6-21-2017 12:00 AM

Essays on Entertainment Analytics

Kyle Maclean, *The University of Western Ontario*

Supervisor: Fredrik Odegaard, *The University of Western Ontario*
A thesis submitted in partial fulfillment of the requirements for the Doctor of Philosophy degree in Business

Abstract

This thesis explores live entertainment analytics and revenue management allocation strategies for live entertainment.

In Chapter two, we look at empirical factors that effect the success of Broadway shows. How well-known actors (stars) effect film revenues has been a recurring question of entertainment producers and academics. Because a film cannot be disentangled from a star involved, researchers have long struggled to rule out “reverse-causality” - that stars have access to higher quality movies. Using a novel data set that includes Broadway show revenues and actor usage, we provide a fixed-effects regression and case studies. We find across multiple specifications that increases in star power in a show improve revenue.

Motivated by social grouping and the associated operational challenges, in Chapter three we formulate and study extensions to the Dynamic Stochastic Knapsack Problem (DSKP). We compartmentalize the knapsack according to predefined reward-to-weight ratios, and incorporate a stochastic interaction between the offered set of open compartments and the item placement. Using a specific interaction function inspired by customer choice in the entertainment industry, we provide an algorithm to determine the optimal solution and obtain insights into structural properties. Given the computational complexity of the dynamic program we also propose and analyze via simulation a heuristic algorithm.

In Chapter four, in a large sequence of simulations, we propose and study practical heuristic algorithms on which seats should be offered to requests. We propose an algorithm that offers revenue improvements from a “naive” policy on the order of 5-10%.

Throughout, we aim for managerial relevance, providing implications to current techniques both in strategy as well as operations.

Keywords: Entertainment, Revenue Management, Markov Decision Process, Simulation, Seating, Operations

Co-Authorship Statement

I declare that this thesis incorporates some material that is a result of joint research. Chapter three "Dynamic Stochastic Knapsack Problem with Adaptive Interaction:Live Entertainment Capacity Based Revenue Management" and Chapter four "A Simulation Analysis of Group Seating Heuristics" are co-authored with Dr. Fredrik Odegaard. As the first author, I was in charge of all aspects of these projects including formulating research questions, literature review, model formulation, simulation programming, and preparing the first and the following complete drafts of the manuscript. With the above exceptions, I certify that this dissertation and the research to which it refers, is fully a product of my own work. Overall, this dissertation includes 3 original papers.

Contents

List of Figures

List of Tables

List of Appendices

Chapter 1

Introduction

The entertainment business is fascinating, containing many unique strands that differentiate it from other industries. In this thesis, we explore these different facets analytically and discuss how managers can face them. In general, we explore two different ideas: the uncertainty of success, and operational allocation of seats.

In virtually all entertainment fields (e.g, films, live theatre, books...etc), there exists a winner-take-all element. Successful entertainment products do tremendously well and capture a large share of revenues, while the majority of products make meager revenues. The problem managers face is that the process of becoming a winner is highly uncertain, and path-dependent (Salganik et al., 2006). William Goldman, a Hollywood screenwriter, once said that "nobody knows anything" about what makes a movie successful (Goldman, 1983).

The entertainment industry is interesting from an operational point of view as well. The industry generally satisfies the criteria of Revenue Management (RM) methodologies; There is a fixed-supply of seats that become perishable after an event, and selling an extra seat incurs low marginal costs. However, live entertainment is different from traditional RM application areas such as airlines. As mentioned above, social factors are more important, and the experience is more uncertain prior to purchase.

The key interlinking idea between these two topics is exploring how businesses should make decisions in an environment where social factors are important and groups of consumers make and engage in purchasing decisions. This creates research questions on an operational level (How to offer seats to incoming groups?) as well as at a strategic level (Should films and live entertainment products focus on "stars"?).

The predominant focus within the RM literature on live entertainment has been regarding pricing (or value) and its effect on demand/sales. For instance, Leslie (2004) examine Broadway data and show that seat price discrimination improved profit by 5% relative to uniform pricing; Hume et al. (2006) examine factors influencing service quality in performing arts; Toma and Meads (2007) empirically explore macro-level determinants of symphony attendance; Veeraraghavan and Vaidyanathan (2012) developed and empirically analyzed location based "seat value" and customer utility of events (specifically baseball); and Dixon and Verma (2013) show how sequence effects play a significant role in influencing season subscription renewals.

A main focus of RM in general is on differential pricing. This comes in the form of *dynamic pricing* and *variable pricing*. In the live entertainment sphere, dynamic pricing refers to changes in prices over time based on changes in realized demand vs predicted demand. For management simplicity, dynamic pricing often uses capacity levels as triggers to alter pricing. For example, to raise prices when 10% of capacity remains.

Variable pricing refers to different prices for different products, where a product is a specific seat to a specific event. Variable pricing is common in live entertainment. For example, a Friday night show may be priced differently than a mid week matinee. The event itself sometimes differs materially too. Sports teams sometimes change prices depending upon who is being played.

The most common type of variable pricing is for seating within a venue however. As far back as Shakespearean days, venues have charged different prices for different seating locations. At that time, the cheapest tickets were for the ground floor where standing was the norm. More expensive upper level seats were available as well however for the upper class. The process of demarcating different zones of different prices on the seating map is known as "scaling

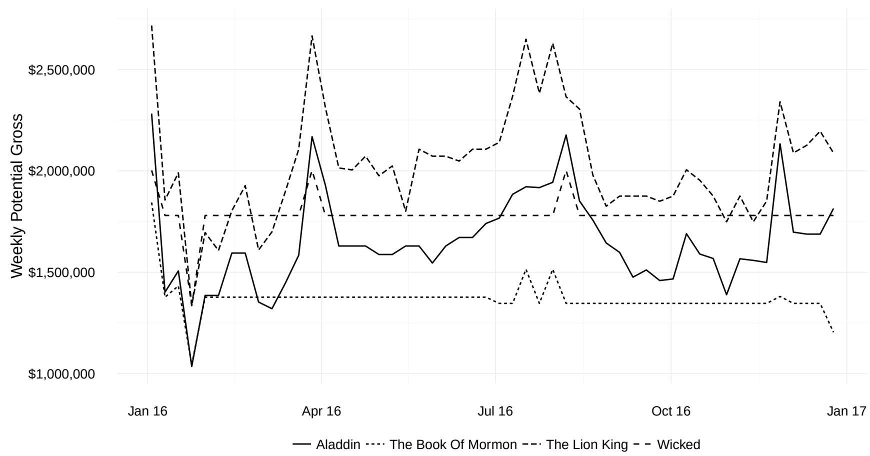

Figure 1.1: Weekly potential gross for selected shows in 2016.

the house".

Dynamic pricing in live entertainment has only begun to be introduced recently but has spread to multiple outlets. Major League Baseball's (MLB) San Francisco Giants introduced dynamic pricing to the sports industry in 2009 (Kemper and Breuer, 2016). In 2012 , 17 MLB organizations were beginning to dynamically price their events in some fashion Dunne (2012).

On Broadway, sophisticated RM pricing methods are often associated with Disney productions, sometimes cited as the "masters" of the process (Healy, 2014). In some of its shows, management sets limits on how high prices can be set, ostensibly to keep prices affordable for families. Other shows on Broadway are generally independently organized and financed, making sophisticated RM processes more unlikely.

To test the extent of these RM techniques on Broadway, we use publicly available data about shows' weekly *potential gross*. Potential gross measures how much a production could possibly make in a given week. The calculation is the summation of each seat, at each performance, at the posted price for that week. I.e. coupons and other discounts are ignored, and so a show can be fully attended without reaching 100% of its potential gross. We hypothesize that an active RM pricing regime would result in changes in prices week to week and thus differing potential gross numbers.

In Figure 1.1 we plot potential grosses for four selected shows on Broadway in 2016. Note the two Disney productions (Aladdin and The Lion King) exhibit changes week to week and throughout the entire year. In fact, it is rare to see two weeks where potential gross does not change. This is consistent with an active variable pricing process. In contrast, we look to Wicked and Book of Mormon, two other successful shows on Broadway. These two shows exhibit very long periods of unchanged prices, even during strong seasonal trends, e.g. the summer tourist season. However, even these shows exhibit temporary bumps during holiday periods - note that all shows have relatively higher potential grosses at the start of January 2016 - the week including the previous New Year's Eve.

In Chapter two of this thesis, we look at empirical factors that effect the success of Broadway shows. How well-known actors (stars) effect film revenues has been a recurring question of entertainment producers and academics. It is critical to the "blockbuster strategy", the idea that stars/budgets/marketing are a successful strategy in entertainment. However, because a film cannot be disentangled from a star involved, researchers have long struggled to rule out "reverse-causality" - that stars have access to higher quality movies. We argue that the commercial live-theatre environment provides a way of controlling for a show's fixed effects. Using a novel data set that includes Broadway show revenues and actor usage, we provide a fixed-effects regression and case studies. We find across multiple specifications that increases in star power in a show improves revenue. We also find that competitive star power has a negative, but not statistically significant effect on revenues. We discuss managerial implications as well as directions for future research based on the Broadway context.

Motivated by social grouping and the associated operational challenges, in Chapter two we formulate and study extensions to the Dynamic Stochastic Knapsack Problem (DSKP). We compartmentalize the knapsack according to predefined reward-to-weight ratios, and incorporate a stochastic interaction between the offered set of open compartments and the item placement. Using a specific interaction function inspired by customer choice in the entertainment industry, we provide an algorithm to determine the optimal solution and obtain insights into structural properties. Given the computational complexity of the dynamic program we also propose and analyze via simulation a heuristic algorithm.

In Chapter three, for tractability, we assume that groups are seated strictly from left to right, arguing that this can be operationalized by managers. In Chapter four, we relax this assumption and allow incoming requests to sit anywhere within a row. In a large sequence of simulations, we propose and study practical heuristic algorithms on which seats should be offered to requests. We find that a "naive" offer-all method is inefficient across a wide variety of conditions. We propose an algorithm that offers revenue improvements on the order of 5-10%, and show how it compares to common industry methods.

Throughout this paper, our aim is to make contributions to both the strategic side of entertainment management as well as the operational. We hope to spur other researchers to look into this field and examine more closely the idiosyncrasies of the industry.

Chapter 2

Mega-Stars and the Blockbuster Strategy: An Empirical Analysis of Broadway Revenue

2.1 Introduction

Entertainment industries (e.g films, live theatre, books, etc.) face a winner-take-all environment. The "winners" grab the vast majority of the revenue while losers make relatively little (Vany and Walls, 1999). Cultural products like films and theatrical productions face high upfront costs with highly uncertain lifetime revenues. *The Blair Witch Project*, with a budget of $60,000[1] eventually made over $200MM at the box office. At the same time, the 2016 summer film *Ben-Hur* cost $100MM to produce, and only made $94MM in domestic revenue.

A constant managerial question is what actions can be taken to increase the probability of success, or mitigate the probability of failure. The screenwriter William Goldman once said that in the movie industry "Nobody knows anything". Similarly, Jeffrey Sellers, the producer behind the mega-hit Broadway musical *Hamilton* suggested that "there is no moneymaking formula on Broadway" (Sokolove, 2016).

One formula suggested has been the "blockbuster" strategy. A "blockbuster" strategy entails using high production values, copious amounts of marketing, and well known actors (stars) to try to create extreme winners (Elberse, 2013). In the movie business, this can imply using known intellectual property (sequels), stars, and high marketing budgets to produce "tent-pole" summer productions. However, academic evidence for this strategy is mixed and popular reaction to it has varied. The media has suggested that the strategy has died off. A 2016 summer NY Times article was titled " Hollywoods Summer of Extremes: Megahits, Superflops and Little Else". A Vox article was titled "Hollywoods go-to formulas [have] stopped working".

One aspect of the blockbuster strategy is using stars to create a hit. Using stars can spur an information cascade, causing herd behavior in consumers (Bikhchandani et al., 1992). Stars also have a high cost though, both financially and logistically. Determining how and whether stars affect financial success is the first step to answering whether these costs are worthwhile. The literature to date on this question has been mixed, and definitive answers are unclear. Some papers find that stars increase film revenue. Others find that no effect or only an indirect effect. Much of this uncertainty in the literature stems from the inability to clearly identify causality. In a film, a star may be correlated with other movie characteristics, such as quality stories, scripts, budgets, and other creative elements. This has been referred to as the "reverse-causality" problem (Liu et al., 2014).

To address this issue, we propose looking at the live theatre environment, where different actors perform in the same show over time. We obtain a novel data set of which actors were present in Broadway productions in a given week, and weekly grosses of those shows. Using a continuous measure of star power, we study how a show's star power drives weekly revenue. We find that increases in star power improve revenues on average, as do major awards, except for Tony awards. Our work builds on a recent effort to correct for endogenous actor selection (Liu et al., 2014) and finds that stars directly, rather than indirectly effect revenues. We also explore competitive effects, finding a negative but not statistically significant effect of other Broadway shows hiring stars. We discuss managerial implications as well as future research directions.

[1] All dollar figures in USD.

2.2 Background and Literature Review

2.2.1 Literature Review

A variety of mechanisms and conceptual models have been proposed for how stars might alter the success of entertainment products. Most literature here is centered in the film industry context. Albert (1998) proposed that stars serve as a "marker" of a successful film type. Similarly, Thomson (2006) and Luo et al. (2010) suggested that stars are a type of human brand. A star might signal the type or quality of a cultural product that a consumer might purchase. For example, a consumer may update their private beliefs about the genre/quality of the movie if the consumer realizes a certain star is involved. In that way, a star would reduce the quality risk for consumers. Cultural products, being experiential goods, have quality that is difficult to predict a priori. It would hold then, that if consumers have a better idea of the type of film, demand may increase.

A rich literature has been developed to empirically test this theory in business environments, with almost all of the literature being based in the film industry. Unfortunately, empirically the literature is mixed. Some find a positive effect from using stars. Using an online prediction market as an event study, Elberse (2007) finds that stars have a statistically significant positive relationship with revenues, and is able to determine the value of specific actors. Nelson and Glotfelty (2012) use IMDB rankings to operationalize star power, and apply a per-country regression, finding a statistically significant positive effect of stardom. Other papers find no statistically significant effect.Prag and Casavant (1994) find no significant effect, after accounting for marketing expenditures, from a star's presence in a movie. But the authors note that the presence of a star also effects marketing expenditures. Ravid (1999) comes to a similar conclusion, noting that when including budgets, stars do not seem to significantly effect revenues. Vany and Walls (1999) use a distributional approach and conclude that stars do not make successful films, the movie itself does. In a fascinating new work, Liu et al. (2014) create an endogenous switching model to control for endogenous matching processes. They find that any effect stars have on revenues is largely indirect, from different movie characteristics as well as increased theatre allocation (i.e. stars cause their movies to be in more screens nationwide).

Other work has researched other factors that stars may effect. Joshi (2015) focused on week-to-week grosses, rather than total grosses, and found that they had less volatility when a film had a star. If the movie business values consistent cash flow, week-to-week riskiness should also be important to them. Liu et al. (2013) finds that stars impact stakeholders in the early project stages (e.g. financing and pre-production) but that the impact on audience may again only be indirect.

When positive results are found, most of the literature has only claimed that their result is an association, rather than causation. This is due to endogeneity in the process of matching star actors with films.

2.2.2 Endogeneity in Actor Selection

In the ideal experiment to test this research question, actors with varying "star power" would be assigned to movies randomly, and the resulting revenues tested for difference in means. This can not occur in reality. Actors are not assigned randomly, but rather are attached to projects through a complex sequence of bargaining and access (Ravid, 1999). This implies that actor selection is endogenous and potentially related to movie characteristics (Liu et al., 2014).

Endogeneity between star power and movie characteristics is bi-directional. Star actors, if they have bargaining power, may be able to influence or seek out specific film characteristics of projects that they are involved in. Given private information about characteristics that go into the quality of a product, star talent may choose projects that are projected to have higher quality. Alternatively, from the producer side, producers may choose to attach stars to specific movies (Eliashberg et al., 2006). They may do so to ensure project financing. Furthermore, they may attach stars only to specific genres.

If these forms of endogeneity exists, it presents a problem for empirically assessing the inclusion of a star in a cultural product. Stars may be *associated* with higher revenues, but only through their association with other movie characteristics. As Vany and Walls (1999) points out, the movie would make the star, instead of the reverse. "Reverse - causality" would remain a potential explanation.

Endogeneity could be dealt with partially through more data. Some movie characteristics are possible to control for by obtaining more data. For example, much of the literature includes budget as an independent variable in the analysis. However, others would be very difficult, if not impossible to quantify and operationalize (e.g the quality of script/story)

Numerous papers have identified this issue as well. Eliashberg et al. (2006) states the need for continuing research into what effect stars have on films. Elberse (2007) points out the reverse-causality issue and attempts to solve it via the event study approach on a public prediction market. Given full information, a prediction market would identify the effect a change in actor has on estimated revenues. Thus, an event study would solve this issue. However, as Ravid (1999) points out, and Elberse (2007) acknowledges, there is assymetric information in the market for films and full information is not available to the prediction market. Script quality, budget, and other movie characteristics may be unknown to the prediction market. With assymetric information, the presence of a star can serve as a signal of these other movie characteristics. Thus, even under the null hypothesis of no effect, the prediction might change as a result of the announcement in stars. Therefore, if the prediction market operates on public information, which the online website does, an event study using prediction markets still suffers from the same issue of a reverse-causality explanation.

Liu et al. (2014) formalizes the notion of endogeneity in the film market and star selection. They attempt to solve the issue via explicitly modelling actor selection and movie characteristics. However, the study only seeks to measure opening week differences. Our study is cleaner in the sense that star power variation occurs naturally over time on Broadway. Furthermore, we notably find different results. Liu et al. (2014) find that stars increase the number of screens a film is on. On Broadway, capacity changes are not possible since a theatre's number of seats is fixed. Thus, we find a different mechanism than them.

To solve the potential reverse-causality problem, we propose looking outside the film industry. For the reasons mentioned above, the film industry is unlikely to be a fruitful place of research for relatively clean data analysis. Instead, we must find a context where it is possible for star actors to be introduced into and out of the underlying cultural product. Only one industry satisfies this requirement to our knowledge: live theatre. Unlike other entertainment where almost all of the production is done up front, live theatre's main component is "produced" live nightly. While the script, direction, choreography and other story and design elements are fixed prior to the show starting[2], the actual performance occurs nightly before a live audience.

2.2.3 Broadway Background

Broadway theatres refers to 41 professional theatres in downtown New York City. A Broadway theatre is classified as any downtown theatre with more than 500 seats. Capacity varies widely though; the smallest theatre has 597 seats, while the largest has 1,935 seats. Shows on Broadway are generally categorized as either plays or musicals.[3] Plays are scripted productions, while musicals incorporate songs with singing and choreography, in an often extravagant manner.

Shows can be categorized as being open ended runs or limited runs. A “run” refers to the amount of time that a show will be playing for. Open ended productions have no specified end date and will run for as long as the show is profitable. A limited run is a show that opens and has a specified end date. Limited runs are common with plays that open with star actors, because of their busy schedule (Reaney, 2014). However, even limited runs can extend their run given heavy demand, depending on the schedule of their cast and on the availability of their theatre. The standard running week of a Broadway show is 8 performances every 7 days, with one “dark” day where no performances are held.

A Broadway project generally progresses according to the following sequence. The show undergoes rehearsal for 2-4 weeks where the cast and creative team rehearse and practice the show. During the final week, a “tech” rehearsal occurs where the lighting and sound elements are brought together with the cast and cues synced. Once public performances begin, a show goes through several weeks of previews. During previews, the general public buys tickets to the show understanding that design elements may change or be a work in progress. Musicals for instance may remove songs, edit dialogue, and alter staging during the preview period. On opening night, the show is locked and no further creative elements are changed [4]. In addition, professional reviewers (e.g. The New York Times, The Washington Post) publish their official review of the show on opening night.

Commercial theatrical productions on Broadway can be long running affairs: The Lion King on Broadway has been running for over 18 years with over 7500 performances. In long running shows, while the creative elements are fixed, the actors themselves change throughout the run of the show. Like in the film industry, actors are chosen based on a negotiated matchmaking process. This would involve both their acting ability as well as their potential star power to draw in an audience. Star actors sometimes headline the show during the beginning of a run in hopes to generate interest. This occured in 2011 when the revival of *How to Succeed in Business Without Really Trying* opened. Daniel Radcliffe, from *Harry Potter* fame, was cast as the lead character and played the role for a little less than a year (Feb 26, 2011 - Jan 01, 2012). Stars however can also be replacements during the middle of the run, sometimes negatively referred to as“stunt-casting”, because of the perceived lack of talent of these actors.[5]

[2]Minor changes to scripts can occur during previews, or sometimes during extended runs however.

[3]Exceptions exist: The magic show Penn & Teller is not obviously described by these labels.

[4]Some small topical references that become outdated can be edited, e.g. after the presidential election in 2008, the musical Avenue Q altered a lyric referencing George W Bush.

[5]More specifically, the term is generally reserved for stars without a live-theatre background.

In the musical *In the Heights*, High School Musical star Corbon Bleu stepped into the lead role partway through its run (Jan 25, 2010 - Aug 17, 2010).

The years in our dataset were years of change on Broadway. Revenue management (RM) techniques like variable pricing became more common, spearheaded by Disney and its productions (e.g. The Lion King) (Healy, 2014). From 2009 to 2015 the average price for sold tickets increased from \$83.02 to \$104.26, a 25% increase that far outpaced inflation. Prices changes like this could be a result of more pervasive RM techniques, or as a result of shifting consumer preferences toward live entertainment.

While the industry has changed, in this chapter we focus solely on the effect of starpower, not necessarily the prediction of grosses in general. Changes in pricing behavior with respect to RM behavior should not create bias in our estimated effects of starpower.

Our research thus also intersects with a literature on the success of live theatre. This literature generally centers around what effects the success of a show, as well as how awards or reviews can effect a shows revenue. Reddy et al. (1998) examine multiple factors that impact the success of a Broadway show including critic reviews and advertising expenditures. They find that critics both predict and influence success. The authors also investigate the effect of actor choices, using as proxies the total number of shows on Broadway the lead cast/director/author had done, as well as the number of awards these individuals received. Boyle and Chiou (2009) examine the value of winning a Tony Award using a discrete choice model as well as a hazard model. They suggest that nominations bring general positive increases to all shows, but that attention shifts to just the winners after the awards are given. They also propose that awards information follows as an information cascade where the main effect is not felt immediately. Simonoff and Ma (2000), using a survival model found mixed results for major paper reviews, but a positive effect from winning awards. Maddison (2004) also performed a survival analysis on Broaday shows, finding that original shows tended to outlast revivals. The author also finds that as shows last longer, the more likely they are to last for a longer time, unlike in the movie industry. Maddison (2005) looked into revivals on Broadway, finding that revivals are becoming more prevalent and that they are less risky, though do less well on average.

Notably, Reddy et al. (1998) investigates how actors affect lifetime revenue on Broadway, however only do so based on the actors who were there on opening night. This is an interesting question but suffers from the same direction of effect questions as the film industry. This analysis was only done for the opening night cast and did not test what happened when these lead actors left the show and thus, similar to the movie industry, this approach is not able to cleanly identify star power effects.

2.3 Model Framework

In Figure 2.1 and Figure 2.2, we provide the conceptual diagrams which guide our research. Broadway productions have two differences from the film industry. Firstly, actors (and therefore the effects of star power) change throughout the run of the show. In film, a movie's cast is unchanged once production is finished. The other change occurs with capacity. In film, capacity can quickly expand and is an important aspect of capturing information cascades/social herding effects (Vany and Walls, 1996). If a film captures the public's interest, it can quickly be distributed to thousands of screens nationwide. In contrast, on Broadway capacity is fixed and

Figure 2.1: Conceptual Model for Films

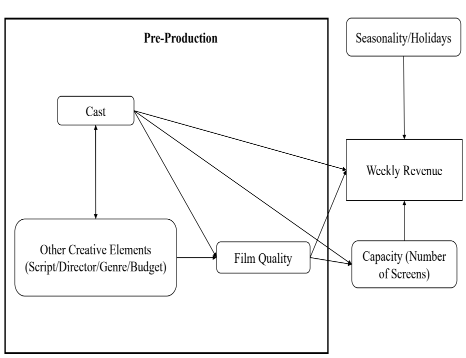

Figure 2.2: Conceptual Model for Broadway

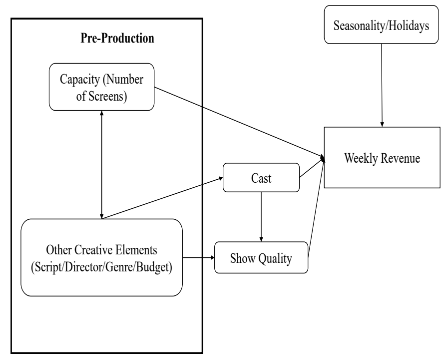

very inflexible. Although it is possible for a show to switch theatres during its run, the switch is very expensive, lengthy, and ultimately rare. Capacity is therefore non-dynamic throughout a show's run.

Quality levels can differ on Broadway throughout a show's run, in contrast to film. By quality, we mean the general artistic ability to provide intended emotions to an audience. Changes in actors may effect the quality of a show through changes in innate acting ability or ability to perform in a live show. For example, it was reported that Bruce Willis in 2015 had to be fed lines through an earpiece on stage [6].

We do not explicitly include changes in quality in our model, largely because data is not available. However, we argue this effect is small. Producers have a vested interest in keeping the quality of their own show at some minimum level. In addition, if quality had high temporal variability we would expect to see new reviews of shows on cast changes. This is rare, indicating that quality, although theoretically effected by cast, is for practical purposes steady. Thus, although in theory quality may be variable over time, in our model quality is captured by the time invariant unobserved effects of a show.

In our model, we index shows by $i = 1, \ldots, N$; where N is the number of shows in our sample. We index time (weeks) by $t = 1, \ldots, T$, where $t = 1$ is the first week in our sample (the week ending May 31st 2009) and $t = T$ is the last week in our sample (the week ending August 1st 2015).Our dependent variable is the log transformed weekly per-performance revenue, y_{it} for show i in week t. $\boldsymbol{x}_{it}$ is a column vector of K explanatory variables for show i in week t.

More details are provided about the specific independent variables in $\boldsymbol{x}_{it}$ in Section 2.3.3. Details about the dependent variable y_{it} are given in Section 2.3.2. Variables are summarized in Table 2.2. Our econometric model is given by

$$\begin{aligned} y_{it} &= \boldsymbol{\beta} \boldsymbol{x}'_{it} + \alpha_i + \epsilon_{it} \\ \text{for } i &= 1 \ldots N, \\ \text{for } t &= T^i_s \ldots T^i_e, \end{aligned}$$

[6]http://nypost.com/2015/11/05/pacinos-not-alone-willis-needs-an-earpiece-to-remember-his-lines-too/

where T_s^i is the starting week for show i and T_e^i is the ending week. $\boldsymbol{\beta}$ is the column vector of parameters/coefficients. α_i is the unobserved individual effect of show i, that is intended to capture time-invariant heterogeneity. ϵ_{it} is the idiosyncratic error term which is assumed to be independently distributed across shows but with a potential dependence across time (We assume that $\epsilon_i \sim i.d.\ (0, \Omega_i)$). That is, we explicitly allow for within show i serial correlation in errors and heteroskedasticity.

We assume that α_i is a fixed effect, rather than a random effect. A random effects would require the assumption that α_i is uncorrelated with the entries of $\boldsymbol{x}_{it}$. In our context, this assumption would imply that star power is uncorrelated with script/budget/genre. This is a strong assumption and one we do not wish to make. Therefore in our analysis we assume that α_i are fixed effects[7]. Each α_i is treated as an unknown parameter to be estimated (Greene, 2007, p. 359).

The theoretical downside to modeling α_i as a fixed effect is that time-invariant covariates would not be able to estimated, since they would be captured by α_i. For our purposes, this is acceptable since we are not interested in time-invariant factors (e.g. genre, budget) and explicitly do not include them in our model. We are specifically focused on variation in starpower over time. However, this does mean that shows who had no replacement in actors are not able to contribute to our estimation of starpower effect.

β and α_i can be estimated using Ordinary Least Squares (OLS). However, since OLS assumes ϵ_{it} are i.i.d., the associated standard errors on β and α_i are generally biased downward (Cameron and Miller, 2015, pg. 3). We instead use robust standard errors following Cameron et al. (2011) to account for any potential serial correlation or heteroskedasticity. This two-step procedure (OLS and robust standard errors) is mentioned by Cameron and Miller (2015) as a more recent method in the literature.

2.3.1 Broadway Data

We collected data from the Internet Broadway Database `IBDB.com`, a website ran by the trade association for the industry, for all shows that were running during the period May 31st 2009-August 1st 2015. This resulted in 290 shows across 327 weeks . The 2009 starting date was chosen due to a change in reporting standards in grosses prior to that date. This data included the "gross gross"[8] of each broadway show, in each week that it was playing. Revenue data is publically available and generally released every Monday, giving grosses for the previous week.

From the same site, we obtained information on which actors played which characters for all shows within the sample period. We refer to character/actor combination as "Roles". Role information included the actors names, the character(s) they played in the show, and the dates they joined and left the show. Figure 2.3 shows an example of the data on the site itself.Since a role is both a character and actor combination, some actors are included in the dataset in multiple roles over different time periods. We index roles by $r = 1, \ldots, R$.

We made significant efforts to ensure the data was cleansed and precise. Role information were removed due to a lack of either a start date, ending date, or imprecise start/ending dates.

[7] A model assuming random effects is presented in the Appendix however for robustness. All qualitative results presented here hold.

[8] Gross Gross takes the total revenue in a given week and does not subtract out transaction fees.

Figure 2.3: Screenshot of example raw role information from IBDB.com

Josh Gad	Elder Cunningham	(Feb 24, 2011 - Jun 06, 2012)
Asmeret Ghebremichael	Ensemble	
Nikki M. James	Nabulungi	(Feb 24, 2011 - Jan 05, 2014)
Clark Johnsen	Ensemble	
Rory O'Malley	Elder McKinley	(Feb 24, 2011 - Jan 27, 2013)
	Moroni	(Feb 24, 2011 - Jan 27, 2013)
John Eric Parker	Guard	
	Ensemble	
Michael Potts	Mafala Hatimbi	(Feb 24, 2011 - Mar 24, 2013)
Andrew Rannells	Elder Price	(Feb 24, 2011 - Jun 10, 2012)

[9] Generally speaking, roles that were removed were either less important characters (so-called "ensemble" members) or non-notable actors. We also ensured that no roles conflicted with each other, e.g. data that specified two actors playing the same character in the same week. For those weeks that showed two actors overlapping the same part, we assumed that the actor with the shorter overall time period was actually in the show. This tended to occur for short vacation periods of the main actor.

After this data cleansing process, we had 928 unique roles in our data set. This consisted of 633 unique actors across 80 shows. The number of shows with role information is less than our entire show sample. This is primarily due to short show runs for which no actors changed and therefore for which IBDB did not provide starting/ending dates. The mean length of a show for which we have actor information was 133 weeks compared to 51 weeks across the entire show sample. In a fixed effects analysis, a show with unchanging actors does not provide information about the effect of those actors. Therefore not having information on these shows did not effect our analysis. On average, an actor in our dataset was in a show for 31.9 weeks before leaving the show.

