have a big impact on daily revenue and therefore, pricing of newly released cards is of major interest, which is explored later in Chapter 5.

4.1 The fundamentals of the ARIMA model

The ARIMA model has been widely used in time-series analysis since Box and Jenkins introduced a method to determine the parameters for the model (1970), and is utilized in economics and social sciences to determine control policies and forecasting (Mccleary and Hay, 1980). This method is also employed by companies and governments for similar purposes, such as user prediction in mobile industry (Ye, 2010), or cruise tickets booking forecast (Sun *et al*, 2011). In our work, we use ARIMA modeling with intervention analysis, the impact of an event on a time series. We are not the first to use ARIMA for this purpose, for example Pearce, Stevenson and Perry (1985) used ARIMA to analyze the effect of merit pay on organization performance.

In our time-series models, we let $\{Y_t\}$ denote the observed time series of the daily revenue. The time series $\{Y_t\}$ is said to follow an ARIMA model if the d^{th} differenced series, $\{W_t\}$, is a stationary *autoregressive moving average* (ARMA) process. For example, if $d = 1$, we have $W_t = Y_t - Y_{t-1}$. ARIMA assumes homogenous non-stationary data, i.e., the mean of the time series changes from time period t to time period $t-1$, and only the difference in means between time periods is constant. ARMA model is a special case of ARIMA model in which the difference in means between time periods is assumed to be zero.

If the observed time series is stationary, then an ARMA model is applied to the sequence. There are two major components to an ARMA model, the *autoregressive* (AR) component and the *moving average* (MA) component, where $AR(p)$ denotes a p^{th} order autoregressive process, and $MA(q)$ denotes and a moving average of order q. Thus, we have:

$$Y_t = \Psi_1 Y_{t-1} + \Psi_2 Y_{t-2} + \cdots + \Psi_p Y_{t-p} + e_t + \Theta_1 e_{t-1} + \Theta_2 e_{t-2} + \cdots + \Theta_q e_{t-q}\ ,$$

where $\Psi_i, i = 1,2, \ldots, p$, and $\Theta_i, i = 1,2, \ldots, q$, are coefficients. $\{e_t\}$ represents a white noise time series, that is, a sequence of uncorrelated, with constant variance, zero-mean random variables. The current value of $\{Y_t\}$ linearly combines the p most recent past values of the time series and the q most recent past values of the white noise (Cryer and Chan, 2008).

Let B denote a *backshift operator* which recalls the previous element in a time series. For example, $BY_t = Y_{t-1}$. Using the introduced notation we know, $ARMA(p, q)$ satisfies the equation:

$$Y_t = \Psi_1 BY_t + \Psi_2 B^2 Y_t + \cdots + \Psi_p B^p Y_t + e_t + \Theta_1 B e_t + \Theta_2 B^2 e_t + \cdots + \Theta_q B^q e_t,$$

thus,

$$\varphi(B) Y_t = \theta(B) e_t,$$

where

$$\varphi(B) = 1 - \Psi_1 B - \Psi_2 B^2 - \cdots - \Psi_p B^p;$$

$$\theta(B) = 1 + \Theta_1 B + \Theta_2 B^2 + \cdots + \Theta_q B^q.$$

Therefore, a time series can be represented with $ARMA(p, q)$ model using polynomials of backshift operator (Wei, 1994) as:

$$Y_t = \frac{\theta(B)}{\varphi(B)} e_t.$$

As we discussed above, if $\{Y_t\}$ is a non-stationary sequence, but the d^{th} difference $\{W_t\}$, $W_t = (1-B)^d Y_t$, is stationary, then the sequence can be represented by an $ARIMA(p,d,q)$ model:

$$W_t = (1-B)^d Y_t = \frac{\theta(B)}{\varphi(B)} e_t.$$

4.2 ARIMA modeling with intervention analysis

Intervention analysis within an ARIMA framework was introduced by Box and Tiao (1975). Interventions affect a time series by changing the trend or the mean. In our case this happens with the release of new cards into the game. Intervention analysis is used to measure the impact of the spikes on the time series. Specifically, intervention analysis is a tool for studying the impact of policy changes or other events on a time series. For example, suppose that a new maximum speed limit is established on a highway and the police want to learn how much the new limit affects accident rates, ARIMA modeling with intervention analysis can be used in this setting.

A single intervention model that consists of a single spike in a time series can be considered as the simplest case. Using the notation of Cryer and Chan (2008), we introduce the dynamic intervention model using a linear equation. The general model for the time series is modeled by process, $\{Y_t\}$, and is given by:

$$Y_t = m_t + N_t,$$

where $\{m_t\}$, is the time series process that represents the effect of the intervention and can be modeled as a *transfer function,* m_t, defined by Box and Tiao (1975). In particular, the transfer function often uses pulse and step functions to measure intervention magnitude and duration. $\{N_t\}$ is the *underlying process* during the intervention, assuming the intervention never occurred. In our analysis, $\{N_t\}$ is defined as the underlying process during the whole time span. It is modeled as an *ARIMA* process and its parameters are estimated using the Box-Jenkins method described by Box, Jenkins and Reinsel (2013). According to Cryer and Chan (2008), the transfer function, $\{m_t\}$, is assumed to be zero before any intervention takes place.

The transfer function, $\{m_t\}$, is a combination of a step function and a pulse function with only auto-regression (*AR*). As an intervention takes place at time T, the transfer function can be modeled as:

$$m_t = \delta_1 B {S_t}^T + \frac{\delta_2 B}{1-\omega B} {P_t}^T,$$

where S_t is a step function with $S_t{}^T = \begin{cases} 1, & if\ t \geq T \\ 0, & otherwise \end{cases}$, P_t is a pulse function with $P_t{}^T = \begin{cases} 1, & if\ t = T \\ 0, & otherwise \end{cases}$. The backshift operator coefficients, δ_1, δ_2, and ω capture the relationship between time periods. δ_1 is the coefficient that describes the change in the mean due to the intervention; δ_2 is the coefficient that describes the magnitude of the spikes resulting from an intervention; ω is the auto-regression coefficient. All three coefficients are estimated from the available data. As m_t describes the intervention on the time series, the functional form of m_t enables us to potentially distinguish identical spikes.

Time-series analysis, allows us to understand how long it takes for the impact of an intervention, such as a new card release, to decay and for the time series to revert back to its mean. The four steps of this analysis are shown in Figure 2.

