

**SEQUENCE EFFECTS IN EVALUATING, SCHEDULING, AND DESIGNING
SERVICE BUNDLES**

A Dissertation

Presented to the Faculty of the Graduate School

of Cornell University

In Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

by
Michael James Dixon
August 2011

© 2011 Michael James Dixon

BIOGRAPHICAL SKETCH

Mike Dixon's interest in service operations management stemmed from multiple jobs in the service sector that have allowed him to see firsthand the impact that operational decision have on customer experiences. As dishwasher, cook, waiter, delivery driver, grocer, baker, and painter Mike worked his way through an undergraduate degree in Hospitality Management from Utah Valley University with a desire to open his own restaurant. By the time he graduated, he was an assistant manager of the university's food court and realized that he had much more to learn about managing a service concept. After graduation he managed a military food service establishment and later was hired as to manage a trendy food-prep start-up. During this time, he took a Quality Management course and was immediately hooked on the ideas of thinking about management from an analytical perspective. He earned his MBA from the University of Utah after which he worked at American Express first in a service-engineering and later a risk-management role. While at American Express, he developed an affinity for data analysis and learned valuable skills to analyze large archival data sets modeling customer behavior. Returning to earn a Ph.D., Mike attended Cornell University's School of Hotel Administration to further his understanding of service management. While at Cornell, his research interests developed to include finding ways to represent behavioral aspects of the design of a service in quantitative and analytical ways.

Mike is a married to Kendra Dixon and they have two small children. His hobbies include cooking, gardening, fly fishing, strumming the ukulele, and tapping drums. After graduation, he will take a position as an Assistant Professor of Operation Management at the Naval Postgraduate School in Monterey, California.

ACKNOWLEDGMENTS

Readers will be quick to notice that throughout this dissertation I used the pronoun “we” in place of “I”. This form of speech was intentional; not simply because it sounds less egotistical, but because I sincerely believe that without the guidance of key individuals this work would not be as it stands today.

Primary in providing guidance is my advisor Rohit Verma; Rohit has taken me by the hand and step-by-step taught me what it means to be an academic. Of his lessons that stand foremost in my mind is his example of respect and love that he has for his students, colleagues, and discipline. I am blessed beyond measure to count Rohit as one of my closest friends. His previous Ph.D. students joke that he is our academic father, but it is a small stretch for any of us to call him father. He is full of patience, persistence, and perspective. Thank you pitaji Rohit.

Second, large portions of this dissertation owe much to the patient personal tutoring of Gary Thompson. Looking back at how specific (and critical) parts of my analysis evolved I attribute much of it to the master teaching of Gary. I learned from him that the master teacher does not just provide the answer to the student — no matter how obvious and simple the solution might seem — but instead leads the student down the path to discovery. While this approach to teaching can be painful for both the student and the teacher, the end result is a more confident, complete student. My dissertation experience has been greatly enhanced by him allowing me to discover on my own what was probably obvious and simple to him. He continues to teach me the value of sharp thinking and consistent work.

Others that have taught me lessons that reveal themselves in this dissertation include Vishal Gaur with his quick understanding of how this research can be enhanced with proper

positioning; Nagesh Gavirneni with his constant confidence in me; Sherri Kimes and Chris Anderson with their sincere interest and feedback; Florio Arguillas and others from CISER providing invaluable help to researchers at Cornell; and David Just — his joint appointment as my econometric professor and ecclesiastical leader not only helped me understand the dangers of endogeneity, but provided me with an example of a balanced father.

My co-conspirator and a best friend in high school Alex Smart acted as a proofreader and editor for this dissertation. He has read my writing since I was 16 and has never told me I suck — everyone needs a friend like that. I am glad he appeared again in my life, not because I needed an editor, but because I needed a friend.

Going back even further in time, my parents never set upper limits on what I could achieve. My mom returned to college while I was in high school and we enjoyed taking calculus together; it was then that I knew that I could only hope to be as smart as her one day. My dad continues to teach me with his example that I can do big things and that big things require being able to do long, arduous, often times hard things.

To conclude, I save my most sincere and heartfelt acknowledgment to the love of my life and closest companion — my wife Kendra. Kendra has followed me across the country and supported through thick and thin (literally) with little thought for what was in it for her. I certainly hope that as we look back on our life there will be many more things that “we” accomplish together – nearly all more important than completing a dissertation. I love you forever Kendra.

We also acknowledge the Vienna Konzerthaus and Booze Allen Hamilton for providing data used in this dissertation.

TABLE OF CONTENTS

Biographical Sketch	iii
Acknowledgements.....	iv
Abstract.....	1
Chapter 1: Behavioral Research As a Major Agenda in Service Operations Management.....	2
Chapter 2: Sequence Effects in Service Bundles	27
Chapter 3: The Role of Scheduling and Bundling Flexibility in Creating Optimal Event Sequences	62
Chapter 4: Future Directions.....	97
Appendix 1: Logistic Regression Model Specifications.....	115
Appendix 2: An Illustration of Recommendations from Chapter 2.....	117
Appendix 3: Mathematical Representation of Interrelated Bundle and Scheduling Problem	122
Appendix 4: Discussion of Simulated Annealing Heuristic Algorithm Used To Solve the Interrelated Bundling and Scheduling Problem	138
Bibliography	156

SEQUENCE EFFECTS IN EVALUATING, SCHEDULING, AND DESIGNING SERVICE BUNDLES

Michael James Dixon, Ph.D.

Cornell University 2011

This dissertation addresses the importance of event sequencing as it impacts the customer experience and design of service bundles. We begin by building a case as to why operations management researchers must transition from historical analytical roots to include behavioral theory and practice in order to fully understand the complexities of operating in a service business. As an example of research that can take an operations management perspective on a behavioral issue, we study the design of event scheduling in the context of performing arts season subscriptions. In our first study, we investigate research in psychology and behavioral economics to develop hypotheses that correlate customer repurchase behavior to event utility sequences. Collectively, we refer to the impact of event utility sequences as *sequence effects*. We use six years of archival data from a renowned performing arts venue to develop an econometric model to test hypotheses; conclusions show that sequence effects are significantly correlated with customer repurchases of season subscription bundles indicating that event planners should consider event schedules and sequence effects as a part of service experience design. We propose a mathematical model that represents a multi-indexed integer programming problem that has an objective to optimize explicitly defined sequence effects across multiple bundles. To solve the problem, we develop a meta-heuristic algorithm that uses local search procedures to find near-optimal schedules. We use the algorithm to test the impact of event scheduling flexibility and bundling flexibility on sequence-effect-based scheduling efforts and find that, in our research design, event scheduling flexibility is more important than bundling flexibility when it comes to event schedule design. Finally, we address future direction of our research and propose that event schedule design, with objective to maximize customer experience, is a sub-discipline of service design with many avenues of available research opportunities.

CHAPTER 1:

BEHAVIORAL RESEARCH AS A MAJOR AGENDA IN SERVICE OPERATIONS MANAGEMENT

Abstract

Academic research in Service Operations Management is split between the traditional, analytical-rooted problems related to service companies and the complexities that human behaviors create in actual operations of a service firm. The purely analytical-quantitative approach, while important, often does not capture the complexities of human behavior. Similarly, behavioral-based literature may not apply appropriate tools for solving complex problems and thus may not provide applicable, prescriptive solutions to complex service operations problems. This chapter argues that by considering *both* behavioral and analytical aspects of a problem, researchers and managers can better understand how to apply behavioral research to a traditional operations problem. We explore the use of both empirical and analytical research methods on one such problem yet to be applied fully to operations practice: sequence effects. We discuss the complexities of doing so and introduce the following two chapters that subsequently utilize both empirical and analytical methodologies to applying behavioral research to service design.

Introduction

The study and practice of Service Operations Management (SOM) can be broadly divided into two approaches. The first approach is to consider a service process much like a manufacturing process and applying the same or altered tools and methodologies to control and improve operations: i.e. reduce costs, errors and rework while improving efficiencies and throughput. This approach is rooted in traditional Operations Management (OM) and Operations Research (OR) utilizing analytical, quantitative, mathematical-modeling tools in both research

and practice; so, we will refer to it as the “analytical” approach to SOM. In an early SOM piece, Chase (1996) proclaims that the “mall is his factory” and that until the 1970s all service operations work was essentially OR applied to a service setting. This early work replaces widgets and products with customers and servers, but does very little else to account for possible differences that a service might entail. This approach has led to improved queuing systems, scheduling, routing, etc. but often lacks the realistic assumptions that make service delivery complex. For example, a traditional bottleneck analysis of a family physician’s office may show that the pace of the doctor defines the speed and number of patients through the office. Traditional advice would be to remove any ancillary tasks from the duties of the doctor through extensive pre-work by other support staff. The result of this approach is that the doctor’s office is able to process more patients and each patient can expect a quicker throughput. However, the office runs the risk of *reduced* patient satisfaction because the actual face-time that each patient has with the doctor is dramatically reduced and patients feel “processed” instead of treated. By ignoring the response of patients to proposed changes in processes, OM might give advice that ignores important human element of services. In one study (Oliva & Sterman, 2001), researchers showed that a pure cost reducing, productivity increasing approach to service management leads to a downward spiral of worse perceptions of service quality, higher employee turnover, less demand and lower revenue. Rust and Bhalla (2010) claim that a cost reducing, efficiency approach to service management at the expense of revenue generation can be detrimental to customer lifetime value, i.e. the future earnings from a customer, thus impacting firm value.

The second approach of SOM is to focus on the human aspects of service encounters with customers. A service concept, according to Goldstein et al., (2002) is made up of *what* is to be done for the customer and *how* it is to be done. They make a case that these two factors need to

be integrated to achieve coherent, consistent service. In many cases, *what* is being done for a customer is emotional or experiential in nature, i.e., intangible. Therefore *how* an intangible concept should be delivered requires that service providers have an understanding of how a human might respond in order to design operational capabilities for that aim. Roth and Menor (2003) refer to the *what* as the “service concept” or the “multidimensional construct that embodies the totality of the service elements” (Roth & Menor, 2003, p. 150) and the *how* as the “service delivery system design” or the total efforts put forth to deliver on the service concept. The difficulty is in linking the design of the delivery system to the realization of expected service concept considering that customers play a vital role in the delivery of a system. Sasser et al. (1978) explain the complexity as follows:

A primary reason for defining the service product in terms of a total service concept is the role the process plays in creating the product. In purchasing a service, the consumer interacts with the workforce, equipment and physical environment that create the service. The process itself is, therefore, one dimension of the product. In contrast, the manufacturing process is isolated from the consumer and has an impact on the consumer only through what effect it has on the product. The elements of the manufacturing process are designed for the effective production of the physical good that is its output. The labor, equipment, and facilities are functionally designed with the cost and quality of the product being the primary criteria for evaluating how effectively these resources are utilized. In contrast, the service delivery system must be designed with the presence of the consumer in mind (Sasser et al., 1978, p. 14).

This simultaneous production and consumption may be further complicated with a need for successful, timely, customized production and distribution of a tangible aspect (e.g., food service, auto repair, dry cleaning, travel, dental work). The success of a service concept is dependent not only on the timely and expected delivery of the tangible expectation, but also on

how it was delivered. For example, in the airline industry it is important to arrive at the desired destination on time. Because of the large number of routes, planes, pilots, crews, and airports, airlines are required to incorporate a complex and dynamic plane routing capability to design and alter schedules. The airline routing process is done out-of-sight of the customer, yet this process is important to how customers perceive the capabilities and qualities of an airline. On the other hand, the interactive customer service that a passenger might receive from an attendant during a flight also leads to perceptions of airline service capabilities and quality. While the tasks of flight attendants and plane routing can both be considered operational capabilities, the approach needed to measure, improve, and research them are very different. The plane routing capability can rely on optimization and simulation algorithms to prescribe solutions to complex and vital scheduling problems while flight attendant effectiveness is largely a function of passenger interactions. Although scheduling routes requires a complex algorithmic approach, handling passengers cannot be considered any less complex because of the human behavior variability of passengers. The current capability of technology has yet to fully map out the complexity of human interaction. Therefore, handling baggage and handling passengers, although both operational functions, require different capabilities.