2.3.2 Dependent Variable: Weekly Gross Revenue

The dependent variable y_{it} is a log transformed per-performance revenue for each show for each week in our sample[10]. Weekly grosses were inflation adjusted to 2015 values using monthly CPI data in the United States[11], and subsequently adjusted to a "Per-Performance" measure by dividing by the number of performances done in a specific week. This change was done to control for any non-standard number of performances held. In a typical week, a Broadway production runs 8 shows per week. However this can vary in previews (prior to a show opening), or because of special holidays throughout the year. For example, during the holiday season in December, it is not uncommon for some shows to have 10 performances on a given week. We refer to this variable as *Gross*. In Figure 2.4, we provide box plots of *Gross* for a sample of shows from our data set. We see that mean *Gross* can vary considerably across shows. We

[9] For example, some data points only indicated a general month of when an actor started with a show.
[10] Monetary variables are particularly suited to being specified in logarithmic form (Wooldridge, 2012, pg. 193).
[11] https://www.bls.gov/cpi/

Figure 2.4: Boxplots of selected shows' weekly grosses for 2009-2015. The mean weekly gross for the entire sample is $102,301.

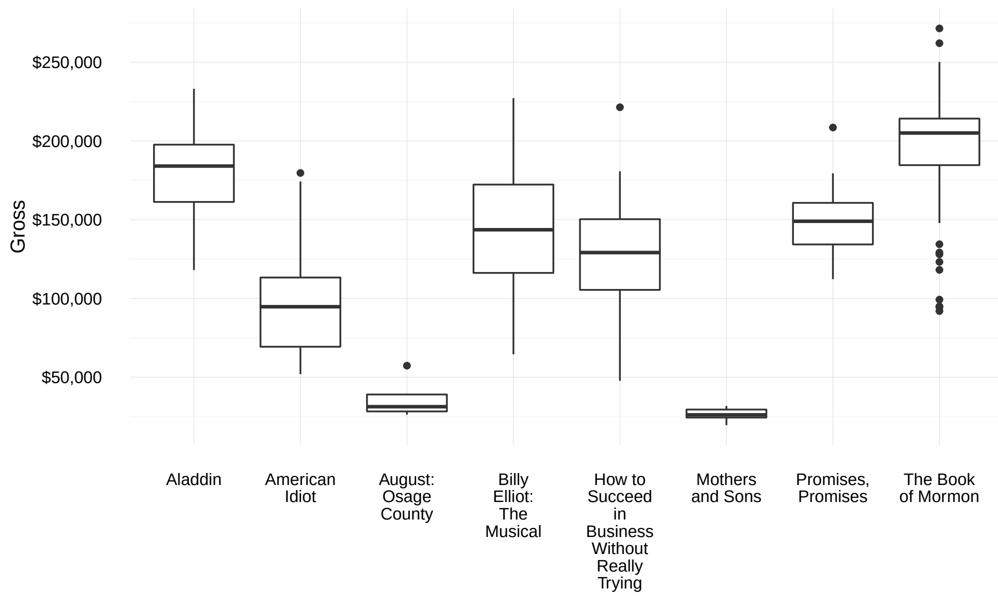

consider this indicative of different fixed effects. This could be capacity, i.e different theatre sizes or different show quality levels. We also see significant variation within each show. We postulate that at least some of this variation is due to changes in actors and starpower.

Our final data set of revenues is an unbalanced panel structure. The panel is unbalanced because of the opening and closing process of the shows, and so some shows do not have data for certain weeks.

2.3.3 Independent Variables: Operationalization of Star Power

To determine how within-show variation in star power explains changes in revenues, we must operationalize the concept of star power. One common approach in the literature is to assign a dummy variable indicating whether an actor was a star. We opted against this approach because we feel it misses the granularity that still exists between actors (Nelson and Glotfelty, 2012). Two stars that might both be classified as a star, might still have different levels of star power. That is, star power is best represented by a continuous variable, rather than a binary one. This is recognized in the media as well, who popularly refer to certain actors as "A-List" stars or "D-List" stars. Dummy variables also have practical operationalization issues. A researcher must still find some underlying variable to operationalize awareness and choose an arbitrary cutoff point. Otherwise, they subjectively assign the dummy variables values such as in Prag and Casavant (1994).

The other common approach is to include two independent variables that combine to form a proxy for star power. Usually this approach uses the number of awards an actor has won, and the quantity/revenue of previous shows the actor was in. For example, in the film industry, a common operationalization is the total revenue of the last five movies an actor was in. This has

Figure 2.5: Screenshot of the STARmeter graph from IMDB.com

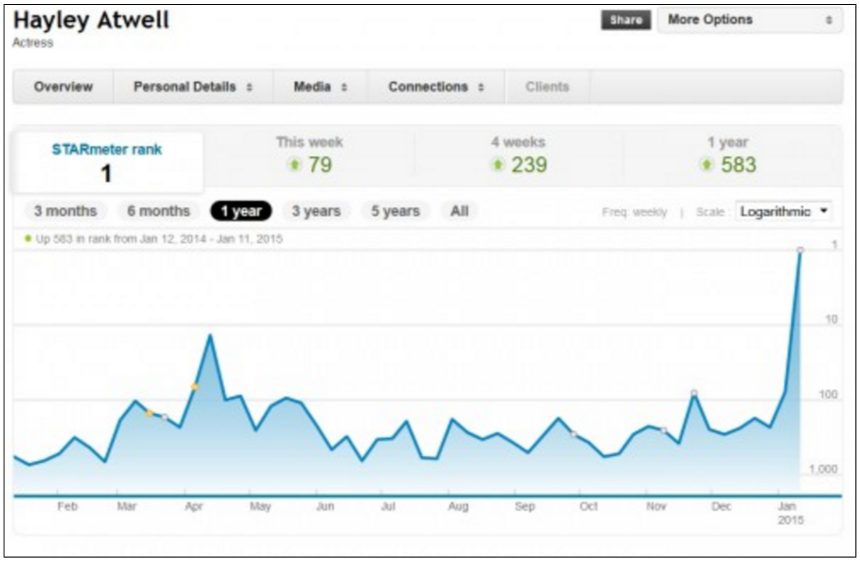

an advantage that such information is publically available and generally easily obtainable.

We feel this is inappropriate for Broadway. Particularly recently, Broadway star actors have come from many industries. They may be successful movie actors (Daniel Radcliffe, Neil Patrick Harris), well known singers (Billie Joe Armstrong, Nick Jonas), television series actors (Matthew Morrison, Sean Hayes), or even well-known actors from Broadway (Kristen Chenoweth, Sutton Foster). Using film box office revenue does not allow us to measure star power across entertainment industries.

Instead we use the recommendation of Nelson and Glotfelty (2012) and use the IMDB STARmeter metric to quantify the popularity and star power of actors.[12] This proprietary ranking system uses search results on the IMDB platform to determine an up-to-date index of which actors are being viewed by the general public. On the website, STARmeter is a rank, such that the top actor has a ranking of 1, and increases in the ranking imply decreases in star power. Figure 2.5 gives an example of how STARmeter appears on the website, and how changes are graphed over time.

For each role, we collected the STARmeter rank of the actor on the date they joined the show[13]. A linear interpolation was performed for start dates that were between STARmeter dates available.

To aid in interpretation, we made two adjustments to the raw data. We assumed that the difference in starpower between low-ranked actors was negligible and truncated it to 20,000. That is, for each STARmeter found, we took the minimum of it and 20,000. 20,000 was chosen because actors beyond this point are not recognizable to us.

After truncating values, we transformed the raw STARmeter ranking into a metric between zero and one, with increases in value implying more star power. We call this variable

[12]Conceptually, it may be important to recognize that the ranking does not measure sentiment. A frequently searched actor may not be generally liked. We do not seek to measure this difference in our study.

[13]In approximately 11% of the roles, we were unable to find a corresponding IMDB profile. We assumed these actors were highly unknown.

Figure 2.6: Histogram of non-zero STARmeter values in our dataset on 6/30/2015. Each bin width is 0.05.

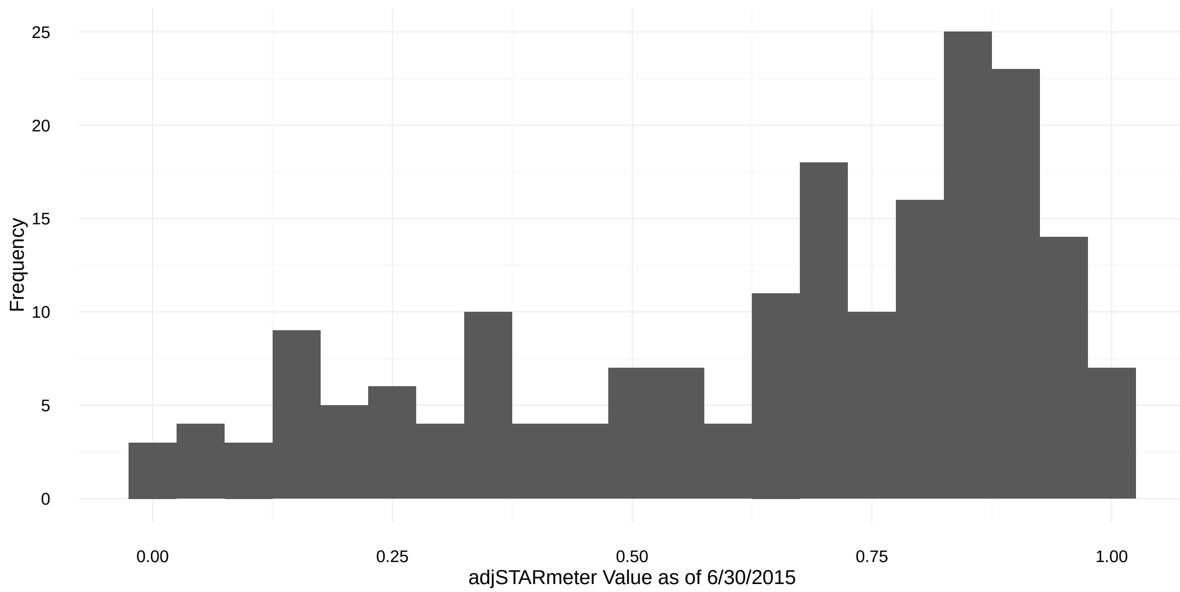

$adjSTARMeter$. We did this by applying the following formula, for role r,

$$adjSTARmeter_r = 1 - \frac{\min(STARmeter_r, 20000) - 1}{19999}. \tag{2.1}$$

Thus, for the highest ranked star in a given week, $adjSTARmeter_r$ returns one. *adjSTARmeter* returns 0 if STARmeter is 20,000 with linear interpolation otherwise and with increases in value implying higher star power.

To provide context and better understand the range of values, Table 2.1 provides ranges of *adjSTARmeter* and example actors from our data set who fit into the ranges. We also provide a histogram of non-zero *adjSTARmeter* values for all actors in our dataset on a given date in Figure 2.6. As can be seen, the distribution of positive values is non-normal with a a large proportion above 0.5.

Table 2.1: Example Actors for adjSTARmeter values

adjSTARmeter Range	0.01-0.50	0.50-0.90	0.90-1.00
	Sutton Foster	Catherine Zeta-Jones	Daniel Radcliffe
	Corbin Bleu	Megan Mullally	Sofia Vergara
	Jason Alexander	Neil Patrick Harris	Emma Stone

Operationalizing star power with STARmeter is attractive because IMDB maintains a history of the ranking with weekly updates. This history allows us to calculate the popularity of a given star at any given moment in the past. This was crucial to rule out an alternative form of reverse causality actors becoming famous because of the success of their show. For example, prior to the television show Glee, Lea Michele was on Broadway in the musical *Spring Awakening*.

The STARmeter is not completely ideal though. An issue with the STARmeter is that it is derived from the actions of users of IMDB. If these users skew differently from the population of Broadway theatregoers, we might expect there to be a bias in the variable. This is a reasonable conjecture. While the IMDB user demographic is not readily available, the average age of a Broadway theatregoer is 44 years compared to a United States median age of 39 (Broadway League, 2015; CIA, 2016). We would further conjecture that being an online platform would skew IMDB demographics even younger than the national median. Thus our measure of star power is likely a measure of how a slightly different demographic views these actors, rather than the demographic we are actually interested in.

Note though that this problem would attenuate the coefficients found in our results. At the theoretical extreme, if demographics are completely different, and if those consumer segments have completely different interests in celebrities, the STARmeter variable would have no correlation with the theoretical quantity we are interested in. In this case, we would expect to find no statistically significant effect from it. Thus, even though this bias is likely present, it is likely to lead to not being able to reject the null hypothesis.

Another issue is that the data we have is an indirect measurement (e.g. web/mobile searches for a given actor), rather than the measurement itself. We only have the ranking of the underlying measurement. This means that linear changes in the rank do not imply linear changes in the underlying data. In fact, this is likely the case. Thus, hidden in the ranking system may be non-linear changes in the underlying data. Unfortunately, that primary data is not directly available.

Finally, for each role, we collected data on the total awards the actor had won for all major awards (Tony, Grammy, Oscar, Emmy), prior to joining the show. To do so, we used information from `imdb.com`, `ibdb.com`, as well as `grammy.com`. We included all individual awards (e.g, Actors starring in a "Best Comedy" were not considered to have won an award).

From the data we collected, we create the following explanatory variables. A reference description of all variables used in our study and the source of data is presented in Table 2.2. Our main research question concerns how star power effects the revenue of a show. We create the variable *Starpower* to test this. This variable returns the max *adjSTARmeter* from all actors who were in the show on that given week. Formally, let S_{it} return the set of roles for show i in week t.

Then we can define

$$Starpower_{it} = \max_{r \in S_{it}} adjSTARmeter_r. \tag{2.2}$$

We hypothesize that the β coefficient for *Starpower* should be positive if increased star power results in higher revenue. The variable returns the max under the hypothesis that a single star actor is what drives revenue. In this case, all that matters would be the star power of the most known actor.

However, Elberse (2007) and Nelson and Glotfelty (2012) point out that there may be ensemble effects in films. Ensemble effects imply that multiple stars have a different effect than a singular top star. It would seem natural that these effects may also exist in live theatre.

On Broadway such casting tactics are common. For example, in 2010 for the revival of the musical *Promises, Promises*, both Sean Hayes (*adjSTARmeter*= 0.90) and Kristen Chenoweth (*adjSTARmeter*= 0.96) were hired as the leads of the show. The *Starpower* variable would only consider Kristen Chenoweth and ignore the combined star power of the two actors. To incorporate ensemble effects, we create the variable *Starpower3* which returns the sum of the three highest *adjSTARmeter* for a given show in a given week.

TonyAwards and *OtherAwards* act similarly for the underlying Tony Award/Other Award data and return the highest number of Tony Awards / Other Awards won by any of the cast who were in the show in the previous week. Note that these need not be from the same actor. For example, *Starpower* could return the *adjSTARmeter* of a well known actor who was in the show that week, but *TonyAward* could return the number of tony awards for a separate actor in the same show that week.

A customer's response to casting choices is not made in a vacuum. A customer makes choices amongst several alternatives. In the film industry, this means customers at a movie theatre make choices among several movies that may be playing. Each film may have different perceived quality by an individual consumer. On Broadway, a similar phenomenon occurs. Since almost all theatres are in the same Manhattan district, search costs are minimal and information on all shows is easy to retrieve.

Consequently, the models above do not take into account competitor actions. If a show adds star power and has a resulting increase in revenue, it is unknown whether Adding a star to a given show could shift purchases from one Broadway show to another. In this case, from an industry perspective, expenses on stars would be inefficient. The alternative is that adding a star entices customers who would otherwise not.

To test for this effect, we create the variable *CompetitionStarpower*, which returns the sum of all other shows *Starpower*. *CompetitionStarpower* thus represents how much star power other shows in the same week have. We would hypothesize a negative coefficient on the variable.

We also include control variables for various time components. To account for seasonality of revenues, we include dummy variables for the month as well as various Holidays. A list of the specific holidays we considered is presented in the appendix. Finally, the variable *DaysOpen* returns how many days the show has been open to control for potential life cycle effects. A final correlation matrix and summary statistics of all non-control variables are presented in Table 2.3 and Table 2.4.

2.4 Results

2.4.1 Graphical Analysis

Our main research question concerns how star power effects revenues on Broadway. We first begin with a graphical analysis. In Figure 2.7 we show scatterplots of *Gross* against *Starpower* and *Starpower3*. In the top row, we provide a best linear fit. For both independent variables, the linear fit suggests a positive correlation between higher *Starpower* and *Gross*. This result could be confounded due to differences in the unobserved individual effects and correlation between those individual effects and star power. For example, a lessor grossing show deciding

Table 2.2: Variable Descriptions

Variable	Description	Data Source
Gross	The inflation adjusted, per performance revenue for a given show in a given week.	`ibdb.com`
Starpower	The max *adjSTARmeter* for the set of actors who were in a given show during a given week.	`imdb.com`
Starpower3	The sum of the three top *adjSTARmeter* for the set of actors who were in a given show during a given week.	`imdb.com`
CompetitionStarpower	The sum of all other shows' $Starpower$ that were running during the same week.	`imdb.com`
TonyAwards	The maximum number of Tony Awards won for the set of actors in show i during week t.	`ibdb.com`
OtherAwards	The maximum number of other major awards for the set of actors in show i during week t.	`imdb.com/grammy.com`
DaysOpen	Variable indicating how many days the show had been open by the end of week t.	N/A
Month	Dummy variables indicating whether the last day of the week was in a given month.	N/A
Holiday	Dummy variables indicating whether a given holiday occured during the week.	N/A

Table 2.3: Correlation Matrix

	Gross	*Starpower*	*Starpower3*	*CompetitionStarpower*	*TonyAwards*
Gross					
Starpower	0.14				
Starpower3	0.12	0.91			
CompetitionStarpower	-0.03	-0.02	-0.02		
TonyAwards	0.07	0.38	0.49	-0.02	
OtherAwards	0.07	0.42	0.5	-0.02	0.28

Table 2.4: Summary Statistics of Variables

Statistic	N	Mean	St. Dev.	Min	Max
Gross	9,420	102,301.100	56,578.950	10,079.130	360,051.900
Starpower	9,420	0.199	0.333	0.000	1.000
Starpower3	9,420	0.289	0.571	0.000	2.831
CompetitionStarpower	9,420	5.682	1.898	0.160	10.786
TonyAward	9,420	0.154	0.526	0	5
OtherAwards	9,420	0.200	0.843	0	18

on less star power.

To partially deal with this issue, in the second row of Figure 2.7 we provide selected show's box plots conditional on if a star was present or not that week. We define a "star week" as any show-week combination where $Starpower \geq C$ where C is some prespecified constant indicating at what level an actor is considered a star. Here, we assumed that $C = 0.90$. In this way we start to control for unobservable individual time-invariant effects.

We see that for the most part there is a difference in behavior between star weeks and no star weeks, with all selected shows reporting higher median gross under a star week than a no star week. Indeed, the effect is visibly apparent.

We can also examine the overall distribution of *Gross* conditional on the presence of a star. We then examine the distribution of *Gross* during star-weeks as compared to non star-weeks. [14]

Our analysis is sensitive to the choice of C, and thus we provide multiple values to test. We try the following values for C: $0.50, 0.70, 0.90, 0.95$. In Figure 2.8, we provide these conditional distributions. Along the x-axis is the specific gross, and along the y-axis is a smoothed estimate of the probability density of a star (green) vs non-star (red) of being at that value. The smoothing was done via a kernel density estimate with a Gaussian smoothing kernel. Bandwidth was selected using Silverman's rule of thumb. In Table 2.5, we provide the associated mean values of star vs non-star weeks at each definition.

Table 2.5: Mean Values for Star Weeks compared to Non Star Weeks

C	Mean / Standard Deviation of *Gross* during Star Weeks	Mean / Standard Deviation of *Gross* during Non-Star Weeks
0.5	\$114,570 (52,394.7) (N=1,993)	\$99,008.9 (57,207.4) (N=7,427)
0.7	\$114,684 (54,913.5) (N=1,498)	\$99,959.7 (56,588) (N=7,922)
0.9	\$99,254.7 (40,467.1) (N=677)	\$102,537 (57,634) (N=8,743)
0.95	\$112,293 (39,590.8) (N=317)	\$101,953 (57,049.6) (N=9,103)

A few points of interest arise. At all star definitions, the distribution of star grosses appears to be have less variance than the distribution of non-star grosses. This graphical intuition

[14] In this analysis, we use *Starpower* however an analysis with *Starpower3* results in qualitatively similar results.

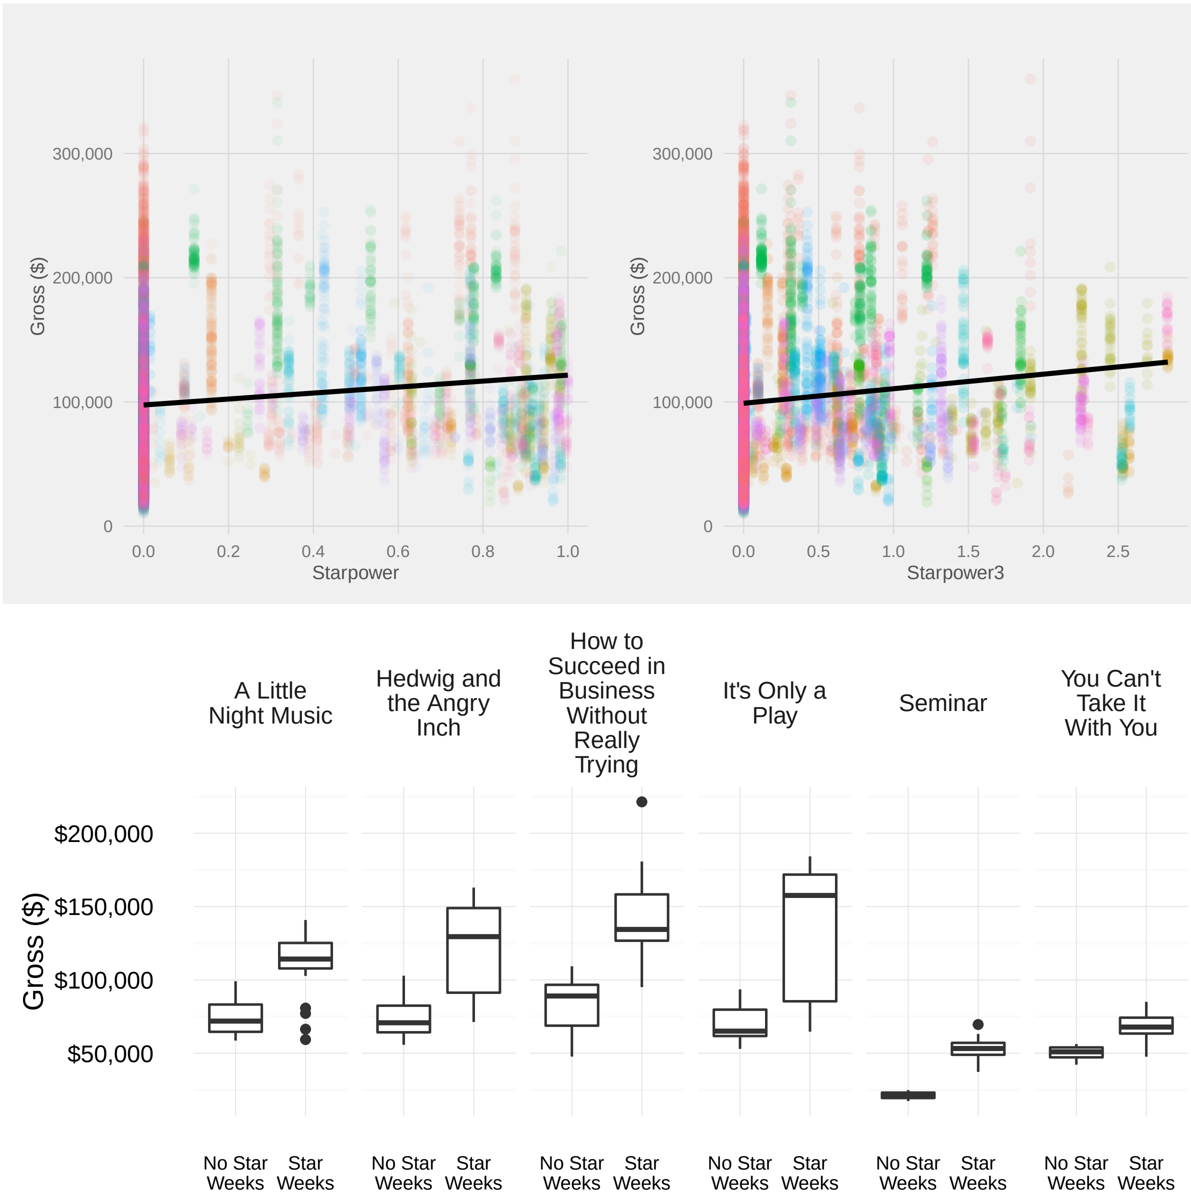

Figure 2.7: Scatterplots of independent variables Starpower and Starpower3 against dependent variable Gross. In all plots, color signifies a specific show. In the first row we provide in black the best fit linear line. In the second row, we provide selected samples. A "star week" is defined as one where Starpower ≥ 0.90.

Figure 2.8: Distribution of Gross conditional on star presence. Densities estimated via a kernal density estimate, with a Gaussian smoothing kernel. Bandwidth was selected using Silverman's rule of thumb.

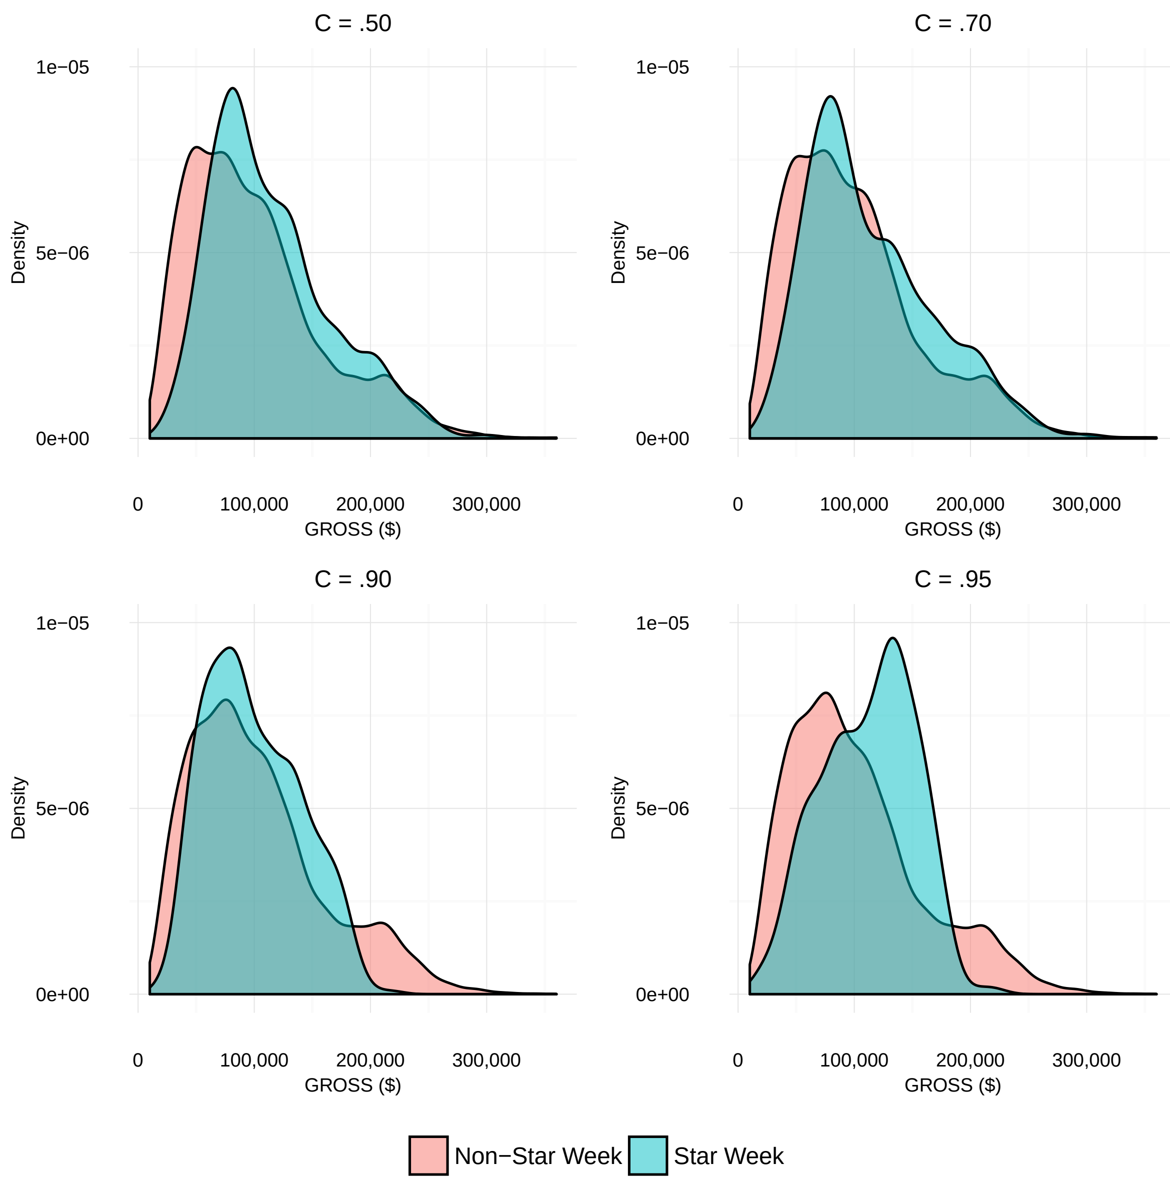

is confirmed by checking the standard deviation of the distributions in Table 2.5. At all C definitions, the standard deviation of star weeks are less than non-star weeks. In addition, in all but one definition, the mean of star weeks is higher than non-star weeks. Finally, under all definitions, the no-star distribution appears to have a higher probability of achieiving relatively low gross amounts. This would seem to bolster the claim that stars help producers reduce risk.

We can also qualitatively examine stochastic dominance of one distribution over the other. In this setting, stochastic dominance exists if, at every potential gross dollar amount, a distribution has at least the probability as the other of achieving a higher gross. Managerially, stochastic dominance would that there is no "downside" to hiring a star as they would have less down-side risk and higher upside risk. At the $C = 0.50$ and $C = 0.70$ definition, there seems to exist stochastic dominance of stars over no-stars. At other definitions however, there is a clear violation in the upper tail. For example, at $C = 0.90$ and $C = 0.95$ the non-star distribution has a higher probability of achieving higher grosses This would imply a tradeoff - stars may reduce downside risk but limit upside risk.

2.4.2 Regression Results

The prior sections provide some insights into how stars play a role in a show's revenue. However, we did not control for show-specific and time-specific effects. Here we provide formal results from multiple fixed-effects regressions and discuss managerial intepretations. Our model aims to predict the log of *Gross* as a function of star power, award covariates, and various controls. We use show-specific dummy variables to account for any unobserved time-invariant show-specific effects that may exist. These effects might include the script quality, song quality, and budgets.

In Models 1 and 2, we use the *Starpower* variable as our measure of star power. In Model 3 and 4 we instead use the ensemble measurement *Starpower3*. In Models 2 and 4, we add *CompetitionStarpower* to test the effects of competition, but leave it out in Models 1 and 3. We attempted squared terms on the *Starpower* and *Starpower3* variable but only found significance with *Starpower3*. We thus only report this variable, but present all attempted models in the appendix.

2.4.3 Model Results

In Table 2.6, we present coefficient estimates (and accompanying standard errors), and R^2 values from these models.

First, we examine our main research question and variable of interest. *Starpower* is positive and statistically significant at traditional levels (5%) across all specifications. The coefficient of *Starpower3* is positive and significant on the squared term but insignificant on the main term. This is consistent with the hypothesis that adding star power improves revenues. The positive and significant squared term would indicate increasing returns to ensembles of stars. Note that the two models using the *Starpower3* variable as a measure of star power have a higher adjusted R^2 value, indicating a better fit.