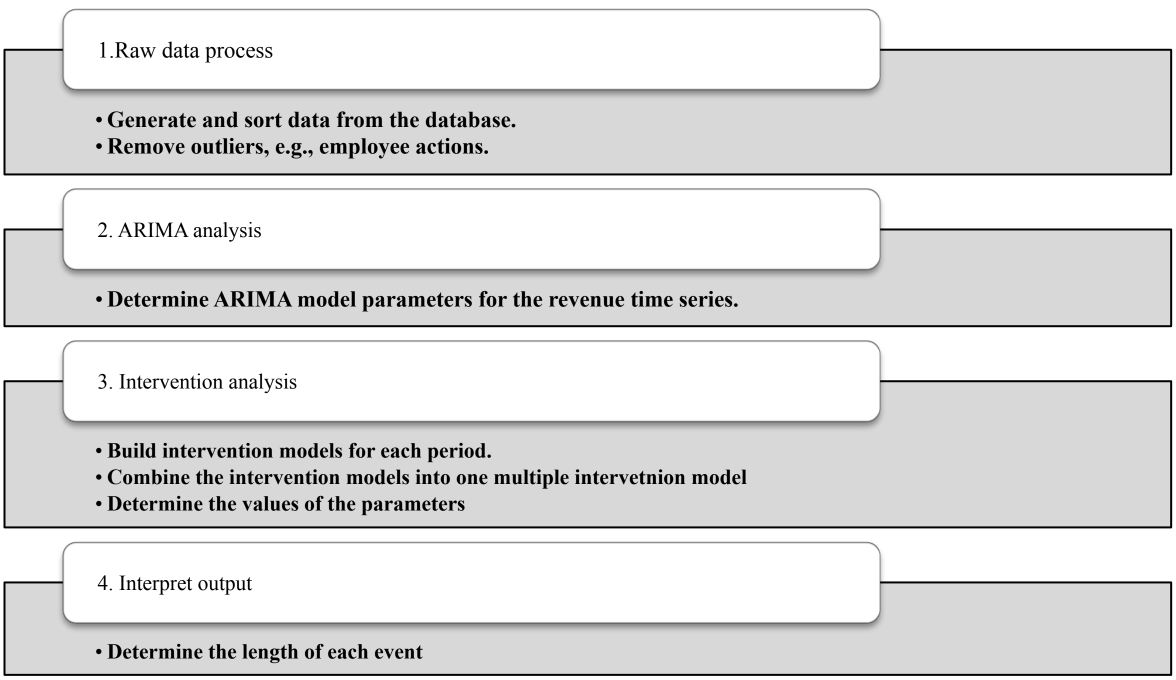

Figure 2: Intervention analysis process

1. **Preprocess Raw Data:** We process the raw data by removing any anomalies that exist such as company employees gifting their personal accounts large sums of premium currency then making a large number of purchases. We also remove all company employees' business accounts from the data and all automated purchases made by computer programs. After

removing all of the outliers, we consider the daily premium currency revenue, from here on referred to as revenue.

2. **Conduct ARIMA Analysis**: We identify the ARIMA models for prediction. First, we decide the order of the difference in ARIMA model, d, by differencing the observed sequence d times until we obtain a stationary time series. Second, we employ the Box-Jenkins method to determine the parameters of autoregressive and moving average in ARIMA model, p and q, respectively.

3. **Apply Intervention Analysis**: We apply intervention analysis to model the transfer function of each spike in the revenue time series. We determine the parameters for all of the transfer functions, using a single ARIMA with intervention analysis model, across all 320 days (Enders, 2008).

4. **Interpret Output**: We interpret the results from the intervention analysis output. When applying ARIMA modeling with intervention analysis to the revenue stream, the parameter values and their associated standard errors are obtained using *Arimax* function in *TSA* package (Chan and Ripley, 2012) from *R* (2008). These estimations are used to determine the coefficients of each intervention function, defined in the intervention model. We plug in the determined parameter values into the original ARIMA model to determine the length of the impact of each intervention on the revenue time series.

4.3 Results of intervention analysis

We employ the *Arimax()* function in R to estimate our ARIMA modeling with intervention analysis coefficients. Various estimations of ARIMA models reveal that based on the smallest Akaike Information Criterion (AIC), $ARIMA(1,0,1)$, i.e. $ARMA(1,1)$, is the best fit model. The estimated coefficients are provided in Table 1: The estimation results.

ARIMA(1,0,1)	Coefficient values	Standard error				
AR(1)	0.69	0.10				
MA(1)	-0.38	0.12				
Mean	0.08	0.01				
Estimations	δ_1	Standard error	δ_2	Standard error	ω	Standard error
Card release 1	0.07	0.02	0.76	0.03	0.30	0.04
Card release 2	-0.03	0.02	0.41	0.03	0.28	0.08
Marketing 1	-0.01	0.01	0.22	0.02	0.91	0.02
Marketing 2	-0.01	0.01	0.11	0.03	0.60	0.17
Marketing 3	-0.02	0.02	0.24	0.03	0.85	0.04
Card release 3	0.08	0.02	0.84	0.03	-0.02	0.04
Card release 4	-0.07	0.02	0.39	0.03	0.18	0.07
Card release 5	-0.01	0.02	0.09	0.03	0.66	0.13
log(likelihood) = 704.06		AIC = -1354.11		BIC = -1248.6		

Table 1: The estimation results

As shown in Table 1, the first part is the estimation on the ARIMA model. With the estimated coefficients, the underlying time series can be presented as a linear equation: $Y_t = 0.69Y_{t-1} + e_t - 0.38e_{t-1} + 0.08$. The coefficients of the transfer functions are displayed in the latter part of the table. For instance, the estimated function of the first intervention, i.e., new card release 1, can be interpreted as $m_1 = 0.07BS_t{}^T + \frac{0.76B}{1-0.3B}P_t{}^T$. The intervention model of each spike in the revenue stream could be determined using the values approximated.

4.4 The duration and magnitude of each intervention

Using the parameters derived from the intervention analysis (Table 1), the number of days, t, it took for the time series to revert back to within 10% of mean after each spike are given in Table 2. The formula that we use to calculate the duration, t, is

$$\delta_2\omega^t = 0.1(\delta_1 + a).$$

Thus,

$$t = \frac{\log \frac{0.1(\delta_1 + a)}{\delta_2}}{\log \omega}.$$

Where δ_2, δ_1, and ω are determined from the intervention analysis, and a denotes the mean before an intervention takes place. Therefore, t can be interpreted as the duration of the intervention. From the second column of Table 2, the impact on the revenue time series from a new card release, Card release 1, 2, 3, 4, and 5, is 5.8 days on average. However, the durations of other interventions, Marketing 1, 2, and 3, can reach as long as 55 days.

Interventions	Days to decay	Ratio of revenue from average intervention day over nominal day
Card release 1	6	3.4
Card release 2	5	2.6
Marketing 1	55	1.5
Marketing 2	10	1.3
Marketing 3	36	1.3
Card release 3	2	6.0
Card release 4	4	2.2
Card release 5	12	1.0

Table 2: Days to decay for each event and ratio of revenue from average intervention day over nominal day

Note that, the spikes resulting from new card releases are quite different from the spikes from other events. The spikes of the new card releases all start with a greater revenue and decay quickly within a short time period. Nevertheless, spikes of other events have lower initial magnitudes, but decay at a slower rate.

Interventions—with an emphasis on release of new cards—stimulate a surge in revenue for a few days. We find that the daily revenue during the interventions is on average 1.4 times that of during nominal days (days with no intervention affect). In particular, for new card releases interventions, even though the number of days impacted by spikes is small, the daily revenue generated during those days is on average more than twice of the revenue during nominal days. We focus on new card releases because for the two types of spikes, new card releases and other events, the daily revenue proportion of the former, is 1.6 times of the latter on average. More specifically, the daily revenue during the first two days of new release is 21.5 times of the daily revenue during nominal days on average.

From these observations, we conclude that initial card prices are of paramount importance during a new card release. The amount of revenue realized per day is higher than during any other day of the year, but the time window during which these high revenues may be realized is very small. Because of this small window, it is crucial that cards are priced correctly when released. The first step towards correct pricing is determining the relative pricing of cards by using the card usage data that we refer to as card quality.

Chapter 5
Initial Pricing of Cards: an Elo Rating Approach

In this chapter we propose an approach based on the Elo rating system to help set initial prices of card as they are released to the market. The first step is determining the cards' relative prices. This is achieved by utilizing the cards' usage data and inferring their relative inherent quality. Specifically, we observe how players use their cards, and presumably determine each card's quality, assuming players play with cards that are of higher quality. One effective method to distinguish quality in a group is the Elo rating system. Traditionally, the Elo rating system is used to evaluate the skills of players in a given game. It was originally used in chess, but is now frequently employed in online gaming to measure the skills of players.