While this dichotomy seems obvious, the field of operations management has largely focused on applying operational knowledge to improve tangible capabilities at the expense of, or perhaps ignoring, the intangible behavioral driven capabilities. Voss et al. (2008) suggest that traditional operations strategy is not sufficient to manage in this environment and suggest a need for greater focus on the management of the intangible:

For the most part, OM research focuses on tangible and functional design from a service provider's perspective — making the delivery system and offerings more efficient

through following best practices. Much less attention has been given to the “intangible” customer side of service design and the dynamics of the service encounter (Voss et al., 2008, p. 252).

They continue that “many sociopsychological aspects of experiences as they pertain to operations and business strategy are not well understood” (Voss et al., 2008, p. 247). And that the service concept deserves both “an infusion of behavioral science *and* systems theory into the technical elements of the design...” (Voss et al., 2008, p. 250; emphasis added).

Even within SOM textbooks and courses the two approaches (analytical vs. behavioral) are at the mercy of the instructor’s bias. A quick review of a popular service operation text (Metters, King-Metters, Pullman, & Walton, 2008) provides the authors’ position in the chapter on Wait Time Management: eleven pages on queuing theory, two pages on psychology of queuing. While a full study of the bias of service operations text and instruction is outside the scope of this dissertation, it may be safe to say that there are differing opinions on what dimensions should be emphasized in considering operations management in a service setting. Still, there is a strong and growing population of scholars that believes human behavior is a topic worth considering in service operations management (Chase & Dasu, 2001; Pullman & M. A. Gross, 2004; Roth & Menor, 2003; Verma, Thompson, & Louviere, 1999; Verma, Thompson, Moore, & Louviere, 2001).

We contend that *both* analytical and behavioral approaches to SOM are necessary and useful in delivering on a service concept, but perhaps neither is sufficient in isolation. While traditionally the analytical approach to SOM has focused on tangible elements of service delivery and design, we posit that OM is aptly positioned to creatively apply analytical approaches to intangible behavioral aspects of service design and delivery. Given the proven

ability for OM researchers to consider large, complex problems (e.g., scheduling, routing, revenue management), we believe that the complexity of managing the intangible is not out of reach. In this paper, we further develop the need for behavioral based research in SOM and describe an example of applying an analytical approach to a behavioral aspect of service design and delivery.

Service Experiences

The linkage between a service concept and service delivery is more important in experience-centric concepts that —according to Pullman and Gross (2004, p. 553)— “occur when a customer has any sensation or knowledge acquisition resulting from some level of interaction with different elements of a context created by a service provider.” Pine and Gilmore (1999) coined the term “Experience Economy” as the next evolution of the service economy. They claim that as services become more and more efficient and effective they become commoditized and indistinguishable in the eyes of consumers. They state that in order to stay competitive, companies have to shift their strategy from cost saving, efficient delivery of a service to providing a significant, memorable, and unique experience. By providing experiences, firms create loyal customers that are eager to share their experience with others.

Pine and Gilmore and others (Grove & Fisk, 2001) compare these experience-centric operations to theater productions, comparing front-line servers to actors, physical surroundings to stages, and customers to audience members. Voss et al. (2008) conclude that OM takes on the role of a choreographer, carefully planning and supervising service delivery in order to evoke in the customer a specific emotional state at a specific time. The emotions that experience-centric

concepts try to convey are not always only ones of satisfaction or delight, but are much more subtle and detailed. About experience-based behaviors, Voss et al.(2008) write:

Experience-based behaviors arise from the uniqueness, knowledge, novelty, memorability, aesthetics, and entertainment that provoke customers' emotions, sensations, imagination, feelings, and perceptions... (Voss et al., 2008, p. 248)

The metaphor of a choreographer is useful in helping researchers understand the role operations management takes in a complex service environment that attempts to produce emotional experiences. Just like a choreographer, OM must deliberately and explicitly consider and define the setting, actions, timings, and sensory elements (sound, light, smell) that lead to a desired experience.

In a recent empirical case study, London Business School researchers explored how experience-centric firms design and manage service experiences (Voss & Zomerdijk, 2007; Zomerdijk & Voss, 2010). They found that these firms often used language to explain their customers "journey" through their experience and would design this journey to include "dramatic structures", i.e., dramatic sensory reactions. A main finding was that these firms spent much of their innovation effort on improving process or journey attributes rather than tangible product/service attributes, i.e., service innovation came not from having a new service product but from creating a new journey to experience. This evokes the cliché "Life is a journey, not a destination," but can illuminate how important the "sensory" aspect of operational capabilities is to service process innovation.

Psychology of Service Delivery

Chase and Apte (2007) recently identified what they termed the “Big Ideas” of service operations management and made a point of leaving out instances of OR applied in a service setting. They identified three main categories: (1) transference of industrial management concepts to services (e.g., McDonald’s production line, Disney’s industrialized fantasy), (2) frameworks for service design and management (e.g., customer contact model, service recovery, service profit chain), and (3) tools and techniques of service operations to improve productivity in services. They divide this last group into capacity planning (revenue management) and service quality tools (e.g., gaps model, poka yoke). After this exhaustive search of service operations research and applications, they conclude by providing future areas for research:

We argue that for service encounter research to be of maximum value, we must approach the *psychological side of services* with at least the same depth and rigor that we have traditionally approached the manufacturing of goods (Chase & Apte, 2007, p. 383; emphasis added).

Similarly, in defining a research agenda for SOM, Roth and Menor (2003) proclaim:

[R]esearch in human behavior is becoming more important. Thus, SOM scholar should investigate how specific strategic design choices can influence *psychological and consumer lifestyle differentiation* (Roth & Menor, 2003, p. 157; emphasis added).

In addition, Schneider and Bowen (2010) make strong points about the need for a behavioral point of view in the new Service Science movement. They claim that focusing on strictly technical and operational aspects of Service Science at the expense of the softer “people” side of services is a mistake:

[I]gnoring the social psychology of the various parties to service delivery and the setting in which they interact is dangerous to the long-term health of the service organization

because that is what can yield *sustainable* competitive advantage (Schneider & Bowen, 2010, p. 33; emphasis in original).

Chase and Dasu (2001, 2008) have led the charge on this attempt to gain a better understanding of the “psychological side of services” starting with their Harvard Business Review article entitled “Want to perfect your company's service? Use behavioral science.” Their research has led them to investigate the psychology and behavioral economics literature to consider applying theories from those fields into the practice of service delivery. Later in this chapter, we discuss their findings in more detail, but it is interesting to note that in the evolution of behavioral-based service operations literature, their cross-disciplinary investigation was unique and highlighted the need for service operations management to fully grasp the amount of human behavior research and theories that are available for application in a service delivery context.

Economics had arguably the first “behavioral” sub-field starting around the time Kahneman and Tversky, two psychologists, published their seminal work on Prospect Theory in *Econometrica* in 1979 (Kahneman & Tversky, 1979). Since that time, trying to understand how human behavior impacts assumptions in the traditionally analytical fields of Finance, Accounting, and more recently Operations Management has gained momentum. Researchers interested in Behavioral Operations Management have traditionally been interested in understanding the rationality of human decisions in traditional operations problems, (e.g., supply chain relationships, news vendor ordering decisions). Typically, the human behaviors being researched are managers or employees, and thus the field has a tight connection to human resources. Although SOM behavioral researchers are often interested in understanding the relationship between employees and service delivery, we are also interested in *consumer*

behavior. For that matter, we tend to be interested in how employee behavior relates to consumer reaction to that behavior. For this reason SOM is more akin to marketing than to human resources in our pursuit of knowledge.

In practice, the functional roles of operations and marketing are certainly blurred more than they are in academic circles. Karmarkar described a trend he noticed as early as 1996:

Traditional functional distinctions are not critical in lean, flat, downsized, reengineered, agile, and virtual organizations. Even if all firms cannot be described by that list of adjectives, it is apparent that boundaries between functions and firms are eroding. (Karmarkar, 1996, p. 125).

We would add experience-centric and customer-service oriented to the list of his original adjectives as the roles of designing, delivering, and describing experiences must cross traditional functional boundaries. Karmarkar (1996) claimed that in spite of the large volume of research in service marketing, operational considerations of simultaneous production and consumption, perishability of inventory, and variability of demand had to be considered in developing a service management research agenda. Further, he stated that operations must consider more than just cost issues in evaluating factors such as design, efficiencies and capacity.

Marketing and SOM Challenges

Chase (1996) claims that, among other thing, you might be a service “junkie” if you have friends in the marketing department. The field of service management has largely been led by marketing scholars because they do not hesitate to draw from behavioral based literature and create theory based on human interactions and perceptions. For example, service quality (Parasuraman, Zeithaml, & Berry, 1985), service blueprinting (Zeithaml, Bitner, & Gremler,

1996), service profit chain (Heskett & Sasser, 2010), and service dominant logic (Vargo & Lusch, 2004) have all been spearheaded by marketing scholars. While OM scholars have certainly made significant contribution towards a better understanding of service management, it is safe to say that the volume of service related work from operations minded faculty is very small compared to that of our colleagues in marketing departments.

Quite unlike operations management, marketing and other management fields are not only accustomed, but *expected* to deal with human behaviors in developing theoretical models that explain business interactions. Compared to a manufacturing process, a service process requires input from a customer in order to progress through a value-added process (Sampson & Froehle, 2006). In a large portion of service businesses, the customer input is the customer themselves (e.g., hospitality, healthcare, retail). Unlike a widget making its way through a factory floor, a customer in a process can add complex variability that is not easily handled. This variability can often be explained by those with an understanding of human behavior, and for this reason, marketing faculty have thrived in developing academic work in service management.

From a managerial standpoint, little is accomplished by delegating the human behavior element of a service experience to the marketing department given that operations will be in charge of executing the experience. The disparity of behavioral based research within operations management faculty should not lead to the conclusion that it is not important to the operations of a service. Menor, Tatikonda, and Sampson (2002) address the possible problems of a marketing centric new service development effort:

[T]he front end (which is classically Marketing-centric) can become isolated from the back-end (which is classically Operations-centric), leading to “over-the-wall” transfer of information and other dysfunctional organizational behavior...

The front and back-ends need to be understood as potentially different processes, but also must be simultaneously coordinated and integrated. The lack of such linkage could lead to inappropriate specification of the service concept; that is, service concepts which are not inherently executable or are resource inefficient (Menor et al., 2002, p. 146).

Appropriate service concept development is critical for designing and delivering customer value (Goldstein et al., 2002; Roth & Menor, 2003), but the integration between concept and delivery is far from seamless. Fynes and Lally (2008) attribute this to first, the difficulty of articulating desired intangible experiences; and second, the difficulty in translating articulations into design and delivery processes. They theorize that firms that focus heavily on articulating the service concept neglect the need for translation into practice and those firms that focus on delivery and processes do not see all processes leading to a holistic service concept, resulting in inappropriate operational decisions. Ideally, the functional roles of service operations and service marketing need to be blurred in order to strike an appropriate balance. Both academic fields of service marketing and service operations have provided tools to this end beginning with the traditional P's of marketing, Product, Price, Place, Promotion (McCarthy, 1972) extended to include People, Physical Evidence, and Processes (Lovelock & Wright, 1999) and Chase's emphasis on process design with customer contact in mind (Chase, 1978, 1981). However, it is our opinion that OM can add greater depth to the discussion by utilizing a higher degree of analytical rigor to integrate behavioral aspects into the design and delivery of a service concept.

The challenge for SOM researchers is to have a traditional understanding of the role, techniques, and common practices of OM *and* a level of understanding of consumer behavior. Additionally, traditional disciplinary methodologies might not be appropriate for SOM researchers; Roth and Menor state the problem as follows:

Service Operations Management research, in contrast to manufacturing [research], tends to focus on different problem types and uses different methodologies that those that are typically found in most POM doctoral curricula. Many service management problems are fuzzy and unstructured; are multidimensional and complex; and are less conducive to normative, analytical modeling (Roth & Menor, 2003, p. 146).