OtherAwards is positive and significant for specifications with *Starpower* but only marginally significant (at the 10% level) for specifications with *Starpower3*. This is partial evidence that, independent of star power, some awards appear to have an association with higher revenues.

Table 2.6: Regression Results

	(1)	(2)	(3)	(4)
Starpower	0.095**		0.096**	
	(0.038)		(0.037)	
Starpower3		−0.055		−0.054
		(0.054)		(0.053)
Starpower3^2		0.118***		0.117***
		(0.028)		(0.028)
CompetitionStarpower			−0.004	−0.004
			(0.006)	(0.006)
TonyAwards	0.052	0.014	0.051	0.014
	(0.033)	(0.035)	(0.032)	(0.034)
OtherAwards	0.040**	0.021*	0.039**	0.020*
	(0.017)	(0.012)	(0.017)	(0.012)
DaysOpen	−0.411***	−0.417***	−0.403***	−0.409***
	(0.103)	(0.098)	(0.102)	(0.097)
DaysOpen^2	0.031***	0.032***	0.031***	0.031***
	(0.010)	(0.010)	(0.010)	(0.010)
Constant	12.801***	12.796***	12.775***	12.770***
	(0.425)	(0.416)	(0.420)	(0.411)
Mean Squared Error	0.03952	0.0381	0.03948	0.03806
Observations	9,420	9,420	9,420	9,420
R^2	0.888	0.892	0.888	0.892
Adjusted R^2	0.884	0.888	0.884	0.889
Residual Std. Error	0.202 (df = 9131)	0.198 (df = 9130)	0.202 (df = 9130)	0.198 (df = 9129)

Note: *p<0.1; **p<0.05; ***p<0.01

Dependent variable of LOG(*Gross*). Control Variables Excluded. Cluster Robust Standard Errors shown in parentheses.

TonyAwards is positive but not statistically different from zero across all specifications. This is a surprising finding because other major awards do appear to have an effect. To the extent the broadway customer base values success on broadway specifically, we would expect that signals of that success to matter. Awards on Broadway, as a signal of quality, would signal to insiders (and maybe outsiders) that a show is good.

Competitor casting choices, and their effects can also be evaluated with our model. *CompetitionStarpower* is negative as predicted, but non-significant across both models. That is, we cannot reject the null hypothesis that competitor casting choices have no effect on a given show. If this hypothesis were true, it would be consistent with the idea that star power attracts consumers who otherwise would not have gone to a Broadway show at all. We discuss the managerial and industry implications of this later in the paper.

Finally, we can estimate the revenue gain from increasing the star power associated with a show. This would be the most managerially relevant outcome from our model, since it would inform whether a blockbuster strategy is likely to pay off. Because of the log transformed dependent variable, to find the average change of *Gross* from a change in *Starpower*, we can evaluate the following expression

$$\frac{\Delta Gross}{\Delta Starpower} = \overline{Gross}\left(e^{\beta \Delta Starpower} - 1\right)$$

where β is the estimated coefficient(s) for *Starpower*.

In our first comparison we examine how our models suggest an average show's *Gross* would improve if the show hired a star actor (*adjSTARmeter*= 0.98) when it previously had no actors with star power(*adjSTARmeter*= 0). This action would imply an increase in *Starpower* and *Starpower3* of 0.98. We summarize these effects in Table 2.7. Estimates of dollar change

Table 2.7: Estimated effects from hiring single actor with adjSTARMeter=0.98

Specification	Average Percent Change	Mean Estimate
Model 1	9.8 %	$9,983 ($79,868 per week.)
Model 2	6.1 %	$6,223 ($49,786 per week.)
Model 3	9.8 %	$10,058 ($80,465 per week.)
Model 4	6.2 %	$6,300 ($50,403 per week.)

Table 2.8: Estimated effects from hiring two actors with adjSTARMeter=0.7 each

Specification	Average Percent Change	Mean Estimate
Model 1	6.9 %	$7,036 ($56,285 per week.)
Model 2	17 %	$16,975 ($135,801 per week.)
Model 3	6.9 %	$7,087 ($56,700 per week.)
Model 4	17 %	$17,073 ($136,585 per week.)

range from $6,223 in Model 2 to $10,058 under Model 3. Generally, the models built on the ensemble variable predict lower increases for a single star hiring.

We can also explore ensemble star hiring. Instead of hiring a single actor, we consider the estimate effects if a show hired two star actors each with an *adjSTARmeter* of 0.7. Here, the increase in *Starpower* is only 0.7, however the increase in *Starpower3* is 1.4. We summarize each model's predicted effects in Table 2.8. Perhaps intuitively, our measure built specifically for ensemble effects predicts a much larger effect than the measure that only measures the top actor. Model 2 predicts a $16,975 average increase in *Gross* as compared to a $7,036 from Model 1.

2.4.4 Model Diagnostics

We must test that our model assumptions are valid. One key model assumption we made is to treat unobserved individual effects as fixed effects. The theoretical reason is that we did not want to assume that the vector of covariates was independent from the unobserved individual effect. However, we also checked for this assumption via a Hausman Test which confirmed a fixed effects specification ($p < 2.2e^{-}16$) [15].

In Figure 2.9, we provide histograms as well as Q-Q plots of the residuals of each of our models. Residuals appear to be normally distributed and centered around zero for all models.

Recall that in our estimation procedure we adjust the standard errors using "cluster robust standard errors". The Cameron et al. (2011) method is robust to serial correlation and heteroskedasticity (Wooldridge, 2002, p.275). Our data shows signs of both issues. A Breusch-Godfrey test for serial correlation in panel models finds statistically significant ($p < 2.2e^{-}16$)

[15] A replication of our analysis assuming random effects is presented in the Appendix however.

evidence of serial correlation in the residuals. A Breusch-Pagan test for heteroskedasticity also finds significant evidence ($p < 2.2e^{-}16$) against homoskedasticity [16].

2.4.5 Case Studies

In the following sections we provide case studies of specific star entries and exits into Broadway shows. These examples are intended to provide understanding of the business environment. They are also intended to show how practitioners can estimate the effect of a specific star entry. This may be of interest for managers interested in determining post-hoc whether a cast change was worthwhile.

American Idiot

American Idiot was a Broadway musical based on the Green Day album of the same name. The show opened in April 2010. On September 27, 2010, the show's producers announced that Green Day lead singer Billie Joe Armstrong would star as the "St. Jimmy" character for only one week, starting only a few days later. After his initial run week run he later played the same characters multiple weeks throughout January and February. On March 10th, 2011 the show announced that it was closing on April 24th and announced even more weeks with the star.

The process of choosing Armstrong for this specific musical was obviously endogenous to the show. The producers chose him due to his connection with the source material. However, the actual weeks chosen was closer to an exogenous process. Unlike most actor contracts on Broadway, the star was in the show in non-contiguous weeks, sometimes with significant breaks. This occured due to the star's previously scheduled tour and logistical availability. Thus, the choice of which specific weeks the star would be in was constrained.

Table 2.9: Summary Statistics of American Idiot Grosses. Star Weeks are defined as those which included Billie Joe Armstrong.

	Star Weeks	Non-Star Weeks
Min	83,190	51,930
1st Quartile	118,300	60,460
Median	129,300	65,940
Mean	132,100	66,310
3rd Quartile	149,300	70,100
Max	179,700	83,090

We first separate weeks with and without the star. We label these star weeks and non-star weeks respectively. Summary statistics of these weeks are presented in Table 2.9. There is a difference in means of $65,829 , a statistically significant difference.

This difference could be explained via three mechanisms. The most obvious mechanism is that the star increased awareness or quality of the show in the eyes of consumers. This would

[16] In the tests, the null hypothesis is no serial correlation and homoskedasticity, respectively.

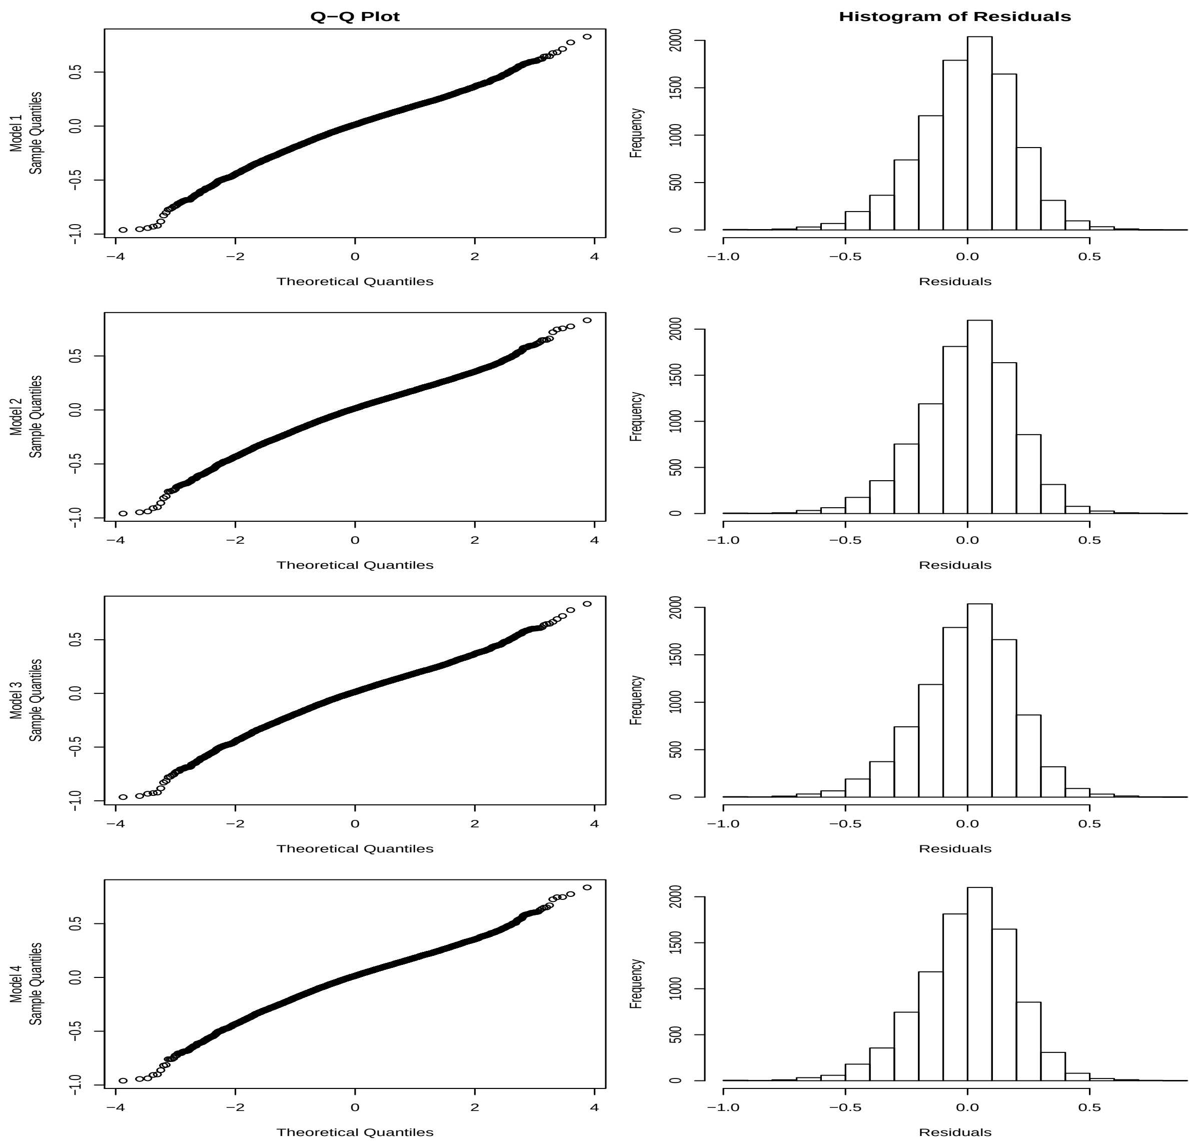

Figure 2.9: Residual diagnostics for four models. Left column is a Q-Q plot which compares theoretical quantiles to sample quantiles in the residual data under a given model. Right column is a histogram of residuals. Histogram bin sizes are chosen using Sturges formula.

Figure 2.10: American Idiot performance relative to control group. Red shading indicates when the star was in the show. Green shading indicates when the public knew there were future performances with the star.

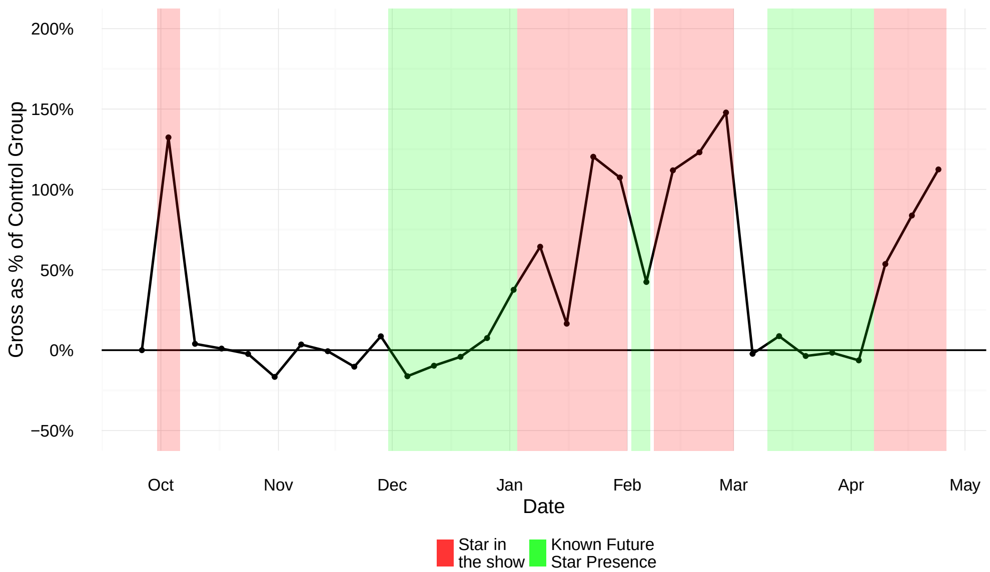

be consistent with the blockbuster strategy and suggest that the star was successful in bringing in additional revenues.

Alternatively, the difference could be due to seasonal time effects. That is, the star being in seasonally "better" weeks, and being associated with an increase without being the cause. In this case study, the likelihood of this mechanism is mitigated because of the constrained process by which weeks were selected. If concert sales and theatre sales are correlated we would expect the star to be in worse weeks, because the "better" dates were taken up by the tour.

Finally, the difference in means could be explained by a temporal altering of purchase by consumers, rather than new consumers. That is, a customer already interested in seeing American Idiot might change the date they see the show to see the star. If so, we would still find a week to week change and association between the star and revenue, but total revenue would be unaltered.

We seek to test whether the second and third explanations are valid in this case study. To do so, we provide a graphical analysis of *American Idiot* performance relative to a control group. The control group consists of all other shows who ran during the entire time period of our analysis. We define performance as the percent change in *Gross* relative to a reference date (i.e. on the reference date both the show and the control group are 100%). The reference date we use is for the week ending September 26th 2010, when it was announced that the star would be appearing in next weeks shows. We use this date because it is the last week under which the star's arrival was not public information. Therefore, it should be unbiased relative to the control group.

We then plot the relative performance of *American Idiot* compared to the control group[17]. Figure 2.10 shows the results of this analysis. The discrete points in the data have been joined by lines for clarity. We shade dates in red to signify when the star was in the show, and green for when it was public knowledge that the star would be in the show in the future.

For the initial week in October with no public notice, the show clearly outperformed the control group. In January, the star enters again and again shows clear outperformance. The same is true for selected weeks in February. Note that the week in February where no-star is present also significantly outperformed the synthetic control. During this week, the singer Melissa Ethridge performed in the show in place of Billie Joe Armstrong. She has won multiple grammies and could be considered a star as well. Finally, there was also a clear outperformance during April while the star was present. This is less useful in establishing star outperformance since customers knew the show was closing soon. Part of the outperformance could be due to closing effects.

We can also use this case study to test the third mechanism, temporal changes in purchasing. To do so, we shift attention to the green shaded sections. Notice that during these times the show appears to have performed similarly to the control group. That is, even when it was public knowledge that the star would be in the show again, revenues do not appear to have underperformed the control group. If temporal changes in purchasing behaviour explained outperformance, we would expect a similar magnitude of underperformance in the immediate weeks before and after the star. We do not see this, and thus this mechanism cannot fully explain the significant difference in revenues we see.

This case study was a unique situation on Broadway in that weeks were largely non-contiguous. In reality, the major variable in our study, *Starpower*, is sluggish in the sense that actors stay for many months/years. In the next section, we look at a case study of this.

How to Succeed in Business Without Really Trying

How to Succeed in Business Without Really Trying (abbreviated How To Succeed for the rest of this section) was a Broadway musical that started previews February 26, 2011, opened on March 27th 2011, and closed May 20, 2012. When the show opened, the main character was played by Daniel Radcliffe, the well-known star of Harry Potter.

Daniel Radcliffe left the show on Janury 1, 2012 and for three weeks the character was played by Darren Criss (a TV actor/songwriter). Afterwards, Nick Jonas (from the band The Jonas Brothers), joined the show for the rest of its run.

We conduct a similar analysis to *American Idiot*. In Figure 2.11, we graph performance of *How to Succeed* relative to the control group of all other shows running during this period. The reference date was chosen as March 6th 2011. Here, the reference date was chosen as a date early in the run. Interpreting this reference date is difficult however as Daniel Radcliffe was in the show from the start. Rather than interpreting performance relative to without the star, performance here can be roughly interpreted as relative to Daniel Radcliffe.

Note that the switch to Darren Criss (red shaded) did not appear to harm revenues. The same relative outperformance appeared to occur during his weeks as Daniel Radcliffe, if not more. Darren Criss had an average outperformance of 51%. In contrast, Daniel Radcliffe had

[17]That is, a 100% value indicates American Idiot doubled what it was expected to make if it had performed like the control group.

Figure 2.11: How to Succeed performance relative to control group.

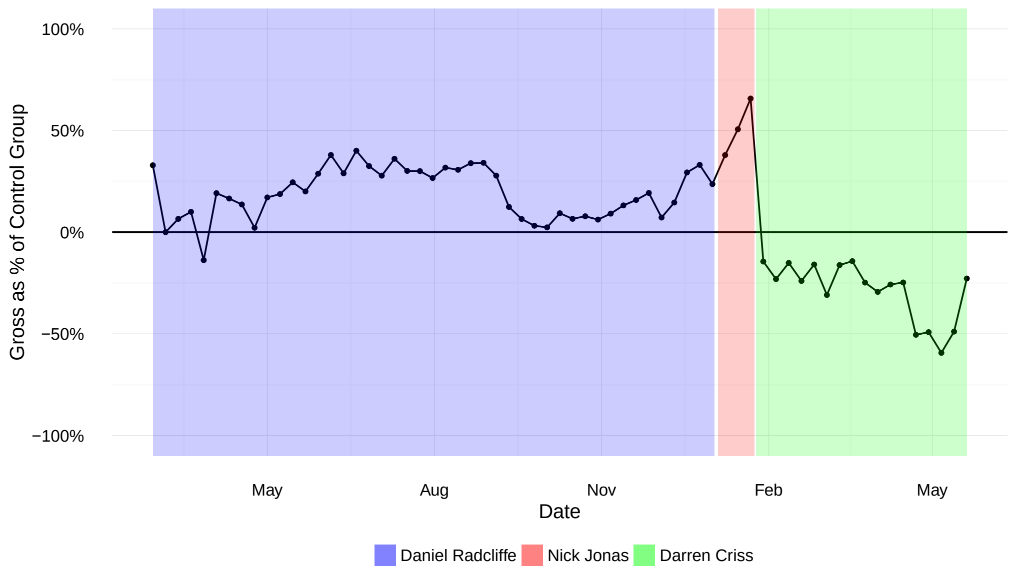

an average outperformance of 19%. A potential explanation is that consumers who were solely interested in seeing him were forced into a smaller number of weeks. The final cast change to Nick Jonas did seem to negatively effect revenues. For the first time, *How to Succeed* began to underperform the control group in a sustained fashion, before finally deciding to close.

It is not immediatly clear how Daniel Radcliffe and Darren Criss differ from Nick Jonas. The qualitative differences between these actors (particularly Criss and Jonas) may appear slight. However, when we look at the corresponding adjSTARmeter of each actor during their run, it is clear that Nick Jonas is the least out of the three. Daniel Radcliffe and Darren Criss had adjSTARmeter values of 0.98 and 0.96 respectively. However Nick Jonas had a adjSTARmeter of only 0.83. In this case, relatively small values of adjSTARmeter were associated with relatively large changes in *Gross*.

2.5 Discussion and Conclusion

Star actors and their effects in entertainment businesses have been, and will continue to be, an important question for entertainment managers and academics studying the industry. The blockbuster strategy is intuitive and popular with practitioners. Previous attempts in the film industry that try to answer these questions have generally failed to identify causality due to potential endogeneity in actor selection. In this paper, we explore a novel dataset in live theatre to answer the research question and use within-show variation in star power to test what changes star power has on revenues. Our paper is unique in that this sort of dataset is unparalleled in the film industry.

Our main finding is that increases in star power are positively associated with increased revenue, even after controlling for show-specific fixed effects. This is consistent with the idea that the actors themselves are causing these revenue increases through increased awareness or increased quality estimates of the show. Note that in the Broadway context, supply is incapable of shifting and thus any effect we find cannot be explained by changes in capacity/supply. Liu et al. (2014) finds that stars in films may increase revenue, but primarily through increases in screen count. Thus, our result suggests a different pathway than what Liu et al. (2014) finds.

Recall from our case study that a potential explanation of our findings is that a temporal shift in purchases occured, rather than an overall increase in revenue. In the case study, we found evidence that this was not occuring, since weeks where the star was not present, and where the audience knew he would be later playing, did not significantly underperform. Note that this is an extreme example, since the star in American Idiot had a chaotic schedule where it would have been very easy for consumers to delay purchases. Contrast this with the general setting, where an average actor was in a show for 31.9 weeks. Consumer behaviour would likely only be able to change at the margin, and the vast majority of weeks would not be effected. Thus, we believe this effect cannot explain our results.

Our results on face value appear to support a blockbuster strategy. Our model suggests that actors with more star power increase revenues. More research is needed though. Deciding whether to use a star depends not just on the revenue they bring in, but on their cost. An in-demand star may fully capture the value they create by demanding higher pay. Unfortunately, without information about salaries on Broadway, it is not yet possible to answer whether this is occurring. Anecdotally, some publications suggest that top stars can make anywhere from $100,000 to $150,000 per week (Simonson, 2010). In this range, our models would suggest that stars fully capture the value they provide, and that hiring them does not increase profits. What this does suggest is contractual parameters for stars should be negotiated with this in mind. Given the high uncertainty, a profit sharing contract might be more feasible than a fixed-payment contract (for a discussion on this see Chisholm (1997)).

Our results also suggest potential industry responses. Broadway is a relatively small local industry, with a small number of theatre owners. This gives the potential for non-competitive strategies to arise. For example, if hiring stars were a zero-sum game, theatre owners could engage in a tit-for-tat type strategy for those who hire them. Our research suggests that stars are not zero sum though. We find that competition that increases star power has a non-significant effect. This is consistent with the conclusion that stars grow the market for Broadway and attract new customers, rather than shift customers.

A natural question that arises is how generalizable our findings are to other cultural products (e.g, the film industry). While our revenue estimates are obviously unique to Broadway, we feel our theoretical results would generalize to most cultural products. The underlying process with which stars are associated with increased revenues is likely to be similar to those in other cultural industries. Live theatre is unique in many respects for cultural products, but the underlying process by which shows become a success is similar - a complex dynamic system of social influences and awareness. If stars contribute to success on Broadway, it is likely that they do in the film industry.

Another potential industry with star effects is sports, where star players have the potential to draw attendance and viewers. In the sports industry, live "performances" are again created in front of a live audience and with potentially different stars. Players can change because of injuries, trades, or strategic rest. In March 2017, NBA commissioner Adam Silver warned that teams could face fines for resting star players without sufficient notice.[18] Future work can work on replicating our results in the sports industry, where demand seemingly is most influenced by the probability of win.

Our paper also has limitations, primarily due to the dataset. While we conclude that reverse causality cannot explain revenue increases that are a result of stars, we also cannot claim full direct causality. The process of choosing stars is still a process endogenous to the show and time. While we have controlled for time factors, as well as show-specific characteristics, we cannot control for other managerial actions. It is plausible that managers change efforts in response to a casting change. Most importantly, they may alter marketing budget in response to changes in star power. In the film context, this explanation has been advanced by (Prag and Casavant, 1994) who found no effect from star presence after accounting for marketing expenditures. On Broadway, marketing costs are private and data is scarce, so our study cannot account for them. This explanation, if true, would suggest that managers constantly under invest in marketing. Given the number of shows running on Broadway at once, and the amount of experience managers have in this industry, we would be surprised if managers were systematically under investing. A single manager who over invests would find it profitable and outperform their peers.

An aspect uncaptured in our paper is the performing history of certain actors, and how that performing history "fits" with a given show. In our case study, it is abundantly clear that Billie Joe Armstrong had a very strong fit with the musical, since he was the original writer of the source material. Yet quantifying chemistry between an actor and a project is difficult and operationalizing the concept is unclear. However, future research should consider these effects. If a star well known for drama does a comedic film, is the effect as large as if an actor already known for comedy was chosen? There could be a positive surprise factor from the move, but also a negative impact from an unknown quality effect. In either case, this would help further dimensionalize our actors.

The literature on stars to date has revolved around the film industry. The theme of stars and star power pervades all cultural industries however, and managers across them face similar dilemmas. We hope that this paper has shown the value of considering theoretical questions in different contexts, and that future research continues to do so.

[18]https://www.usatoday.com/story/sports/nba/2017/03/20/adam-silver-sends-memo-nba-owners-warns-teams-not-rest-star-players/99436202/

Chapter 3

Dynamic Stochastic Knapsack Problem with Adaptive Interaction: Live Entertainment Capacity Based Revenue Management

3.1 Introduction

The Knapsack problem is a classic combinatorial Operations Research (OR) problem. Stated informally, it asks how to select a set of items with varying associated weights and rewards into a knapsack with finite weight capacity. Given the combinatorial complexity, and despite "good" approximations and algorithms, even *deterministic* versions of the Knapsack problem pose great computational challenges; the decision and optimization versions are NP-complete and NP-hard, respectively. Extensions to incorporate *stochastic* weights and rewards, and *dynamic* arrival of items adds further complexity, but in a sense also a "simpler" structure to the original problem. In this paper we consider an extension to the Dynamic Stochastic Knapsack Problem (DSKP) by incorporating an *adaptive interaction* with the items regarding their placement in the knapsack. We assume the knapsack decision maker can only make certain *compartments* within the knapsack available but ultimately the items themselves select their placement. This autonomy includes a *balk* option not to enter the knapsack. Our motivation to consider such extensions stems from Revenue Management (RM) applications of group seatings in live entertainment venues.

Live entertainment venues such as theatres, concert halls, and sport arenas represent a multi-billion dollar industry in the United States. To illustrate, the combined National Football League teams have been estimated to be worth upwards of 62 billion dollars (Ozanian, 2015); Broadway commercial theatre productions make more than one billion dollars in revenue per year (in a 6 square block radius of downtown New York city!); and Taylor Swift's "1989" tour, grossed over 250MM dollars (Lewis, 2015). Furthermore, live entertainment possess some of the key RM characteristics: a fixed supply of perishable inventory (i.e. expired tickets are worthless); a predominantly fixed cost structure; and highly uncertain demand both in the short and long term. In the short term, an entertainment venue must be able to predict booking curves and demand levels for performances at different times of day, for different days of the week, in different times of the year. In the longer term, live entertainment is a notoriously tricky business to predict. In cultural products, success in the long term is a complex mix of social dynamics and quality that is highly path dependent (Salganik et al., 2006). However, despite all the suitable traits and large economic value, live entertainment represents a relatively new context for RM approaches to pricing and inventory management.

The "first generation" of live entertainment RM consisted of moving away from *general admission* type of pricing to *variable* pricing based on seat or zone location; colloquially referred to as *scaling the house* (Courty, 2000; Phumchusri and Swann, 2014). In addition, many professional sports teams and commercial theaters have recently extended variable pricing to account for desirability of specific games/shows, opposing teams/cast members, day of the week, seasonal effects etc., and introduced *dynamic* pricing by adjusting prices based on time and demand (Xu et al., 2016). Oddly though, little attention has been given to the inventory management or capacity control aspects of RM.

One reason why RM has been lagging in the live entertainment industry is that models traditionally developed in the airline context do not include some of the fundamental consumer dynamics reflected in live entertainment; the most important being the experiential attribute of the service consumption. Purchasing a flight is fundamentally done as a means to another ultimate goal, e.g. personal travelers fly to visit a vacation destination or family, while business travelers fly to attend meetings. The act of flying commercial airlines is rarely done as

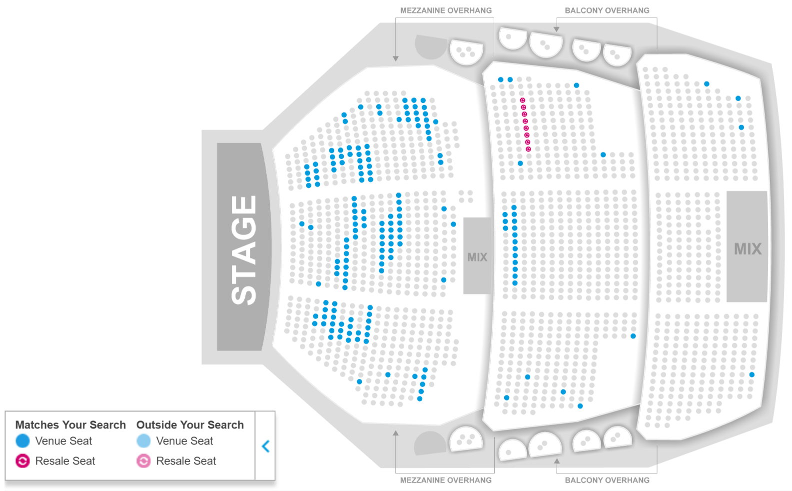

Figure 3.1: Screenshot of the Broadway Musical Aladdin's seating chart. Source: Ticketmaster.com

an enjoyable activity for its own sake. Contrast this with live entertainment where the experience is the goal. This is known as experiential consumption (Gilovich et al., 2015; Holbrook and Hirschman, 1982). The experience is composed not just of the production or sports game but also of the people you attend with. Gilovich et al. (2015) noted that "We . . . go to restaurants, concerts, and sporting events with fellow foodies, music lovers, and sports enthusiasts." This perhaps obvious observation changes an underlying assumption in many RM models. In particular, that customers are not separable. A group of n customers cannot be modeled or approximated as n individual customers. Experience at the box office suggests that if n contiguous seats are not available customers may balk. This also suggests that n individual seats are not worth the same amount as n contiguous seats.

There is real-world evidence that suggests this may be hurting revenue. Figure 3.1 shows the seating map for the Broadway show *Aladdin* from `ticketmaster.com` for a Tuesday evening performance in June 2016. This seating map is from a couple of days prior to the performance and shows in dark which seats are still available for purchase. Across the whole theatre there are many single seats still available, but sparsely situated across the venue. With two days until show time, these are arguably more difficult to sell as individual customers are rare. On the other hand, had customers been offered and purchased seats differently these remaining seats could have been less sparse and easier to sell.

The purpose of this paper is to propose a dynamic method for the optimal capacity control. Specifically, our objective is to analyze how venues should dynamically update the set of seats made available based on time, remaining capacity and customer group size - which we refer to as *stacking the house*.