The Elo rating system is named after Arpad Elo, who invented the method to rate chess players (Elo, 1978). The purpose of this system is to measure chess players' skill by updating their ratings as players proceed with the game over time. Reid and Nixon (2011) employ Elo ratings to reduce human inaccuracies when obtaining human descriptions of soft biometric traits, such as height, hair color, hair length, etc. Similar to Reid and Nixon (2011), we use Elo ratings to infer the relative performance of objects. Namely, we use Elo ratings to infer players' relative preference of one card to another using historical data. To understand our method, we first discuss how the Elo rating system is traditionally used.

5.1 The Elo ratings specifics

The Elo ratings measure players' strengths based on observable match outcomes, where a player either wins, loses or draws the match. The outcomes reflect players' abilities. Based on past outcomes, and assuming players' skills are normally distributed, we compute players' ratings and probability of winning upcoming matches.

A player's Elo rating increases or decreases according to the outcomes of previous games between other rated players. After every match, points transfer from the losing player to the winning player. The amount of points transferred is a function of players' ratings, with the amount determined by the difference in players' ratings. The difference in the amount of points transferred makes the Elo rating system *self-correcting*: players' true rating is found through the course of gameplay. Another advantage of the system allows new players to join at any time without having to adjust existing players' ratings.

The Elo rating system predicts match outcomes. Let $R_A$represent the Elo rating of player A before a match with player B, and let R'_A be the Elo rating of player A after the match with player B. E_A is player A's probability of winning the match against player B, and S_A is player A's actual score after the match. If player A won, $S_A = 1$; otherwise, $S_A = 0$. Ties are excluded in our study and hence are not discussed herein.

Let $Q_A = 10^{\frac{R_A}{400}}$ (Elo, 1978) be a transformation of R_A serving as a proxy of the likelihood of player A winning a match against a random opponent. The expected score of player A from playing against player B is defined as follows:

$$E_A = \frac{Q_A}{Q_A + Q_B}.$$

R'_A is updated using the following formula:

$$R'_A = R_A + K(S_A - E_A),$$

where K is the maximum possible adjustment per match, and is any positive constant set a priori. Essentially, K as specifying how much the most recent match outcome affects players' ratings.

5.2 Set-based Elo rating method

In this section we formalize our proposed method, which we refer to as the *set-based Elo rating* method. In the game we analyze, players select a subset of i cards—a deck—out of their total number of cards, k, which may very well exceed i. When playing a battle, players select their strongest deck. As we do not know how players select their decks, we assume that they use a single Round-Robin tournament to select cards. That is, within their card collections, they compare every two cards head-to-head once and select the best i cards.

However, in the available data, we only observe the top i cards, 40 in our case, selected by each player. We assume that these 40 cards are the 40 highest scoring cards in the single Round-Robin tournament. Therefore, each card individually in the 40 selected cards is said to win against each of the unselected cards in a head-to-head comparison.

The initial Elo ratings are equally assigned to each card. Once the selected card set and the unselected card set in a deck are identified, we update the Elo ratings assuming each card in the selected cards wins against each of the unselected cards in a head-to-head comparison. Since there are millions of players' decks available in the data, a higher value card will be identified by the Elo ratings determined from the described set-based comparison method across all players. We also acknowledge that cards may have synergistic properties, for example card 1 and card 2 individually may not be preferred, but together they are strongly preferred. We do not account for such synergies in our analysis, and make this simplifying assumption in order to ease our study given the available data.

We next formulate our set-based Elo rating method. Assume each player has a set of cards he owns, $\{S\}$, and each player enters a match with a deck, $\{D\}$, that is a subset of $\{S\}$. Then all cards in $\{D\}$ are considered to be the most preferred cards of $\{S\}$ and thus beat all cards that are not selected, $\{S - D\}$. We then update the Elo ratings of all cards in $\{D\}$ and $\{S - D\}$. We repeat this process for every game played and as new cards are introduced in the game.

In our set-based Elo rating system, it is assumed that all the cards in the selected set $\{D\}$ are better than the cards in the unselected set $\{S - D\}$, thus every single card in $\{D\}$ can beat every single card in $\{S - D\}$, in a head-to-head comparison as:

$$\{card_i\} \succ \{card_j\}, \forall\, card_i \in D, \forall\, card_j \in \{S - D\}.$$

We assume that cards in $\{D\}$ and $\{S - D\}$ are not compared to other cards in $\{D\}$ and $\{S - D\}$ respectively.

The set-based Elo rating is updated as:

$$R'_{card_i} = R_{card_i} + K(S_{card_i} - E_{card_i}),$$

$$R'_{card_j} = R_{card_j} + K(S_{card_j} - E_{card_j}),$$

$$\forall\, card_i \in D, \forall\, card_j \in \{S - D\},$$

where $S_{card_i} = 1$, and $S_{card_j} = 0$, and $E_{card_i} = \frac{Q_{card_i}}{Q_{card_i} + Q_{card_j}}$, and $Q_{card_i} = 10^{\frac{R_{card_i}}{400}}$, E_{card_j} and Q_{card_j} are analogously defined.

5.3 Comparing head-to-head Elo ratings to set-based Elo ratings

In this section, we show that as the number of selected decks tends to infinity, the set-based Elo ratings of a card collection will approach the traditional Elo ratings that we refer to as the head-to-head Elo ratings.

Below we show the results of two representative simulations carried out to examine the performance of the set-based Elo ratings method as compared to the head-to-head Elo ratings. In these simulations, we assume there are 10 cards to be considered in total with their global preference ordering set a pirori. In the set-based Elo ratings simulation, each participant, i.e. player, is assumed to own 5 cards out of the total 10 cards. Each time, a participant is required to select a random number, uniformly distributed from 1 to 4, of cards from his card collection. Using the assumption that we discussed in the previous section, any selected card is said to be better than any unselected card, we proceed to determine the Elo ratings according to participants' selections.

In contrast, since the Elo rating system is traditionally developed for paired comparisons, we also determine cards' ratings using the head-to-head selection. In this simulation, we compare only two of the 10 cards at a time, with the preferred one beating the other card. The Elo ratings of these 10 cards are calculated based on the outcomes of the paired comparisons.

We proceed to compare the ratings obtained from these two simulations when $K = 1$, we qualitatively observe the same results for different values of K. The number of the observations are 1000 and 1,000,000, respectively. As shown in Figure 3, we find that as the number of observations increases, the ratings of the set based Elo method approach the ratings of the head-to-head method. The correlation coefficient between the set-based Elo method and the head-to-head Elo method is 0.9987 with 1,000,000 observations. This implies that the relative ranking of the cards can be estimated to a high degree of precision as the number of observations increases.