Because of the focus on behavioral issues and the non-traditional methodologies, traditional OM researchers often label SOM research as belonging to the marketing field. This labeling may be somewhat appropriate given our discussion above about the need for greater integration between the two fields, but it also creates an apparent neglect of service research in OM (Metters & Marucheck, 2007) and makes SOM research more difficult to publish in mainstream OM journals (Metters, 2010). According to Metters (2010), SOM research is held to a higher standard by editors and reviewers because services are so familiar to reviewers that assumptions are often not allowed that would be acceptable in a manufacturing-centered project.

What then is an SOM researcher to do? We propose that SOM researchers are uniquely positioned to make a substantial academic contribution distinct from our marketing colleagues and acceptable to the traditional OM population by applying the strengths of our analytical heritage to provide practical, prescriptive managerial insights about the application of behavioral leanings in a complex service concept.

Analytical Approach — Prescriptive Results

While the strict analytical work of early service operation management scholars often assumes away the complex human behaviors inherent in a service system, the behaviorist approach to service management driven by marketing academics often ignores the complexities of applying theory into the design and management of an intricate service process. The tools and

concepts that make up the practice of OM and OR are more than capable of handling extreme cases of complexities, but the behavioral research and theories are just beginning to be modeled in the analytical, mathematical ways accustomed to these researchers. The difficulty lies in quantifying behaviors and responses in order to fit them into a mathematical representation of a service system. An additional challenge is blending the behaviorist theories into the traditional OM and OR practices in order to approach complex problems in their entirety, i.e., to include the theoretic behavioral responses of humans into the realistic complexities that make up modern service processes.

The sister fields of OM and OR are not easily differentiated, but the major difference is that OR researchers seem to be more interested in the development of tools to solve problems, while OM researchers are more interested in the problems and subsequent solutions themselves, i.e., the application of the solutions. For example, while OR researchers are interested in thinking about how to solve complex routing problems faster and more efficiently, OM researchers might consider if the right problem is being solved. For this reason, OM researchers often use a more empirical data-driven methodology versus the analytical, mathematical-modeling methodology of OR researchers. The division between empirical and modeling researchers is apparent after just a short time in the community, but there is a need for both in developing and validating theory in operations management research. The argument for the need of empiricists versus rationalists is an old one that has been argued in depth in the philosophy field (Markie, 2008). Empiricism claims that knowledge is acquired from experience, while rationalists rely on reason. Since human behavior is known not to follow rational reason and since the field of OM is rooted in OR reason-based modeling, it is easy to understand the bias the OM field has against incorporating behavior aspects into theory and research. Most behavior based research is

empirically validated through the use of experiments designed to systematically observe changes in behavior under controlled conditions. Even the most creative and complex of experiments, however, pale in complexity to the design of a complex service concept. The rationalist modeling-based methodologies of OR, however, can lead to developing systems, designs, or ideas that are not easily observable because they do not yet exist, are too difficult to observe, or are very unlikely to occur. Therefore, a modeling approach to service design can capture and compute extreme complexity and, we posit, can be used to incorporate empirically validated behavioral results into a complex service delivery design.

Behavior Based Empirical and Analytical Research in SOM

Recent published research provides evidence that service operations management researchers can and should address consumer behavior issues. As early as 1999, Verma and Thompson (1999) introduced the importance of customer choices as it pertains to developing service attributes. Verma and his co-authors continued with a series of research projects addressing various aspects of customer choice and market utility as a tool to design services: e.g., customer choice as it pertains to capacity (Pullman, Goodale, & Verma, 2000) , multicultural customer segments choices as it pertains to service attributes (Pullman, Verma, & Goodale, 2001), the tradeoff between operational difficulty and customer choice (Verma et al., 2001) customer choice of self-service technology (Ding, Verma, & Iqbal, 2007) and more recently, customer choices as it pertains to revenue management pricing practices (MacDonald, C. K. Anderson, & Verma, 2010) and flow experiences in online financial services (Ding, Hu, Verma, & Wardell, 2010). Most of this stream of research has utilized discrete choice modeling as the methodology used to determine service attribute utility (Verma & Plaschka, 2003, 2005).

An analytical-modeling based stream of behavioral SOM research has begun recently to emerge within operations management journals. For example, Veeraraghavan and Debo have modeled the behavior of customers in choosing which queues to join if service quality could be determined by the queue length (Veeraraghavan & Debo, 2008). They utilize “herding behavior”, the tendency for humans to do what others are doing, as a way to explain customer choices between service offerings of unknown quality (Debo & Veeraraghavan, 2009). Others have modeled the response of strategic consumers in the trade-off of speed and quality in customer-intensive services (Anand, Paç, & Veeraraghavan, 2011), similar to the example given in the beginning of the chapter neglecting processing time in a healthcare setting. There are recent modeling-based research on the response of strategic consumers on retail assortment (Caro & Martínez-de-Albéeniz, 2009), communication of inventory levels to customers (Allon & Bassamboo, 2009), and revenue management practices (Jerath, Netessine, & Veeraraghavan, 2009, 2010; Liu & van Ryzin, 2008; Shen & Su, 2007). Most of these recent analytical papers have modeled their problem as a queuing model solving for a game-theoretic equilibrium to determine optimal design principles assuming specific customer behavior.

Another emerging line of consumer behavior focused SOM literature utilizes experiment design as a methodology. Victorino (2008) has begun to investigate customers’ perception of the degree of scripting across several types of services encounters by utilizing a unique video experiment design. Buell and Norton (2011) found that even with unfavorable results, the visual progress, or illusion of work, improves customer perception of online search results. McGuire and Kimes (2006) investigated perceived fairness of different wait-list management techniques using survey and experiment design.

These examples of recent empirical and analytical research points to the realization among researchers that customer behavior cannot be considered exogenous to the design of the service offering. Additionally, customer reaction to specific designs might be contrary to traditional OM practice and so therefore need to be more fully investigated and better understood by service designers. We present the current dissertation as a case study of research that lies at the intersection of empirical and modeling methodologies and analytical and behavioral theories. This dissertation attempts to borrow from the strengths of each camp. Behavioral aspects of a complex problem are empirically validated in order to develop and define theory. These behavioral complexities are mathematically modeled and included in a realistic problem in order to capture and discuss the difficulties of including them in an actual application. Operations research methodologies are used to solve the problem, and data driven experiments are tested to investigate possible managerial decisions. The problem will be well defined in later sections, but the major contributions of this dissertation are threefold: first, developing and testing theory concerning event sequencing as it affects human perception of a series of services; second, incorporating this theory into a mathematical representation of a complex problem; and third, exploring managerial insights that can be considered by solving the complex problem. While the research relies heavily upon the work of psychologists, marketers, behavioral economists, and other behavioral researchers in theory development, the outcome of the work focuses on the operational decision of scheduling and attempts to tie into the practical operations management research of scheduling, planning, service design, process design, and quality. Furthermore, the dissertation uses an equal amount of empirical and modeling methodologies as appropriate in order to address research questions.

The strength of this dissertation is its ability to incorporate behavioral aspects of a customer experience into an operations management framework, complete with analytical modeling methodologies. Incorporating behavior aspects into a proposed schedule allows service providers an operational advantage based on assumptions of expected psychological effects. This is accomplished by defining, quantifying, and translating expected psychological effects out of the language of the psychologist and into the language of a scheduler – mainly mathematical notation. After a proposed behavioral effect can be quantified, the OR techniques of optimization, simulation, mathematical modeling, etc. can be incorporated to investigate ramifications and prescribe solutions. This dissertation thus attempts to bridge gaps between behavioral theory and operational practice.

Sequence Effects

Chase and Dasu (2001, 2008) identified five principles of behavioral science theory that they determined should be applied to service design of delivery: (1) finish strong, (2) get bad parts over with early on, (3) segment the pleasure and combine the pain, (4) allow choices, and (5) build rituals. Of the five principles, three are specifically related to the timing of service delivery, and one (build rituals) might also relate to timing. These timing principles, which will be referred to as “sequence effects” throughout this dissertation, originate from research from psychologists and behavioral economists interested in understanding the importance of the ordering of pleasure, pain and other hedonic moments. Much of this work was initiated by the Nobel laureate Daniel Kahneman (Kahneman, Fredrickson, Schreiber, & Redelmeier, 1993; Redelmeier & Kahneman, 1996; Redelmeier, Katz, & Kahneman, 2003) and has been followed up by other famous names in psychology and behavior economics such as Dan Ariely (1998; Ariely & Carmon, 2000) and George Lowenstein (1987; 1993; Loewenstein & Sicherman,

1991). Sequence effects are exhaustively reviewed in the following chapter; however, the basic premise of sequence effects is that the order of the events during an encounter impacts participant's evaluation of the encounter. Humans don't place equal weight on all time periods of an encounter; most notably, we tend to place heavier weight on the end of an encounter and on parts that are more hedonically extreme, i.e., the most painful or most pleasurable part. Similarly, humans prefer upward trends and favor separating highly pleasurable events from one another.

The importance of these ideas in the design of service is twofold: first, an encounter can be planned and scripted to realize a specific sequence; and second, a series of encounters that a customer will experience over time can be scheduled to anticipate a specific sequence. This second design consideration is best characterized by time-elapsing¹ service bundles for which customers will have multiple encounters that are similar in form but may be different in utility or value. For example, season subscription ticket packages, a terms-worth of classroom lectures making up a course, or a multi-day tournament of sporting events. Compared to the disparate parts of an individual service encounter (e.g., a visit to a doctor includes checking-in, waiting, seeing a nurse, seeing a doctor, checking out — all very different in form), time-elapsed service bundles provide multiple encounters of similar form that we refer to as events (e.g., a season of baseball tickets includes a number of games that are all similar in form: i.e., all baseball games). Although similar in form, each event carries different value or utility and thus provides a natural palette to express the sequence effect theories of the behavioral researchers within the realms of service design. While sequence effects are certainly applicable to specific encounter design, this

¹ The term “time-elapsing” service bundle is used to define those bundles made up of individual parts that will be experience discretely over time as opposed to bundles made up of parts that might be used concurrently. For example, a telecommunications company might market a service bundle that includes telephone, internet, and cable services – this type of bundle is not “time-elapsing” as all three services could be used at once.

dissertation focuses on the sequence of discrete, similar-form events or encounters separated by time.

Television provides a practical example of sequence effect in action. Writers for television series have an inherent sense for sequence effect and use them liberally in creating tension in plot development and anticipation for further segments. A typical crime mystery plot will begin with a highly dramatic crime scene for which there is no apparent answer for the protagonists. Throughout the episode, detectives discover evidence, some of which is shocking, and commercial breaks are strategically placed at highly dramatic points in the episode to encourage viewers to sit through the commercials out of fear of missing something important. Finally, the tension builds and the protagonists solve the mystery just in the nick of time, that is, just in the final few minutes. The hook at the beginning is interesting enough to keep viewers from channel surfing, and the ebb and flow of the plot tension keeps us guessing until the very end— at which time tension is released and justice served.

Since service experiences are often compared to theater or story telling (Grove & Fisk, 2001; Pine & Gilmore, 1999) an analogy of a television series as a service experience is useful in order to consider the ways that service designers can better hold the attention of customers. In this analogy, we could consider one episode as an event and a season worth of episodes a bundle of events. Television producers hope that an episode is interesting enough to encourage viewers to return the following week. The writers try to create some sort of upward trend that will climax at the season-ending episode in order to encourage watchers to return to watch their show the following season. Similarly, the delivery of a service encounter or series of encounters could have a similar design, with hopes that the sequence will increase customer loyalty, spending, and perceived quality.

Although there is considerable research done in the area of sequence effects in psychology and other behavioral fields, very little work has tried to apply or test the theories in an actual service concept, except in the specific context of pain management in a healthcare setting. As noted before, the complexities of a service design will most likely far outweigh any complexities found in a controlled observed experiment. Additionally, applying the principles of the sequence effects into a complex service design might not be simple in the face of complex service design. The next two chapters address these concerns: Chapter 2 validates empirically—through econometric estimation—the correlation of sequence effects and customer repurchase decisions; and Chapter 3 proposes to incorporate sequence effects into in a realistic, complex scheduling problem.