To study this problem in the context of RM, we model live entertainment sales as a Dynamic Stochastic Knapsack Problem: future demand is uncertain, arrivals have stochastic varying weights and rewards, and irrevocable decisions are made for each item. In order to reflect important attributes of live entertainment we include two extensions. First, we assume that the finite capacity of the knapsack is *compartmentalized* according to a predefined reward-to-weight ratio scheme (e.g. a theatre with rows priced differently). Second, to reflect the fact that customers choose seats, we extend the traditional DSKP, where item placement into the knapsack is deterministic, with a *post-decision* stochastic item placement. In other words, we model items as having autonomy in selecting their compartment placement. In particular, since customers may choose *not* to purchase any seats we permit items to "reject" entering the knapsack altogether. We label these extensions as "DSKP with *adaptive interaction*."

3.1.1 Literature Review

The Knapsack problem dates back to some of the original and classic OR papers, such as Dantzig (1957), and has been applied in a wide and rich context, including transportation and manufacturing, e.g. Cheng et al. (1996); Cherbaka and Meller (2008); Sharma and Dubey (2010). In the original formulation the problem is to select a subset of items, with varying but known weights and rewards, such that the total reward is maximized subject to staying under a total weight capacity (of the knapsack). This formulation is also known as a 0-1 knapsack problem since each item can only be selected and placed once. Over time several extensions have been formulated, including: *unbounded* - where items can be selected as many times as desired, i.e. sampling with replacement (Andonov et al., 2000); *multiple choice* - where items are of different types and there exists a constraint that at least one item from each type must be placed (Kellerer et al., 2004); and *multiple knapsack* - where there are several knapsacks, each with different capacity, and each item can be assigned to only one knapsack (Martello and Toth, 1980). Although our problem formulation resembles a dynamic and stochastic version of the multiple knapsack problem, there is a key difference. We assume a single knapsack is compartmentalized according to predefined reward differences, e.g. rows or zones at a venue that are priced differently. In the multiple knapsack problem it is assumed an item's reward is independent of the knapsack it is assigned to. However, in the special case of a venue with all rows (or zones) priced equally our model formulation can be regarded as a dynamic stochastic version of the multiple knapsack problem.

The knapsack problem is related to another classic OR problem. Namely, the bin-packing problem, e.g. Johnson (1974). In the bin-packing problem, a set of items must be placed into exogenously sized bins, so as to minimize the number of bins used. In this sense, the (multiple) Knapsack and Bin-packing problems have reversed objectives and constraints. In the bin-packing problem, it is assumed that all items must be placed, and the objective is to minimize the amount of bins (knapsacks) used. In the knapsack problem, the amount/capacity of knapsacks (bins) is fixed, and the objective is to choose which items to maximize the reward. For our purposes of live entertainment sales where capacity (seats) is fixed the knapsack problem is more suitable.

Traditionally, the original knapsack problem as well as variants have been *static* in time or the number of items. That is, even when weights and rewards are allowed to (stochastically) vary, the total number of items is known at the start of the problem. In contrast, Papastavrou et al. (1996) defined a *Dynamic Stochastic* Knapsack Problem (with deadlines). By dynamic, the authors extended the original formulation to include that the entire set of items is unknown and irrevocable accept/reject decisions are made for sequentially arriving items. By stochastic, the authors mean that the weights and rewards of arriving items are drawn from a distribution. Papastavrou et al. (1996) show that a threshold policy in the reward is optimal for the discrete time version of the problem. That is, that at each state there is a critical reward threshold, such that items with rewards above the threshold are accepted. The authors also examined extensions to some special cases (equal weights, equal rewards, proportional weight to reward). Kleywegt and Papastavrou (1998) examine and define the DSKP for fixed weights but in continuous time. They introduce a per-time waiting cost and give an option to "stop" the arrival of items, showing that for finite horizons a threshold type policy is still optimal. van Slyke and Young (2000) and Kleywegt and Papastavrou (2001) both extend the continuous time DSKP to stochastic item weights. Notably, van Slyke and Young (2000) extended the model into higher dimensional capacity settings and noted that this could apply to group seating dynamics.

Another related paper is Dizdar et al. (2011), who also study a dynamic knapsack problem with incoming agents of uncertain weight and reward. Similar to our problem, the reward is defined by an underlying *willingness-to-pay* (wtp). Their decision is the payment (price) and allocation of the (single) knapsack's capacity. Framed in the economics literature, their work is primarily focused on discussing concavity properties of the objective function. In contrast, our research focus is to analyze properties regarding optimal policies, and in particular to establish conditions on managerially relevant threshold type policies.

In the RM literature, although little has been analyzed with regard to group seating, a related paper in the restaurant context is Bertsimas and Shioda (2003). They explore group seatings when reservations can be made and waiting-lines develop. The authors develop an Integer Program as an approximation to the Dynamic Program (DP) and show via simulation how it improves on traditional first-come-first-serve heuristics. Our paper shares some of their DP model formulation but overall is fundamentally different.

Group booking is a common theme in the hotel and airline RM literature (e.g. Vinod, 2013; **?**). Group booking means a large group (for example, a conference) requesting large amounts of rooms for a specific product (e.g. night stay or airline leg). Group booking in live entertainment is similar - large number of seats (> 10) for a specific night, but who do not necessarily have to be seated together. Elements of this problem are similar to ours. The non-linear benefits that accrue only when the entire request can be accepted for example. However, where our problem differs is that the capacity requested must be contiguous and in some sense linked. For traditional group booking, although it is an all or nothing proposition, there is no allocation requirement like this. In this and the following chapter we focus on the process of allocating and optimally offering individual requests (e.g. 1-6 seats).

An alternative approach that initially might seem appealing for our purpose, and that is common within customer choice based RM, is to analyze the problem in the context of Network RM, e.g. Zhang and Adelman (2009). The underlying premise in a Network RM model is that there are a set of products (or services) that make use of a set of shared and finite capacity-based resources. To model the network of which products use which resources, an *incidence matrix* is constructed. However, for live entertainment seating under the assumptions we provide the incidence matrix would simply be the identity matrix. Therefore, the main benefits of Network RM seem unlikely to translate or lead to any improved managerial insights over the proposed knapsack-based approach.

3.1.2 Technical and Managerial Contributions

Our paper contributes to the current literature by extending the DSKP to a setting where the knapsack is compartmentalized and includes an adaptive interaction with the items. The classic

DSKP becomes a special case of our general formulation. Allowing for stochastic placement *post-decision* is relevant for RM models for customer choice, as well as classic knapsack implementations in manufacturing and transportation.

Choosing to focus on the entertainment context here, we provide an optimal algorithm to control the allocation of seat offers to incoming customers and provide structural properties of the optimal decisions. We show conditions under which managerially friendly threshold policies can be shown to work, and the implications of these conditions. Importantly, we show that under reasonable conditions, there exists a threshold in time, such that managers will close a row early in a sales process, opening it closer to the event date.

We also detail how model parameters could efficiently be estimated from transactional box-office data or managerial judgment - or ideally from a combination of both. Finally, due to the computational complexity of the DP we also propose and analyze a heuristic based approach. Through simulation, we show that significant revenue gains are possible, and that the gains are driven by an unequal distribution of incoming groups.

Note that while we formulate the problem in discrete time, most of our formulation holds when generalized to a continuous time model. For ease of flow, unless specified otherwise, terms are used in a non-strict sense (e.g. decreasing instead of strictly decreasing). All proofs are provided in the Appendix.

3.2 Knapsack Model Formulation

We formulate a DSKP with adaptive interaction. The problem is *dynamic* in the sense that items arrive individually, without prior knowledge about the entire set of items, and that irrevocable decisions are made for each arrival. The problem is *stochastic* in that both the weights and rewards are random. The items' weight is realized and known upon arrival, but the reward is only realized after the decision. Finally, we introduce the notion of *adaptive interaction*. By *interaction*, we mean that there is a stochastic item placement that occurs post-decision. By *adaptive*, we mean that this stochastic process depends upon which compartments are opened. This problem's most important contribution to the literature is allowing for a process of random compartment placement that depends on the decisions of which compartment to "open". In previous DSKP literature the decision has been to accept or reject a given item with consequences being deterministic. Here, we consider a knapsack with multiple compartments, of different capacity, where the specific compartment placement and potential balk being stochastic.

We assume a stationary arrival process of items over T time periods, with sufficiently small intervals such that at most one item will arrive during each period. We index time backwards so that $t = T$ is the first time period and at $t = 0$ the event happens. Consider a knapsack with K finite compartments and where compartment $k \in \mathbb{K} \equiv \{1, 2, \ldots K\}$ has finite, discrete capacity κ_k. Let the remaining capacity in compartment k be c_k and the remaining capacity for all compartments in period t be the column vector $C_t = [c_1, c_2, \ldots, c_K]$. Thus, $C_T = [\kappa_1, \kappa_2, \ldots, \kappa_K]$. We denote a knapsack with no capacity remaining in any compartments as $C_O = [0, 0, \ldots, 0]$. Further consider an arriving item being of type $i \in \mathbb{I} \equiv \{0, 1, 2, \ldots, N\}$; with $N < \infty$. We define $i = 0$ to indicate no arrival. All items of type $i \geq 1$ have the same fixed, discrete and finite weight $w^i \in \mathbb{N}$, while $w^0 = 0$. The probability of item type i arriving in period t is stationary and denoted by $\alpha^i \in [0, 1]$, and defined such that $\sum_{i=0}^{N} \alpha^i = 1$. Extensions to consider, for instance, the common RM assumption of Poisson arrivals is straight forward. To simplify notation, we reserve the symbols i, k, t for item types, compartments, and time, respectively. Furthermore, throughout we use superscript for item types, and subscripts for time and compartments.

Upon a request being made and the knapsack decision maker observing item type i, they must decide which compartments to open. Denote the decision to open the k^{th} compartment for item type i by d_k^i. For $i = 1, \ldots, N$, $k \in \mathbb{K}$,

$$d_k^i = \begin{cases} 1 & \text{if the } k^{\text{th}} \text{ compartment is opened to item type } i, \\ 0 & \text{otherwise,} \end{cases} \tag{3.1}$$

while for $i = 0$, $k \in \mathbb{K}$, $d_k^i = 0$. Note again that the decision is which compartments to open up, as compared to an accept/reject decision. Specifically, it is possible for the knapsack decision maker to open multiple compartments. Let D_t^i denote the column vector of compartment decisions for item i in time period t, for $i \in \mathbb{I}, t = 1, 2 \ldots, T$,

$$D_t^i = [d_1^i, d_2^i, \ldots, d_K^i]. \tag{3.2}$$

The set of feasible decisions $\mathbb{D}(C, i)$ is defined as the set of all decisions D_t^i where the resource constraint holds. Note that the set of feasible decisions does not depend on t, i.e. for all C, i,

$$\mathbb{D}(C, i) \equiv \{[d_1^i, \ldots, d_K^i] : d_k^i w^i \leq c_k, \forall k\}. \tag{3.3}$$

Opening up multiple compartments implies there must be some underlying process by which an item is placed into one of the open compartments. Let $\gamma_k(D_t^i) \in [0, 1]$ denote the probability of item i "choosing" compartment k given a decision vector D_t^i. For instance, this is the probability that a customer will choose row k, given offer set D_t^i. We restrict $\gamma(\cdot)$ to functions that meet the following criteria: If $d_k^i = 0$ then $\gamma_k(D_t^i) = 0$, $\forall i, k, t$, and $\sum_{k=1}^{K} \gamma_k(D_t^i) \leq 1$, $\forall i, t$. The first criteria states that if a specific compartment is not open, there is zero probability that an item will be placed there. The second states that the total probability of being placed in any compartment must be less than or equal to one. If the total probability is less than one, it implies there exists a chance the item will not choose any compartment. That is, unique to our problem, there exists a chance that items balk.

Let $r_k(i)$ denote the real-valued reward *if* item type i enters the k^{th} compartment, $r_k(i)$: $\mathbb{K} \times \mathbb{I} \to \mathbb{R}^+$, $i = 1, 2, \ldots, N$, and $r_k(0) = 0$, $\forall k$. We restrict attention specifically to positive and time-stationary rewards. If a reward were negative then (of course) the knapsack decision maker would never open that compartment. The restriction to time-stationary rewards is mainly for mathematical tractability (and ease of notation). If an item balks then neither reward nor cost is incurred. Note that if the reward function $r_k(i)$ is independent of the compartment k then the model formulation can be considered as a dynamic, stochastic version of the multiple knapsack problem.

3.2.1 Optimality Equations

Given the above model formulation we formulate the DSKP with adaptive interaction as a stochastic dynamic program. The value-to-go for a given state (C, i) and period $t = 1, 2, \ldots T$, is given by,

$$V_t(C,i) = \max_{D_t^i \in \mathbb{D}(C,i)} \left\{ \sum_{k=1}^{K} \gamma_k(D_t^i)\Big(r_k(i) + \sum_{j=0}^{N} \alpha^j V_{t-1}(C_{[k]}, j)\Big) + \Big(1 - \sum_{k=1}^{K} \gamma_k(D_t^i)\Big) \sum_{j=0}^{N} \alpha^j V_{t-1}(C, j) \right\}, \tag{3.4}$$

with boundary conditions $V_0(C,i) = V_t(C_O,i) = 0$, $\forall t, i$, and where $C_{[k]} \equiv C - w^i \lambda_k$, with λ_k a column vector with an entry of 1 in the k^{th} row and zeroes elsewhere; i.e. $C_{[k]}$ represents the new capacity vector after subtracting w^i from compartment k. We denote the optimal decision vector for time period t as D_t^{i*}, and the optimal decision for compartment k as d_k^{i*}.

Note that $V_t(C,i)$ is the expected value over two sources of uncertainty: compartment choice (γ) and arrival type (α). In the first summation, we are adding the reward if a given compartment is chosen plus the expected value-to-go from the next time period, given the new capacity. Recall that the reward itself is realized post-decision, and that the decision only influences the probability of a given reward. The second term is the probability of the item choosing no compartment and moving into the next time period with the same capacity configuration (and no reward incurred).

To simplify the value-to-go, it is convenient to let $W_t(C)$ be the expected value-to-go over all item types, for $t = 1, \ldots, T$,

$$W_t(C) = \sum_{i=0}^{N} \alpha^i V_t(C,i) \tag{3.5}$$

This allows us to more succinctly rewrite the value-to-go as, for (C,i) and $t = 1, 2, \ldots T$,

$$V_t(C,i) = \max_{D_t^i \in \mathbb{D}(C,i)} \left\{ \sum_{k=1}^{K} \gamma_k(D_t^i)\Big(r_k(i) + W_{t-1}(C_{[k]})\Big) + \Big(1 - \sum_{k=1}^{K} \gamma_k(D_t^i)\Big) W_{t-1}(C) \right\}, \tag{3.6}$$

with boundary conditions $V_0(C,i) = V_t(C_O,i) = 0$, $\forall t, i$. Using this formulation allows us to characterize some properties of the value function we are dealing with. For asymptotic properties, similar to van Slyke and Young (2000), it is helpful to refer to the (optimal) value function of the unbounded (static) knapsack problem with an abundance of all item types. Let $S^*(C)$ refer to this (optimal) value for a given capacity C. We first summarize some of the basic properties of the value-to-go function.

Proposition 3.2.1 *The value-to-go function (3.6) has the following properties:*

1. $V_t(C,i)$ *is a positive finite function.*
2. $V_t(C,i)$ *is increasing in* t.
3. $V_t(C,i)$ *is increasing in* c_k, *for* $k = 1, 2, \ldots, K, t = 1, \ldots, T$.
4. $\lim_{T\to\infty} V_T(C,i) = S^*(C)$.

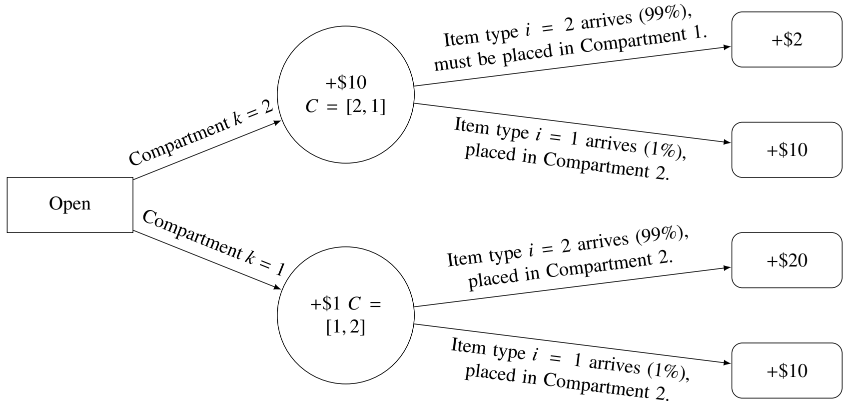

Figure 3.2: Decision Tree of Numerical Example

For the most part, these claims are straightforward and analogues to previous literature; e.g. van Slyke and Young (2000). To see the first three, you must only recognize that adding time or capacity only expands the feasible decisions. The last property comes about from the fact that the upper limit is if all capacity was utilized, and that as t approaches infinity, we would expect for all capacity to eventually be used. All of these properties trivially extend to $W_t(C)$.

Let us now note how this problem presents counter-examples to "common-sense" statements. For instance, we might initially think that it would be optimal to open the compartment with the highest reward. This is not always the case due to the combinatorial nature of the problem and the dynamic arrival of random item weights. If placing an item leaves a remaining capacity that is difficult to sell, it may make sense not to offer the compartment with the highest reward, or in some cases even reject the item entirely (by closing all compartments). Recall the issues illustrated by the seating chart in Figure 3.1.

To further elaborate on the trade-off and complexities involved we present a stylized example. Consider a two compartment knapsack, $\mathbb{K} = \{1, 2\}$, where an item will take the highest compartment offered. Suppose also that we only have two item types to consider, $\mathbb{I} = \{1, 2\}$, with $w^1 = 1, w^2 = 2$, and $\alpha^1 = 1\%, \alpha^2 = 99\%$. Thus item type two arrives almost all the time. Suppose that $r_1(i) = w^i, r_2(i) = 10w^i$, $i = 1, 2$. Let both rows have only two open units remaining, $c_1 = c_2 = 2$. For our current decision, imagine that we are in the second to last time period, $t = 2$, and that an item type $i = 1$ has arrived. The decision can be analyzed via a decision tree; see Figure 3.2.

Here working backwards, in the last time period $t = 1$ the trivial decision is always to open the available row with the highest reward available. If compartment 1 is opened in period $t = 2$, any incoming item in period $t = 1$ can be placed into compartment 2. On the other hand, if compartment 2 is opened in period $t = 2$, then there is insufficient capacity in compartment 2 should an item of type 2 arrive in period $t = 1$. Based on the decision tree, it is clear, for $t = 2$, $i = 1$, to only open row 1, $d_1^{1\star} = 1, d_2^{1\star} = 0$; $1 + .99 * 20 + .01 * 10 > 10 + .99 * 2 + .01 * 10$.

Placing the group in the highest valued row would have a high opportunity cost, since the only way to sell the other seat would be if a single person arrives. That leftover seat worth \$10 would go unsold 99% of the time. Instead, it is optimal to place the incoming customer in the lessor priced row, and save the two seats in the higher priced row for the next period.

This stylized example is meant only to illustrate the trade-off of higher immediate reward and worse capacity configurations. The problem gets more complex when stochastic placement is considered.

To capture these trade-offs between higher immediate reward and future state values, let B_k be the *marginal benefit* of an item i in period t being placed in compartment k, compared to not being placed at all. Although the B_k value implicitly depends on the current capacity remaining C_t and request type i, we drop this for notational convenience. Specifically, for a given t, k, i,

$$B_k \equiv r_k(i) - (W_{t-1}(C) - W_{t-1}(C_{[k]})). \tag{3.7}$$

This is similar to how van Slyke and Young (2000) structure their problem. The main difference is that their decision space is $\{accept, reject\}$. Thus their benefit function calculates the difference between accepting and rejecting an item. In our model, the marginal benefit is the difference between deterministically placing an item *in a particular compartment k* and rejecting the item completely. This is an important difference, since in our problem items retain autonomy over their placement after they have been "accepted". Since B_k is the marginal benefit of a compartment receiving an item, it can only be calculated over compartments for which there is sufficient capacity to place. Therefore, in any compartment k where $c_k < w^i$, we define $B_k \equiv 0$. Using this definition, and rearranging (3.6), allows the following result.

Proposition 3.2.2 *The value-to-go for a given (C,i) and period $t = 1, 2, \ldots T$, is given by,*

$$V_t(C, i) = \max_{D_t^i \in \mathbb{D}(C,i)} \left\{ \sum_{k=1}^{K} \gamma_k(D_t^i) B_k + W_{t-1}(C) \right\}, \tag{3.8}$$

with boundary conditions $V_0(C, i) = V_t(C_O, i) = 0$, $\forall t, i$.

Rewriting the value-to-go in this manner is helpful because the latter term is constant with respect to the decision and B_k is unaffected by the decision. If items had no autonomy post-decision, they would be placed in the compartment with the highest B_k. Instead though, the problem can be conceptualized as maximizing the weighted average of a series of values. Notice that the decisions of which compartments to open only influences $\gamma(\cdot)$.

Solving the value-to-go for the optimal decisions can be done through standard backward induction. Starting from the last period, the optimal decisions and corresponding values can be found for all state spaces. Saving these values and moving back a time period allows us to determine optimal decisions for previous time periods. Note that in the process of backward induction all B_k values will also be calculated. A downside of the proposed DP is that it exhibits the usual "curse of dimensionality" since the state space grows exponentially as compartments, item types, and periods are added.

3.2.2 Monotonicity Properties

It is often of interest to provide properties of the dynamic program with respect to underlying ordering and monotonicity. In order to do so we first define $\Delta W_t(C)$ to be the average marginal benefit of an extra period of arrivals to capacity configuration C at time t, i.e. $\Delta W_t(C) \equiv W_t(C) - W_{t-1}(C)$. Second, we define $\Delta W_t(C)$ to have a "time-invariant ordering" if between states the order is unaffected by the time remaining, i.e. for *all* t and *any* two different capacity configurations C and C', either: (i) $\Delta W_t(C') = \Delta W_t(C)$, (ii) $\Delta W_t(C') \geq \Delta W_t(C)$ (and at least one strict inequality), or (iii) $\Delta W_t(C') \leq \Delta W_t(C)$ (and at least one strict inequality). In other words, if at an "early" stage the gain over one period is larger for capacity configuration C' versus capacity configuration C, then this order will also hold at a "later" time period. Given this setup we can derive the following property.

Lemma 3.2.3 *If $\Delta W_t(C)$ has time-invariant ordering then B_k is decreasing in t, $\forall k, i, C$.*

A time-invariant ordering of $\Delta W_t(C)$ implies that if extending the sales horizon in one capacity configuration is preferred to another capacity configuration in period t, then the same holds true for any other time period. However, this is not necessarily true and in particular not in live entertainment. Consider cases where, asymptotically speaking, the value function of a capacity configuration is greater than another, but with few time periods remaining, may be less valued. For example, we might consider a capacity configuration that is very difficult to fill with one period to go, but that - given enough time - would have a higher expected reward. For instance, if a given row k had a very small $\gamma_k(\cdot)$ but a very high $r_k(i)$.

Nevertheless, this property has appealing implications as seen in Proposition 3.2.3. Specifically, it implies B_k to be decreasing in t. This is reasonable in that we might expect the marginal benefit to generally be decreasing in time periods remaining. With few periods to go, the benefit of placing an item compared to rejecting it should be relatively large, since there exists few remaining potential arrivals.

In addition to monotonicity regarding the value-to-go function, it is also of interest to provide structural properties regarding the optimal policy, i.e threshold-based rules on either time, item, reward, or capacity. Providing a structure on one of these variables enables more efficient computation and is easier to communicate to managers. In this case, providing a structure to the optimal policy is not obvious as the problem is still relatively undefined. Without a functional form on $\gamma(\cdot)$, it is not obvious how to maximize even the expected current reward, except by numerically searching all possible combinations.

For example, opening *only* the row with the highest B_k is not necessarily optimal, if there is a high chance the item balks. Similarly, opening *all* compartments with positive benefits is not optimal if it reduces the chance of a high B_k being chosen. Even more oddly, opening a compartment with $B_k < 0$ might be advantageous if it increases the probability of going into some other compartment with a very high marginal benefit. Finally, for certain functional forms of $\gamma(\cdot)$, it could be the case that even in the *last* period, it would not be optimal to open all available compartments. This would occur if an item becomes less likely to choose a higher reward compartment when a compartment with a lower reward is chosen, e.g. customers buying down.

Note too that B_k has not been defined as the marginal value from an optimal decision, and hence $\lim_{T\to\infty} B_k$ is not necessarily zero. If an item i placement in compartment k is not part of an optimal static solution, then $\lim_{T\to\infty} B_k < 0$. However, if the item placement is in the optimal static solution then $\lim_{T\to\infty} B_k = 0$; similar to the corollary in van Slyke and Young (2000). Additional factors that influence the asymptotic limit of B_k is the structure of $r_k(i)$. If, for a given k, $r_k(i)$ is structured such that rewards are proportional to the weights of all item types i, then B_k would converge to zero for any capacity C and item type i, e.g. rewards based on price multiplied by quantity purchased. However, it is not clear if this convergence occurs from above or below.

3.3 Entertainment Knapsack Problem

We now place the DSKP extension specifically into an entertainment RM context and illustrate how the problem could be implemented and solved. We consider the sales process of single ticket sales for a specific performance at a venue. That is, we assume that season ticket or large group sales (e.g. 10+) have previously been sold, and the available capacity adjusted accordingly. The compartments correspond to the venue's rows (or zones) and are numbered such that increases in the row imply rows closer to the stage, i.e. $k = K$ is the front row. We assume each row has an exogenously fixed price per seat p_k, and are such that $p_1 \leq p_2 \leq \cdots \leq p_K$. Requests for random group sizes arrive and are represented by the item type i, $i \in \mathbb{I}$. The weight of each group type w^i corresponds to the size of the group, i.e. $w^0 = 0, w^1 = 1, \ldots, w^N = N$. The reward of request type i being placed into the k^{th} row is therefore given by $r_k(i) = p_k w^i$. Note that in this context there is a fixed ratio of reward-to-weight that differs per row.

Recall that our challenge in providing structural results for this problem largely hinge on how general we have left $\gamma(\cdot)$. In the following sections, we discuss how numerical estimations of $\gamma(\cdot)$ could be performed using historical transactional data, as well as providing a model-theoretical basis for a specific function.

3.3.1 Transactional Data-Driven Estimation of Row Selection

Implementing our model means attaching specific values to the arrival probability α and row selection probability $\gamma(\cdot)$. First, let us focus on how to estimate α. Traditionally, sales at entertainment venues are made via multiple channels, e.g. online, phone, and in person at a box office. Each of these channels would certainly possess data on transaction times and the seats sold. However, recall that α^i is the arrival probability of a request and therefore using only sales data would underestimate the arrival probability since we would not record customers who arrive and do not purchase (either because of lack of capacity, or because they did not like the offered rows). The sales data thus censor demand.

However, censoring is only an issue via the physical channels. Using online sales transaction data to estimate α^i is relatively straightforward. Users going through the process of requesting seats are logged in the IT system and with IP addresses can be uniquely identified as having looked into purchasing, whether or not they eventually purchased. Incoming phone calls, if they are logged, could be done the same way. In person requests are more difficult as it is unlikely that data currently is recorded for no-purchase requests. To the extent that it is reasonable to assume different customers types use the different channels equally, box office and phone sales data could be scaled up by the same factor as seen in the online based transactions. Even if this assumption is not completely reasonable (online sales might be skewed toward the young), it provides an approximation.

Using online transaction data (that contain both successful sales and unsuccessful no-buy transactions), row selection probabilities across item types could be estimated by multinomial logistic regression,

$$\gamma_k(D_t) = \begin{cases} \frac{e^{\beta_k D_t}}{1+\sum_{m=1}^{K-1} e^{\beta_m D_t}} & \text{for } k = 0, 1, \ldots, K-1, \\ \frac{1}{1+\sum_{m=1}^{K-1} e^{\beta_m D_t}} & \text{for } k = K, \end{cases} \tag{3.9}$$

where $k = 0$ denotes the "no-buy" option. The β-coefficients are estimated through maximum likelihood. Here we assumed row selection do not depend on the request type i but only on which rows were offered, i.e. row preferences are identical across request types and no need for super-script i on D_t and d_k. However, extensions to estimate row selection probabilities across item types is straight forward. Note too, that additional covariates can easily be incorporated, e.g. customer demographics, holidays, etc.

In some situations, it may still be infeasible to produce data that includes no-buy transactions. In these cases equation (3.9) would represent the *conditional* probability of choosing row k given a purchase was made, e.g. $Pr\{\text{select row } k|\text{purchase}\} = (e^{\beta_k D_t})/(1 + \sum_{m=1}^{K-1} e^{\beta_m D_t})$.

In order to complement the analysis and produce *unconditional* row selection probabilities $\gamma_k(\cdot)$, purchase/no-purchase probabilities must somehow be elicited or estimated. One solution would be for the manager to manually estimate $Pr\{\text{purchase}|D_t\}$; with $Pr\{\text{no purchase}|D_t\} = 1 - Pr\{\text{purchase}|D_t\}$. Given the large permutations of possible D_t's, a tractable and reasonable approach is to assume that $Pr\{\text{purchase}|D_t\}$ only depends on the lowest (cheapest) row offered. If so, only K values would have to be estimated. The resulting row selection probabilities would then be given by, for $k \in \mathbb{K}$,

$$\gamma_k(D_t) = Pr\{\text{select row } k|\text{purchase}\} \times Pr\{\text{purchase}|D_t\}. \tag{3.10}$$

3.3.2 Customer Choice Model Based Row Selection

If transactional data is completely lacking, row selection probabilities could be derived using management judgment. Recall that $\gamma(\cdot)$ must return a value for all possible decision vectors D_t^i, $\forall i, t$. This results on the order of $\sum_{x=1}^{K} \binom{K}{x}$ values to estimate; an infeasible task for any reasonable value of K. We propose an alternative: asking managers for different customer segments' preference orderings and taking a weighted average.

Suppose that all requests belong to a customer segment $y \in \{1, 2, \ldots, Y\}$, that is distinct from the N groups defined above; e.g seniors, students, tourists, etc. Let $Pr\{\text{Customer is } y\}$ be the probability that a request is from customer segment y; $\sum_{y=1}^{Y} Pr\{\text{Customer is } y\} = 1$. Suppose also that each customer type has a strict preference ordering over the rows and "no-buy" option. This works since prices are exogenously defined. For example, "students" might have the following preference ordering for a two row theatre: Row 2 $\prec$ No Buy $\prec$ Row 1. Row 2 (the row closest to the stage) may be priced sufficiently high such that a student will never purchase in Row 2. Let $I_k^y(D_t^i)$ be an indicator function that returns 1 if a customer type y would purchase row k given the offered rows D_t^i, i.e. if row k is preferred to any other opened rows as well as the no-buy option. Given D_t^i and k,

$$\gamma_k(D_t^i) = \sum_{y=1}^{Y} Pr\{\text{Customer is } y\} I_k^y(D_t^i). \tag{3.11}$$

Given a specific customer segment, we would know with certainty which row they would choose given an offered set. However, in many cases it might not be possible to know which customer segment is making a request. So, we take the weighted average over all the customer segments.

The advantage of this approach is the relative ease at which a manager could estimate aggregate preferences for specific customer segments and then disaggregate the individual row selection probabilities. Splitting customers into different segments should be natural and already completed for marketing purposes. Estimating or using data to ascertain the relative sizes of these segments should also be straightforward. The last step involves asking for managerial judgment as to the preference ordering of the rows, for each customer segment. We believe it would be reasonable for a manager to be able to give preference orderings for each customer segment. In the case of managers being uncertain, finding preference orderings can be done in a less-direct fashion using pointed questions, e.g. "Would a student prefer to buy in the orchestra or mezzanine?"