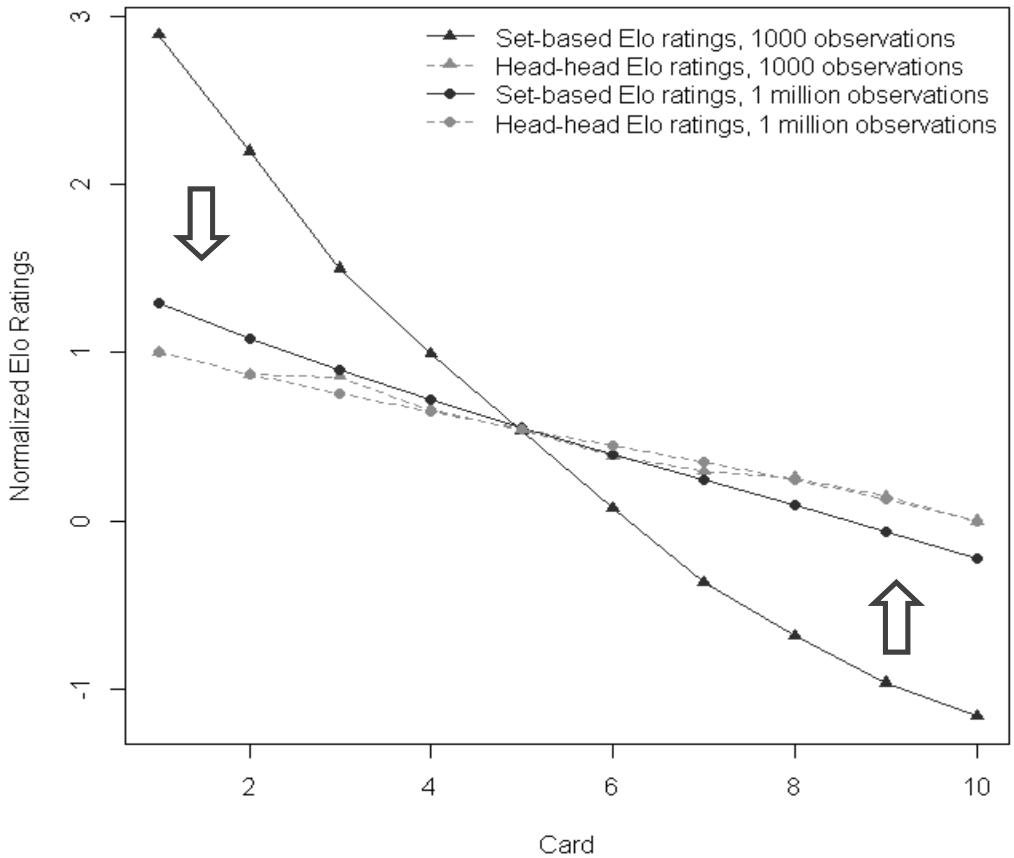

Figure 3: The ratings from set-based Elo method, uniformly choosing 1 to 4 cards within 5 cards, and the head-to-head Elo method, randomly choosing any two cards, with different observations

We now formally show the relationship between the set based and head-to-head Elo rating systems.

Proposition 1. *In a set-based Elo rating method, the number of times that any two cards are compared tends to infinity as the number of observations tends to infinity.*

Proof. Assuming we have n cards in total, the probability of any two cards being randomly selected for head-to-head comparison is $\frac{1}{\binom{n}{2}}$. Now consider the set-based Elo rating method. Of the n cards, k cards, $k \geq 2$, form a player's card collection, $\{S\}$, and out of these k cards, the best i cards are selected to be the player's deck $\{D\}$. Assuming uniform distribution over the cards in the set, the probability of any one card being selected in $\{D\}$ and another one being left in $\{S - D\}$ is p, and p satisfies the following inequality:

$$p \geq \frac{1}{\binom{n-2}{k-2}} \cdot \frac{1}{k-1} > 0,$$

where $\frac{1}{\binom{n-2}{k-2}}$ is the probability of selecting any two specific cards, A and B, into $\{S\}$, and $\frac{1}{k-1}$ is a lower bound on the probability of keeping A and not B in $\{D\}$. As p is always non-zero, therefore if the number of observations tends to infinity, the number of times any two cards are compared tends to infinity. ■

According to the proof described above, if we have a sufficiently large number of observations, all card comparisons are observed. Before implementing the set-based Elo rating method to the data in our study, we test if the data has enough observations to generate accurate set-based Elo ratings. This data set is tested by simulating the card selection process in the game. We calculate the rank correlation of the set-based Elo ratings to a pre-determined card ranking. The simulation stops when we reach a rank correlation of 0.99. Afterwards, we verify that the number of observations in the data exceeds this threshold.

Simulating the game, we have $n=500$, $i=40$, and k follows a normal distribution with mean of 150 and standard deviation of 30, bounded between 41 and 500. Once Spearman's rank correlation reaches a coefficient of 0.99 (Sheskin, 2003), the simulation is stopped. Figure 4 demonstrates the occurrence of the event after 200,000 iterations. In the game we study there are 3,774,517 observations in total, which exceeds the minimum threshold of 200,000, and hence allows us to use the set-based Elo rating method.

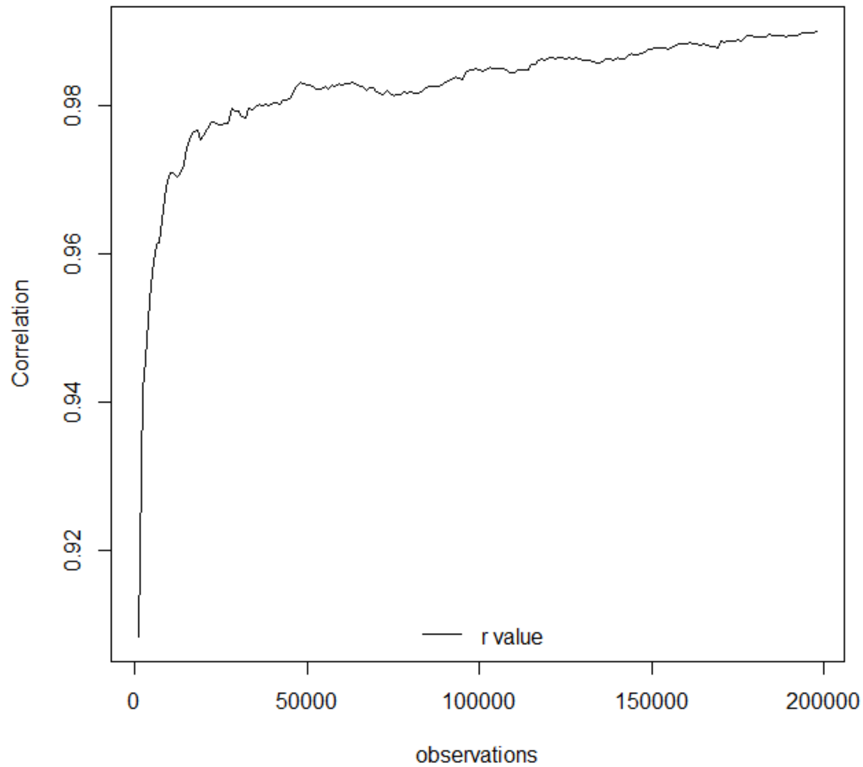

Figure 4: Correlation tracked with number of observations

5.4 Implementing the rating system

As the number of observations in our data is sufficient to compare any two cards using the set-based Elo rating method, we now proceed to implement this method in order to obtain the relative card quality, and ultimately to set initial card prices.

Any new card has an Elo rating of 1500 and this rating is updated as more and more deck information is observed. Remember that our goal is to determine the in-game market prices of cards when they are first released. Our industry partner makes new cards available in a pre-release sale during which cards may only be acquired in card packs. We focus on the Elo ratings of new cards after the pre-release sale and before the market release date to help determine their initial prices. However, we still update all cards' Elo ratings before and after the pre-release sale to measure all cards' quality.

We use Elo ratings to measure the card quality, hence, we expect that ratings are linked with prices, as well as the number of units sold for all cards. Once the Elo ratings of the cards are computed, we use Spearman’s rank and Person’s correlation to measure the two relationships between cards’ Elo ratings and card revenue, and between cards’ Elo ratings and units sold.

Since the Pearson’s correlation measure the linear dependence between two variables, we use $Q_A = 10^{\frac{R_A}{400}}$ to linearize the normal-scaled ratings before calculating Pearson’s correlations, where we refer to Q_A as the *modified Elo rating*. The modified Elo rating is used when calculating the Pearson’s correlation coefficient to determine the linear correlation between cards’ Elo ratings and card revenue, as well as the units sold.

5.5 Using Spearman’s and Pearson’s correlations to measure performance

We use Spearman's and Pearson's correlations to measure how well our set-based Elo method predicts aggregate revenue and aggregate number of units sold across all available cards in twelve days from the new card release day. The reason we use 12 days is that this is the maximum duration of any of the card release interventions, see Table 2.