Chapter 2 — Sequence Effects in Service Bundles

Sequence effects are valuable in the context of event sequencing if the order of events is correlated to or causes an increase in loyalty, revenue, demand, or perceived quality. Therefore, an early step in applying behavior sequence literature to service bundle scheduling is to validate that such a relationship exists; i.e., we must find evidence that customers behave differently having experienced different sequences. This step requires the researcher to be able to identify the behavior or attitudes over time. The first of our studies is an empirical investigation of the sequence effects in the context of a season subscription of performing arts events. The study develops theory that specific sequence of events within a season subscription should yield different repurchase rates while controlling for many other bundle, events and customer attributes. The expectation is to see that the placement of high and low utility events should make a difference in whether patrons are likely to repurchase a season subscription. Data for the project was provided by the Vienna Konzerthaus, one of the largest concert venues in Europe.

Six years of tickets sales and over one million transactions were used to create a binary choice model that estimates the significance of variables derived to represent various psychological effects. The complexities inherent in customer choice of repurchase are considered and controlled for and econometric parameter estimation provides evidence that the sequence of event utility significantly influences customer repurchase behavior.

Appropriately applied econometric estimation can help researchers understand the relationship of observed characteristics on choice outcomes; therefore, this project required us to translate expected behavioral effects developed by psychologists into operational, measurable comparisons from within an extensive and complex archival data source. We synthesized past — mostly experimentally tested — research results into service design theory that could be observed and compared given a sufficiently diverse and complex set of service offerings. The Vienna Konzerthaus provided us with an ideal dataset that included multiple subscription bundles that can be tracked over the six years with varying event sequences across bundles and across years. Although we do not claim the project to be flawless, it does provide a good case study in translating behavioral theory into observable, service delivery design (i.e., operational) characteristics of an archival dataset. Given the seemingly endless capacity of data collection capabilities across modern service companies, this type of translation should continue to prove useful to empirical researchers attempting to test behavioral theories in the context of varying service concepts.

Chapter 3 — Optimal Event Sequencing

The empirical investigation of Chapter 2 provides us with ample evidence that event scheduling considering sequence effects is an appropriate method to increase customer

repurchase and loyalty in the context of performing arts season subscription bundle design. However, as with most complex service designs, the Vienna Konzerthaus service offering is complex enough that applying sequence effects into a scheduling effort may not be trivial. The re-ordering of events within a bundle impacts a master schedule of events across all bundles and a more proper sequence might be created by considering the events within the bundle. There are reasonable restrictions placed on when events can be scheduled and with the bundles to which they can belong. The purpose of Chapter 3 then is to define a way that this sequence-effect based scheduling is to be done and to explore various research questions that doing so might encourage.

Wherein Chapter 2 required us to translate behavioral theories into measurable and observable characteristics of data, in Chapter 3 we explicitly define an “optimal sequence” mathematically requiring all aspects of sequence effects to be mathematically represented in terms of event date schedules, bundle membership and venue placement. A non-linear integer model is developed to schedule events into bundles, date / time slots, and venues (event halls) with an objective of maximizing explicitly defined sequence effects across all bundles within a season. The result is an analytic scheduling and bundling model with a behavior theory driven objective. The complexity of the problem becomes apparent as notation allows for events to be in multiple bundles and as constraints across multiple dimensions are developed to ensure realistic representation of the complexity of bundle and delivery design.

Due to its complexity, the problem is approached using a meta-heuristic search procedure that can robustly and quickly solve realistic sized (200 events, 50 bundles, 300 dates) in a matter of minutes. With an efficient way to solve this complex problem, we are able to ask questions that will provide managers with an understanding of the application of sequence effects across

different scenarios. We specifically explore and compare the ramifications of flexibility in event bundle membership and date/time constraints. Our findings lead to us to and understanding of how different event flexibilities can impact sequence effect optimality, thus helping event planners make decisions with holistic service concept design in mind.

Similar to the project in Chapter 2, we believe that this project serves as a good example of how to translate behavioral theory into operational decision making, this time via mathematical modeling and application of OR tools. The project provides service designers with a prescriptive understanding of how to apply behavioral theory into service design and how service design decisions can impact the service concept. Because our problem is complex, solving it gives us a measure of confidence that sequence effects can be applied appropriately in other, less complex contexts. Furthermore, the development of the solution will allow us to explore many other aspects of service design.

Chapter 4 — Future Directions

The final chapter of this dissertation discusses the areas that naturally lead from our findings. Most notably, there are a number of different design issues that can be addressed by testing different types of problems in the meta-heuristic algorithm developed in Chapter 3. We discuss the weaknesses of both our studies and propose characteristics of an ideal study to continue exploration on the topic of sequence effects.

Conclusion

The objective of this chapter was to position the following dissertation in the context of current SOM research; mainly, an attempt to apply and further our understanding of a set of

behavioral theories in the context of service concept, design, and delivery. We believe that SOM is well positioned to translate behavioral aspects into operational representations and provide guidance on how to appropriately prescribe solutions. In so doing, SOM researchers are able to address managerially relevant complexities and provide researchers and managers alike with an understanding of the ramifications of the complexities in service design and delivery. The following two chapters provide a case study of how SOM can achieve this end using both empirical and analytical methodologies.

CHAPTER 2

SEQUENCE EFFECTS IN SERVICE BUNDLES

Abstract

Past research in psychology and behavioral economics have shown that the sequence of events or interactions play an important role in the way individuals evaluate experiences. It has been shown that the pleasure or pain associated with the peak event, the last event, the general trend, and spread of the events are important in predicting overall memory of an experience. In this paper we investigate whether the sequence effects within a service bundle impacts customer repurchase behavior. Using an extensive archival database provided by a renowned performing arts venue, we build and test an econometric model to predict season ticket subscription repurchase and determine if the temporal placement of events impacts repurchase. We find evidence of peak, end, trend, and spread effects and discuss the importance of sequence in determining service design and scheduling. These results have implications for effective service design and capacity planning for a wide range of service industries.

Introduction

Scholars have suggested that the sequence of events within a service encounter can influence customer's overall perception of the quality and satisfaction associated with the service (Chase & Dasu, 2001, 2008; Cook et al., 2002). Specifically, Chase and Dasu (2001) suggest various strategies for service sequencing including placing the lowest point or bad news at the beginning of the encounter, ending the service on a high note, and improving the experience over time. While these ideas have intuitive appeal, to our knowledge they have not been empirically

validated. The value of service sequencing on future customer behavior (e.g., repurchase) or operations (e.g., scheduling of events; capacity planning) have also not been explored.

In a related research stream, it has been shown that different attributes of a service are not equally important to the customers, i.e., customers place different weights or utility on various elements of a service (Verma et al., 1999). However, the past research has not explored if the temporal aspects of an element of a service (i.e., sequence) also have unequal utility. Consumer behavior scholars have theorized that the underlying values of consumption of a service or product not only include functional values, but also conditional, and emotional values (Sheth, Newman, & B. L. Gross, 1991). Behavioral research suggests that temporal sequencing influences these non-functional attributes in such a way as to significantly influence perception.

Effective service design involves developing a service concept that appeals to end users considering operational constraints (Verma et al., 2001). Furthermore, past research has emphasized that operations management's role in designing a service concept involves understanding "what" should be done and "how" it should be done (e.g., Goldstein et al., 2002). While the methods and frameworks to accomplish the "how" of a service concept are in abundance, the often unasked questions within "how" is "when" i.e., does the delivery sequence of the service concept have an impact on customers' experiences?

In this paper, we investigate how customer repurchase behavior is impacted by the temporal placement or sequence of events within a service bundle. Using a comprehensive multi-year ticket purchase database from a world-renowned performing arts venue, we test the impact of event sequence on customer repurchase of subscription packages. Specifically we identify that the placement of high-utility events and the trend of the event utilities impact the probability of

subscription repurchases. Furthermore, we illustrate how the estimated weights for sequence parameters can be used to make better operational and marketing decisions.

The rest of the paper is organized in the following manner: first, we provide a review of literature related to service bundling and sequence-related behavioral research; second, we present our theoretical framework and hypotheses; third, we describe our research design and analysis approach; fourth, we present our results and associated discussion; and finally we discuss theoretical and managerial implications of this research.

Service Bundling

In this paper, we address the temporal sequence of events within the context of a *service bundle*, i.e., a combination of a number of different services sold in one package. Product and service bundling is a heavily researched topic in marketing (Gaeth, I. P. Levin, Chakraborty, & A. M. Levin, 1991; Guiltinan, 1987; Harlam, Krishna, Lehmann, & Mela, 1995; Stremersch & Tellis, 2002). The practice is common across many service industries, for example fast food industries offer meal packages, telecommunications and cable companies offer packages with several different services at one price, and performing arts venues sell season subscriptions that include tickets to a number of events. Some service bundles are created by bundling a number of different services that are intended to be used simultaneously, or concurrently. For example, for one monthly charge telecommunication firms provide internet, cable television, and home telephone service as a service bundle that is typically used concurrently. Other service bundles are created by placing similar discrete services together in a way that they have to be experienced across time or sequentially; we refer to these bundles as time-elapsing. For example, a course taught over 12 weeks may have 12 separate class sessions, a cruise ship package includes 5 days

of separate experiences to different locations, or season ticket sales for performing arts or sporting events includes a number of different events experienced across a season. Within this second type of service bundles, the event sequence of some bundles is constrained, e.g., the 5 day cruise typically visits islands in a physically linear fashion. However, in other service bundles the sequence is not assumed fixed or at least not entirely fixed, e.g., the schedule of performances within a performing arts season subscription can be altered. Time-elapsing service bundles provide ideal testing grounds for applying sequence related behavioral research in the context of service design and scheduling because the sequence of the discrete segments can be changed. Further, time-elapsing service bundles tend to lead to service relationships, i.e., customers form relationships with service providers because there are several different encounters over time.

Different hierarchical levels of bundling effectively act as a pricing rate fence, for example a cell phone company that bundles phone, IM, and internet access can charge different prices for different combinations of bundles. Thus, operations management researchers to date have primarily concerned themselves with revenue management or pricing issues surrounding product and service bundling (Aydin & Ziya, 2008; Bitran & Caldentey, 2003; Bitran & Ferrer, 2007) and supply chain issues of supplier bundling or product mix purchasing (Rosenthal, Zydiak, & Chaudhry, 1995; Schoenherr & Mabert, 2008). From an economic perspective, customers purchase bundles because their reservation prices for all individual elements are met, i.e., the actual price for highly demanded elements is lower than the reservation price so the surplus is transferred to the less desired element of the bundle. Revenue management principles suggest that in order to optimize revenue on bundled services, the bundle should include both high demand and low demand services. To reach overall capacity maximization, managers would

do best to separate the most popular events into different bundles reaching a higher capacity for the less popular elements. Leveraging highly popular elements is at the cornerstone of revenue management with bundled services.

In a related research stream, a number of procedures to find “optimal” product and service attribute profiles have been developed to find an attribute mix that maximizes sales, market share (Green & Krieger, 1989; Ho & Zheng, 2004; Shocker & Srinivasan, 1979), or profit (Green & Krieger, 1991; Moore, Louviere, & Verma, 1999; Morgan, Daniels, & Kouvelis, 2001; Raman & Chhajed, 1995). Other researchers have developed attribute mix optimization models while considered operating constraints such as capacity (Pullman & Moore, 1999), production costs (Moore et al., 1999), waiting time and labor scheduling (Pullman et al., 2000), and operational difficulty (Verma et al., 2001). This stream of research has contributed to an understanding of consumers’ choice of product and service attributes; however, to our knowledge, none of the optimization models have considered the sequence related attributes of service delivery.

Sequence-Related Behavioral Research

Based on a review of past behavioral research, Chase and Dasu (2001) proposed that among over things, customers remember three aspects of a service experience:

- 1.) The trend in the sequence of pleasure and pain*
- 2.) The high and low points*
- 3.) The ending*

These three aspects have been researched heavily in psychology and behavioral economics and are called Trend Effects, Peak Effects, and End Effects respectively.

Trend Effects

Generally speaking, individuals prefer a sequence of events that improves over time (Loewenstein & Prelec, 1993). For example, Ross and Simonson (1991) demonstrated that gamblers prefer to first lose \$15.00 then subsequently win \$85.00 over first winning \$85.00 then losing \$15.00. Although the net gain is the same, the trend in the sequence of winning seems to impact the utility of the overall win.