Finally, we present a theoretical alternative that is based on a customer choice model. Consider that each request arrives with a stochastic *willingness-to-pay* (wtp) per seat X. Let the cumulative distribution function $F(x) = Pr\{X \leq x\}$, be defined on $[p_1, P]$; with $P > p_K$. The lower support p_1 is set such that all arrivals are willing to purchase a seat, while the upper support is such that there exists at least some proportion of customers willing to purchase front row seats. Given X and D_t^i, a group chooses the "best" row that is priced below their wtp, i.e. we assume all requests prefer rows closer to the event. Formally, an arriving item type i with realized wtp x chooses row,

$$\hat{k} = \max_k\{k | p_k \leq x, d_k^i = 1\}. \tag{3.12}$$

If no row k satisfies these conditions then the group balks and makes no purchase. Note that this process is not necessarily the same as maximizing consumer surplus, since we have not modeled consumers' wtp for different rows. Furthermore, we are assuming *myopic* customers that make a single "buy" or "leave" decision. The extension to consider *strategic* customer behavior requires further model assumption and left for future studies.

There exists asymmetric information as to the value of X and so row preferences are uncertain from the decision maker's point of view. Let u_k represent the probability a customer is willing to purchase row k (given the opportunity), then $u_k = F(p_{k-1}) - F(p_k)$, $1 \leq k < K$, and $u_K = 1 - F(p_K)$. Figure 3.3 illustrates how u_k can be represented on a graph of the cumulative distribution function. We can then define the probability of purchase in row k given decision vector D_t^i as,

$$\gamma_k(D_t^i) = d_k^i\left(u_k + \sum_{x=k+1}^{K} u_x \prod_{j=k+1}^{x} (1 - d_j^i)\right). \tag{3.13}$$

Defining $\gamma(\cdot)$ in this way produces a function with some nice and reasonable properties. First, the overall probability of a purchase in *any* row is affected only by the lowest priced row offered. In particular, if $d_1^i = 1$, then $\sum_{k=1}^{K} \gamma_k(D_t^i) = 1$, i.e. no balking. This assumption seems like a reasonable approximation to reality. In our experience, the vast majority of importance is assigned not to specific seat locations, but to the event itself. On the other hand, it is easily

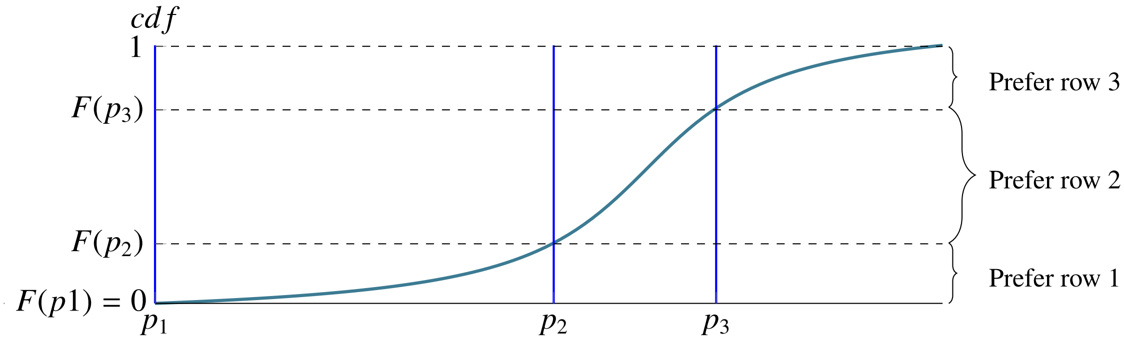

Figure 3.3: CDF of wtp and resulting row preferences.

conceivable that that same group, even if they were willing to pay for front row seats, would purchase back-section seats if front row tickets were not available. In other words, most of the utility value of attending an entertainment event is in the event itself, with a relatively small amount dependent on the specific seat.

Another nice property is that the probability of purchase in a specific row k is altered only by the decisions of rows *above* it, i.e. $\gamma_k(D_t^i)$ does not depend on $[d_1^i, d_2^i, \ldots, d_{k-1}^i]$. This is due to our assumption that customers choose the best row less than their wtp. If a row is already open, and a row beneath is opened, there is no fraction of customers who opt to buy down.

Before proceeding to optimal row allocation it is informative to reflect on why traditional RM approaches (that assume groups can be split up) are not suitable and results in trivial solutions. This is because if arriving groups were (deterministic) groups of size one, then the obvious solution would be to open all possible rows since customers would self-select into the row closest to their wtp - and no combinatorial effects would occur. The non-trivial part of the problem is in considering the combinatorial and stochastic effects of future arrivals' weights and rewards.

In general, it is still difficult to derive threshold policies on any single attribute of an incoming item or of the state space, i.e. on i, t, k, or C. However, the marginal benefit B_k can be utilized in deciding which rows to open. Specifically, for each period t and row k there is a threshold $\tilde{B}$ such that if $B_k \geq \tilde{B}$ then it is optimal to open up the row (and else close it). The advantage in doing so is that B_k incorporate state space information, as well as the incoming item information. This forms the basis for the proposed algorithm which we label as DSKPAI.

Algorithm DSKPAI Row Allocation for $t = 1, 2, \ldots, T$

Set $\tilde{B} = 0$.
Set $d_k^i = 0, \forall k$.
for $k \in \mathbb{K}$ **do**
  if $B_k \geq \tilde{B}$ and $c_k \geq w_i$ **then**,
    set $d_k^i = 1$,
    set $\tilde{B} = B_k$.
  end if
end for

DSKPAI starts by evaluating the lowest price row and moves "forward" to the most expensive row by continuously updating the threshold value $\tilde{B}$. Although the marginal benefit B_k only considers the value of (deterministically) placing an arriving request into row k versus not placing it at all, the proposed algorithm produces an optimal row allocation. In other words, and counter-intuitively, there is no need to consider the combinatorial effect of marginal values between rows. We formally summarize this result.

Theorem 3.3.1 *A solution $D_t^{i\star}$ found using* DSKPAI *is a globally optimal solution.*

To see why DSKPAI works, it helps to note that the probability of any purchase depends only on the lowest row opened up. Let k_{min} denote this lowest row. Given that k_{min} has been chosen and is the sole row opened, the expected benefit is $(1 - F(p_{k_{min}}))B_{k_{min}}$. Opening up a row $k > k_{min}$ is beneficial only if $B_k > B_{k_{min}}$. This logic holds as we move up by each row. Opening up the next row does not change the overall probability of purchase, but means that some requests which would choose row k_{min} now choose row k. However, the firm only wants this if the benefit of the k^{th} row exceeds the benefit from the k_{min} row. This procedure can be conceptualized as taking steps that increase the weighted average. Note that the solution from DSKPAI is not necessarily the *unique* optimal solution. For instance, if two rows have the same benefit, the knapsack decision maker is indifferent between opening and closing the higher row.

An immediate consequence of Thereom 3.3.1 is the following result.

Corollary 3.3.2 *In the final period it is optimal to open all rows, $D_1^{i\star} = [1, 1, \ldots, 1]$, $\forall i$.*

The intuition behind this result is that there is no longer any opportunity cost of selling a "bad" seating arrangement. In the last period, a firm should sell the highest priced seat that a customer is willing and self-selects to purchase. Note this was not the case for the general formulation in Section 3.2.

This idea can be extended somewhat into providing conditions on when a "naive" policy like this could be optimal for all time periods. The time-invariant ordering condition that we set out and defined before, provides this sufficient condition.

Proposition 3.3.3 *If $\Delta W_t(C)$ have a time-invariant ordering then it is optimal to always open all rows, $D_t^{i\star} = [1, 1, \ldots, 1]$, $\forall t, i$.*

The condition for $\Delta W_t(C)$ to have a time-invariant ordering turns out to be relatively strict. First, it requires conditions on all capacity configurations and in essence "smooths" out the combinatorial complexity - the crux of all knapsack problems. Second, it ensures the B_k to be positive, decreasing in time (Proposition 3.2.3), and increasing in k. While these attributes provides sufficient conditions for a naive policy to be optimal, it is unlikely they would hold across different arrival processes or willingness-to-pay distributions. On the other hand, it tells us that for any non-trivial policy, there must be time-invariance across at least some capacity configurations in $W_t(C)$.

Threshold Based Policies

An ancillary benefit of DSKPAI is that it enables us to investigate conditions regarding threshold-based policies in time, row, and row-capacity. However, despite the particular underlying row selection process which provides structure for $\gamma(\cdot)$, the problem is still quite general. Therefore, in order to ensure monotone optimal policies further conditions are necessary.

Time-Period Threshold - We first establish conditions on when a threshold in time exists. It is helpful to explore the behaviour of differences in the weighted average value of capacity configurations. For t, w^i, and rows l, k $(0 \leq l < k)$, we define $\delta_t^W(l,k)$ as the marginal value of weight w^i placed in row l over row k, $\delta_t^W(l,k) = W_t(C_{[l]}) - W_t(C_{[k]})$; with $\delta_t^W(0,k) \equiv W_t(C) - W_t(C_{[k]})$. Note that while $\delta_t^W(0,k) \geq 0$, $\forall t$, by Proposition 3.2.1 (iii), the same is not necessarily true for $\delta_t^W(l,k)$, $l > 0$, due to the combinatorial complexity of the knapsack problem. The problem of establishing time-based thresholds is to ensure how, over time $\delta_t^W(l,k)$ behave with respect to $w^i(p_k - p_l)$; where for $l = 0$, $p_l \equiv 0$. Sufficient conditions are summarized in the following result.

Proposition 3.3.4 *If for any C, i, and $0 \leq l < k$, either*
(A1) $\delta_t^W(l,k)$ is monotone in t;
(A2) $\delta_t^W(l,k)$ has a single local minimum on $[1,T]$;
(A3) $\delta_t^W(l,k)$ has a single local maximum on $[1,T]$, and $\delta_T^W(l,k) > w^i(p_k - p_l)$;
then there exists $\tilde{t} \in \{1, 2, \ldots, T\}$, such that for i and c_k,

$$d_k^{i*} = \begin{cases} 1 & \text{if } t \leq \tilde{t}, \\ 0 & \text{if } t > \tilde{t}. \end{cases} \tag{3.14}$$

Proposition 3.3.4 presents an appealing result. Firstly, it is intuitive. With many time periods to go, a decision maker can afford to be picky about which rows are opened, because there are likely to be many more arrivals in the future. With relatively few time periods though, a decision maker chooses to open the row to maximize their immediate reward. In particular, from Corollary 3.3.2, in the final period we will always open the row and hence at an extreme case $\tilde{t} = 1$. Second, although there is no closed-form solution for $\tilde{t}$, the result is helpful because it enables more efficient computation. If a threshold in time is known to occur, or assumed to occur, calculating all time periods becomes unnecessary. If, in the process of backward induction, it is realized a row is closed, then the result can be applied for all time periods prior (without explicitly calculating the value functions).

The technical intuition of the result is that conditions $A1$, $A2$, and $A3$ are sufficient to guarantee there is *at most* one time-period t for which $\delta_t^W(l,k) \leq w^i(p_k - p_l)$, and $\delta_{t+1}^W(l,k) > w^i(p_k - p_l)$ (i.e. possibly there is none and $\delta_t^W(l,k) \leq w^i(p_k - p_l)$, $\forall t$). Note that conditions $A2$ and $A3$ are less restrictive than concavity or convexity. In fact, concavity (covexity) is generally not sufficient for a time threshold, since it does not guarantee $\delta_t^W(l,k)$ has the "single-crossing" property with respect to $w^i(p_k - p_l)$. Observe too that condition $A1$ does not specify the monotonicity direction. A specific example of the required property is when, for $l = 0, 1, \ldots, k-1$, and for any $t \leq T$, $\delta_t^W(l,k) \leq w^i(p_k - p_l)$, in which case row k should always be opened. In this example, it is not sufficient that $\delta_t^W(l,k) \leq w^i(p_k - p_l)$, for a given $l < k$, since there may be a row l' for which $\delta_t^W(l,k) > w^i(p_k - p_l)$ (and hence row k should be closed).

Another important point is that the three conditions are stated for a fixed and finite T. Since $\lim_{T\to\infty} \delta_T^W(l,k) = w^i(p_k - p_l)$, we would not expect all three conditions to hold asymptotically, e.g. $\delta_t^W(l,k)$ cannot be monotonically decreasing in t as the sales horizon is extended. In our numerical analyses, we generally observe $\delta_t^W(l,k)$ to be monotonically increasing in t. On the other hand, in the cases when $\tilde{t} \neq 1$ or T, $\delta_t^W(l,k)$ generally behaves according to $A3$, with differences reaching some global maximum point and decreasing towards the limit $w^i(p_k - p_l)$ as T increases. See discussion below and Figures 3.4 (b), and (d).

Finally, note that we only need to assume these conditions on a limited subset of capacity configurations, and not for any given C and C'. This is important because in general, we can imagine situations where a general C and C' have multiple local extrema - when different rows have vastly different underlying selection probabilities and different prices. Here, we are considering only capacity configurations that differ at most in two rows.

Row-Index Threshold - The insight from DSKPAI also enables statements on when certain row-index threshold policies are optimal. Clearly these will require further conditions on the price-difference between rows. First though, let a row-index $\tilde{k}$ be defined by, for i, C, t,

$$\tilde{k} = \begin{cases} \min\{k|B_k \geq 0\} & \text{if } \exists B_k \geq 0 \\ K+1 & \text{if } \nexists B_k \geq 0. \end{cases} \tag{3.15}$$

We then have the following result.

Proposition 3.3.5 *For a given i, C, t, if* $(p_k - p_{k-1})w^i \geq W_{t-1}(C_{[k-1]}) - W_{t-1}(C_{[k]})$, $\forall k$, *then*

$$d_k^{i*} = \begin{cases} 1 & \text{if } k \geq \tilde{k}, \\ 0 & \text{if } k < \tilde{k}, \end{cases} \tag{3.16}$$

where $\tilde{k}$ *is the given by (3.15).*

Proposition 3.3.5 provides the condition under which a row-based threshold type policy is optimal. The condition implies that above the lowest row with a non-negative B_k, the remaining B_k are increasing. Notably, with this property, *if* $B_1 \geq 0$, then it is optimal to open all rows. Otherwise, there is a strict threshold with which it is optimal to open all available rows above the threshold.

Row-Capacity Threshold - It should be clear that closing a row only occurs when the opportunity cost of accepting a customer is relatively high. This can occur in two ways: if placing the customer means not being able to place larger groups of customers in the same row later, or if a decision-maker would prefer the group to go to some lower row. Therefore, in cases where a row *would be* filled entirely by an incoming group there is little opportunity cost. In other words, with groups self-selecting the highest-priced row and if all available capacity in the row would be filled, the optimal and managerially intuitive policy is to open the row, i.e. if $w^i = c_k$, then $d_k^\star = 1, \forall i, t$.

It may therefore be tempting to consider a row-capacity based threshold $\tilde{c_k}$ policy of the form, for i and t,

$$d_k^{i\star} = \begin{cases} 1 & if\ c_k \geq \tilde{c_k}, \\ 0 & if\ c_k < \tilde{c_k}. \end{cases}$$

This type of threshold would state that a row should be open if it has sufficiently large capacity remaining. However, this form of capacity threshold is unlikely to hold in general. For instance,

i	0	1	2	3	4	5
α^i	6/40	1/40	15/40	6/40	8/40	4/40

Table 3.1: Table of Numerical Illustration parameters.

from above we reasoned that a row should be opened when $w^i = c_k$, which implies that $\tilde{c_k} = w^i$, i.e. all rows with sufficient capacity should be opened. Intuitively, this is unlikely to hold as illustrated by the counter example in Figure 3.2. On the other hand, it would seem reasonable to expect that any optimal policy would always contain an "upper row-capacity threshold" $\tilde{c_k}$ such that, for i, if $c_k \geq \tilde{c_k}$, then $d_k^{i\star} = 1$; but for $c_k < \tilde{c_k}$, the optimal decision is not necessarily monotone.

3.3.3 Numerical Illustration

Next we illustrate how the algorithm works and provide further insight into why the problem can be difficult. We consider a two row venue, $K = 2$, along with a maximum group size $N = 5$, and $T = 26$ time-periods. To mirror the entertainment reality we set the arrival probabilities such that requests for single seats occur infrequently relative to most other request types; see Table 3.1. We normalize prices such that the second row is twice as expensive, $p_1 = 1, p_2 = 2$, and assume $u_1 = u_2 = 1/2$, which corresponds to a uniform wtp, $X \sim U[1, 3]$. Figure 3.4 illustrates some of the results. Note that the problem is discrete in time but that curves have been joined for presentation purposes.

First, in Figure 3.4 (a) we show how $W_t(C)$ changes with respect to t. This figure helps to illustrate the structural properties in Proposition 3.2.1. Note how the value-to-go is increasing in t and in overall capacity. Also see how at some point all of the curves flatten. This is when they approach the optimal static knapsack value $S^*(C) = p_1c_1 + p_2c_2$. The last feature of interest is that with relatively few time periods remaining, the difference between high capacity value functions is minimal. Intuitively, if there are few periods remaining, differences between high capacity states should be negligible since incoming requests could fill neither; i.e. past a certain point there are decreasing benefits to increased capacity.

In Figure 3.4 (b) we illustrate some properties and complexities of the problem by graphing the benefit of placing an arrival item type $i = 3$ into either of the rows of a venue with $C = [4, 4]$. Recall how DSKPAI works: A row should be opened if it's benefit is greater than the best option below it. Here,when $t \leq 5$, $d_1^{3\star} = 1$ and $d_2^{3\star} = 1$, since the benefits are strictly ordered. This is also an example of when Proposition 3.3.5 is valid and when a time-based threshold policy exists over the rows. When $t \geq 6$, it is no longer optimal to open row 1 since its benefit is now negative. However, B_2 is still positive for $t = 6, 7$, and so the optimal decision is to (solely) open the second row then, while for $t \geq 8$, $B_2 < 0$ and so both rows exhibit negative benefits and both should be closed. Intuitively, this can be explained as the cost of leaving one seat open. Since we assumed the available capacity is 4 for both rows, accepting an item $i = 3$ means leaving only one seat available that may be difficult to sell. Later in the sales process, it becomes preferable to take the reward rather than rejecting the request. However, earlier in the sales process, it is preferable to wait for a request that will either fill up the row, or that will leave a more desirable number of seats remaining. Interestingly, the higher priced row is

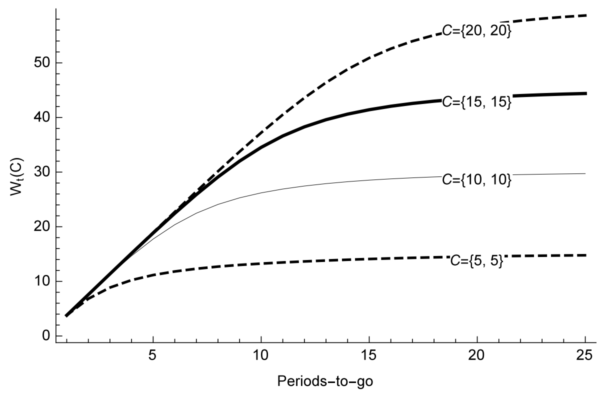

(a) $W_t(C)$ for various capacity configurations and periods-to-go.

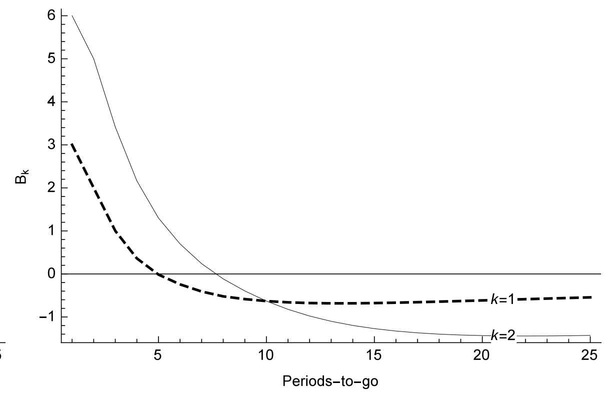

(b) The benefit of placing item type $i = 3$ in either row with $C = [4, 4]$ with different periods-to-go.

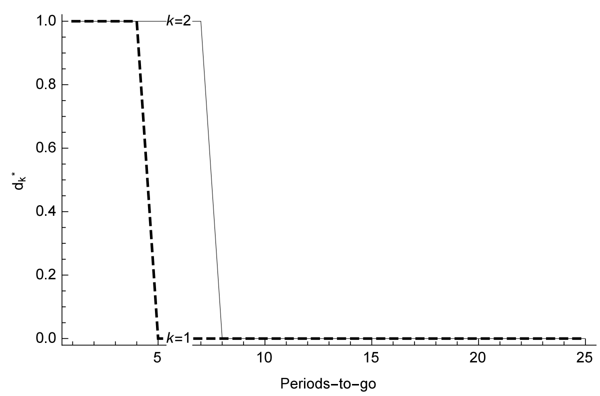

(c) The optimal decisions graphed as a function of periods-to-go.

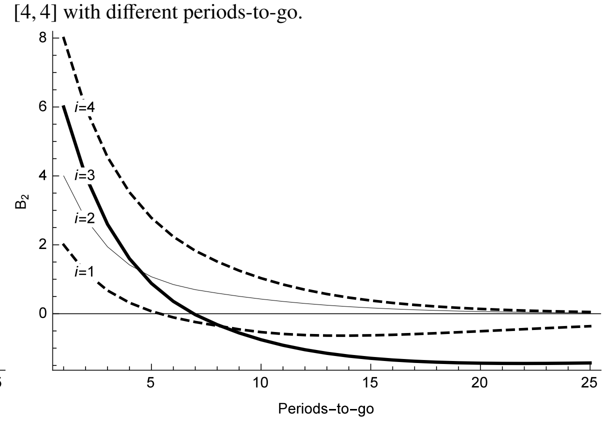

(d) The benefit of placing different group sizes in row $k = 2$ with $C = [0, 4]$ at different periods-to-go.

Figure 3.4: results from Numerical Illustration.

opened prior to the lessor priced row. The corresponding optimal decisions for $i = 3$ are shown in Figure 3.4 (c).

Finally in Figure 3.4 (d) we examine the benefits of placing different item types in a specific row over time. We assume a capacity configuration of $C = [0, 4]$ so that we can focus on only one row. The various lines correspond to the benefit of placing different group sizes. Since there is only one row remaining, the decision maker must only verify if $B_k \geq 0$, to decide whether to open the row to a given item type. As can be seen, for the most part the functions remain above zero, and so the row should be opened. Again, the complexity in attempting to use this information is that even for simple examples, the benefits do not retain order over time. Instead, it depends on the time remaining. This information allows us to answer which group sizes would be preferred at a given t. In the last period it is obvious that you would like to sell as much as possible and the graph shows this, with the sizes being strictly ordered. However, in earlier periods this is not necessarily the case. Specifically we see that the benefits for group sizes of 2 and 3 are not always ordered.

3.4 Single-Row DSKP Heuristic

As mentioned earlier, similar to most Dynamic Programs the proposed model formulation suffers from the "curse of dimensionality". Even when values are stored, the initial calculation involves calculating large amount of future states. The explosion in state space is primarily driven through the number of rows in the problem, because each row adds capacity configurations to evaluate. Thus, to solve the problem efficiently, a heuristic or approximation is likely necessary for decision makers.

3.4.1 SINGLE-DSKP Heuristic Algorithm

To make the problem tractable for realistically sized problems, we propose that on a row-by-row basis (assuming sufficient capacity), the single-row DSKP *without* adaptive interaction is solved. That is, the DSKP with a single compartment and no option for an item to balk; this is similar to Papastavrou et al. (1996). In the single-row DSKP, we solve the DP with the same w^i, t, and row-capacity c_k. Price is a simple scaling factor and so we can set it to 1 without loss of generality; c.f. Papastavrou et al. (1996). Once the single-row DSKP is solved and an optimal decision for the row has been determined, we use this decision as a heuristic for the row in a multi-compartment DSKP *with* adaptive interaction. The reasoning behind the heuristic is as follows. If it is not optimal to open a row when no other rows exist, it is unlikely it would be optimal to open it in a multi-row setting. In a multi-row setting, the opportunity cost of closing a row is diminished somewhat, because there is a chance the customer could choose a different row. We name this heuristic algorithm SINGLE-DSKP.

Numerically, SINGLE-DSKP is more efficient because each single-row problem only scales with the individual c_k, and not the number of rows K. In reality, even for large venues, the capacity of a single row is bounded and relatively small. While a large sports arena might have thousands of seats, each individual row is likely less than 30 seats wide, for example. The single compartment problem thus remains the same size, even with increases in K. In addition, once solved, finding the decision at a smaller c_k becomes simple, since smaller versions would

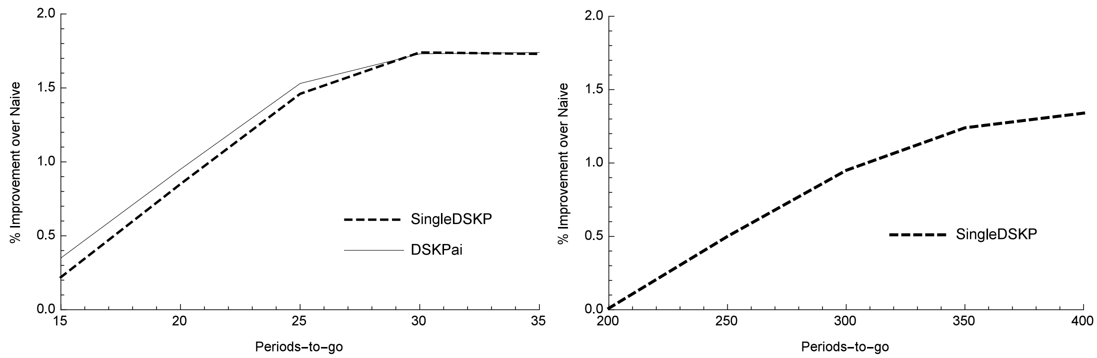

(a) Study 1 percent improvement over naive policy for a 2 Row theatre with different time periods remaining. Grey line indicates optimal policy. Dashed line indicates heuristic. Mean from 2000 simulations.

(b) Study 2 percent improvement over naive policy for a 40 Row theatre with different time periods remaining. Mean from 2000 simulations.

Figure 3.5: Summary of Simulation Study.

have been solved during the backwards induction process. So, once solved to completion, the optimal decisions can be found via a lookup table.

3.4.2 Simulation Study

We provide two simulation studies to test how well the SINGLE-DSKP heuristic performs compared to the optimal DSKPAI algorithm. We compare three row allocation policies: a “naive” policy of opening all feasible rows, the SINGLE-DSKP heuristic, and the optimal policy based on DSKPAI. In each study we vary the amount of time periods as a proxy for different demand levels, and run 2000 trials based on common arrival probabilities (given by Table 3.1) and wtp uniformly distributed between 1 and 3.

While we use a simulation to evaluate heuristic performance, an alternative would be to generate upper-bounds on the optimal result and compare heuristic performance to these upper bounds. One potential upper bound on the objective function is to decompose groups into single arrivals. In this case, the upper bound on the value function would be the minimum of expected total future arrivals and capacity.

In Study 1, we consider a $K = 2$ configuration with 15 seats per row, prices $p_1 = 1, p_2 = 2$, and time periods $t = 15, 20, 25, 30, 35$. In Study 2, we consider a $K = 40$ configuration with 15 seats per row, prices $p_1 = \cdots = p_{20} = 1$, and $p_{21} = \cdots = p_{40} = 2$ (i.e. customers are equally likely to prefer the first and last 20 rows), and time periods remaining of $t = 200, 250, 300, 350, 400$. Note that although there are only two price-points, Study 2 is fundamentally different (and more complex) than a problem with two rows of 300 seats. Within the same price-point, the decisions in each row are still unique and can be different depending on the remaining row-capacity. For Study 2 the curse of dimensionality meant that we were unable to provide exact optimal results. Figure 3.5 summarizes the results by showing the net gain of DSKPAI and SINGLE-DSKP over the naive policy. In each graph, the x-axis represents the periods-to-go, and the y-axis the percentage net gain over the naive policy. The grey lines indicate DSKPAI and the dashed line indicates SINGLE-DSKP.

There are a couple of prominent and interesting results that stand out. First, in both studies we observe a small but non-trivial revenue gain of about 1%. Second, for Study 1 (in panel a), we see that the heuristic policy appears to get close to the actual optimal results. This gives us hope that the same is true for larger capacity configurations. Third, as one might expect, at low demand levels (low t) there is little gain from any of our alternative policies. With low demand, there is excess capacity such that almost all requests can be accepted, and hence the "naive" policy works quite well. It would perhaps be natural to assume that at high demand levels (high t) there would be excess demand to fill any gaps, such that the "naive" policy would perform well. However, this is clearly not the case. Across both studies, the highest demand setting results in highest levels of improvement for the optimal/heuristic policies. The reason being, at high demand levels existing gaps may be difficult to sell because of how rare certain group sizes are. Offering the rows in an optimal fashion, however, makes it more likely that gaps are filled, i.e. that the venue does not end up in situations like Figure 3.1.

The core of the Knapsack problem (both static and dynamic), is of course the combinatorial effect of the items' weights. Consequently, the main reason for the observed revenue gains might stem from the assumption that single seat requests were rare. To test this hypothesis, we reproduced the two studies but assumed equal arrival probabilities, i.e. we set $\alpha_i = 1/6$, $\forall i \in \mathbb{I}$. The results were indeed that there now was essentially no gain from DSKPAI and SINGLE-DSKP over a naive policy. This is because there are very few "bad" state spaces. Unlike the original α settings (Table 3.1), single seats are now "easy" to sell since requests for them are relatively common.

3.5 Conclusion

In this paper we have extended the DSKP to bridge the problem with a burgeoning literature of customer choice models in RM. The novel extensions allow incoming items to select their own knapsack compartment, and to have rewards vary based on their choice. Furthermore, the choice probabilities were assumed to depend on the decision maker's offered set of compartments. We provided the model formulation and discussed the complexities inherent in solving it in full generality and why control policies are unlikely to exist. Motivated by live entertainment and customer choice we then discussed how the proposed model could be implemented and applied for RM purposes. We provided an optimal row-allocation algorithm and conditions under which various threshold-type policies would apply. Due to the inherent "curse of dimensionality" we also provided a heuristic that achieves close to optimal results.

Live entertainment is a multi-billion dollar industry and possess many of the key RM attributes. Despite this there has been limited model development designed specifically for the entertainment industry - both research and industry application. One reason is the mathematical complexity of dealing with (small) group requests. Traditional RM models, motivated by air transport, hotel and car rental, tend to rely on disaggregating the small group request into multiple individual requests. Furthermore, in traditional RM industries, management can usually control or adjust the allocation of customers to specific inventory (capacity) *after* the purchase, e.g. move a customer from one seat in the plane to another or offering a different category of car. In live entertainment, such as theater, this is rarely the case, and hence ensuring optimal dynamic allocation policies as requests arrive are of bigger importance. Finally, there is a simple legacy in live entertainment of dichotomizing bookings between (large) group-sales and individual ticket sales.

From our analysis we showed how, with some reasonable assumptions, a time-threshold exists in deciding when to open or close rows. Specifically, we established sufficient conditions such that closer to the event, the manager becomes less restrictive and offers more choices. Using this knowledge, management can closer to the event, and without "much" loss of revenue, implement the simple *open-all* policy. This can be advantageous in certain situations, such as on the day of the event, if the box-office is busy and implementing dynamic allocation policies is difficult. Prior to the event day however, management may be more strategic in their choice of which rows to open.