Current card prices are updated based on normalized daily revenue of the cards, an approximation of the gradient of the price function. We use these current price values for comparison to our Elo rating method. As the card price has a direct relationship with revenue, we also use the number of units sold for each card when comparing the performance of the proposed Elo method to update card prices. The number of units sold is a proxy for potential revenue, as we observed some high selling cards, did not have their prices updated due to the price update method used by our industry partner.

We carry out a similar calculation for the same cards' Elo ratings, available through card usage data available from pre-release sales. Only when we determine Pearson’s correlation coefficients, we employ the modified Elo ratings, and otherwise we use the Elo ratings of the cards. As we know, Pearson’s correlation coefficient measures the strength of a linear association between two variables, while Spearman’s rank correlation measures the strength of association between the ranks of the two variables and identifies whether the two variables relate in a monotonic fashion.

As shown in Figure 5 and Figure 6, the Spearman’s rank correlation and the Pearson’s correlation both yield a strong correlation between cards' Elo ratings and the daily revenue.

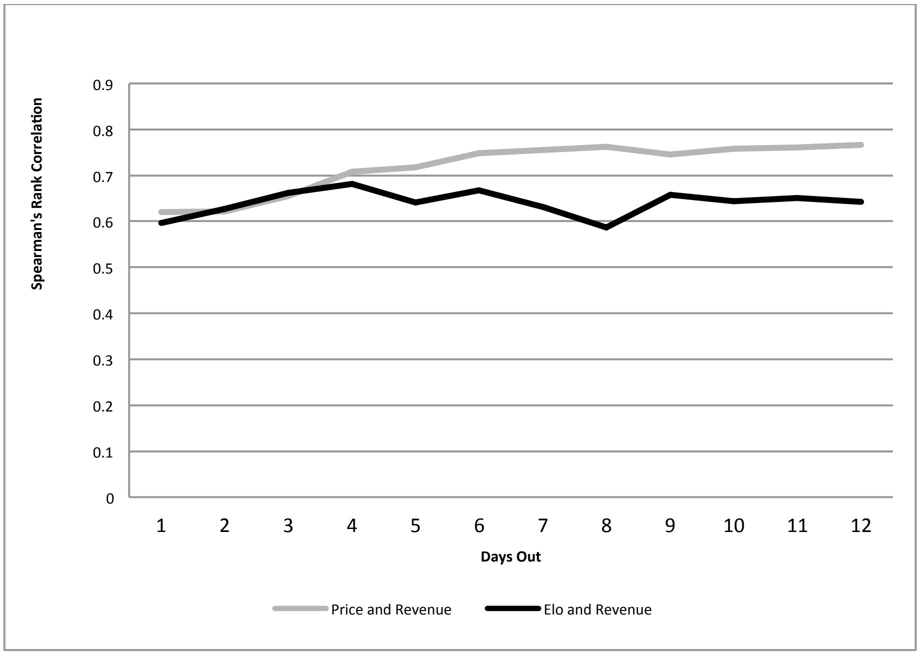

Figure 5: Spearman's rank correlation between prices and revenue, and between Elo ratings and revenue

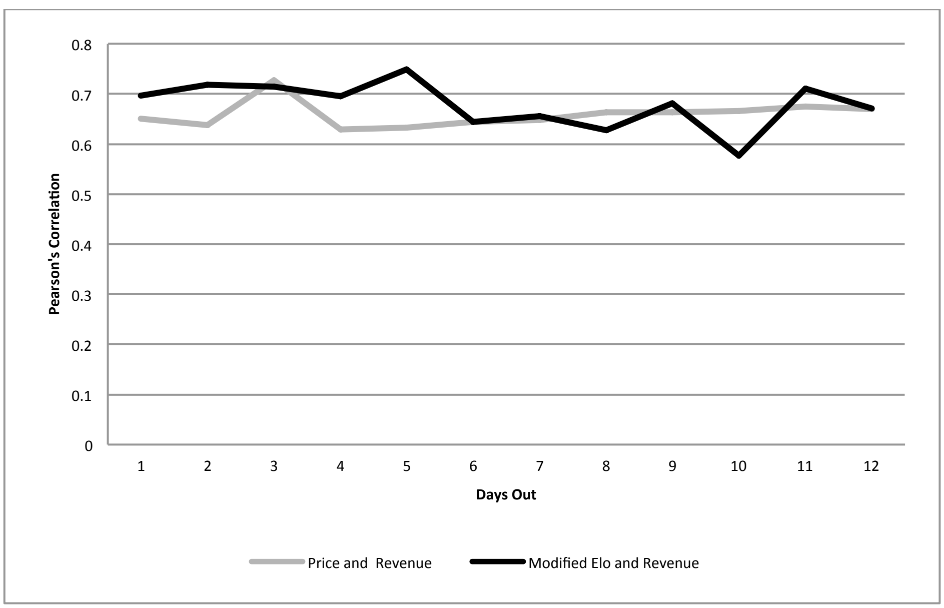

Figure 6: Pearson's correlation between prices and revenue, and between modified Elo and revenue

Using rank correlation, price starts out having a similar level of correlation as the Elo ratings and then progressively increases with time, while Elo fluctuates around the value of 0.65. This suggests that in terms of revenue prediction on the basis of rank, Elo does not have an advantage over the current price update method.

Using Pearson's correlation, which also takes the actual size of variables, rather than their order, into account, we can see that modified Elo ratings may offer an improved prediction of future revenue on certain days, in particular, the first day, in which a large proportion of the revenue is made. This suggests that modified Elo ratings can be one option when setting the initial prices on the first day (Figure 6).

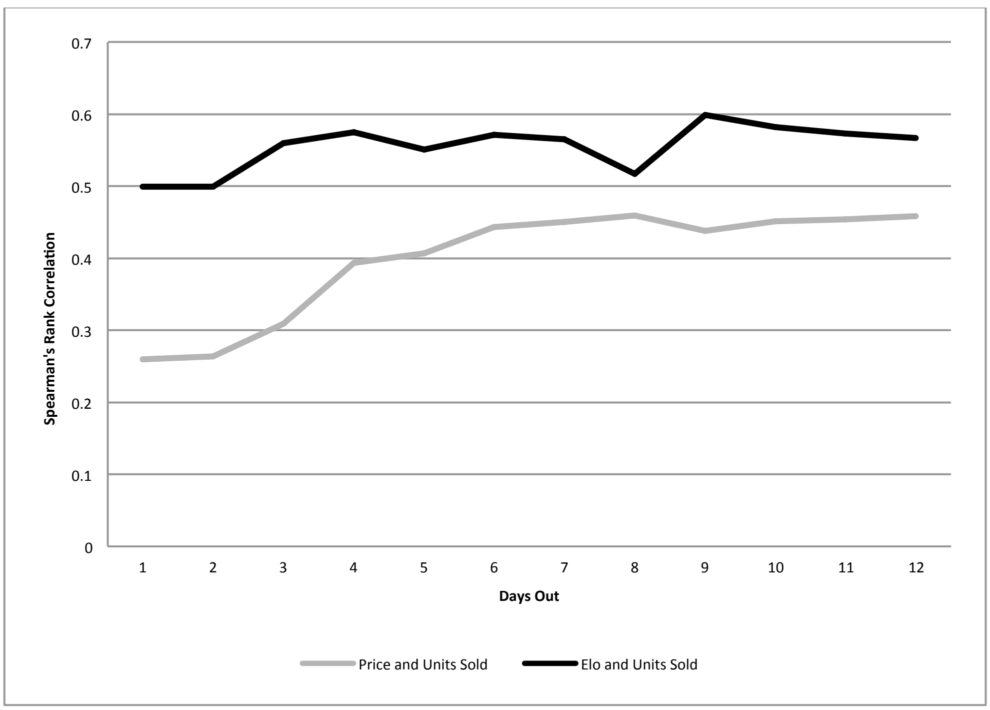

Figure 7: Spearman's rank correlation between prices and units sold, and between Elo ratings and units sold