In a legal research article (Walker, Thibaut, & Andreoli, 1972) researchers found that the presentation sequence of different pieces of evidence impacts the overall judgment. The sequences that start with weak evidence and ends with strong evidence generally yield the most favorable judgments. In another study, Loewenstein and Prelec (1993) describe an experiment which asks participants to choose between visiting a good friend one weekend and an abrasive aunt another weekend. A majority choose to postpone the (good) friend and visit the (abrasive) aunt first. They explain this behavior as a tendency to want to *savor* good outcomes by postponing them and quickly get through bad outcomes to eliminate a feeling of *dread*.

Similarly, other studies have shown that, all else being equal, an increasing wage profile is preferred to a declining or flat one (Loewenstein & Sicherman, 1991). Ariely (1998) describes an experiment in which participants were asked to rate their pain under different pain sequence profiles inflicted with the aid of a calibrated vice squeezing the participant's hand at different pressures. Profiles that started with high pressure and decreased over time rated much lower (56 out of 100) than those that started low and increased over time (75 out of 100). As we begin to

adapt to the most recent stimulus, an improvement feels like a gain while a worsening move can feel like a loss. According to prospect theory, we tend to be more sensitive to a loss than a gain hence we exhibit a preference for upward trends (Kahneman & Tversky, 1979).

Peak and Spreading Effects

Researchers studying memory have found that human minds are more prone to selectively capturing and remembering the snapshot of extreme high or low points (i.e., peaks) from a past experience rather than recording every detail of their lives (Burt, Mitchell, Raggatt, Jones, & Cowan, 1995; Nguyen & Belk, 2007). Furthermore, the intensity and sequence of an experience seem to be more important than the duration of the experience. For example, Redelmeier and Kahneman (1996) discovered that the overall pain experienced by a patient is highly correlated with the highest degree of pain for patients during colonoscopies regardless of duration, e.g., patients whose colonoscopy lasted 1 hour compared to those whose colonoscopy lasted 15 minutes experience similar overall pain highly correlated to the *peak* pain felt.

In situations with multiple high points, Loewenstein (1987) identified a *spreading effect* explained by a preference to spread out preferred outcomes in a sequence. When participants were asked to choose between sequences with two good outcomes and one mediocre outcome (two fancy dinners and one dinner at home), a majority choose to separate the good with the mediocre. In a follow up study, Loewenstein and Prelec (1993) asked subjects to schedule 2 future weekends to use a pair of hypothetical \$100 coupons to a restaurant. When subjects were told they had two years to use the coupons they spread out their plans using the first, on average, at week 8 and the second on week 31. Thaler and Johnson (1990) showed that people think they will be more happy if two positive events are temporally separated than if the same two events

are temporally close (winning 2 lotteries on the same day vs. separated by a week). Chase and Dasu (2001) recommend that service businesses consider segmenting pleasurable aspects of an encounter and combine the painful segments. They state that most people would prefer to win two \$5 gambles as opposed to one \$10 gamble essentially spreading out the winning episodes. The spreading effect ensures that a sequence is well “covered” by positive events.

End Effects

In a clinical trial Redelmeier et al. (2003) prolonged the less painful, yet still uncomfortable *end* of colonoscopy procedure for some patients and compared the assessment of pain for these patients against those of other patients. The results showed that the overall pain assessment was lower for the experiential group. Similarly, those patients whose most *intense pain* (peak) was near the end of the procedure reported higher overall pain. Similarly, in his calibrated vice experiment, Ariely (1998) found that pressure profiles that started low and ended high, resulting in lower total pressure, had statistically equivalent ratings as a control group that had a consistent high pressure. This *end effect* reveals that the end of an experience impacts remembered utility (Kahneman, Wakker, & Sarin, 1997).

Serial position effects explain that the presentation sequence impacts memory (Ebbinghaus, 1902). Researchers have shown that when presented with a list of nonsense words to memorize, subjects displayed two types of serial position effects: *primacy*, or the ability to better recall the first items, and *recency*, the ability to better recall the last items. Primacy and recency have been found to form impressions and influence decision making (N. H. Anderson & Barrios, 1961; Asch, 1946). More recently, researchers have found that subjects rely heavily on their initial reference point in decision making. This effect has been termed an *anchoring* effect

because the initial reference acts as an anchor that is not often or easily adjusted (Ariely, Loewenstein, & Prelec, 2003; Tversky & Kahneman, 1974). Recency, primacy, and anchoring suggest that what is remembered and used to form impression is at the beginning or the end.

Marketing researchers have used the above ideas in explaining how customer expectations are formed and how satisfaction with a product or service is expressed (e.g., Oliver, 1980; Parasuraman et al., 1985). Within the operations management literature, sequence effects have been less researched. In their seminal book *Service Breakthroughs: Changing the rules of the game*, Heskett, Sasser and Hart (1990) discuss the idea of “service bookend” and emphasize the need for services to provide not only a strong ending, but also a strong beginning mirroring the ideas of primacy, recency, anchoring, and spreading effects. Similarly, Johnson (1995) proposes that exceeding customer’s expectation early in an encounter is more likely to delight customers throughout the service encounter because customers are primed to see good service. As stated earlier, Chase and Dasu (2001, 2008) are the pioneering operations management scholars to suggest that behavioral research ought to be considered in service design; however, they do not provide any additional empirical evidence. They, however, propose that an upward trend and a strong ending are more important than a strong beginning (Chase, 2004). Other researchers have shown through experimentation (Hansen & Danaher, 1999) and service content analysis (Verhoef, Antonides, & de Hoog, 2004) that an upward trend of sequence performance leads to higher perception of quality and satisfaction; however, these studies only tested for a change in performance level across a fixed sequence, not for changes in the sequence of the process itself, i.e., the service process remained unchanged and only the performance levels changed. Other scholars (Bolton, Lemon, & Bramlett, 2006) have shown that more recent service encounters as well as “extra mile” or extremely favorable experiences influence system support

service contract renewals. More recently, Bitran, Ferrer, and Oliveira (2008) further refine a conceptual framework of duration in a service encounter and how it applies to profitability. They cite behavioral literature as it applies to duration and the sequence of an encounter and conclude by calling for more varying techniques of empirical based evidence across different industries and context.

Our research adds to the past multi-disciplinary literature by testing the presence of sequence effects by econometric modeling. Furthermore, we are interested in temporal event placement within a service bundle i.e., we hope to uncover the effect that a change in the *sequence* of events might have on customers, not just the change of the *performance levels* over time of a fixed process. Finally, we provide insight on how sequence effects may be used in event scheduling by searching for sequence effects in a service bundle that elapses over a long period of time.

Theory and Hypotheses

The sequence literature reviewed above suggests that the sequence of events ought to impact customer evaluation of a service. The decision to repurchase a repeating service bundle is based largely on the evaluation of the previous experience with the service (LaBarbera & Mazursky, 1983; Rust & Zahorik, 1993). The evaluation of a product or service is in large part a function of evaluating multiple attributes that make up the offering (Gensch & Recker, 1979). In this study we are attempting to find evidence to support the theory that sequence effects, as attributes of a service bundle, impacts future customer behavior. At a highest level, we propose that sequence effects influence customer evaluations of service bundles which, in turn influences customer behavior.

Sequence (End, Peak, Spreading, Trend) Effects → Evaluation of Service Bundle → Future Customer Repurchase Behavior

The link from sequence effects to customer evaluations may be through customer satisfaction, service quality, or by adding value, but we leave the complete causal model to future research. However, we posit that sequence effects should be a proxy for customer evaluations and so should influence customer behaviors in predictable ways that will be outlined in our hypotheses.

Within consumer behavior research there is a proposed model of consumption that identifies independent product and service attributes constructs that lead to utility (Sheth et al., 1991). These include functional, conditional, and emotional attributes. Functional value is gained from the functional or utilitarian aspects of the physical attributes. Conditional value is added when the conditions are right for consumption of the product or service. Finally emotional value is added by arousing feeling or affective states. Sequence effects may increase value through conditional means by scheduling the right event at the right time, i.e., an event in a non-ideal time slot will not add value as much as it would in a more appropriate time. Similarly, sequence effects can influence emotional values by more positively influencing the affective state of an individual through considering trend, peak and end effects.

Random utility theory and the corresponding discrete choice modeling approach suggest the utility of an alternative (e.g., service) is based on the characteristics of individual decision-maker (e.g., customer) and the attributes (e.g., price, quality, brand name, etc.) of the alternatives (Ben-Akiva & Lerman, 1985; Luce, 1959; McFadden, 1980). In particular, research suggests that after acquiring information and learning about possible alternatives, decision-makers define a set of determinant attributes to use to compare and evaluate alternatives. After comparing available

alternatives with respect to each attribute, decision-makers eliminate some alternatives and form a final choice set containing a few alternatives. They then form impressions of each alternatives' position on the determinant attributes, value these attribute positions vis-à-vis one another (i.e., make tradeoffs), and combine the attribute information to form overall impressions of each alternative.

The random utility theory assumes that individuals' choice behavior is generated by maximization of preferences or utility. Louviere (1988) defines utility as "judgments, impressions, or evaluations that decision makers form of products or services, taking all the determinant attribute information into account." The idea of utility maximization and its relation to human choice behavior is not new. For example, McFadden (1986) quotes from a 1912 economics text by Taussig: "An object can have no value unless it has utility. No one will give anything for an article unless it yield him satisfaction. Doubtless people are sometimes foolish, and buy things, as children do, to please a moment's fancy; but at least they think at the moment that there is a wish to be gratified."

The typical customer choice model can be expressed as: $Prob_i = f(\text{customer's attributes, alternative's functional attributes})$. In words, the probability of choosing alternative i is a function of the customer's attributes and the alternatives (or product and service functional) attributes (Ben-Akiva & Lerman, 1985). In this manner, service design researchers have shown that service providers can gain valuable information about what type of customers are drawn to their offering and what these customers prefer.

When predicting a repurchase of a service, we assume that including the evaluation of the sequence effects in addition to customer attributes and product and service attributes will

improve the prediction. Since we believe that sequence effects should proxy customer evaluations we theorize that including sequence effect attributes into a probabilistic choice model will improve the model's fit. We propose that an econometric prediction of customer repurchase will be improved with the inclusion of variables that represent the sequence of the utility of discrete events within the bundle. This hypothesis can be tested by comparing nested models, i.e., comparing the fit of an estimated model that lacks sequence related variables against a model that includes them.

***H1:** Prediction of customer repurchase will improve significantly by considering sequence attributes above and beyond just considering customer characteristics and product (goods and/or service) features.*

This first hypothesis is our primary concern, simply put; we hope to find the sequence of events matters in the repurchase decision. We follow up this primary research question with less general hypotheses that investigate specific sequence theories.

The colonoscopy related research (Redelmeier & Kahneman, 1996) found that a patient's perception of overall pain was influenced significantly by the peak pain level suggesting that as the peak event increases in utility, its impact is more pronounced. These lines of research suggest that the peaks are more remembered and influence customer's evaluation more than other events which leads to our next hypothesis. Note that we consider the event with the highest utility to be the peak, i.e., our notion of "peak" is positive (utility) instead of negative (pain), but we expect similar results.

***H2a:** Customers are more likely to repurchase as the peak event utility increases.*

Similarly, the colonoscopy research found that the pain level at the end of the procedure also significantly impacted overall perception. The end effect suggests that the last event of a sequence impacts customer evaluation, and so we predict that as the utility of the last event increases its effect will remain salient in customer's minds and will result in a higher overall assessment.

***H2b:** Customers are more likely to repurchase as the last event utility increases.*

Combining the peak and end effects, colonoscopy research found that by extending the end of the procedure so that the peak pain was further from the end led to improved pain evaluations. Similarly, procedures that ended shortly after the peak pain had worse evaluations. In our case, we expect that the placement of the peak event near the end of a sequence should positively influence assessment. A bundle ending with a peak should result in a higher overall assessment of the bundle, e.g., if the last event in a season subscription package includes an all-star cast of musicians performing traditionally crowd pleasing pieces, then the patrons will remember the event and give high marks to the entire subscription. On the other hand, if the all-star performance occurs further from the end of the season, the subsequent, less exciting events may diminish the utility of the peak experience thus lowering the overall assessment.