Focusing on the individual ticket sales part of the business, we show that small group size requests should not be assumed away or ignored. Contrariwise, our analysis indicates significant revenue enhancements are possible when incorporating the dynamics of group requests. A key driver of the potential gain is the distribution of the arriving group sizes. The more unequal the distribution of incoming sizes, the higher revenue enhancement that can be achieved and consequently the more attention that should be paid. Note that high-demand events are not exempt from this problem. Revenue enhancements are possible even when high-demand is present because spaces in the capacity configuration persist until exact group sizes arrive. This last feature perfectly embodies the crux of the knapsack problem.

Although we have kept prices exogenous, the proposed model can be used to derive insight regarding operational pricing issues, e.g. through shadow prices. A motivating example is online ticketing where entirely closing a row may not be possible. Instead, it may make sense to implement dynamic pricing to make the venue indifferent between which rows a customer chooses. This would entail setting prices such that the marginal benefit are equal between different rows. Another possible application would be to offer a *post-purchase* discount or buy-up offer for customer to move from one row to another. While there may initially be some behavioural scepticism or even backlash to these proposed applications, they have been accepted and are now well-established in traditional RM industries like air transportation, hotel, casino and car rental.

Chapter 4

A Simulation Analysis of Group Seating Heuristics

4.1 Introduction

In the entertainment industry, it is rare to consume a service individually. Individuals attend theatrical performances, sporting events, and concerts in groups. These types of businesses, which we will we refer to as entertainment venues, serve as experiential goods. Experiential goods are a type of good in which the utility gained is through the experience of consuming the good. The experience itself is heightened through through experiencing it socially.

The operational problem this creates is that discrete groups of customers purchase tickets assuming they will be seated contiguously. Without being seated in groups, the experience itself is seen as significantly less valuable, and the consumer is less likely to buy. This is corroborated by how ticket buying occurs in reality. On the major entertainment ticket buying websites (e.g. Ticketmaster, Telecharge, Stubhub), the number of tickets the customer chooses to buy will automatically be arranged next to each other.

Assuming discrete groups sit together is atypical from typical Revenue Management (RM) settings. In a commerical airline environment, consumers do not require that they are seated together and tickets are often purchased prior to knowing seat locations. Indeed, seat locations are relatively fluid in airlines because many airlines do not assign seats until check-in, leaving them flexibility with how seats are assigned. In contrast, specific seats are reserved upon buying for entertainment events.

This presents a problem. When groups must be seated together, demand cannot necessarily be seated up until capacity. Depending on where groups are currently seated, there may be small pockets of available seats that can only be sold to small groups. Without demand from groups of that size or smaller, the empty pockets will go unsold, regardless of how many seats are deanded overall.

The problem is made more complicated by the fact that customers still make the final decision of where to sit. Venue managers can limit their options however by not allowing certain seat assignments. To the best of our knowledge, most entertainment venues do not use sophisticated analytical techniques in deciding which seats to offer. The venue simply offers every available seat to a prospective customer.

In this paper, we seek to investigate the effect that group seating has on the effective capacity of a venue, and under what conditions. Because an exact approach is not tractable, we propose three heuristics. We explore revenue enhancements under different assumptions on customer behavior, group distribution, and demand settings.

4.2 Literature Review

One reason that group requests have received little attention is that they present challenging combinatorial complexities and often "ruin" the structural properties of existing RM models (e.g. concavity and monotonicity). An early paper that noted this issue is Lee and Hersh (1993). They showed that the expected marginal value of a seat no longer has to be monotone when group requests are considered. Talluri and van Ryzin (2004) noted similar issues - that the structure of optimal allocation policies is "profoundly" impacted by group requests. Few studies have explicitly modeled the dynamics of group bookings.

One example, although motivated by restaurants, is Bertsimas and Shioda (2003). Like entertainment venues, restaurants must consider discrete groups of stochastically arriving consumers, who also must be seated together. Bertsimas and Shioda (2003) develop multiple classes of models to assist restaurant managers in deciding when/where to place restaurant groups. The authors find that significant revenue gains were possible over the simple First-Come First-Serve approach.

The restaurant context is similar to our problem in that if small groups are placed in larger tables, they take up space for the larger groups that should produce higher revenue. Therefore group size and placement matters. However, unlike in the restaurant problem, an entertainment venue manager cannot create a queue and must decide dynamically where to offer seats. Another important difference is that entertainment venues exhibit customer choice. The most important difference however is the nature of the inventory. In a restaurant, tables are (relatively) fixed and can sit a certain number of people, and thus can accept any group less than that amount. Once it is taken, the unused seats are unavailable. In an entertainment venue with rows of seats, there are an enormous amount of possible configuration of groups seated. It is not trivial to view the seats as discrete inventory like in restaurants.

It is worth highlighting two classic operations research problems that our problem may appear similar to: The bin-packing problem, and the knapsack problem (Dantzig, 1957). The bin-packing problem is described as the following. Given bins of a certain size, and items of different sizes, how can we place items into bins such that we use the least number of bins? The bin-packing problem differs from our problem in that we need not fit all groups like the bin packing problem assumes. In fact, the implicit decision is whether it is worth "rejecting" groups sometimes. In addition, the bin packing problem assumes contiguous placement and no item autonomy.

The knapsack problem suffers from the same assumptions. It assumes contiguous placement of items, and it assumes deterministic placement. In Chapter three we relax the deterministic placement, by giving items autonomy, which we called "adaptive interaction". However, we still assumed contiguous placement. In this paper, we relax both of those assumptions. Due to the complexity of the problem, we take a simulation based approach to evaluate computationally efficient heursitics.

4.3 Venue Seat-Selection Process

We consider an entertainment venue with a rectangular seating map. We denote S as the seating map matrix with dimensions R by C with individual elements denoted by $s_{r,c}$. We define $s_{r,c}$ to be one if the seat at row r, column c is filled or otherwise unavailable[1]. For completeness, we also define $s_{r,c}$ to be one if the row/column index are outside of the seating map. Let $S_{r,c}(i)$ refer to the i seats starting at row r, column c and extending to row r, column $c+i-1$. $S_{r,c}(i)$ is defined as zero if the entire set of seats is available, and one otherwise. Formally, for $r = 1, \ldots, R$ and $c = 1, \ldots, C - i + 1$

$$S_{r,c}(i) = \begin{cases} 1, & \text{if } \sum_{n=c}^{c+i-1} s_{r,n} \geq 1, \\ 0, & \text{otherwise.} \end{cases} \tag{4.1}$$

[1]For instance, if the venue is not completely rectangular, these variables could be set accordingly to accomodate that situation.

For $r = 1, \ldots, R$ and $c = C - i + 1, \ldots, C$, $S_{r,c}(i) \equiv 1$ as it extends "off the map". We let $p > 0$ be the price per seat, and assume it is homogenous across seats and customer type. That is, we assume no discounts exist for large groups. Our model focuses on maximizing the quantity of customers able to be seated, rather than engaging in pricing decisions. Thus, this model may be considered as a model of a particular section of a venue. For example, we could separately model the Orchestra section from the Mezzanine section of seats at a theatre, where prices within each section are very similar, even though between them they are different. Without loss of generality, we can let $p = 1$.

The sales horizon is discretized into T periods, with each period sufficiently short such that at most one customer arrives. We then number the periods in reverse order, with $t = 1$ representing the last period before the event starts. Customers of group size $i = 1, \ldots, I; I \leq C$ arrive and request i consecutive seats in the venue. The (stationary) probability of group size i arriving is denoted by α_i; $\sum_{i=0}^{I} \alpha_i = 1$. α_0 is reserved for the probability of no arrival.

At each time period, the venue must decide which seats to offer the incoming request. Let $x_{r,c}(i)$ be the decision to offer the i consecutive seats to the right of (and including) seat [r,c]. $x_{r,c}(i) = 1$ if the seat-set (from [r,c] to [r,c+i-1]) is opened, and zero otherwise.For a given request size of i, the entire decision vector is denoted by X_i and given by $X_i = [x_{1,1}(i), x_{1,2}(i), \ldots, x_{R,C}(i)]$. A resource constraint exists on the decision. A seat-set can only be opened if all seats within it are available, i.e. if $S_{r,c}(i) = 0$.

If no seat-sets for group size i are offered (i.e. if $X_(i) = [0, \ldots, 0]$), the customer does not purchase. If at least one seat-set is offered, we assume that customers choose probabilistically according to the following process. Let $u_{r,c}$ be the utility of seat [r,c]. We assume that the front center point $[R, C/2]$ has a utility value of one and that there is an exponential decay function based on the euclidean distance from this point.[2] The utility of a specific seat is then given by

$$u_{r,c} \equiv e^{-\beta\left(\sqrt{((r-1)^2+(c-C/2)^2)}\right)}. \tag{4.2}$$

$\beta > 0$ is an exogenous parameter that determines how utilities scale across the venue. As β increases, utility differences between seats increase and front-center seat-sets become more likely to be chosen. We provide an example heatmap of individual seat utilities when $\beta = 1$ in Figure 4.1. The front of the stage is toward the bottom of the figure. As can be seen, the utility values nicely form a decreasing ring around the front center point.

Note that in the special case of $\beta = 0$, all seats have a common utility of 1. This would represent an assumption that customers choose from seat-set offers randomly with uniform probability. Although an aggressive assumption, this provides a worst-case scenario from the point of view of the venue. With customer's choosing completely randomly, many small pockets are likely to arise.

Let $U_{r,c}(i)$ be the utility value of seat-set $S_{r,c}(i)$. We assume that utilities of seat-sets are given by the sum of the individual seat utilities. Then, $U_{r,c}(i) = \sum_{n=c}^{c+i-1} u_{r,n}$. We then assume that customers choose from offered seat-sets probabilistically according to proportion of utility made up in the offer sets. The probability that a customer of group size i chooses a specific seat

[2]Note that in the case of an even valued C, the front-center point would not correspond to a specific seat. This does not pose a problem for our analysis.

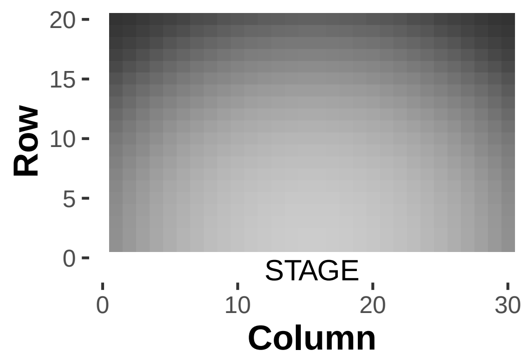

Figure 4.1: A 20 row, 30 column venue with heat map of $u_{r,c}$ values when $\beta = 1$. Front of the stage is the bottom of the figure. $s_{1,15}$ is the best seat in the venue.

set is given by, for $X_i \neq [0, \ldots, 0]$,

$$Pr\{S_{r,c}(i) \text{ chosen} \mid X\} = \frac{x_{r,c}(i) U_{r,c}(i)}{\sum_{m=1}^{R} \sum_{n=1}^{C} x_{m,n}(i) U_{m,n}(i)}. \tag{4.3}$$

This expression is a well defined probability distribution. The probability of a given $S_{r,c}(i)$ being chosen is less than or equal to one, because the denominator includes the numerator. In addition, the sum of all probabilities over the entire outcome space is equal to one.

To explore how β influences customer selection, we simulate 100 arrivals of groups of size $i = 1$. Using an "offer-all" policy, we plot how the map looks after a single run of 100 arrivals. These plots are shown in figure 4.2. A black spot indicates the seat is taken, while a grey spot means the seat is available. For $\beta = 0$ and $\beta = 1$, the seating map looks roughly uniform. However, there is a distinct grouping of requests for $\beta = 5$ towards the front/center. This would seem to replicate an assumption that customers prefer those areas and select seats accordingly.

We choose to model seat choice probabilistically to capture heterogenuous customer preferences. The Mirvish organization, a Toronto theatre producer, notes on their FAQ that

> *[M]any people insist on sitting in the first two or three rows - wouldn't dream of sitting anywhere else - while others find the first rows too close for comfort (or for a full view of the action); some people love the box seats for their roominess and their aura of exclusivity, while others say they don't like the box seats' angled view of the stage*" [3].

This problem can be formulated as a Markov Decision Process (MDP) and solved exactly via backwards induction. This approach however suffers from the so-called "curse of dimensionality". A dynamic program would have to explicitly consider the probability of a consumer choosing to sit in a specific spot, for all future stages. For commercial organizations with hundreds or thousands of seats, this is not tractable, particularly in the times needed to make online set decisions.

Similarly, a Deterministic Linear Program could be developed. Taking *expected* future demand for each group type, the optimization problem would allocate seat-sets to expected

[3]Retrieved from https://www.mirvish.com/ticket-info/faq

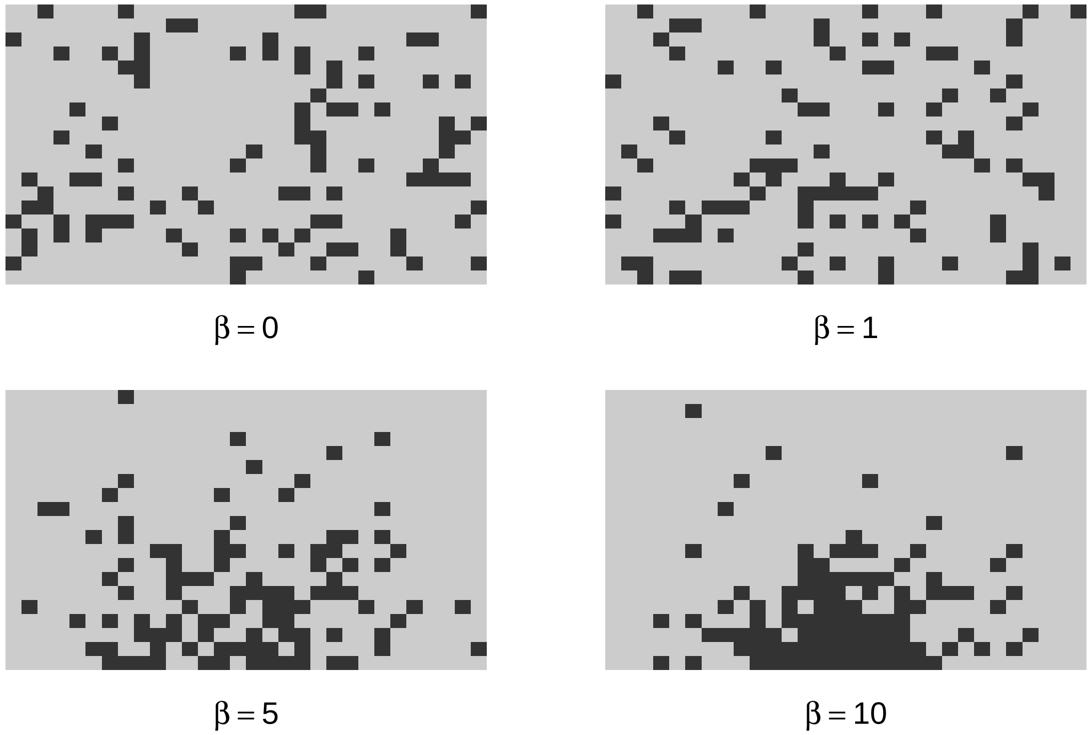

Figure 4.2: A 20 row, 30 column venue after 100 arrivals of group type $i = 1$, with differing β parameter. Dark indicates the seat is filled, and grey indicates the seat is available.

demands across the seating map. In the Appendix we provide a full formulation of this optimization. Nevertheless, even for a modest sized theatre, the optimization problem is quite extensives. For a 600 seat theatre, one constraint matrix has 18,000,000 elements.

Since both of these alternatives potentially are not able to run in an online fashion, we explore alternative algorithms based on heuristics. All algorithms we develop are compared to an“offer-all” method, which we call the NAIVE method. The NAIVE method is, for $r = 1, \ldots, R$ and $c = 1, \ldots, C - i + 1$,

$$x_{r,c}(i) = \begin{cases} 1, & \text{if } S_{r,c}(i) = 0 \\ 0, & \text{otherwise.} \end{cases} \tag{4.4}$$

4.4 Seat Offering Heuristics

No-Single-Seat Heuristic

The heuristic currently implemented in the Ticketmaster platform restricts choices to seat-sets that, if chosen, would not leave a single seat open. For example, in an empty-segment of four seats, an $i = 2$ request would not be able to choose the middle two. The customer could select two seats at either end of the empty-segment.

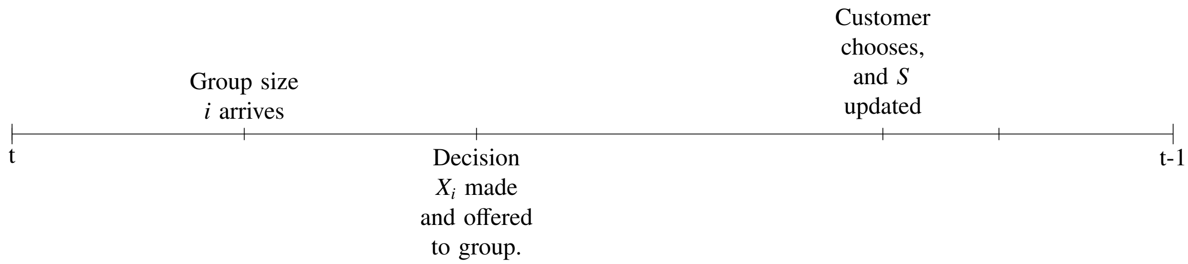

Figure 4.3: Sequence of Events

We test two variants of the algorithm. In the SINGLEa heuristic, seat-sets that would leave a single seat open are strictly not offered. SINGLEa is defined as the following; for $r = 1, \ldots, R$ and $c = 1, \ldots, C - i + 1$,

$$x_{r,c}(i) = \begin{cases} 1, & \text{if } S_{r,c}(i) = 0 \wedge (S_{r,c-2}(2) = 0 \vee s_{r,c-1} = 1) \wedge (S_{r,c+1}(2) = 0 \vee s_{r,c+1} = 1) \\ 0, & \text{otherwise.} \end{cases} \tag{4.5}$$

In the SINGLEb heuristic, seat-sets that would leave a single seat open are not permitted, *unless* it would force the customer to be rejected. The same definition as SINGLEa is used, but if $X_i = [0, \ldots, 0]$, then the NAIVE heuristic is used. In other words, SINGLEb will permit a single seat remaining if it is the only option. It is not clear a priori which variant is preferred and so we include both in our analysis. SINGLEa may be too selective and reject customer unnecessarily while SINGLEb may open seat-sets unnecessarily.

Greedy Heuristic

The third heurstic we are going to test is a GREEDY heuristic. At a high level, the heuristic works as follows. At each time period, the seating map is "optimally" allocated for *expected* future demands of each group size. A group request for i' seats can then be offered the seats already allocated for its group size. We provide details of the allocation process below.

Let Y_i be the expected number of arrivals of group size i from future periods. In period t, Y_i, for $i = 1, \ldots, I$ is given by

$$Y_i = \begin{cases} (t-1) \times \alpha_i + 1, & \text{if } i' = i \\ (t-1) \times \alpha_i, & \text{otherwise.} \end{cases} \tag{4.6}$$

Note that Y_i depends on t but we suppress this notation for convenience.

Let E be the set of empty-segments. An empty-segment is defined as a seat-set where all seats are empty, and in which both sides are bordered by filled seats (or the end of the row). Therefore E is given by

$$E = \left\{ S_{r,c}(i) \mid S_{r,c}(i) = 0 \wedge s_{r,c-1} = 1 \wedge s_{r,c+1} = 1 \right\}. \tag{4.7}$$

We index the set E by e. Let e_r and e_c refer to the row and column index, respectively, of given empty-segment e. Let $k(e)$ be the length of empty-segment e, e.g. the number of seats. Finally, for the allocation, we let $d_{r,c}(i) = 1$ represent the allocation of i seats to group size i staring from row r, column c, and zero otherwise.

The initial allocation algorithm works as follows. The largest group size j that still has unfulfilled expected future demand (e.g. $Y_j > 0$) and which fits (e.g. $j \leq k(e)$) is “placed” to the furthest left and Y_j decreased by one. The current empty-segment length is then updated and this process is repeated until either the empty set is fully allocated (i.e. $k(e) = 0$) or there is no group j which satisfies the requirements. This process repeats for all empty-segments. This process is given in Algorithm 2 and is repeated in each period that a request arrives.

Algorithm 2 Initial Allocation

Set $d_{r,c}(i) = 0, \forall r = 1, 2, \ldots, R \ \ \forall c = 1, 2, \ldots, C \ \ \forall i = 1, 2, \ldots, I$
for $e \in E$ **do**
  Set $m = k(e)$
  while $m > 0$ **do**
    if $\exists j | Y_j \geq 1 \wedge j \leq m$ **then**
      Set $j = \max \ j | Y_j > 0 \wedge j \leq m$
      Set $d_{e_r, e_c + k(e) - m}(j) = 1$
      Set $m = m - j$
      Set $Y_j = Y_j - 1$
    else
      Set $m = 0$
    end if
  end while
end for

For each period t, Algorithm 2 takes expected future demands Y_i and allocates them seats on the seating map. The Algorithm returns the allocation D. This procedure is not necessarily optimal. In an empty segment of 5 seats for example, the heuristic would allocate a group of 4 and 1, rather than allocating groups of 2 and 3 which may be better.

Note that a request of size i' has at most $Y_{i'}$ allocations using Algorithm 2. This implies that relatively rare request sizes would have few selections. We provide two additional algorithms that modify the initial allocation, and in the process expand the offer set. We label these modifications the reverse procedure and switching procedure respectively. We define D to be the vector of all allocations; $D = [d_{1,1}(1), \ldots, d_{r,c}(I)]$.

The two additional algorithm find equivalent allocations D' and D'', given an initial allocation D. In both algorithms, we consider only the current request i'. It is unnecessary to consider all i because this heuristic is performed in each time period and we only wish to expand the offer set for the current request i'.

The motivation for the reverse procedure is that the ordering of groups within an empty segment is irrelevant. Therefore, for each empty-segment, the ordering of groups can be reversed. This algorithm is presented in Algorithm 3. A simple example is used to illustrate. Suppose an empty-segment $S_{1,1}(4)$ exists and the initial allocation is $d_{1,1}(1) = 1$, and $d_{1,2}(3) = 1$. Reversing this segment for a request $i' = 1$ would imply setting $d'_{1,4}(1) = 1$.

Algorithm 3 Reverse Allocation

Set $d'_{r,c}(i) = 0, \forall r = 1, 2, \ldots, R \ \ \forall c = 1, 2, \ldots, C \ \ \forall i = 1, 2, \ldots, I$
for $e \in E$ **do**
  for $l = 0, 1, \ldots k(e) - 1$ **do**
    if $d_{e_r, e_c+l}(i') = 1$ **then** Then
      Set $d'_{e_r, e_c+k(e)-l-i'}(i') = 1$
    end if
  end for
end for

Algorithm 3 takes the initial allocation d and provides the reversed allocation for request size i'.

The motivation for the switching procedure is that the order of filling empty-segments is arbitrary. Thus, we can switch allocations between equally sized empty-segments. Thus, the second procedure we propose to find equivalent allocations is outlined in Algorithm 4. Again, we illustrate. Suppose there exist two empty segments: $S_{1,1}(2)$ and $S_{2,1}(2)$. Suppose that we are considering request size $i' = 1$ and that $d_{1,1}(1) = 1$, $d_{1,2}(1) = 1$, and $d_{2,1}(2) = 1$. Switching allocations between segments for request size $i' = 1$ solution would imply setting $d''_{2,1}(1) = 1$, $d''_{2,2}(1) = 1$.

Algorithm 4 Switching Allocation

Set $d''_{r,c}(i') = 0, \forall r = 1, 2, \ldots, R \ \ \forall c = 1, 2, \ldots, C \ \ \forall i = 1, 2, \ldots, I$
for $e \in E$ **do**
  for $e' \in E$ **do**
    if $k(e') = k(e)$ **then**
      for $l = 0, 1, \ldots k(e) - 1$ **do**
        if $d_{e_r, e_c+l}(i') = 1 \lor d'_{e_r, e_c+l}(i') = 1$ **then**
          Set $d''_{e'_r, e'_c+l}(i') = 1$
        end if
      end for
    end if
  end for
end for

Algorithm 4 takes allocations D and D' and returns new allocation D''. Note that although we present the expansion in this order, e.g. reverse and then switch, the order is not consequential. The final decision of what to offer a request size i' is to set $x_{r,c}(i') = \max\left(d_{r,c}(i'), d'_{r,c}(i'), d''_{r,c}(i')\right)$. Thus, as long as a seat-set $S_{r,c}(i)$ is opened in *any* allocation, it is offered. Because this heuristic is ran for every time period, it is unnecessary to track which specific allocation the offer is derived from.

To summarize, the objective of the GREEDY heuristic is to determine what seat sets to offer to a request for i' seats in period t. It works by using three nested algorithms. An initial allocation is done via Algorithm 2. Nested Algorithm 3 expands the offer set via a reverse procedure, and Nested Algorithm 4 expands the offer set further via a switching procedure.

A Numerical Illustration of Offered Set

We start our analysis by comparing the offer sets the algorithms make under the same capacity configuration. We randomly fill a 20 row, 30 column venue to three capacity levels (0 seats, 100 seats, 300 seats) and assume that an $i = 3$ request has arrived. In Figure 4.4 we present a figure to show the offered seat-set various algorithms make under the three capacity configuration. The lightest color indicates the seat-set has been offered. The second lightest indicates the seat-set was not offered. Black indicates the seat was filled and thus unavailable. Blue indicates the set is open but not offered, and grey indicates the seat has been taken. We do not consider SINGLEb as the offer set would be identical to SINGLEa in this example.

In general we see that the GREEDY heuristic is more selective of which seat sets it offers. Even when the venue is completely empty, the GREEDY heuristic offers less choice to consumers. However, a pattern emerges when the venue is empty: GREEDY offers a pattern such that seats chosen will leave spots that will be more easily filled by predicted groups. Similarly, SINGLEa removes the second to last seat-set in the row, in order to not leave a single spot blank, but leaves the very end of the row.

4.5 Simulation Study

4.5.1 Simulation Parameters

We run a computational experiment to test the relative performance of the proposed algorithms vs a NAIVE method. We simulate a 20 Row by 30 column theatre, resulting in 600 seats. This theatre size replicates a small Broadway theatre, such as the Helen Hayes theatre, which has 597 seats. It could also represent the orchestra level of a larger venue. We assume a max group size of $I = 5$. Each period a random request arrives. Using the heuristics above, seat-sets are offered and customers choose randomly based on the probabilistic model in Section 4.3. Arrivals are held constant across the heuristics, however the filled seat maps evolve independently. At the end of the T periods, we find how many seats were filled in total for each heuristic and compare against the NAIVE method.

We assume two sets of α values that are given in 4.1 and that are based on the author's best estimate. The important differences between the two are the α_1 probabilities and therefore we also use this to denote the two sets (e.g. $\alpha_1 = 5\%$ and $\alpha_1 = 10\%$). The second case has double the frequency of α_1. We would hypothesize that as single seats become more common (e.g. as α_1 increases), alternative seating algorithms become less effective, because smaller pockets become easier to fill naturally.

We also test differences in demand levels. Given a total number of periods T, and α values, we can calculate the expected number of seats demanded. This is given by $T \times \sum_{i=1}^{I} \alpha_i i$. We vary T to allow for different demand levels. We try 80%, 100% and 120% demand levels, where percentages are of total capacity.

The combination of parameter values results in 18 unique configurations of parameters, e.g. α, T and β. 500 trials are ran for each configuration.

4.5.2 Simulation Results

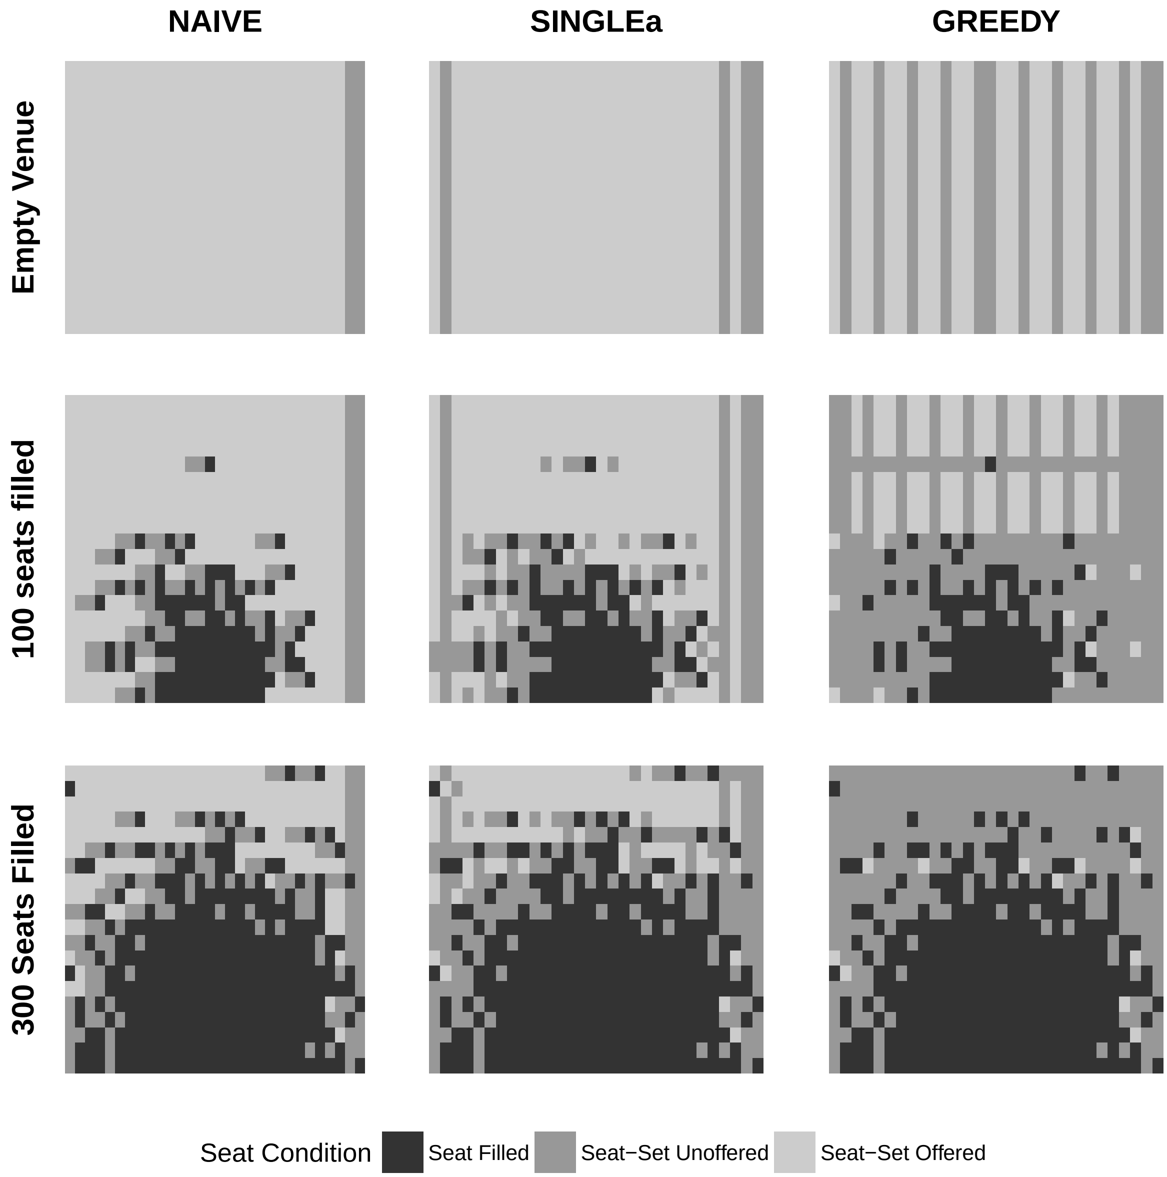

Figure 4.4: A comparison of offered sets for three heuristics at a 20 row, 30 column venue with an $i = 3$ group arriving.