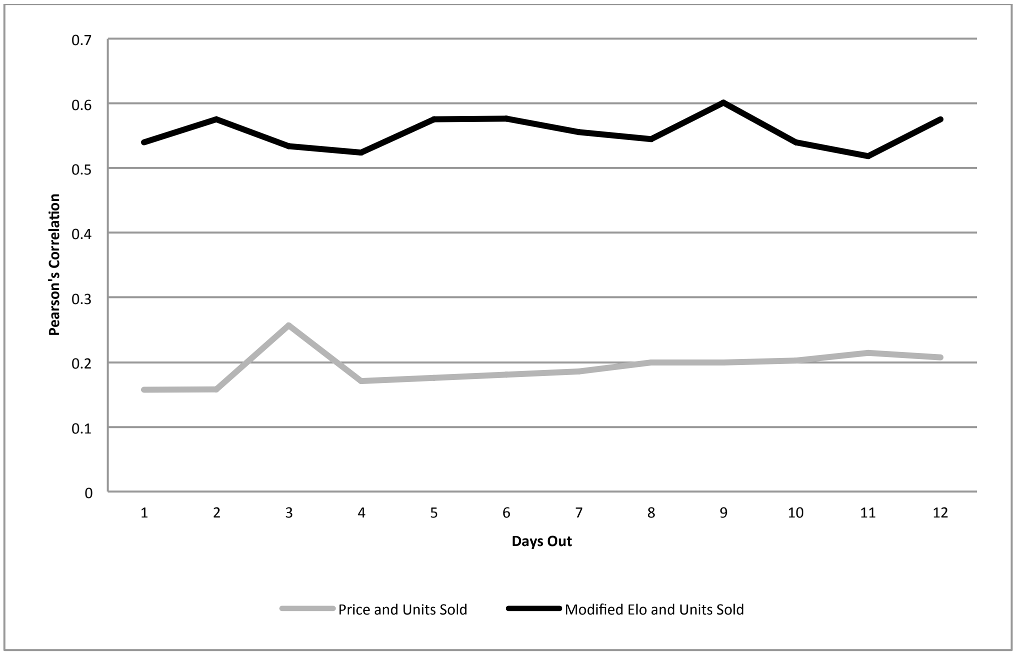

Figure 8: Pearson's correlation between prices and units sold, and between modified Elo and units sold

From **Figure 7** and **Figure 8**, we find that the correlation between price and units sold is lower than the correlation between Elo ratings and the units sold for all days. In **Figure 7**, Spearman’s rank correlation of price would approach Elo ratings when predicting units sold. However, Elo ratings always have a better correlation. The fact that the price correlation curve approaches the Elo curve in **Figure 7** after approximately two weeks suggests that to capture the same information about the number of units sold, it will take the price updating mechanism used by our industry partner more than two weeks after cards are released. Conversely, card Elo ratings give as much information about the number of cards sold as card prices will after approximately two weeks. In **Figure 8**, the Elo ratings can predict the number of units sold with Pearson’s correlation coefficient value of approximately 0.55. As a comparison, the correlation coefficient between the number of units sold and the current price is at most 0.25. One would think card quality is related to price, as higher quality goods may be expected to have higher demand not accounting for budgetary constraints. However, our study shows that card quality, as measured by cards' Elo ratings, is a better indicator of number of units sold than daily price. It is therefore concluded that cards may be mispriced on any given day.

The figures above show that the cards are mispriced. Therefore, the revenue can be increased by setting the prices of higher quality cards higher than their initial prices. In particular, as card prices have a Pearson's correlation value of at most 0.25 with the number of units sold, it follows that some participants are willing to buy high priced cards. However, as Elo ratings suggest, with a correlation of at least 0.52 with the number of units sold, some cards may need to be priced higher or lower than their initial values in order to increase sales and potentially increase revenue. Note, that this is hard to test from the available data as current price and revenue are closely related.

Chapter 6
Concluding Remarks

In this thesis we applied multiple intervention time-series analysis, ARIMA modeling with intervention analysis, to assess the observed effects of a company's operational events on sales. This time-series analysis includes the selection of suitable models, the estimation of model parameters, and the application of these models to actual sales data. Specific attention was given to clustering similar events, such as new card releases, and to the initial daily sales as new sets were released. This approach can generally be employed to analyze sales data with multiple interventions, and forecast the features of similar events in the future. Moreover, since the magnitude and the duration of any identical event are predictable, a company can adopt this approach to determine not only the appropriate timing to release new items to the market, but also the number of new item varieties to release in one event as the company maximizes its revenue. In addition, this thesis can also be extended to customer relationship management. For example, if a company were to permanently alter its prices of existing goods, or change the makeup of its goods, such as decrease size of food products, a similar study can measure the impact of such decisions on a company's existing customer base.

In addition to investigating time-series data from revenue, customer relations, or other, we show that the Elo rating system may be used to evaluate perceived quality of a product line. We also showed that the Elo rating system can be used to not only compare items head-to-head but also a set of items to another set of items. In our study, the one set of items are said to be "selected" while another set of items are said to be "not selected." We showed that if we update items' Elo ratings by assuming items in the selected set "beat" items in the not selected set, then the Elo ratings of all items converge to the same Elo ratings obtained from head-to-head. This, more general, set-based form of Elo comparisons may be employed by retailers to determine shelf space of a product collection, websites to determine the collection of links on each page, and organizations to compare departments of one another. Therefore, the trend of customer preference can be tracked using the set-based Elo rating method; thus, the retailer can make decisions to withdraw or provide an item. Finally, we believe that Elo ratings can help organizations evaluate their portfolio of items as they are added or removed. In particular, as Elo ratings are history-independent and self-correcting, they allow organizations to use Elo ratings for new and old items alike.

Application of the ARIMA modeling with intervention analysis and the Elo rating system to the sales data provided by our industry partner confirmed the reliability of our method. Therefore, our method may be widely applied to sales data with similar attributes to help in pricing decisions.

Appendix

Study on Card Distribution Channels

1. Revenue Discrepancy Between Two Distribution Channels

The two current distribution channels—single card market and card pack bundles—are used by our industry partner in their online game. The single card market displays all cards used in the game. If a player is after a particular card, then the player may purchase the card in the single card market. A card pack bundle is a collection of unknown cards draw from the single card market according to a known probability. Both distribution channels can provide all the cards for players.

We pre-process the data and group all the purchased items into three categories: packs, single cards, and miscellaneous. The packs category and the card category include all the packs and cards purchased, respectively. The miscellaneous category includes items purchased to upgrade the hero's ability or to enlarge the deck size. We proceed to calculate the revenues of each category accordingly. The packs' revenue and the single card market revenue contribute to approximate 57% and 23% of the total revenue, respectively. It is noted that the packs' revenue is more than twice of the single card revenue, which inspires us to explore the motivation that drives players to buy packs rather than single cards in the market.

Since a player only knows the probability of getting a particular card, purchasing a pack essentially reflects the player's gambling behaviour. However, there are other players who do not like taking chances to get a particular card, but prefer the single card market. In this appendix, we will discuss the revenue discrepancy and the players' risk attitudes. One explanation for this revenue discrepancy between the two distribution channels is that players make a trade-off between the currency spent and the value gained from the purchase. The result implies that given the same amount of in-game currency spent, a player would get more value from buying packs than from buying single cards.