***H2c:** Customers are more likely to repurchase as the peak event nears the end of the sequence.*

Finally, Chase and Dasu (2001) suggest that as a sequence improves over time the feeling of loss is avoided and customer evaluations improve. We predict that an upward trend of event utility should impact customer evaluation positively.

H2d: Customers are more likely to repurchase as the trend of the events utility over time increases.

Research Design

In order to test the proposed hypotheses we estimate a series of econometric models that predict the probability that a customer who had purchased a given service bundle for a given time period, would again purchase a bundle from the same cycle the subsequent time period. (In this paper we refer to purchase of each service bundle as a “cycle” and a time-period as a “season”).

Specifically, for the set of customers C who bought cycle j season t , we are interested in predicting whether or not each customer will buy cycle j season $t+1$, i.e., the same cycle the subsequent season. The unit of analysis is individual customers who purchased a given cycle the previous season and our dependent variable is binomial: 1 if the customer purchased the same cycle the subsequent year, 0 if they did not. Since our dependent variable is binary, we have chosen to model the data using logistic regression. Our econometric model uses the following form,

$$\ln\left(\frac{P(Y_{cjt+1} = 1)}{1 - P(Y_{cjt+1} = 1)}\right) = \beta X + \varepsilon \quad | \quad Y_{cjt} = 1 \quad (1)$$

where $Y_{cjt+1} = 1$ represents a repurchase of bundle j in season $t+1$ (the next season’s bundle for the same cycle) by customer c , X is a vector of predictors, β is the vector of coefficients including an intercept, and ε are the errors. This model is estimated across all customers i who purchased bundle j in season t . The model predicts the log-odds of repurchase given the set of independent variables using a maximum likelihood estimator assuming the distribution of errors

follow a logit distribution. Described in more detail below, the independent variables include customer characteristics, service attributes, and sequence-related variables.

Data Description

To test the proposed hypotheses we use a multi-year subscription ticket purchase database for an internationally renowned performing arts venue (location disguised in blind format of the manuscript). This concert venue houses 5 concert halls that can be used simultaneously. The venue hosts approximately 300 events per year and offers over 40 different subscriptions to its customers. The database includes 6 years (*seasons*) of ticket sales data from 2001 to 2007 including over 1 million individual ticket sales transactions for more than 2,400 events purchased by over 50,000 unique customers. The database includes the date and time of the ticket purchase, the price paid, membership status of the customer during time of purchase, general seating category (based on price category), and whether the ticket was purchased as a part of subscription. Additionally, we are given details about all the events such as the date and time of the event, the genre of the event (out of 16 possible genres), and the specific concert hall used for the event. Finally, we have limited customer specific information that is optional when creating an account with the venue: gender, title, degree held, postal code, etc.

The subscriptions offered by the venue are theme based cycles offered year over year. Most cycles are based either on a certain genre or are specific to a particular ensemble. Themes based on genre alone include Jazz, Classical Symphony, Music and Film, Piano, Children's Music, etc. Other themes include Rising Stars, International Orchestras, International Quartets, Beethoven, Original compositions, etc. The cycling nature of the subscriptions allow us to link subscription bundles year to year to determine if a given customer repurchased the same cycle

the next season. For the purpose of terminology we will now refer to a subscription *cycle* as a theme based subscription that can be tracked year over year over several *seasons* and a subscription *bundle* as a specific season in a subscription cycle. In the six years of data we find 41 subscription *cycles* that can be tracked for all six *seasons* for a total of 246 subscription *bundles*. There are other subscriptions cycles that do not span over all six seasons, but for reasons forthcoming, they are left out of the analysis.

Customer Specific Variables

In predicting repurchase, three general sets of variables are considered: first customer specific attributes, second bundle specific attributes, and finally sequence specific variables. We are not primarily interest in customer and bundle specific attributes, but they are included in the model to act as control variables. Additionally, our main hypothesis states that by including sequence attributes our model should improve; therefore, we compare models that include sequence variables with those that do not.

Customer specific attributes include gender, seating category of tickets (seat placement), number of bundles purchased (for a given bundle, not across all bundles) total number of unique bundles purchased for the season, days from purchase date to first event in the bundle (measure of how early a bundle was purchased), and membership status. Since we are predicting the purchase of cycle j season $t+1$, we will derive the above mentioned variables from ticket sales data for season t . Additionally; we have created a variable to determine the customer's loyalty with the bundle. We have classified customers into four groups and subsequently predict that the groups can be thought of as ordinal in their likelihood to repurchase. The first group consists of those customers who have purchased the given subscription cycle for the past 3 seasons; we

Table 2.1: Descriptive Statistics of Customer Attributes Variables

	Non - Repurchasers	Repurchasers	All Customers
Days from purchase to first event*	115.72	130.70	126.41
Bundles Purchased* (for bundle j)	1.74	1.76	1.75
Unique Bundles Purchased* (in season t)	1.94	2.14	2.08
Gender			
Male	38%	41%	40%
Female	8%	7%	7%
Unknown	55%	52%	53%
Member	20%	10%	13%
Non-Member	80%	90%	87%
Experience			
Loyal	24%	69%	56%
Potential	4%	2%	2%
Fickle	16%	13%	14%
New	57%	17%	28%
Seating Category 1**	35%	27%	29%
Seating Category 2**	16%	17%	17%
Seating Category 3**	21%	18%	19%
Seating Category 4**	17%	17%	17%
Seating Category 5**	12%	12%	12%
Seating Category 6**	11%	12%	12%
Seating Category 7**	5%	5%	5%
Seating Category 8**	1.2%	2.2%	1.9%
Seating Category 9**	0.3%	0.4%	0.3%
Total	9108	22708	31816

* averages reported

** Percentage of customers who purchased subscriptions from a given price category.

** Percentage do not sum to 100% because some customer purchased from multiple price categories.

** Seating Categories start with lowest priced seats (category 1) and ascend to highest price seats (Ca

named these customers *Loyal*. The second group consists of customers who have purchased a given cycle for the past 2 seasons, but not 3 seasons; we name these *Potential* as in “Potentially Loyal”. The third group is named *Fickle* and is made up of customers who have purchased a given cycle one season ago and three seasons ago, but not two seasons ago. They are fickle because they are not consistent in repurchasing. Finally the last group is called *New* and is made up of those customers who have purchased the cycle for only one season. By calculating the loyalty variable we set a limit on the data that can be used in the model. We begin with predicting the fourth season ($t=4$) since season 1 would be season $t-3$, season 2 would be season

$t-2$ and season 3 would be season $t-1$. Still, with this restriction we are left with data for seasons 4, 5, and 6 for which have 44 cycles giving us 128 bundles (40 cycles with 3 seasons + 4 cycles with 2 seasons). Within those 128 bundles we find a total sample size of $n = 31,816$ customers who had purchased a given cycle the previous season. Given the total size of the dataset and the resulting sample size for the model, we are satisfied with reducing the data in order to derive the loyalty variables.

In our final model estimation, we excluded a random 10% of the observations to use to validate the accuracy of the model. Further, we identified and excluded 1 outlier observation that proved to be a significant influence on the model estimation; it is described in more detail in the Appendix section. Table 2.1 shows a summary of the customer specific variables.

Bundle Specific Variables

Both marketing and operations management researchers consider product and service mix as an important aspect of customer satisfaction, perception, intention, and subsequent choice processing. Product and service mix is the set of attributes for a given product and service, e.g., a hotel property might include an exercise facility, a pool, a restaurant, wireless internet, and concierge service; a credit card might have fraud protection, online account access, automatic bill pay and cash back rewards; a car might have good gas mileage, five cup holders, moon roof and Bluetooth capability. Service providers have to choose what attributes to include in their offering in order to entice the right customer to purchase. In the case of the concert venue, management must create bundles of subscriptions that include attributes such as the number of events in the bundle, the genre mix of the events, and the percent of events on weekend (Friday – Sunday) vs. weekday, and the percentage of non-matinee events vs. matinee (before 5 pm).

Adding to the list of bundle specific variables, we include a measure of *total bundle utility* calculated as the sum of all the individual event utilities - event utility calculations are described in the next section. This variable can be thought of as a measure of the total number of events within the bundle as well as the relative popularity of the subscription as a whole.

Determining the Utility of an Event

Researchers have dissected a service's utility into the individual parts and attributes of the service. In our context, we would like to determine the utility that customers receive from each event within a subscription, e.g., if there are eight events within a subscription we want to determine utility of each event. Perhaps the most appropriate measure of utility would come by asking each customer to rate the performance at the end of each show, but unfortunately we do not have access to such data. Instead, we have formed two measures that we use separately to represent event utility at an aggregate level.

The first measure of event utility that we use is the event's *average ticket price*. As a part of the essential "P's" of marketing, price communicates to customers the company's intended value for its product or service (Kotler, 2002). From a customer's perspective, the price of the product must exceed or at least equal any expected value derived from the product or service. Therefore, it is in the company's best interest to set its price in relation to the value delivered and perceived by its customers. If the price is set too high the customer will not buy and if it is set too low the company is foregoing potential profits. It is for this reason, we assume that the concert venue does a fair job at pricing the events such that its price close to the value customers expect to experience from the event.

Kotler (2002: page 487) explains that as customers get a feel for the actual quality as opposed to perceived quality of a product the price plays a smaller role representing quality. However, as an actual measure of quality is unknown to a customer, the price is the primary signal used to determine expected quality. In the case of concert events, it may be difficult for customer to assess the quality of a specific event especially if it is an artist or performance that is unknown to the customer.

The second measure of utility is a measure of both seat occupancy and ticket price:

Revenue per Available Seat (REVPAS). REVPAS is calculated by dividing the total revenue for each event by the total number of available seats for the event.

ticket price_{ce} = the price that customer *c* paid for event *e*.

available seats_e = the number of available seats for event *e*.

$$REVPAS_e = \frac{\sum_c ticket\ price_{ce}}{available\ seats_e} \quad (2)$$

REVPAS is adapted from the revenue management field for which some measure of revenue per available unit is maximized. For example, revenue per available room (REVPAR) is used widely in the hotel industry and has been shown to be highly correlated to customer satisfaction (Davidson, 2003; Davidson, Manning, Brosnan, & Timo, 2001), service quality (Kimes, 1999, 2001), and brand loyalty (H. B. Kim & W. G. Kim, 2005; H. B. Kim, W. G. Kim, & An, 2003). Consumers are often uncertain about the quality of a hotel property (as they are with performing arts events) and use price as a signal for expected quality. Consistently demanding higher prices and filling more rooms suggests that hotel with high REVPAR are able to deliver on the expectation set by the price signal resulting in repeat business, loyalty, and

positive word of mouth. Not unlike hotel brands that have proven their value, some performers or performances can demand higher prices and fill more seats because they provide a higher valued event.

Our data does not provide us with the means to derive an individual customer level

Table 2.2: Descriptive Statistics for Bundle Attributes Variables

Total Subscription Bundles Considered		128
	Average Bundle	Standard Deviation
Genre Mix*		
Ancient Music	3%	15%
New Music	4%	19%
Jazz	2%	15%
World Music	5%	21%
Children's Music	19%	39%
Literature	2%	15%
Organ Music	2%	14%
Piano Music	8%	26%
Chamber Music	23%	41%
Vocals	7%	21%
Choral Music	0%	1%
Orchestra	21%	37%
Film	3%	15%
Number of Genres in Bundle	1.33	0.74
Percentage of Events on Weekends	52%	33%
Percentage of Events in the Evening	80%	37%
Total Bundle Utility	REVPAS 93.23 (72)**	Average Ticket Price 133.82 (75)**

* represents percentage of genre in all bundles

** average(standard deviation)

$$\text{Total Bundle Utility} = \sum_{\text{event} \in \text{subscription}} \text{Utility}_e$$

where $\text{Utility}_e =$

$$\text{REVPAS}_e = \frac{\sum \text{ticket price}_{ce}}{\text{available seats}_e}$$

$$\text{Avg ticket price}_e = \frac{1}{N_e} \sum_c \text{ticket price}_{ce}$$

measure and so we choose to test our hypotheses with aggregate measures. Certainly this is a weakness of our model from an individual customer's perspective and is not ideal in deriving a choice model; however, from the stand point of the service provider, an aggregate measure is needed to implement a scheduling methodology based on our results. We assume that event schedulers forecast aggregate demand for each event and set prices accordingly. The forecasts are based on a combination of past attendance data and industry trend knowledge. Because the forecasts are derived from the same data we use in deriving event utility, they can then be used to sequence the events according to the results of our model.