Table 4.1: Group Distribution Configurations

Variable	$\alpha_1 = 5\%$	$\alpha_1 = 10\%$
α_0	20%	20%
α_1	5%	10%
α_2	35%	33.75%
α_3	10%	8.75%
α_4	25%	23.75%
α_5	5%	3.75%

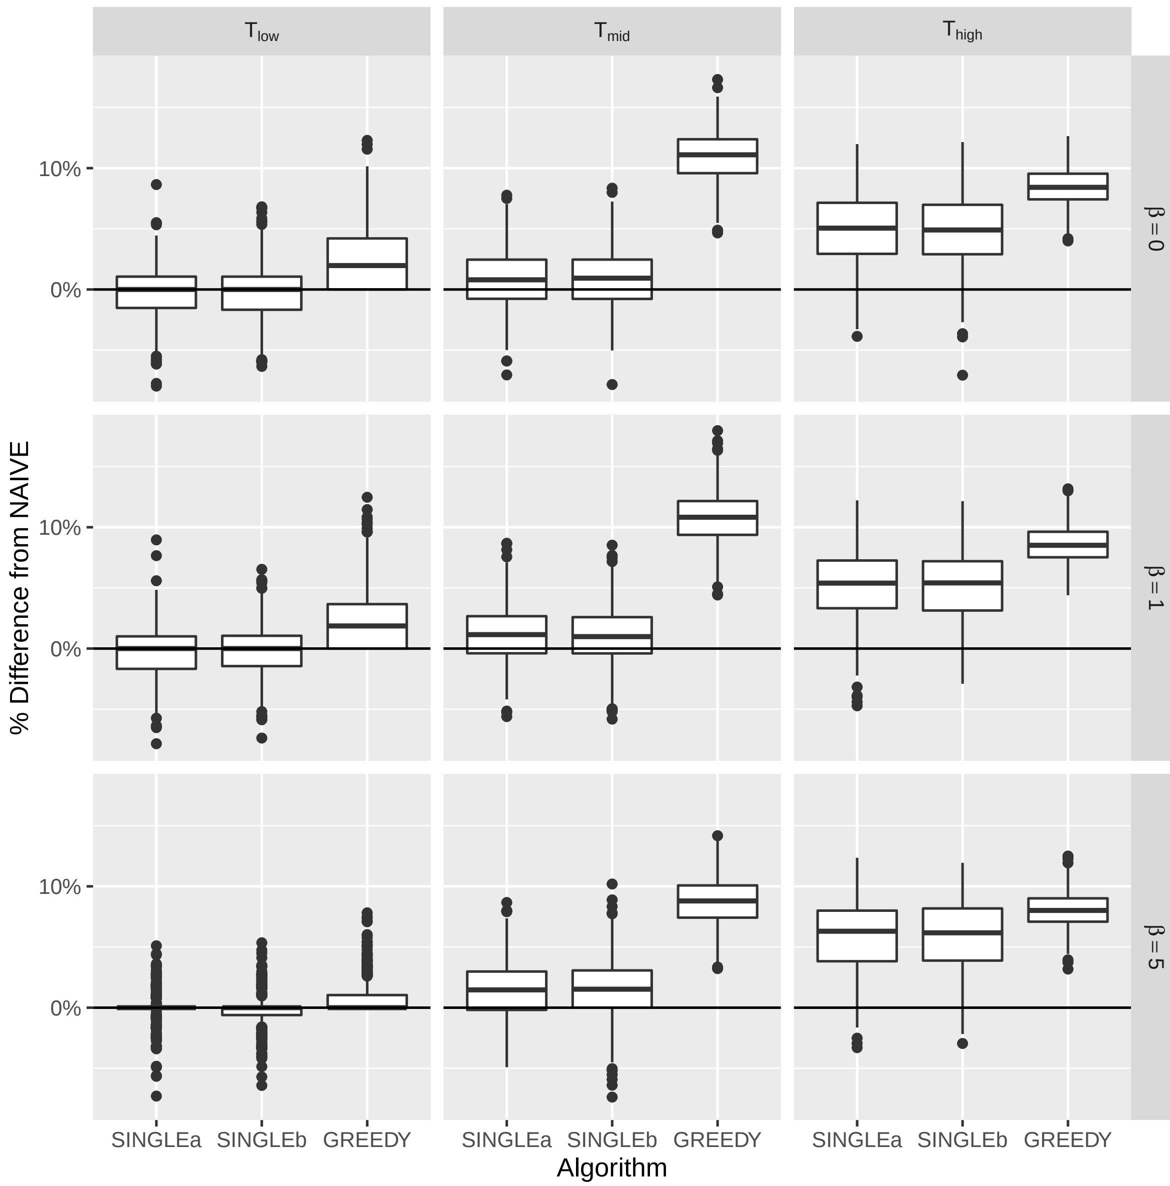

Figure 4.5: Box plots of percent difference between given algorithm and NAIVE in a given trial under the $\alpha_1 = 5\%$ condition

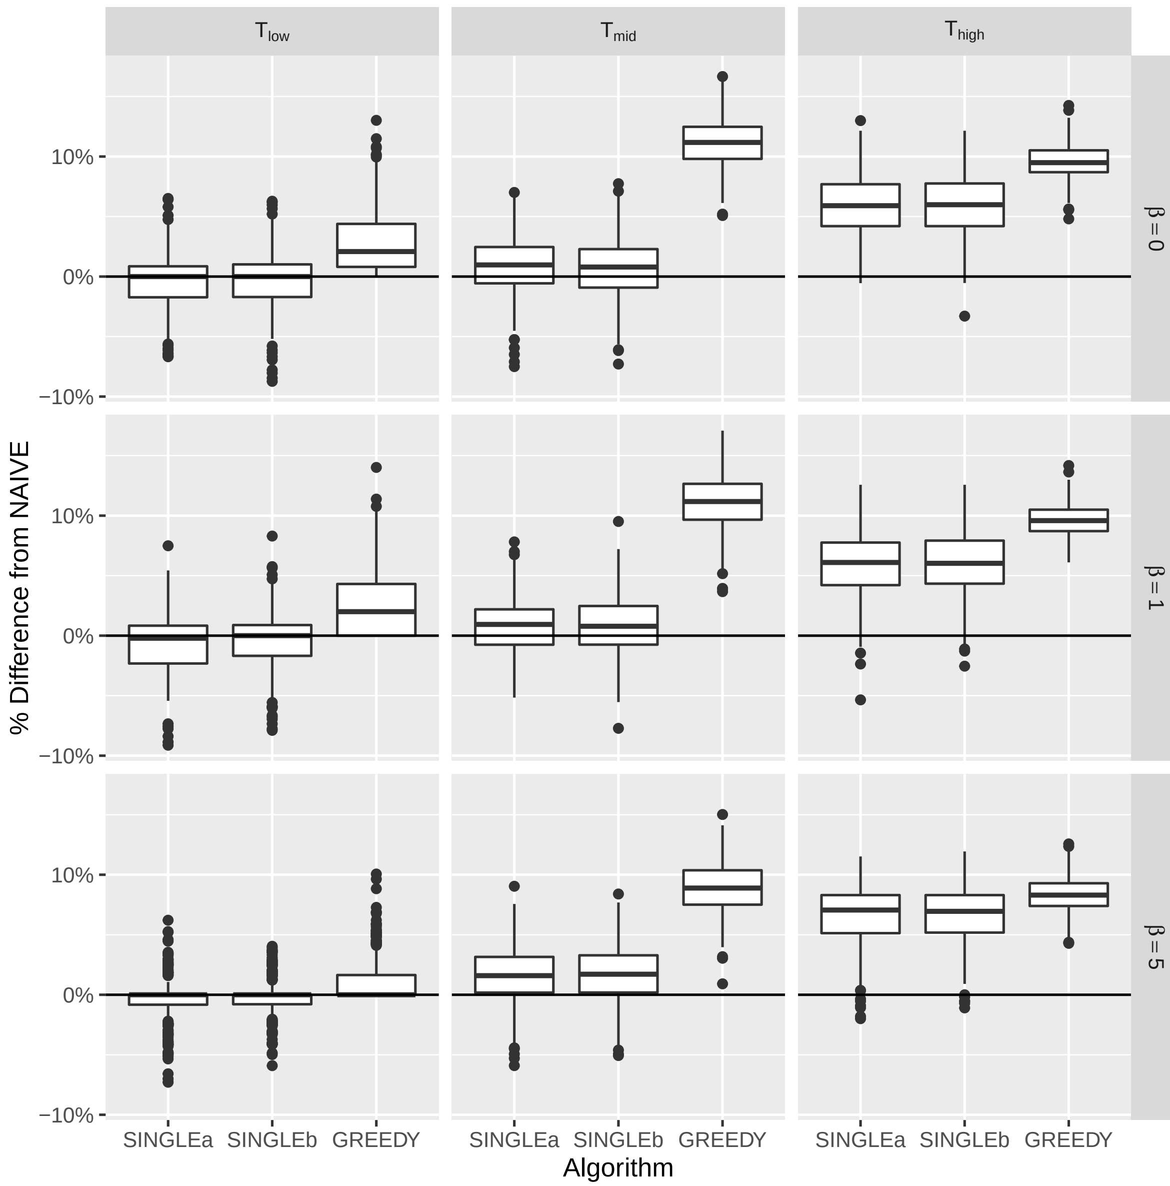

Figure 4.6: Box plots of percent difference between given algorithm and NAIVE in a given trial under the $\alpha_1 = 10\%$ condition

In Figure 4.5 and 4.6 we present boxplots for results for all configurations with $\alpha_1 = 5\%$ and $\alpha_1 = 10\%$ respectively. The boxplot displays the percent difference in seats filled between the heuristic and the NAIVE method. Positive percentages indicate the heuristic performed better than NAIVE. Accompanying parentheses indicate the standard deviation of results.

Our results suggest that the choice of SINGLEa or SINGLEb is unclear. Across almost all configurations the two heuristics have statistically indisguishable results from each other. This finding points to the fact that the differences in procedure occur infrequently and that differences are minor. In addition, even when they make different offer sets, it is not clear when it is beneficial or not. Because of this, we focus our attention solely on SINGLEa and GREEDY below.

Our results suggest that possible revenue enhancement is most sensitive to demand. All values given are the average performance across 500 trials (with standard deviation given in parentheses). At the low demand level configuration (T_{low}), SINGLEa does not produce outperformance. GREEDY however outperforms at the $\beta = 0$ and $\beta = 1$ configurations. The magnitude of outperformance is relatively small though. GREEDY outperforms by 2.4%(2.43) and 2.7%(2.55) respectively on average.

In contrast, at higher demand levels (T_{mid}, T_{high}), SINGLEa also exhibits outperformance, but is still bested by GREEDY. In the base case, the mean outperformance of SINGLEa is 1.1%(2.28), as compared to 10.8%(2.24) for GREEDY. At T_{high}, the mean outperformance of SINGLEa is 5.2%(2.97), as compared to 8.6%(1.56) for GREEDY.

It was unexpected for gains to hold at higher demand levels. We theorize that these gains may be holding because, at demand levels higher than 100%, the GREEDY method relatively easily reaches capacity. However, because single seats are the only inventory remaining, the NAIVE method can still only increase its total through more single seat groups. Increases in other group sizes do not help increase the average as they are unable to be filled anyways. Thus, even with more overall demand, the method cannot capture many of those gains.

As predicted, customer choices parameter β had an effect on performance. While similar at the $\beta = 0$ and $\beta = 1$ condition, at the $\beta = 5$ level, we see a drop off in performance of GREEDY. In the base case, GREEDY had a mean outperformance of 10.8%(2.24) compared to 8.8%(1.99) at $\beta = 5$. We theorize that this occurs because higher β encourages customers to select seats close to each other, limiting the problematic phenomenon.

We also test the effect of changes in group size distribution. Recall that we predicted that $\alpha_1 = 10\%$ would have smaller benefits. We see that our hypothesis is not true and effects persist even across group distributions. In the $\alpha_1 = 5\%$ condition, GREEDY had a mean outperformance of 10.8%(2.24) compared to 11.1%(2.27) in the $\alpha_1 = 10\%$ condition. Clearly substantial benefits remain even when the relative magnitude of single-seat requests is doubled.

4.6 Discussion and Conclusion

In this paper we explore group seating effects and how managers can use relatively simple allocation methods to produce revenue enhancement. Overall, we find evidence that the predominant method used to offer seats is sub-optimal. Significant revenue enhancements, on the order of 2-10%, are possible by using other methods. As well, these gains are relatively robust to a wide variety of parameters, including group distribution, customer choice parameters, and overall demand levels. We find that revenue enhancement is possible even at high demand levels.

Our model explicitly assumes homogenous prices across the venue. This is helpful in cleanly demonstrating the effect, but does not hold in reality. Entertainment venues are well-known for pricing seats differently depending on perceived customer valuation of different locations. Future research should incorporate variable pricing and consumer "buy-up" behavior. Another fruitful extension could consider a scenario based optimization, rather than the mean-value approach we consider here. In a mean-value problem, we do not consider the tradeoff involved in uncertainties. An optimization model incorporating this tradeoff may advance our GREEDY heuristic and provide insights into how to improve it.

The problem and model described above lay a framework for further research and raise a series of further questions. The insight that current seating methods may forego revenue leads into obvious questions about pricing and marketing. Instead of looking into the operational problem of allocating seats, should managers adjust prices on different seats because of these effects? For example, rather than not allow a group to sit in a specific seat, perhaps the price should be raised on them using a bid-price model. Our research also makes clear the need for more experimental research into what impacts consumer seating choices. Modeling a consumer's seating choice is, and will be, a difficult exercise without foundational research driving it. Research into the factors impacting seating choice would improve this line of research, but would also be of practical importance to managers increasingly looking for more granular pricing.

Chapter 5

Conclusion

In this thesis, we explore how entertainment managers can use analytics to make strategic and operational decisions. We dive into unique aspects of the live entertainment business and show that investigating these problems can produce fruitful research opportunities. Throughout the previous chapters, we have noted specific avenues for future research. We would encourage further research to explore in these areas.

Research topics in the entertainment revenue management space are plentiful though. Secondary markets exist for live entertainment tickets, e.g. Stubhub and SeatGeek. These online markets offer the ability to sell unwanted tickets for whatever price the market will bear. This includes the ability to sell tickets for above face-value, i.e. above the original market price.

These markets thus allow for *speculation* on the market price of future tickets. Speculators may purchase tickets with the intention of selling them later at a higher price. Market prices may change over time because of a host of reasons. Consumer uncertainty around scheduling reduces as events draw nearer. To the extent that uncertainty would reduce expected utility, we would expect this effect to lead to increases in price. In addition, underlying event demand can change as performers/teams become more or less sought after. This price change is unclear. Changes in the event itself are possible as well, e.g. forecasts for bad weather. All of these produce potential changes in the value of a ticket.

The fact that this marketplace dynamic exists at all is interesting - if tickets persistently sell above face-value, why do firms not raise prices? This phenomenon also exists in restaurants, where long lines or waiting lists exist, yet prices stay (relatively) low. A number of explanations have been raised for this peculiar fact. Some suggest that producers are risk-averse (Swofford, 1999) or that demand is positively correlated with demand from other consumers (Becker, 1991).

Ticket speculation and selling tickets above face-value has produced public outrage. New York State signed a law in 2016 banning the use of "bots" (online robots) to buy tickets online, responding to concerns that individual consumers were being pushed out of online markets. [1] In February 2017, the province of Ontario in Canada announced that it too was developing legislation to ensure that people had a "fair shot" at receiving tickets.[2] Research into these markets thus is of high interest to policy makers and entertainment firms.

Theoretical work on consumer fads, social herding, and network effects is also of high relevance to live entertainment managers. The business is one where winner-take-all effects dominate. Foundational work on modelling these issues is relevant to understanding how to take advantage of them. For example, a seminal paper on fads/herding behavior suggested that consumers may rationally ignore their private beliefs if sufficient public signals exist (Bikhchandani et al., 1992). Pricing that takes into account the social learning nature of the industry would be an interesting research avenue.

We also hold a broader goal for this thesis. We hope that we have shown the value in exploring new RM application areas deeply. We encourage this process to continue and echo the call from Currie (2015) for specific models for new RM industries. Other literature has started this process. Cruise ships used to be viewed as "floating hotels" (Talluri and van Ryzin, 2004). Today, a new line of research has specifically modeled lifeboat capacity constraints,

[1] https://www.governor.ny.gov/news/governor-cuomo-signs-legislation-combating-unfair-ticket-purchasing-and-reselling-practices

[2] https://news.ontario.ca/mag/en/2017/02/putting-fans-first-in-ontario.html

among other unique cruise ship dynamics (Biehn, 2006; Maddah et al., 2010). We hope that this continues, and that fundamental research is not viewed as merely applications, but as fundamental to expanding RM's use in the real world.

In this thesis, we have continued down this path. Group effects have been considered before, but not specifically within live entertainment. They could be modeled using a Network RM model. However, this is unlikely to lead to insights for the specific problem. Instead, we choose to explicitly model group distribution and customer choice, and study the resulting problem. This was done intentionally to lead to specific insights to the specific group problem.

Of course, applying RM to new industries requires domain level knowledge. Researchers should be humble in modeling new industries, recognizing that hotels and airlines may not accurately approximate them. In fact, we should work under the *assumption* that hotels/airlines do not approximate them. Practitioners are thus critical to the process. Work in new RM applications should include practitioners as much as is practical.

Bibliography

Albert, S., 1998. Movie Stars and the Distribution of Financially Successful Films in the Motion Picture Industry. Journal of Cultural Economics 22 (4), 249–270.

Andonov, R., Poirriez, V., Rajopadhye, S., Jun. 2000. Unbounded knapsack problem: Dynamic programming revisited. European Journal of Operational Research 123 (2), 394–407.

Becker, G. S., 1991. A Note on Restaurant Pricing and Other Examples of Social Influences on Price. Journal of Political Economy 99 (5), 1109–1116.

Bertsimas, D., Shioda, R., 2003. Restaurant Revenue Management. Operations Research 51 (3), 472–486.

Biehn, N., Jul. 2006. A cruise ship is not a floating hotel. Journal of Revenue and Pricing Management; Basingstoke 5 (2), 135–142.

Bikhchandani, S., Hirshleifer, D., Welch, I., 1992. A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades. Journal of Political Economy 100 (5), 992–1026.

Boyle, M., Chiou, L., Feb. 2009. Broadway productions and the value of a Tony Award. Journal of Cultural Economics 33 (1), 49–68.

Broadway League, Sep. 2015. The Demographics of the Broadway Audience 2014-2015. `https://www.broadwayleague.com/research/research-reports/`.

Cameron, A. C., Gelbach, J. B., Miller, D. L., 2011. Robust Inference With Multiway Clustering. Journal of Business & Economic Statistics 29 (2), 238–249.

Cameron, A. C., Miller, D. L., Mar. 2015. A Practitioner's Guide to Cluster-Robust Inference. Journal of Human Resources 50 (2), 317–372.

Cheng, C. H., Madan, M. S., Motwani, J., 1996. Knapsack-Based Algorithms for Designing Cellular Manufacturing Systems. The Journal of the Operational Research Society 47 (12), 1468–1476.

Cherbaka, N. S., Meller, R. D., 2008. Single-plant sourcing decisions with a multidimensional knapsack model. Journal of Manufacturing Systems 27 (1), 7–18.

Chisholm, D. C., 1997. Profit-Sharing versus Fixed-Payment Contracts: Evidence from the Motion Pictures Industry. Journal of Law, Economics, & Organization 13 (1), 169–201.

CIA, 2016. The World Factbook. `https://www.cia.gov/Library/publications/the-world-factbook/fields/2177.html`.

Courty, P., 2000. An economic guide to ticket pricing in the entertainment industry. Recherches Économiques de Louvain / Louvain Economic Review 66 (2), 167–192.

Currie, C. S. M., Mar. 2015. How far will the airline model stretch? Journal of Revenue and Pricing Management; Basingstoke 14 (2), 120–122.

Dantzig, G. B., 1957. Discrete-Variable Extremum Problems. Operations Research 5 (2), 266–277.

Dixon, M., Verma, R., Mar. 2013. Sequence effects in service bundles: Implications for service design and scheduling. Journal of Operations Management 31 (3), 138–152.

Dizdar, D., Gershkov, A., Moldovanu, B., May 2011. Revenue maximization in the dynamic knapsack problem. Theoretical Economics 6 (2), 157–184.

Dunne, P., Feb. 2012. Dynamic pricing trend sweeps across Major League Baseball.

Elberse, A., Oct. 2007. The Power of Stars: Do Star Actors Drive the Success of Movies? Journal of Marketing 71 (4), 102–120.

Elberse, A., 2013. Blockbusters: Hit-Making, Risk-Taking, and the Big Business of Entertainment, first edition Edition. Henry Holt and Company, New York.

Eliashberg, J., Elberse, a., Leenders, M., 2006. The Motion Picture Industry: Critical Issues in Practice, Current Research, and New Research Directions. Vol. 25.

Gilovich, T., Kumar, A., Jampol, L., Jan. 2015. A wonderful life: Experiential consumption and the pursuit of happiness. Journal of Consumer Psychology 25 (1), 152–165.

Goldman, W., 1983. Adventures in the Film Trade. Warner Books.

Greene, W. H., 2007. Econometric Analysis, 6th Edition.

Healy, P., Mar. 2014. Ticket Pricing Puts 'Lion King' Atop Broadway's Circle of Life. The New York Times.

Holbrook, M. B., Hirschman, E. C., 1982. The Experiential Aspects of Consumption: Consumer Fantasies, Feelings, and Fun. Journal of Consumer Research 9 (2), 132–140.

Hume, M., Sullivan Mort, G., Liesch, P. W., Winzar, H., Jun. 2006. Understanding service experience in non-profit performing arts: Implications for operations and service management. Journal of Operations Management 24 (4), 304–324.

Johnson, D. S., Jun. 1974. Fast algorithms for bin packing. Journal of Computer and System Sciences 8 (3), 272–314.

Joshi, A., 2015. Movie Stars and the Volatility of Movie Revenues. Journal of Media Economics 28 (4), 246–267.

Kellerer, P. H., Pferschy, P. U., Pisinger, P. D., 2004. The Multiple-Choice Knapsack Problem. In: Knapsack Problems. Springer Berlin Heidelberg, pp. 317–347.

Kemper, C., Breuer, C., Feb. 2016. How Efficient is Dynamic Pricing for Sport Events? Designing a Dynamic Pricing Model for Bayern Munich. International Journal of Sport Finance; Morgantown 11 (1), 4–25.

Kleywegt, A. J., Papastavrou, J. D., 1998. The Dynamic and Stochastic Knapsack Problem. Operations Research 46 (1), 17–35.

Kleywegt, A. J., Papastavrou, J. D., 2001. The Dynamic and Stochastic Knapsack Problem with Random Sized Items. Operations Research 49 (1), 26–41.

Lee, T. C., Hersh, M., Aug. 1993. A Model for Dynamic Airline Seat Inventory Control with Multiple Seat Bookings. Transportation Science 27 (3), 252–265.

Leslie, P., 2004. Price Discrimination in Broadway Theater. The RAND Journal of Economics 35 (3), 520–541.

Lewis, R., 2015. Taylor Swift's '1989' is 2015's highest grossing concert tour by far. `http://www.latimes.com/entertainment/music/posts/la-et-ms-taylor-swift-1989-tour-highest-grossing-concerts-pollstar-20151230-story.html`.

Liu, A., Liu, Y., Mazumdar, T., Aug. 2013. Star power in the eye of the beholder: A study of the influence of stars in the movie industry. Marketing Letters 25 (4), 385–396.

Liu, A. X., Mazumdar, T., Li, B., Aug. 2014. Counterfactual Decomposition of Movie Star Effects with Star Selection. Management Science 61 (7), 1704–1721.

Luo, L., Chen, X., Han, J., Park, C. W., 2010. Dilution and Enhancement of Celebrity Brands Through Sequential Movie Releases. Journal of Marketing Research 47 (6), 1114–1128.

Maddah, B., Moussawi-Haidar, L., El-Taha, M., Rida, H., Nov. 2010. Dynamic cruise ship revenue management. European Journal of Operational Research 207 (1), 445–455.

Maddison, D., Sep. 2004. Increasing returns to information and the survival of broadway theatre productions. Applied Economics Letters 11 (10), 639–643.

Maddison, D., Nov. 2005. Are There Too Many Revivals on Broadway? A Stochastic Dominance Approach. Journal of Cultural Economics 29 (4), 325–334.

Martello, S., Toth, P., 1980. Solution of the zero-one multiple knapsack problem. European Journal of Operational Research 4 (4), 276–283.

Nelson, R. A., Glotfelty, R., Feb. 2012. Movie stars and box office revenues: An empirical analysis. Journal of Cultural Economics 36 (2), 141–166.

Ozanian, M., 2015. The Most Valuable Teams In The NFL. `http://www.forbes.com/sites/mikeozanian/2015/09/14/the-most-valuable-teams-in-the-nfl/`.

Papastavrou, J. D., Rajagopalan, S., Kleywegt, A. J., 1996. The Dynamic and Stochastic Knapsack Problem with Deadlines. Management Science 42 (12), 1706–1718.

Phumchusri, N., Swann, J. L., Jan. 2014. Scaling the house: Optimal seating zones for entertainment venues when location of seats affects demand. International Journal of Revenue Management 8 (1), 56–98.

Prag, J., Casavant, J., Sep. 1994. An empirical study of the determinants of revenues and marketing expenditures in the motion picture industry. Journal of Cultural Economics 18 (3), 217–235.

Ravid, S. A., Oct. 1999. Information, Blockbusters, and Stars: A Study of the Film Industry. The Journal of Business 72 (4), 463–492.

Reaney, P., 2014. Hollywood stars set Broadway aglow in limited-run plays. Reuters.

Reddy, S. K., Swaminathan, V., Motley, C. M., 1998. Exploring the Determinants of Broadway Show Success. Journal of Marketing Research.

Salganik, M., Dodds, P., Watts, D., 2006. Experimental study of inequality and unpredictability in an artificial cultural market. Science 311 (5762), 854–856.

Sharma, S., Dubey, D., 2010. Multiple sourcing decisions using integrated AHP and knapsack model: A case on carton sourcing. The International Journal of Advanced Manufacturing Technology 51 (9-12), 1171–1178.

Simonoff, J., Ma, L., 2000. An empirical study of factors relating to the success of Broadway shows. The Journal of Business 76 (1), 135–150.

Simonson, R., Aug. 2010. Actors' Salaries & The Levels of Stage Management. `http://www.playbill.com/article/ask-playbillcom-actors-salaries-the-levels-of-stage-management-com-170516`.

Sokolove, M., Apr. 2016. The C.E.O. of 'Hamilton' Inc. The New York Times.

Swofford, J. L., Sep. 1999. Arbitrage, Speculation, and Public Policy Toward Ticket Scalping. Public Finance Review 27 (5), 531–540.

Talluri, K., van Ryzin, G., 2004. The Theory and Practice of Revenue Management.

Thomson, M., 2006. Human Brands: Investigating Antecedents to Consumers' Strong Attachments to Celebrities. Journal of Marketing 70 (3), 104–119.

Toma, M., Meads, H., Sep. 2007. Recent evidence on the determinants of concert attendance for midsize symphonies. Journal of Economics and Finance 31 (3), 412–421.

van Slyke, R., Young, Y., 2000. Finite Horizon Stochastic Knapsacks with Applications to Yield Management. Operations Research 48 (1), 155–172.

Vany, A. D., Walls, D. W., 1996. Bose-Einstein Dynamics and Adaptive Contracting in the Motion Picture Industry. The Economic Journal 106 (439), 1493–1514.

Vany, A. D., Walls, D. W., 1999. Uncertainty in the movie industry: Does star power reduce the terror of the box office? Journal of Cultural Economics, 285–318.

Veeraraghavan, S., Vaidyanathan, R., Jan. 2012. Measuring Seat Value in Stadiums and Theaters. Production and Operations Management 21 (1), 49–68.

Vinod, B., Jul. 2013. Revenue management for the optimal control of group traffic. Journal of Revenue and Pricing Management; Basingstoke 12 (4), 295–304.

Wooldridge, J. M., 2002. Econometric Analysis of Cross Section and Panel Data. MIT Press.

Wooldridge, J. M., 2012. Introductory Econometrics: A Modern Approach, fifth edition Edition. South-Western.

Xu, J. J., Fader, P., Veeraraghavan, S. K., Jun. 2016. The Revenue Impact of Dynamic Pricing Policies in Major League Baseball Ticket Sales.

Zhang, D., Adelman, D., Aug. 2009. An Approximate Dynamic Programming Approach to Network Revenue Management with Customer Choice. Transportation Science 43 (3), 381–394.

Appendix A

Other Data and Models

Table A.1: List of Holidays Controlled for in Chapter Two

Holiday Name
Easter
Independence Day
Memorial Day
Labor Day
Federal Election
Thanksgiving
New Years Eve
Christmas

Table A.2: Regression Results from all models in Chapter Two

	(1)	(2)	(3)	(4)	(5)	(6)
Starpower	0.095**	0.086		0.096**	0.086	
	(0.038)	(0.173)		(0.037)	(0.171)	
Starpower^2		0.011			0.013	
		(0.212)			(0.210)	
Starpower3			−0.055			−0.054
			(0.054)			(0.053)
Starpower3^2			0.118***			0.117***
			(0.028)			(0.028)
CompetitionStarpower				−0.004	−0.004	−0.004
				(0.006)	(0.006)	(0.006)
TonyAwards	0.052	0.052	0.014	0.051	0.051	0.014
	(0.033)	(0.033)	(0.035)	(0.032)	(0.032)	(0.034)
OtherAwards	0.040**	0.039**	0.021*	0.039**	0.039**	0.020*
	(0.017)	(0.017)	(0.012)	(0.017)	(0.017)	(0.012)
DaysOpen	−0.411***	−0.411***	−0.417***	−0.403***	−0.403***	−0.409***
	(0.103)	(0.104)	(0.098)	(0.102)	(0.104)	(0.097)
DaysOpen^2	0.031***	0.031***	0.032***	0.031***	0.031***	0.031***
	(0.010)	(0.011)	(0.010)	(0.010)	(0.010)	(0.010)
Constant	12.801***	12.800***	12.796***	12.775***	12.774***	12.770***
	(0.425)	(0.423)	(0.416)	(0.420)	(0.418)	(0.411)
Observations	9,420	9,420	9,420	9,420	9,420	9,420
R^2	0.888	0.888	0.892	0.888	0.888	0.892
Adjusted R^2	0.884	0.884	0.888	0.884	0.884	0.889
Residual Std. Error	0.202 (df = 9131)	0.202 (df = 9130)	0.198 (df = 9130)	0.202 (df = 9130)	0.202 (df = 9129)	0.198 (df = 9129)

Note: *p<0.1; **p<0.05; ***p<0.01
Dependent variable of *LOG(Gross)*. Control Variables Excluded. Clustered Standard Errors shown in parentheses.

A random effects treatment and assumptions can be found in (Wooldridge, 2002, pg. 257).

Table A.3: Regression Results from Chapter Two assuming random effects on unobserved individual characteristics.