2. The Single Card Market Revenue Summary

We first show the current revenue of the single cards in the market in Figure A-1 before comparing the two distribution channels. We calculate the revenue from the sales records and then sort the revenue from the largest to the smallest. Figure A- 1 reveals that 40% of the cards attribute to 90% of the total revenue in the single card market, while 25% of the cards attribute to 80% of the total revenue. The revenue of the top 10% cards is significantly higher than the rest of the cards, which suggests that several cards have significantly higher demands than other cards.

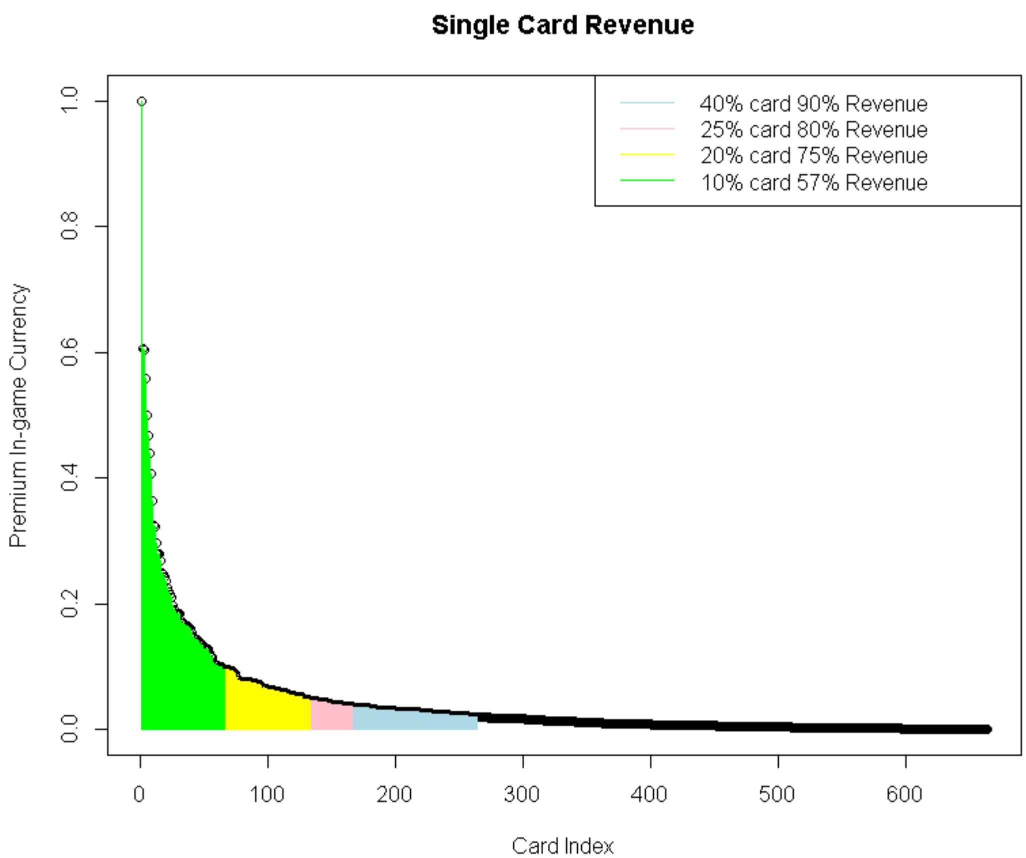

Figure A- 1: Proportion of total revenue from cards

3. Thresholds of Players Purchasing Decision

Since players prefer popular cards and avoid getting the same cards they already owns, choosing a proper distribution channel is to weigh values gained from the two distribution channels according to the players' card collections. We introduce threshold, T, to determine whether a player would buy a card in the single card market or buy pack bundles. This threshold indicates which players gain a positive utility from purchasing a particular card or a particular set of cards.

Consider this following example. There are two players, player A and player B, both are interested in owning card X. Player A currently owns 20 cards in his card collection, while player B owns 300 cards. The price of the card X in the single card market is 46 in-game currency, while the pack may contents this card is priced at 110 in-game currency. As a new player, player A's priority is to expand the card collection. Player A would choose to buy the pack, he might receive card X from the pack along with other cards. He would be satisfied with getting other popular but valuable cards that are not in his card collection. Player B, on the other hand, would choose to buy card X from the single card market. Given player B's large card collection, if he buys the pack, it is very likely that he will get some cards he already owns, which are not valuable at all to him. In this example, we show that each player has a certain threshold to value cards based on their card collections.

In order to determine the threshold, T, of each player, we introduce a criterion, D, which we refer to as *distance*. As discussed in the previous example, more discerning players would consider a card valuable if the card's price is closer to the card's upper bounds price, while less discerning players may accept a larger range. Given D, the threshold is determined using the formula:

$$T = LB + D * (UB - LB),$$

where

$LB: lower\ bound\ price\ for\ the\ card,$

$UB: upper\ bound\ price\ for\ the\ card,$

$D \in [0,1].$

LB and UB are prices that set up by our industry partner in their current pricing mechanism. Namely, LB and UB are constraints imposed on the prices of the cards. D is a threshold parameter. More discerning players have a value of D closer to one; while less discerning players have a lower D closer to zero. For example, a value of $D = 0$ means all cards are welcome by a player. Hence, the cards are considered to have their original values in the single card market. However, a value of $D = 1$ means the player only wants those cards sold at the upper bound price, and all other cards are assigned a value of zero. After having a better understanding of the threshold, we calculate the expected values of the packs with respect to the single cards' prices. When comparing expected values of the packs with their current pack prices, rational players will choose the one with higher value.

4. Expected Value of Packs for Risk-neutral and Risk-averse Players

As mentioned above, the other factor that affects players' purchasing decision is the player's risk attitude. The simplest scenario is that players are risk-neutral, meaning a player is indifferent between choices with equal expected payoffs even if one choice is riskier. For example, if two alternatives are offered: (i) \$50 with 100% certainty or (ii) a 50% chance each of \$100 and \$0, then a risk neutral player would have no preference since the expected value of the latter case is also \$50. However, a risk seeking (respectively, averse) player might prefer the latter (respectively, former) alternative.

There are several different packs available in the online game. We randomly picked one of them to carry out the D-criterion analysis. Let us refer to this pack as Pack A. Table A- 1 provides the expected value of Pack A on the basis of single card market prices with different D values. According to the first column of Table A- 1, if the price of Pack A is 10 in-game currency, players pursuing cards with $D = 0.7$ would buy the specific cards in the single card market. Otherwise, players would prefer seeking more value from buying packs. The number of players with D above 0.7 is relatively small among all the players, which explains the observation that the revenue from selling packs is almost twice of the revenue from selling the single cards.

Even though it is unlikely that all players are exactly risk neutral, players who are risk-seeking would buy packs even if the received expected value is less than the cost of the pack. However, risk-averse players would prefer buying packs if the expected value that could receive from the pack is greater than the pack's price.

There was some work on investigating how risk-averse a player is. The bounds for the risk aversion parameter is obtained based on an study conducted by Taylor (2013) by assuming that utility is described using the *constant relative risk aversion* (CRRA) utility function,

$$u(W) = \frac{W^{(1-r)}}{1-r},$$

where $u(W)$ represents the consumer's utility with payoff of W. r is the risk averse coefficient of CRRA. $r = 0$ indicates risk neutrality, while values of $r > 0$ indicate different magnitudes of risk aversion. As r increases the utility decreases, and when W increases the utility increases. The utility function says, given a value of r, as the payoff becomes greater, the less a participant would care about an increment in revenue. For example, an increase of \$100 is more substantial when a player is endowed with \$200 than having another \$100 when the player endowed with \$2 million.