Similarly, an individual level utility measure can also be rolled up to make aggregate forecasts, and its advantages include being able to create bundles that target a specific segment for which the aggregate utility represent poorly. We leave for future research target market bundle creation based on individual utility measures, but with this exploratory research we are content with assuming that the aggregate measure reflects a fair starting point in investigating the presence of sequence effects in our context. Table 2.2 show descriptive statistics for bundle variables.

Sequence Variables

The sequence variables are of primary interest in this model as they will be used to test our hypotheses. Recall from the previous section that we have two measures of utility that will be calculated for each event. We identify the event with of the highest utility within a subscription and capture its utility as the *peak event utility*. Additionally, the *last event's utility* is considered. Additionally, we measure the number of *days from the peak event to the last event*. To consider the trend of the sequence of events, we calculate the *utility slope* for the line fit in

ordinary least squares regression through event utilities and the number of days from the beginning of the bundle.

Finally, we have created variables to indicate if bundles include a true peak, a valley or if the events in the bundle are relatively homogenous in utility. To determine these categories we plotted the event utilities across time for each subscription and coded bundles that appeared to have a peak, a valley or neither. After coding, we observed that those with a peak or a valley had a range of utility that was at least greater than 10 (with REVPAS); 10 corresponded closely to the 75th percentile of ranges for all bundles. Those bundles with ranges less than the 75th percentile were then coded *Flat*. For the remaining bundles, we calculated the average utility within a bundle and compared it to the peak event utility and the valley event utility. If the difference from the peak to the average was greater than that from the valley to the average, then the bundle was coded as *Peak*. If the opposite was true, the bundle was coded *Valley*. Figure 2.1 shows an example of each of the three categories and Table 2.3 shows the summary of the Sequence Attributes.

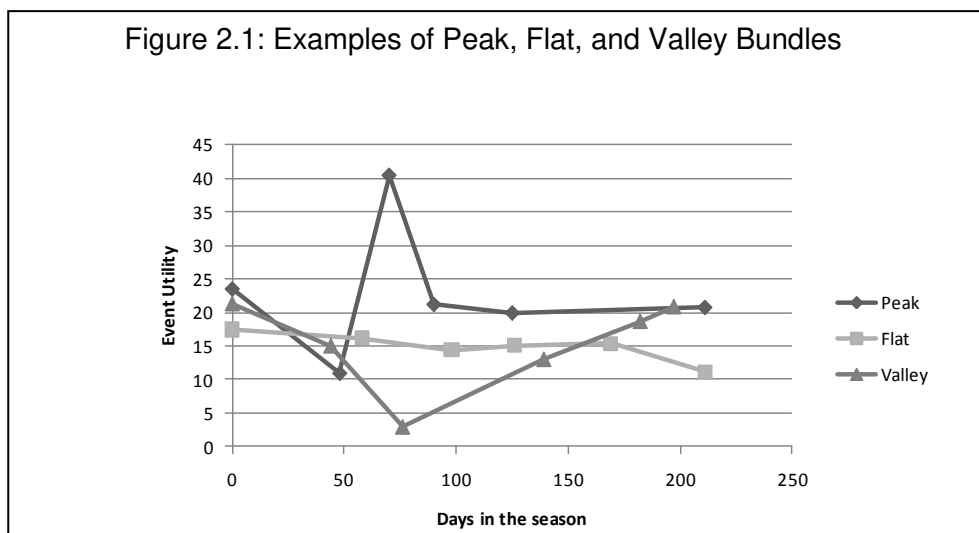


Table 2.3: Descriptive Statistics of Sequence Attributes

	REVPAS	Average Ticket Price
	mean(std dev)	
Peak Event Utility	19.26 (13.9)	25.48 (14)
End Event Utility	14.49 (10)	21.07 (9)
Slope of Sequence	0.0011 (0.02)	-0.0003 (0.02)
Days from Peak Event to End Event	104.39 (77.4)	93.86 (77.7)
Count of Subscriptions		
Days from Peak Event to End Event	REVPAS	Average Ticket Price
0 Days	18	23
1 to 50 days	20	25
50 to 100 days	18	18
100 to 150 days	23	18
150 to 200 days	20	19
200 to 250 days	29	25
Flat Subscription	95	96
Peaked Subscriptions	28	22
Valley Subscriptions	5	10
Total Subscriptions	128	128

$$Slope\ of\ Bundle_b = \frac{\sum_{e \in E_b} (utility_e - avg_utility_b)(daysfromfirst_e - avg_daysfromfirst_b)}{\sum_{e \in E_b} (utility_e - avg_utility_b)^2}$$

Results

Due to the large number of variables available to predict repurchase, we have chosen to create 3 models nesting the three main variables types: customer specific, bundle specific, and sequence specific. Nested model comparisons can be used to determine if adding additional variables leads to an improved fit. The following three models are estimated:

Model 1: Customer Specific Variables

Model 2: Customer Specific and Bundle Specific Variables

Model 3: Customer Specific, Bundle Specific, and Sequence Variables.

Since we use two measures of event utility, and variables derived from event utility are found both in bundle specific and sequence specific variables, models 2 and 3 are estimated twice – once for each utility calculation. Because of the panel nature of our data (same customer over several time periods) we include a fixed effect for season by adding two dummy variables for seasons 4 and 5. This will control for unobserved homogeneity within each season. To control for unobserved homogeneity within customers, we estimated the model by adjusting for standard errors using Huber White Robust (sandwich) errors clustered on customer ID. The results of the models are shown in Table 2.4 and the model specifications are discussed in the appendix.

Recall that the customer and the bundle attributes are not the primary concern for this study. We are interested in Model 1 and 2 primarily in comparison to Model 3. Therefore, we will only briefly discuss their results. The customer attribute model shows intuitive results:

- The coefficient for the number of days from purchase to the first event is positive indicating that customers are more likely to repurchase if they buy their tickets early.
- The more subscriptions purchased (both within the subscription and across the season) the more likely the customer is to repurchase.
- Males are more likely to repurchase compared to females.
- Customers that are also Members are more likely to repurchase than non-members.
- Compared to *New* customers, *Loyal*, *Potentially Loyal*, and *Fickle* customers are all more likely to repurchase. Surprisingly *Fickle* customers are more likely to repurchase than *Potentially Loyal* customers.

Table 2.4: Logistic Regression Results

n=28708 Standard errors adjusted for 9312 clusters in customer ID		Model 1: Customer Attributes Model	Model 2: Customer and Bundle Attributes Model		Model 3: Customer, Bundle, and Sequence Attributes Model	
	Parameter	All	REVPAS	Average Ticket Price	REVPAS	Average Ticket Price
	Intercept	-0.702**	-0.082	0.083	-0.187	-0.296~
	Season 4	-0.028	-0.017	-0.026	0.007	0.012
	Season 5	0.014	0	-0.003	-0.013	0.031
Customer Attributes	Days from purchase to first event	0.002**	0.003**	0.003**	0.002**	0.002**
	Subscriptions Purchased	0.02	0.098**	0.098**	0.102**	0.103**
	Total Subscriptions purchased in the season	0.018	0.046**	0.045**	0.048**	0.046**
	Gender M vs. F	0.061~	0.028	0.028	0.028	0.03
	Gender Unknown v F	-0.014	-0.034	-0.033	-0.03	-0.029
	Membership	0.351**	0.203**	0.216**	0.183**	0.188**
	Loyal vs. New	2.175**	2.022**	2.033**	2.015**	2.028**
	Potential vs. New	0.386**	0.358**	0.359**	0.376**	0.368**
	Fickle vs. New	0.974**	0.902**	0.908**	0.898**	0.911**
	Seating Category 1	-0.296**	-0.063	-0.078	-0.073	-0.053
	Seating Category 2	-0.112*	-0.086	-0.095~	-0.099~	-0.074
	Seating Category 3	-0.186**	-0.145**	-0.152**	-0.163**	-0.122*
	Seating Category 4	-0.092~	-0.115*	-0.123*	-0.121*	-0.103~
	Seating Category 5	-0.046	-0.032	-0.037	-0.044	-0.013
	Seating Category 6	0.019	0.006	0.001	-0.012	0.015
	Seating Category 7	0.047	-0.014	-0.024	-0.028	-0.01
Seating Category 8	0.391**	0.365**	0.355**	0.363**	0.351**	
Seating Category 9	-0.062	-0.016	-0.017	-0.001	-0.016	
Bundle Attributes	Ancient Music vs Orchestra		0.044	0.009	.0276**	0.168
	New Music vs Orchestra		0.141	-0.007	0.413**	0.159
	Jazz vs Orchestra		-0.249**	-0.334**	0.024	-0.266**
	World Music vs Orchestra		-0.286**	-0.395**	-0.169~	-0.226*
	Children's Music vs Orchestra		-1.539**	-1.697**	-1.46**	-1.479**
	Literature vs Orchestra		-0.478**	-0.578**	-0.37**	-0.422**
	Organ Music vs Orchestra		-1.007**	-1.174**	-0.681**	-0.948**
	Piano Music vs Orchestra		-0.454**	-0.515**	-0.297**	-0.451**
	Chamber Music vs Orchestra		-0.156*	-0.241**	-0.142*	-0.157*
	Vocals vs Orchestra		-0.021	-0.071	-0.061	0.021
	Choral Music vs Orchestra		-0.288	-0.217	-0.957	-0.771
	Film vs Orchestra		-0.13	-0.27*	0.064	0.026
	Total Number of Genres in the Subscription		-0.204**	-0.222**	-0.204**	-0.219**
	Percent of Events on Weekend		0.029	0.047	0.119	-0.023
Percent of Events in the Evening		-0.613**	-0.59**	-0.806**	-0.72**	
Sum of all Events Utility		0.001**	.001~	0	0.001~	
Sequence Attributes	Peak Event Utility				0.012**	0.011**
	Last Event Utility				0.009*	0.003
	Days from Peak event to Last Event				0.001*	0.001*
	Slope of Event Utility Over Time				2.545**	0.694
	Peak subscriptions vs. Flat Subscriptions				-0.404**	-0.443**
	Valley subscriptions vs. Flat Subscriptions				0.092	0.017
		20 DF	36 DF	42 DF		
Coefficients of Determination						
Pseudo R-Squared		0.171	0.184	0.184	0.186	0.185
Nested Model Comparison Statistics						
-2 log Pseudo Likelihood		28,536	28,083	28,096	28,017	28,052
Predictive Accuracy - Calculated with observations excluded from model estimation: n = 3107						
Brier Score		0.1649	0.1617	0.1620	0.1616	0.1621

** Significant at alpha < 0.01

* Significant at alpha < 0.05

~ Significant at alpha < 0.10

- Customers who purchase higher priced seats (Seat Categories) have a higher likelihood of repurchase.

The customer attributes in the second and third model retain their sign and general magnitude. The new variables introduced in the customer and bundle model show the following results:

- Compared to Orchestra, nearly all genres have negative estimated coefficients indicating lower likelihood of repurchase.
- As the number of genres in a bundle increases, repurchase likelihood decreases indicating that on average, mixed genre bundles do not fare as well as single genre bundles.
- As the percentage of weekend events in a bundle increases, repurchase is more likely.
- As the percentage of evening events in a bundle increases repurchase is less likely.
- Total bundle utility (sum of all event utility) is significant (with REVPAS) and positive (with both measures of utility) in the 2nd model, indicating that as the total bundle utility increases, repurchase likelihood increases. However, when the sequence variables are introduced in the 3rd model, the total bundle utility variable loses significance, indicating that total bundle utility can be better explained with the sequence variables.