	(1)	(2)	(3)	(4)
Starpower	0.101***		0.101***	
	(0.037)		(0.036)	
Starpower3		−0.048		−0.047
		(0.052)		(0.051)
Starpower3^2		0.114***		0.114***
		(0.028)		(0.028)
CompetitionStarpower			−0.005	−0.005
			(0.006)	(0.006)
TonyAwards	0.054*	0.016	0.053*	0.015
	(0.032)	(0.034)	(0.032)	(0.034)
OtherAwards	0.040**	0.021*	0.039**	0.021*
	(0.017)	(0.012)	(0.017)	(0.012)
DaysOpen	−0.391***	−0.398***	−0.382***	−0.389***
	(0.099)	(0.095)	(0.099)	(0.094)
DaysOpen^2	0.030***	0.031***	0.029***	0.030***
	(0.010)	(0.010)	(0.010)	(0.010)
Constant	11.346***	11.342***	11.359***	11.356***
	(0.111)	(0.107)	(0.113)	(0.109)
Observations	9,420	9,420	9,420	9,420
R^2	0.803	0.806	0.804	0.807
Adjusted R^2	0.801	0.804	0.801	0.804

Note: *p<0.1; **p<0.05; ***p<0.01
Dependent variable of *LOG(Gross)*. Control Variables Excluded. Clustered Standard Errors shown in parentheses.

Formulation of a Deterministic Linear Program in Chapter Four

A Deterministic Linear Program (DLP) takes *expected* future demands deterministically and makes operational decisions accordingly (Talluri and van Ryzin, 2004). The DLP takes the expected group demands, and allocates them locations into the venue.

$$
\begin{aligned}
\underset{X}{\text{maximize}} \quad & \sum_{i=1}^{I}\sum_{r=1}^{R}\sum_{c=1}^{C}(i)x_{r,c}(i) \\
\text{subject to} \quad & \sum_{r=1}^{R}\sum_{c=1}^{C}X_{r,c}(i) \leq Y_i,\ i = 1,\dots,I. \\
& \sum_{\max(1,c-i+1)}^{c}\ \sum_{i=1}^{I} x_{r,c,i} + s_{r,c} \leq 1,\ r = 1,\dots,R,\ c = 1,\dots,C. \\
& x_{r,c}(i) \leq 0,\ r = 1,\dots,R,\ i = 2,\dots,I,\ c = C - i + 1,\dots,C. \\
& x_{r,c}(i) \in \{0,1\},\ r = 1,\dots,R,\ i = 1,\dots,I,\ c = 1,\dots,C.
\end{aligned}
$$

The objective function seeks to maximize the expected number of filled seats by reserving groups of seats in the seating layout for specific group sizes.

Constraint (1) ensures that the amount of group type i placed do not exceed the expected demand. Constraint (2) ensures that for every seat in the venue, a unique group is partitioned for that seat. Without this constraint, a single seat could be reserved for separate groups, or equivalently, groups could overlap. Constraint (3) ensures that we do not allocate demand to seating spots which do not exist. Finally constraint (4) is to ensure decision variables are binary.

Appendix B

Proofs of Theorems

Proof for Proposition 3.2.1: (i) To see that the function is positive, note that a feasible solution includes to always close all compartments, resulting in zero reward. There is finite capacity in finite compartments, thus the value is clearly bounded.
(ii) Proof by induction in time. In period $t = 1$, $V_t(C, i) \geq V_{t-1}(C, i) = 0$, by the boundry condition and because the value-to-go function is non negative by (i). Assume $V_t(C, i)$ is increasing in t, for $t = 1, 2, \ldots, m-1$. Let $t = m$, and the state be given by (C, i), then $V_t(C, i) \geq V_{t-1}(C, i)$,

$$\begin{aligned}\iff \max_{D_t^{i'} \in \mathbb{D}(C,i)} &\left\{ \sum_{k=1}^{K} \gamma_k(D_t^{i'}) r_k(i) + \sum_{k=1}^{K} \gamma_k(D_t^{i'}) W_{t-1}(C_{[k]}) + (1 - \sum_{k=1}^{K} \gamma_k(D_t^{i'})) W_{t-1}(C) \right\} \geq \\ \max_{D_{t-1}^{i} \in \mathbb{D}(C,i)} &\left\{ \sum_{k=1}^{K} \gamma_k(D_{t-1}^{i}) r_k(i) + \sum_{k=1}^{K} \gamma_k(D_{t-1}^{i}) W_{t-2}(C_{[k]}) + (1 - \sum_{k=1}^{K} \gamma_k(D_{t-1}^{i})) W_{t-2}(C) \right\}\end{aligned}$$

Let D^* denote the arg-max of the LHS and D° denote the arg-max of the RHS, $V_t(C, i) \geq V_{t-1}(C, i)$,

$$\begin{aligned}\iff &\sum_{k=1}^{K} \gamma_k(D^*) r_k(i) + \sum_{k=1}^{K} \gamma_k(D^*) W_{t-1}(C_{[k]}) + (1 - \sum_{k=1}^{K} \gamma_k(D^*)) W_{t-1}(C) \geq \\ &\sum_{k=1}^{K} \gamma_k(D^\circ) r_k(i) + \sum_{k=1}^{K} \gamma_k(D^\circ) W_{t-2}(C_{[k]}) + (1 - \sum_{k=1}^{K} \gamma_k(D^\circ)) W_{t-2}(C)\end{aligned}$$

We subtract $\sum_{k=1}^{K} \gamma_k(D^\circ) r_k(i) + \sum_{k=1}^{K} \gamma_k(D^\circ) W_{t-1}(C_{[k]}) + (1 - \sum_{k=1}^{K} \gamma_k(D^\circ)) W_{t-1}(C)$ from both sides.

$$
\begin{aligned}
\iff & \sum_{k=1}^{K} \gamma_k(D^*) r_k(i) + \sum_{k=1}^{K} \gamma_k(D^*) W_{t-1}(C_{[k]}) + (1 - \sum_{k=1}^{K} \gamma_k(D^*)) W_{t-1}(C) \\
& - \sum_{k=1}^{K} \gamma_k(D^\circ) r_k(i) - \sum_{k=1}^{K} \gamma_k(D^\circ) W_{t-1}(C_{[k]}) - (1 - \sum_{k=1}^{K} \gamma_k(D^\circ)) W_{t-1}(C) \geq \\
& \sum_{k=1}^{K} \gamma_k(D^\circ) r_k(i) + \sum_{k=1}^{K} \gamma_k(D^\circ) W_{t-2}(C_{[k]}) + (1 - \sum_{k=1}^{K} \gamma_k(D^\circ)) W_{t-2}(C) \\
& - \sum_{k=1}^{K} \gamma_k(D^\circ) r_k(i) - \sum_{k=1}^{K} \gamma_k(D^\circ) W_{t-1}(C_{[k]}) - (1 - \sum_{k=1}^{K} \gamma_k(D^\circ)) W_{t-1}(C)
\end{aligned}
$$

$$
\begin{aligned}
\iff & \sum_{k=1}^{K} \gamma_k(D^*) r_k(i) + \sum_{k=1}^{K} \gamma_k(D^*) W_{t-1}(C_{[k]}) + (1 - \sum_{k=1}^{K} \gamma_k(D^*)) W_{t-1}(C) \\
& - \sum_{k=1}^{K} \gamma_k(D^\circ) r_k(i) - \sum_{k=1}^{K} \gamma_k(D^\circ) W_{t-1}(C_{[k]}) - (1 - \sum_{k=1}^{K} \gamma_k(D^\circ)) W_{t-1}(C) \geq \\
& \sum_{k=1}^{K} \gamma_k(D^\circ) \Big\{ W_{t-2}(C_{[k]}) - W_{t-1}(C_{[k]}) \Big\} + (1 - \sum_{k=1}^{K} \gamma_k(D^\circ)) \Big\{ W_{t-2}(C) - W_{t-1}(C) \Big\}
\end{aligned}
$$

Note however that the LHS of the inequality is non-negative since D° is a feasible but not necessarily optimal decision in period t. Note also that the RHS of the inequality is non-positive.

To see this, recall that $\gamma_k(\cdot)$ are probabilities and so $\sum_{k=1}^{K} \gamma_k(D^\circ)$ and $(1 - \sum_{k=1}^{K} \gamma_k(D^\circ))$ are both positive, while both inner terms are non-positive by the induction assumption: $W_{t-2}(C_{[k]}) - W_{t-1}(C_{[k]}) = \sum_{j=0}^{N} \alpha^i (V_{t-2}(C_{[k]}, j) - V_{t-1}(C_{[k]}, j)) \leq 0$,

and $W_{t-2}(C) - W_{t-1}(C) = \sum_{j=0}^{N} \alpha^i (V_{t-2}(C, j) - V_{t-1}(C, j)) \leq 0$. Therefore the inequality holds and $V_t(C, i)$ is increasing in t.

(iii) We seek to show, for $t = 1, 2, \ldots, T$, $i \in \mathbb{I}$, capacity configuration C, and $C' = C + c^+ \lambda^{k'}$, for a compartment $k' \in \mathbb{K}$, and scalar c^+, that $V_t(C', i) \geq V_t(C, i)$. Proof by induction in t. At time period $t = 1$, $V_t(C', i) \geq V_t(C, i)$,

$$
\begin{aligned}
\iff & \max_{D_t^i \in \mathbb{D}(C', i)} \Big\{ \sum_{k=1}^{K} \gamma_k(D_t^i) r_k(i) + \sum_{k=1}^{K} \gamma_k(D_t^i) W_{t-1}(C'_{[k]}) + (1 - \sum_{k=1}^{K} \gamma_k(D_t^i)) W_{t-1}(C') \Big\} \\
& \geq \max_{D_t^i \in \mathbb{D}(C, i)} \Big\{ \sum_{k=1}^{K} \gamma_k(D_t^i) r_k(i) + \sum_{k=1}^{K} \gamma_k(D_t^i) W_{t-1}(C_{[k]}) + (1 - \sum_{k=1}^{K} \gamma_k(D_t^i)) W_{t-1}(C) \Big\} \\
\iff & \max_{D_t^i \in \mathbb{D}(C', i)} \sum_{k=1}^{K} \gamma_k(D_t^i) r_k(i) \geq \max_{D_t^i \in \mathbb{D}(C, i)} \sum_{k=1}^{K} \gamma_k(D_t^i) r_k(i)
\end{aligned}
$$

Since the feasible set on the RHS is a subset of the one on the LHS, $\mathbb{D}(C', i) \supseteq \mathbb{D}(C, i)$, and $\gamma_k(\cdot)$ do not change with C, the above inequality holds trivially. Assume the result holds for

$t = 1, 2, \ldots, m-1$. Let $t = m$, then we show that for C and C', $V_t(C', i) \geq V_t(C, i)$. Let D^* and D° be the optimal decisions for the capacity configuration C' and C, respectively, then $V_t(C', i) \geq V_t(C, i)$

$$\begin{aligned}
\iff & \sum_{k=1}^{K} \gamma_k(D^*) r_k(i) + \sum_{k=1}^{K} \gamma_k(D^*) W_{t-1}(C'_{[k]}) + (1 - \sum_{k=1}^{K} \gamma_k(D^*)) W_{t-1}(C') \\
& \geq \sum_{k=1}^{K} \gamma_k(D^\circ) r_k(i) + \sum_{k=1}^{K} \gamma_k(D^\circ) W_{t-1}(C_{[k]}) + (1 - \sum_{k=1}^{K} \gamma_k(D^\circ)) W_{t-1}(C)
\end{aligned}$$

We can subtract $\sum_{k=1}^{K} \gamma_k(D^\circ) r_k(i) + \sum_{k=1}^{K} \gamma_k(D^\circ) W_{t-1}(C'_{[k]}) + (1 - \sum_{k=1}^{K} \gamma_k(D^\circ)) W_{t-1}(C')$ from both sides.

$$\begin{aligned}
\iff & \sum_{k=1}^{K} \gamma_k(D^*) r_k(i) + \sum_{k=1}^{K} \gamma_k(D^*) W_{t-1}(C'_{[k]}) + (1 - \sum_{k=1}^{K} \gamma_k(D^*)) W_{t-1}(C') \\
& - \sum_{k=1}^{K} \gamma_k(D^\circ) r_k(i) - \sum_{k=1}^{K} \gamma_k(D^\circ) W_{t-1}(C'_{[k]}) - (1 - \sum_{k=1}^{K} \gamma_k(D^\circ)) W_{t-1}(C') \\
& \geq \sum_{k=1}^{K} \gamma_k(D^\circ) r_k(i) + \sum_{k=1}^{K} \gamma_k(D^\circ) W_{t-1}(C_{[k]}) + (1 - \sum_{k=1}^{K} \gamma_k(D^\circ)) W_{t-1}(C) \\
& - \sum_{k=1}^{K} \gamma_k(D^\circ) r_k(i) - \sum_{k=1}^{K} \gamma_k(D^\circ) W_{t-1}(C'_{[k]}) - (1 - \sum_{k=1}^{K} \gamma_k(D^\circ)) W_{t-1}(C')
\end{aligned}$$

$$\begin{aligned}
\iff & \sum_{k=1}^{K} \gamma_k(D^*) r_k(i) + \sum_{k=1}^{K} \gamma_k(D^*) W_{t-1}(C'_{[k]}) + (1 - \sum_{k=1}^{K} \gamma_k(D^*)) W_{t-1}(C') \\
& - \sum_{k=1}^{K} \gamma_k(D^\circ) r_k(i) - \sum_{k=1}^{K} \gamma_k(D^\circ) W_{t-1}(C'_{[k]}) - (1 - \sum_{k=1}^{K} \gamma_k(D^\circ)) W_{t-1}(C') \\
& \geq \sum_{k=1}^{K} \gamma_k(D^\circ) \Big\{ W_{t-1}(C_{[k]}) - W_{t-1}(C'_{[k]}) \Big\} + (1 - \sum_{k=1}^{K} \gamma_k(D^\circ)) \Big\{ W_{t-1}(C) - W_{t-1}(C') \Big\}
\end{aligned}$$

Note that the LHS is non-negative, since D° is not necessarily the optimal decision at capacity C', and that the RHS is non-positive. To see the latter, recall that $\gamma_k(\cdot)$ are probabilities and so $\sum_{k=1}^{K} \gamma_k(D)$ and $(1 - \sum_{k=1}^{K} \gamma_k(D))$ are both positive, while both inner terms are non-positive by the induction assumption: $W_{t-1}(C_{[k]}) - W_{t-1}(C'_{[k]}) = \sum_{j=0}^{N} \alpha^i (V_{t-1}(C_{[k]}, j) - V_{t-1}(C'_{[k]}, j)) \leq 0$, and $W_{t-1}(C) - W_{t-1}(C') = \sum_{j=0}^{N} \alpha^i (V_{t-1}(C, j) - V_{t-1}(C', j)) \leq 0$. Therefore the inequality holds and $V_t(C, i)$ is increasing in c_k.

(iv) For any $i \in \mathbb{I}$ and positive integer n, the probability of n arrivals of item i approaches 1 as the remaining time approaches infinity. Thus, items become unconstrained and the problem is equivalent to the unbounded problem, with the same value function. $\square$

Proof for Proposition 3.2.2: From definition (3.7), we can replace $r_k(i) + W_{t-1}(C_{[k]}) - W_{t-1}(C)$ with B_k,

$$\begin{aligned} V_t(C,i) &= \max_{D_t^i \in \mathbb{D}(C,i)} \left\{ \sum_{k=1}^{K} \gamma_k(D_t^i)\big(r_k(i) + W_{t-1}(C_{[k]}) - W_{t-1}(C)\big) + W_{t-1}(C) \right\} \\ &= \max_{D_t^i \in \mathbb{D}(C,i)} \left\{ \sum_{k=1}^{K} \gamma_k(D_t^i) B_k + W_{t-1}(C) \right\} \end{aligned}$$

□

Proof for Lemma 3.2.3: We seek to show that if $\Delta W_t(C)$ has time-invariant ordering then B_k are decreasing in time for any k, t, i, and C. Recall that $B_k = r_k(i) - \big(W_{t-1}(C) - W_{t-1}(C_{[k]})\big)$. Decreasing in time implies

$$\begin{aligned} & r_k(i) - (W_t(C) - W_t(C_{[k]})) \geq r_k(i) - (W_{t+1}(C) - W_{t+1}(C_{[k]})) \\ \iff & W_{t+1}(C) - W_t(C) \geq W_{t+1}(C_{[k]}) - W_t(C_{[k]})) \\ \iff & \Delta W_{t+1}(C) \geq \Delta W_{t+1}(C_{[k]}) \end{aligned}$$

From Proposition 3.2.1 (iii) it follows that $W_t(C) \geq W_t(C_{[k]})$, $\forall t$, and noting that $W_t(C) = \sum_{t'=1}^{t} \Delta W_{t'}(C)$ holds for all t, and assuming a time-invariant ordering on $\Delta W_t(C)$, the inequality and result holds. □

Proof for Theorem 1: The algorithm starts with all compartments $k \in \mathbb{K}$ closed. As a part of the if statement, compartments are only opened if $c_k \geq w_i$. Therefore, the algorithm must produce a feasible result. Note that,

$$\begin{aligned} V_t(C,i) &= \sum_{k=1}^{K} \gamma_k(D^{i*}) B_k + W_{t-1}(C) \\ &= \sum_{k=1}^{K} d_k^{i*}\Big(u_k + \sum_{x=k+1}^{K} u_x \prod_{l=k+1}^{x} (1 - d_l^{i*})\Big) B_k + W_{t-1}(C) \\ &= \sum_{k=1}^{K} d_k^{i*} u_k B_k + \sum_{k=1}^{K} \sum_{x=k+1}^{K} d_k^{i*} u_x B_k \Big(\prod_{l=k+1}^{x} (1 - d_l^{i*})\Big) + W_{t-1}(C) \\ &= \sum_{k=1}^{K} d_k^{i*} u_k B_k + \sum_{k=2}^{K} \sum_{x=1}^{k-1} d_x^{i*} u_k B_x \Big(\prod_{l=x+1}^{k} (1 - d_l^{i*})\Big) + W_{t-1}(C) \\ &= \sum_{k=1}^{K} d_k^{i*} u_k B_k + \sum_{k=2}^{K} u_k \Big(\sum_{x=1}^{k-1} d_x^{i*} B_x \prod_{l=x+1}^{k} (1 - d_l^{i*})\Big) + W_{t-1}(C) \\ &= u_1 d_1^{i*} B_1 + \sum_{k=2}^{K} u_k \Big(d_k^{i*} B_k + \sum_{x=1}^{k-1} d_x^{i*} B_x \prod_{l=x+1}^{k} (1 - d_l^{i*})\Big) + W_{t-1}(C) \end{aligned}$$

Notice that $d_1^{i*} B_1 = 0$, if $d_1^{i*} = 0$, and $d_1^{i*} B_1 = B_1$, if $d_1^{i*} = 1$. Similarly, for a given $k = 2, 3, \ldots, K$, $d_k^{i*} B_k + \sum_{x=1}^{k-1} d_x^{i*} B_x (\prod_{l=x+1}^{k} (1 - d_l^{i*}))$ evaluates to 0, if $d_x = 0\ \forall x = 1, 2 \ldots, k$.

Otherwise, if at least one row is opened, the equation returns the benefit B_x of the *highest* row x opened at or below k.

Let $L(k, D)$ return the row-index of the highest row opened at or below compartment k given a decision vector D: $L(k, D) = \max\{k' : d^i_{k'} = 1 \text{ and } k' \leq k\}$. Therefore, for $k \geq 2$,

$$B_{L(k,D)} = d^i_k B_k + \sum_{x=1}^{k-1} d^i_x B_x (\prod_{l=x+1}^{k} (1 - d^i_l))$$

and $B_{L(1,D)} = d_1 B_1$; for convenience, we define $B_0 = 0$. Using this definition,

$$V_t(C, i) = \sum_{k=1}^{K} u_k B_{L(k,D^*)} + W_{t-1}(C).$$

Note that an upper-bound to $V_t(C, i)$ is $\sum_{k=1}^{K} u_k \max_{k' \leq k}\{B_{k'}\} + W_{t-1}(C)$. That is, the value-to-go is less than if you chould arbitrarily choose from *any* B_k below it, rather than only opened rows. This is an upper-bound since we have removed a constraint.
Let D be a solution achieved by performing DSKPAI. We show via contradiction that it must be the case that the algorithm produces the upper-bound value. It is sufficient to show that $B_{L(k,D)} = \max_{k' \leq k}\{B_{k'}\}, \forall k \in \mathbb{K}$. Suppose $\exists k \in K$, such that $L(k, D) \neq \arg\max_{k' \leq k}\{B_{k'}\} \equiv \tilde{k}$. If $\tilde{k} > L(k, D)$, then this contradicts the algorithm, since if $B_{\tilde{k}} > B_{L(k,D)}$, then row $\tilde{k}$ would be opened by the algorithm and $L(k, D) = \tilde{k}$. If $\tilde{k} < L(k, D)$, then this also contradicts the algorithm, since if $B_{\tilde{k}} > B_{L(k,D)}$, the algorithm would not have opened row $L(k, D)$. Hence, we know that $B_{L(k,D)} = \max_{k' \leq k}\{B_{k'}\}, \forall k \in \mathbb{K}$. Since the algorithm solution results in an upper bound on the value-to-go, $D = D^*$. □

Proof for Corollary 3.3.2:

We must show that in $t = 1$ for any C and i, that $D^{i\star}_1 = [1, 1, \ldots, 1]$. From Thereom 1, this is equivalent to showing that $B_K \geq B_{K-1} \geq, \ldots, \geq B_1 \geq 0$. $B_k = w^i p_k - (W_{t-1}(C) - W_{t-1}(C_{[k]}))$ and by the boundary condition, $W_{t-1}(C) = 0$ and $W_{t-1}(C_{[k]}) = 0$. Thus in $t = 1$, $B_k = w^i p_k$. Thus, this condition reduces to

$$\begin{aligned} p_K w^i \geq p_{K-1} w^i \geq, \ldots, \geq p_1 w^i \geq 0 \\ \Rightarrow p_K \geq p_{K-1} \geq, \ldots, \geq p_1 \geq 0 \end{aligned}$$

which is true since we have assumed ordered p_k. □

Proof of Proposition 3.3.3: We must show that given a time-invariant property on $\Delta W_t(C)$, for all t, C and i, that $D^{i\star}_1 = [1, 1, \ldots, 1]$. From Thereom 1, this is equivalent to showing that $B_K \geq B_{K-1} \geq, \ldots, \geq B_1 \geq 0$. First observe from the definition of the time-invariant property, for two rows l, k $(l < k)$ and an arrival item with weight w^i, either:
(1) $W_t(C_{[l]}) = W_t(C_{[k]}), \forall t$
(2) $W_t(C_{[l]}) \geq W_t(C_{[k]}), \forall t$ (with at least one strict), and $W_t(C_{[l]}) - W_t(C_{[k]})$ is increasing in t.
(3) $W_t(C_{[l]}) \leq W_t(C_{[k]}), \forall t$ (with at least one strict), and $W_t(C_{[k]}) - W_t(C_{[l]})$ is increasing in t.

To see the above cases, note that $W_t(C_{[l]}) = \sum_{j=1}^{t} \Delta W_j(C_{[l]})$, and $W_t(C_{[k]}) = \sum_{j=1}^{t} \Delta W_j(C_{[k]})$, and that from time-invarance definition we may have, for all t: (1) $\sum_{j=1}^{t} \Delta W_j(C_{[l]}) = \sum_{j=1}^{t} \Delta W_j(C_{[k]})$;

(2) $\sum_{j=1}^{t} \Delta W_j(C_{[l]}) \leq \sum_{j=1}^{t} \Delta W_j(C_{[k]})$; or (3) $\sum_{j=1}^{t} \Delta W_j(C_{[l]}) \geq \sum_{j=1}^{t} \Delta W_j(C_{[k]})$. Where for the two inequality cases there is at least one strict inequality. In addition, for the two inequality cases, we also can observe: $\Delta W_t(C_{[l]}) \geq (\leq) \Delta W_t(C_{[k]}) \Rightarrow W_t(C_{[l]}) - W_{t-1}(C_{[l]}) \geq (\leq) W_t(C_{[k]}) - W_{t-1}(C_{[k]}) \Rightarrow W_t(C_{[l]}) - W_t(C_{[k]}) \geq (\leq) W_{t-1}(C_{[l]}) - W_{t-1}(C_{[k]})$.

Next, we show that $B_k \geq 0, \forall k, t$. At $t = 1$, $B_k \geq 0$

$$\begin{aligned} &\iff w^i p_k - \Big(W_{t-1}(C) - W_{t-1}(C_{[k]})\Big) \geq 0 \\ &\iff w^i p_k \geq 0 \end{aligned}$$

and since w^i and p_k are positive, the statement is true at $t = 1$. Asymptotically in the time-horizon T, we have

$$\begin{aligned} \lim_{T\to\infty} W_T(C) &= \sum_{m=1}^{K} c_m p_m \\ \lim_{t\to\infty} W_T(C_{[k]}) &= \sum_{m=1}^{k-1} c_m p_m + (c_k - w^i) p_k + \sum_{m=k+1}^{K} c_m p_m \end{aligned}$$

and hence asymptotically in the time-horizon T, B_k approaches zero,

$$\lim_{T\to\infty} B_k = \lim_{T\to\infty} (w^i p_k - (W_T(C) - W_T(C_{[k]}))) = w^i p_k - w^i p_k = 0.$$

Finally, recall that from Proposition 3.2.3, we established that B_k is decreasing in time and so $B_k \geq 0, \forall t$. Next we show that for $l < k$, $\forall t$: $B_l \leq B_k$

$$\begin{aligned} &\iff w^i p_l - \Big(W_{t-1}(C) - W_{t-1}(C_{[l]})\Big) \leq w^i p_k - \Big(W_{t-1}(C) - W_{t-1}(C_{[k]})\Big) \\ &\iff w^i (p_l - p_k) \leq \Big(W_{t-1}(C_{[k]}) - W_{t-1}(C_{[l]})\Big) \end{aligned}$$

Clearly the left hand side is non-positive since $p_l \leq p_k$. If the above cases (1) *or* (3) holds, then $W_{t-1}(C_{[k]}) - W_{t-1}(C_{[l]}) \geq 0$ and the result holds.

If the above case (2) holds, then $W_{t-1}(C_{[k]}) - W_{t-1}(C_{[l]})$ is non-positive and decreasing in t. Thus $W_{t-1}(C_{[k]}) - W_{t-1}(C_{[l]}) \geq \lim_{T\to\infty} W_T(C_{[k]}) - W_T(C_{[l]})$. Since,

$$\begin{aligned} \lim_{T\to\infty} W_T(C_{[k]}) &= \sum_{m=1}^{k-1} c_m p_m + \sum_{m=k+1}^{K} c_m p_m + (c_k - w^i) p_k \\ \lim_{T\to\infty} W_T(C_{[l]}) &= \sum_{m=1}^{l-1} c_m p_m + \sum_{m=l+1}^{K} c_m p_m + (c_l - w^i) p_l \end{aligned}$$

it follows that $w^i (p_l - p_k) \leq W_{t-1}(C_{[k]}) - W_{t-1}(C_{[l]})$ and the result holds. □

Proof for Proposition 3.3.4: From Lemma 3.3.2 we have, for $t = 1$, $d_k^{i*} = 1$, for all i, k. Let $\hat{t}$ be the largest time-period $t \in [1, T)$, such that, for $t = \hat{t}$, $d_k^{i*} = 1$, and for $t = \hat{t}+1$, $d_k^{i*} = 0$; if $d_k^{i*} = 1$, for all t, then $\hat{t} \equiv T$, and the result holds and we are done. Observe that $\hat{t}$ is well-defined. For the remaining proof we assume $\hat{t} < T$. Thus, for $t \leq \hat{t}$, $B_K \geq \max\{0, B_1, \ldots, B_{k-1}\}$, and for $t = \hat{t} + 1$, $B_K < \max\{0, B_1, \ldots, B_{k-1}\}$. For $t = \hat{t} + 1$, define $l_{\hat{t}+1}$ to be the row with the largest marginal value, i.e. $l_{\hat{t}+1} = \arg\max\{0, B_1, \ldots, B_{k-1}\}$, where $l_{\hat{t}+1} = 0$ indicates $B_k < 0$. In order for the result to hold it is sufficient to show that for all $t > \hat{t}$, $B_k < B_{l_{\hat{t}+1}}$, which is equivalent to, for $t > \hat{t}$, $\delta_t^W(l_{\hat{t}+1}, k) > w^i(p_k - p_{l_{\hat{t}+1}})$. Recall we are assuming three properties for $\delta_t^W(l_{\hat{t}+1}, k)$. For case $(A1)$, $\delta_t^W(l_{\hat{t}+1}, k)$ is increasing in t and hence the result holds. For case $(A2)$, if the sequence has a single local minimum, it implies that $\delta_t^W(l_{\hat{t}+1}, k)$ decreases in time until some time-period and then increases. Since $\delta_0^W(l_{\hat{t}+1}, k) \leq w^i(p_k - p_{l_{\hat{t}+1}})$, and $\delta_{\hat{t}+1}^W(l_{\hat{t}+1}, k) > w^i(p_k - p_{l_{\hat{t}+1}})$, we know that the sequence must be increasing for all $t > \hat{t}$, and hence the result holds. For case $(A3)$, if the sequence has a single local maximum and $\delta_T^W(l_{\hat{t}+1}, k) > w^i(p_k - p_{l_{\hat{t}+1}})$, then it implies that $\delta_t^W(l_{\hat{t}+1}, k)$ increases in time until some period t' and then decreases. Since $\delta_{t'}^W(l_{\hat{t}+1}, k) \geq \delta_T^W(l_{\hat{t}+1}, k)$, and $\delta_{t'}^W(l_{\hat{t}+1}, k) \geq \delta_{\hat{t}+1}^W(l_{\hat{t}+1}, k)$, and there is at most a single local maximum we know, for all $t > \hat{t}$, $\delta_t^W(l_{\hat{t}+1}, k) > w^i(p_k - p_{l_{\hat{t}+1}})$ and the result holds. Hence $\tilde{t} = \hat{t}$. □

Proof of Proposition 3.3.5:
From DSKPAI the result holds if $B_K \geq B_{K-1} \geq \cdots \geq B_{\tilde{k}} \geq 0 > B_{\tilde{k}-1} \geq \cdots \geq B_1$. By the first case of the definition for $\tilde{k}$ we have $B_{\tilde{k}} \geq 0$, and $d_{\tilde{k}}^{i*} = 1$, and $d_k^{i*} = 0$, $\forall k < \tilde{k}$. Need to show that B_k is increasing in k: $B_k \geq B_{k-1}$

$$\begin{aligned}
&\iff w^i(p_k) - (W_{t-1}(C) - W_{t-1}(C_{[k]})) \geq w^i p_{k-1} - (W_{t-1}(C) - W_{t-1}(C_{[k-1]})) \\
&\iff w^i(p_k - p_{k-1}) \geq W_{t-1}(C_{[k-1]}) - W_{t-1}(C_{[k]})
\end{aligned}$$

which is by assumption and hence the result holds; if the second case of the definition for $\tilde{k}$ applies then the result holds trivially. □

Curriculum Vitae

Name:	Kyle Douglas Stevens Maclean
Post-Secondary Education and Degrees:	Ivey Business School,Western University London, ON 2008-2012 H.B.A.
Publications:	Maclean, K.D.S., Wilson, J.G., Krishnamoorthy, S. Forthcoming. Pricing of Excess Inventory on Groupon. International Journal of Revenue Management.
Teaching Experience:	Revenue Management 4520 2014, 2016 Ivey Business School Mean Overall Instructor Effectiveness: 6.2/7, 6.62/7.
Honours and Awards:	Ontario Graduate Scholarship 2014-2015
Related Work Experience:	Teaching Assistant The University of Western Ontario 2012 - 2017 Summer Investment Banking Analyst Citibank 2011 Workflow Consultant Thomson Reuters 2010 - 2012