The expected utility of Pack A for players with different thresholds and different risk-averse degrees is shown in Table A- 1. After comparing the utility gained from each distribution channel, players in the grey area will find buying packs more valuable than buying single cards in the market. As can be seen from the table, given a certain r, players with a higher threshold tend to be customers of the single card market. The reason is that, with a cost of the pack, they would get less than the value of the pack's price. Given a certain threshold, however, players who are more risk-averse would find the single card market more appealing since there is no risk in the single card market.

		r					
		0	0.15	0.29575	0.4415	0.55925	0.677
D	0	7.3	5.4	4.3	3.7	3.5	3.6
	0.1	4.9	3.6	2.7	2.2	2.0	2.0
	0.2	3.9	2.8	2.0	1.6	1.4	1.3
	0.3	2.7	1.9	1.3	1.0	0.9	0.8
	0.4	2.1	1.5	1.0	0.8	0.7	0.6
	0.5	1.5	1.0	0.7	0.5	0.4	0.4
	0.6	1.1	0.7	0.5	0.4	0.3	0.3
	0.7	0.8	0.5	0.4	0.3	0.2	0.2
	0.8	0.7	0.4	0.3	0.2	0.2	0.1
	0.9	0.7	0.4	0.3	0.2	0.2	0.1
	1.0	0.7	0.4	0.3	0.2	0.2	0.1
Pack utility		1.0	0.8	0.7	0.6	0.6	0.7

Table A- 1: Expected utility of Pack A with respect to D (the criterion) and r (the risk aversion coefficient). The values of utility are normalized and rounded. The r values were selected based on Taylor (2013). Shaded cells reflect preference for individual card; otherwise the player prefers a pack.

5. Suggestions of balancing two distribution channels

We already comprehend how the players compare the expected utilities between the two distribution channels. The explanation of the large discrepancy in revenue is that, players find that buying packs brings more utility than buying cards individually. This imbalanced expected utility is coming from our industry partner's different pricing mechanisms for each distribution channel. The pack prices are fixed; while card prices in the market, however, are adjusted along with the demand, which suggest that the card prices reflect the inherent value better than the pack prices. The observation that more players buy packs instead of single cards implies that the packs are relatively underpriced. Therefore, from the company's perspective, more welfare can be obtained to maximize the revenue.

The basic principle to improve the current pricing system is to eliminate the imbalance between the two distribution channels. Since the prices for packs are fixed currently, the main purpose of our following suggestions is to make the packs' prices become dynamic, for example, as a function of the demand, replacing with a new draw distribution, or expand the range between the upper bound price and lower bound price.

First of all, pack prices should fluctuate as a function of market prices. If the demand changes, the prices of packs change as well. For example, if a card's price increases, the price of the pack that may content this card should also increase. Therefore, players who want this card from the single card market will not turn to take their risks to buy the pack.

Second, the probability of receiving a card from a draw should be set as a function of the card's market price. Currently, the draw is uniformly distributed, which is the probability of getting each card is the same despite of price differences. We suggest that, the higher a card's price is, the lower the probability of receiving that card in a pack will be.

Last but not least, a dynamic upper bound change policy can be implemented in the pack pricing system, which is, if the price of a card reaches the upper bound price, then raise the card's upper bound price. The reason why we update upper bound and lower bound prices is that upper bound and lower bound prices will truncate the demand and thus hide the real demand.

Bibliography

Bala, R. (2012) Pricing online subscription services under competition*, Journal of Revenue & Pricing Management*, vol. 11, no. 3, pp. 258-273.

Bala, R. & Carr, S. (2010) Usage-based pricing of software services under competition, *Journal of Revenue & Pricing Management,* vol. 9, no. 3, pp. 204-216.

Box, G.E., Jenkins, G.M. & Reinsel, G.C. (2013) *Time series analysis: forecasting and control*, John Wiley & Sons.

Box, G.E. & Tiao, G.C. (1975) Intervention analysis with applications to economic and environmental problems*, Journal of the American Statistical Association*, vol. 70, no. 349, pp. 70-79.

Chan, K., & Ripley, B., (2012) *TSA: Time Series Analysis.* R package version 1.01. http://CRAN.R-project.org/package=TSA, accessed 27 March 2014.

Cryer, J.D. & Chan, K. (2008) *Time series analysis: with applications in R*, Springer.

DFC Intelligence, (2013) DFC Intelligence Forecasts Worldwide Online Game Market to Reach $79 Billion by 2017. DFC Intelligence report, 4 June, http://www.dfcint.com/wp/?p=353, accessed 27 March 2014.

Elmaghraby, W. & Keskinocak, P. (2003) Dynamic pricing in the presence of inventory considerations: Research overview, current practices, and future directions, *Management Science*, vol. 49, no. 10, pp. 1287-1309.

Elo, A.E. (1978) *The rating of chessplayers, past and present,* Batsford, London.

Enders, W. (2008) *Applied econometric time series*, John Wiley & Sons.

Hamari, J. (2009) Virtual goods sales: new requirements for business modelling?. Graduate thesis, University of Jyväskylä, Finland.

Hamari, J. & Lehdonvirta, V. (2010) Game design as marketing: How game mechanics create demand for virtual goods. *International Journal of Business Science & Applied Management*, vol. 5, no. 1.

Lehdonvirta, V. (2005) Real-money trade of virtual assets: new strategies for virtual world operators. *Proceedings of Future Play, Michigan State University*, vol. 13, pp. 15.

Lehdonvirta, V. (2009) Virtual item sales as a revenue model: identifying attributes that drive purchase decisions*, Electronic Commerce Research*, vol. 9, no. 1-2, pp. 97-113.

Mccleary, R. & Hay, R. (1980) *Applied Time Series Analysis for the Social Sciences*, Beverly Hills, CA: Sage Publications; London.

Meagher, K. & Teo, E.G.S. (2005) Two-part tariffs in the online gaming industry: The role of creative destruction and network externalities*, Information Economics and Policy*, vol. 17, no. 4, pp. 457-470.

Nojima, M. (2007) Pricing models and motivations for MMO play. *Proceedings of DiGRA 2007*, pp. 24-28.

Oh, G. & Ryu, T. (2007) Game design on item-selling based payment model in Korean online games. *Proceedings of DiGRA 2007*, pp. 650-657.

Pearce, J.L., Stevenson, W.B. & Perry, J.L. (1985) Managerial Compensation Based on Organizational Performance: A Time Series Analysis of the Effects of Merit Pay. *The Academy of Management Journal*, vol. 28, no. 2, pp. 261-278.

R Core Team, (2013) R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. http://www.R-project.org/, accessed 27 March 2014.

Reid, D.A. & Nixon, M.S. (2011) Using comparative human descriptions for soft biometrics. *Biometrics (IJCB), 2011 International Joint Conference on IEEE*, pp. 1.

Sheskin, D.J. (2003) *Handbook of parametric and nonparametric statistical procedures*, crc Press.

Sun, X., Gauri, D.K. & Webster, S. (2011) Forecasting for cruise line revenue management. *Journal of Revenue & Pricing Management*, vol. 10, no. 4, pp. 306-324.

Taylor, M. P. 2013, Bias and brains: Risk aversion and cognitive ability across real and hypothetical settings. *Journal of Risk and Uncertainty*, vol. 46, no. 3, pp. 299-320.

Yao, D., Wang, Z., Mukhopadhyay, S.K. & Cong, Y. (2012) Dynamic pricing strategy for subscription-based information goods. *Journal of Revenue & Pricing Management*, vol. 11, no. 2, pp. 210-224.

Ye, X. (2010) The application of arima model in chinese mobile user prediction. *Granular Computing (GrC), 2010 IEEE International Conference on IEEE*, pp. 586.