Hypotheses Testing and Discussion

Our primary hypothesis, by including sequence variables the model will improve, can be tested by comparing nested model comparison statistics. We can see that the models improve as they progress as the -2 pseudo log likelihood are decreasing as more variables are added. Using the difference in degrees of freedom across the models we can create a hypothesis test to determine if the added variables in the model significantly add to the fit of the model. Table 2.5

shows that comparing Model 2 to Model 1 there is evidence that the added variables improved the model ($p < 0.00001$). Similarly, going from Model 2 to Model 3 (within utility type) there is evidence that the sequence variables also improve the model's fit significantly ($p < 0.00001$) providing support for H1.

Table 2.5: Likelihood Ratio Test for Nested Model Comparison

$$(-2 \text{ Log Likelihood}_{\text{model 1}}) - (-2 \text{ Log Likelihood}_{\text{model 2}}) \sim \chi^2$$

$$\text{df} = \text{df}_{\text{model 2}} - \text{df}_{\text{model 1}}$$

	Model 1: Customer Attributes Model	Model 2: + Bundle Attributes		Model 3: +Sequence Attributes	
	All	Revaps	Average Ticket Price	REVPAS	Average Ticket Price
-2 Log Psuedo Likelihood	28,536	28,083	28,096	28,017	28,052
DF	20	36	36	42	42
Likelihood Ratio		453	440	66	44
Degrees of Freedom		16	16	6	6
Pr > ChiSq		<.00001*	<.00001*	<.00001**	<.00001**

* comparing Model 2 with Model 1
** comparing Model 3 with Model 2 within utility type

This conclusion indicates that the sequence variables, as a whole, significantly impact the repurchase behavior of the customers in our dataset regardless of the utility measure that is used. As discussed in the hypotheses development section, we believe that the sequence of a service will contribute to its value in turn impacting its evaluation and ultimately the behavior of the customer. In our dataset, we had no direct measure of customer evaluation, but this result implies that by impacting repurchase decisions, the sequence variables can act, at some level, as a predictor of perceived value and customer evaluations and hence are an important aspect of service design.

The remaining hypotheses can be tested by considering the estimated parameters of the sequence variables. The coefficient for the Peak Event Utility is significant ($p < .01$) and positive

under both REVPAS and average ticket price, indicating that H2a (customers are more likely to repurchase as the peak event utility increases) is supported. The coefficient for the Last Event is significant for REVPAS ($p < .05$) and positive, indicating that H2b (customers are more likely to repurchase as the last event utility increases) is supported. The coefficient for Days from Peak Event to Last Event is significant ($p < .05$) and positive for both REVPAS and average ticket price, indicating that as the peak event is further from the last event, repurchase is more likely. This result contradicts H2c (Customers are more likely to repurchase a subscription as the time from the peak event to the time of last event shortens) and is discussed in detail in the next section. The coefficient for the bundle slope is significant with REVPAS ($p < .001$) and positive, indicating that as the utility of events improve over time (positive upward slope) repurchase probability increases, providing support for H2d.

Managerial Implications

As predicted, the effects of the peak event utility and the last event utility play a significant role in predicting repurchase; as the utility of the peak event and last event increase, so too does repurchase. The non-trivial finding is that the likelihood of repurchase does not increase as the peak event nears the end, rather the likelihood increases as the peak event gets *further* from the end (or closer to the beginning). This can be explained by considering the spreading effect that has been found to be present in sequences with more than one desirable event. When asked to plan two nights in the future to eat at a fancy restaurant, most people would prefer to separate the dinners over time (Loewenstein, 1987). The same seems to be occurring in our case, mainly that the last event should have high utility and that the peak event (highest utility event) should be near the beginning, spreading out the two desirable events. In the sequence of pain literature, the peak or highest amount of pain was best placed further from

the end as well; the explanation was that the highest pain would be less-remembered if it was further from the end. In our case, the peak is best suited further from the end for reasons of spreading pleasure across a longer period of time.

The spreading effect that we have found is bounded by the positive coefficient for the slope, i.e., if an early peak placement creates a negative slope, the probability of repurchase will decrease instead of increase. This means that a peak event should be placed as far away from the end so as to maximize marginal benefit from the spreading effect without decreasing the benefit derived from the trend (slope) effect. In our data, this usually means that the ideal spot for the peak event is rarely the first event, but can often be the second or third event (see illustration of scheduling optimization in appendix for an example). This suggests two things, first that the trend or momentum that is inherent in progressively increasing event utility lead to improved sequences, and second that the first event and hence first impressions do not need to be (perhaps should not be) the climax of the bundle. Intuitively it would be ideal to set the initial expectation as low as possible by placing a low utility event first followed closely by the peak event with a drop in utility after building up to a high ending. This structure allows the customer's initial expectations to be surpassed, it allows for the spreading of high events (peak and end), and allows for an overall positive trend.

Of interesting and unexpected note, the coefficient for *Peak Subscriptions* vs. *Flat Subscriptions* is significant ($p < .01$) and negative, indicating that customers are more likely to repurchase a bundle that is homogenous as opposed to one that has a peak. However, there is no evidence that repurchase probability differs between flat subscriptions and valley subscriptions. By changing the reference category from *Flat* to *Valley* the coefficient for *Peaked* vs. *Valley* is significant ($p < .01$) and negative (-0.49 for REVPAS), indicating that repurchase decreases for

peak subscription compared to valley subscription. The result that “peaked” bundles fared worse than “flat” bundles adds further complexity to our findings. Just as the number of days from peak was bounded by a preference for positive slopes, the preference for high peak and end utilities is bounded by the preference for bundles made up of relatively homogenous events in terms of utility. This suggests that customers are not impressed by highly leveraged bundles that include one highly-popular event bundled with less-popular events. This strategy may improve short term occupancy and ticket sales numbers; however, it does not appear to lead to sustainable repurchase rate. To the extent that highly popular events occur, we provide two suggestions: 1.) Bundle all high utility events together; or, 2.) Spread out high utility events to bundles two at a time and spread them out in order to optimize the peak, end, trend, and spreading effects.

The scheduling implication for subscriptions that have a clear peak should be to simply schedule the events from worst to best with the last event being the peak. In this manner the end effect and trend effects are maximized. However, if there appears to be multiple high utility events or if all the events are homogenous in utility, the best approach would be to rank-order all events from worst to best and then move the peak event to the second or third position. In this manner, the spreading effect is maximized while the end and slope are very close to optimal. See the appendix for an example of optimizing these two types of schedules.

Conclusions, Limitations, and Future Research

At the highest level, this research has provided a degree of empirical support for the peak, end, trend and spreading effect theories set forth by previous researchers. Uniquely, we find evidence that these effects can be found in long sequences that elapse over an entire subscription

season while past research has been focused on single interactions. Additionally, the model shows that scheduling sequence decisions may impact repurchase behavior of customers.

Although our research may not be completely generalizable, we believe that the effect of utility based scheduling can be realized outside the context of performing arts; certainly scheduling sporting events, conferences, courses, and tour packages have similar bundling attributes that make them akin to scheduling based on estimated utilities. Profit-maximizing managers price higher demanded events accordingly and so already have a feel for what events have higher utility. The unsophisticated approaches of scheduling that we proposed suggests that a simple scheduling solution can lead to an improved customer experience leading to higher repurchase rates and higher profits. We estimated an average of 2% increase of repurchase probability across all customers in one season (see appendix). The specific recommendation that we have made above may not be applicable to all service bundles, but designers of all services should consider the sequence an attribute worth considering.

In reality creating subscription bundles and scheduling an entire season of events is not as trivial as moving one event to a different place in time. Some events have constraints placed on them by the performers (e.g., a guest artist in town) and others may be seasonal by nature (e.g., a Christmas show). We could easily find the local optimum given a set of events, but the more challenging problem is to solve a global optimum across all the bundles and events given that bundles can be made up of a much larger set of events across many different days. The problem becomes much more challenging and interesting if events are not only scheduled, but also put into the appropriate bundle. This problem and its insights are left for future research, most likely solved with heuristic optimization methods.

Our model is limited in that it predicts only one year of repurchase given the attributes of the previous year's bundle. Instead, it may be important to consider the entire lifecycle of a cycle over many years and consider how sequence effect may impact an even longer view of the cycle. Although we found little evidence for primacy or anchoring, research on expectations shows that once an expectation is set, it is difficult for a service provider to lower its standard again. Does this imply that if a season ends on a high note, but the next season begins on a much lower note, customers will experience more disconfirmation because expectations are very high? Also, does the impact of the first show provide an anchor for which all other shows are judged? If so, then does a peak need to be more or less intense in order to be effective?

We acknowledge that our model suffers from self-selection bias since customers choose which subscription to purchase and only customers that buy are modeled for repurchase. Although this bias may make it difficult to discern causality (vs. correlation) for many of our independent variables, we are concerned primarily about the sequence variables. It may be the case that high repurchases customers use the future sequence of a subscription as an attribute of initial choice modeling, but that in itself is also a very interesting finding. Since we have controlled for as many other product and customer attributes, we feel that the effects of the sequence variables are distinct and whether they cause a repurchase or if customer who tend to repurchase prefer a specific sequence, the results still support the managerial implications of sequencing events in a specific way. We call for future research to test the causality of sequence effects by using controlled experimentation techniques.

In considering peak and end effects it is not a stretch to think of an offering as a series of nested sequences all of which include some sort of peak and end effect. A subscription consisting of 8 different musical performances can be considered as a series of nested sequences. At the

highest level we have the events scheduling within the subscription as addressed in this paper, next the musical song sequence (how do you choose the order of the songs within a given concert), and finally each piece of music itself invokes peaks and valleys by creating tension and dissonance and release and harmony. Composers use volume, rhythm, tempo, timbre, and chord progression to evoke a sense of movement in which a peak is found and the end is achieved. At the highest level is the subscription cycle over its lifetime across several seasons for which the general trend of each season may benefit from an upward trend building up to a peak.

We do not consider this study to be a “peak” in this field of research and certainly not its “end”. Business scholars can most certainly learn much from those experiences around us that perhaps unknowingly evoke the power of the sequence effects. Drawing from other industry practices, business scholars can begin to understand how genuine peaks are created and how the lowest most level of the nested sequences may drive the peaks of the higher levels.

CHAPTER 3

THE ROLE OF EVENT SCHEDULING AND BUNDLING FLEXIBILITY IN CREATING SEQUENCE-BASED EVENT SCHEDULES

Abstract

When a service provider schedules and creates bundles of events, flexibility in when an event can be scheduled and which bundles it may be part of can affect customer perceptions of both the event and the bundle. The placement of an event within a bundle and across time may strongly influence the evaluation of the entire bundle by altering the expected sequence effects created by purposefully ordering events in a psychologically pleasing way. In this study we propose to test the relative importance of both scheduling flexibility and bundling flexibility to determine which factor may be more important when considering a complex set of service bundles. Using a complex problem faced by a renowned performing arts venue, we propose strategies to schedule and bundle events in a way that incorporates peak, end, trend and spreading effects shown to influence repurchase of season subscriptions. We present a realistic mathematical representation of an event scheduling and assignment problem that aims to maximize the psychological sequence effects. The problem is solved using a meta-heuristic and a series of hypothetical scenarios are tested to provide managerial insights and direction to service providers about how best to evaluate the importance of event flexibility as it relates to maximizing expected sequence effects.

Note: *This chapter is being prepared for submission to publication and it builds heavily off the work of chapter 2. References to chapter 2 are referred to as “Dixon and Verma 2011” instead of “Chapter 2” throughout this chapter.*

Introduction

Many event venues combine multiple events together to sell as a subscription package or bundle. Such bundles are common for sporting events and the performing arts, among others. Sales of these bundles, often called season subscriptions, are important to such organizations because they provide a large part of the capital needed to cover upfront production costs associated with providing a season of events. However, over the years, season subscription purchases have decreased (Pogrebin, 2002) and the question of subscription utility maximization is raised, i.e., what can be done to increase the value of the subscription in the eyes of the patrons? One answer might be to book acts that provide higher utility to customers. At first blush this approach might seem appropriate; however, more popular acts require higher upfront costs and including them in a bundle may not necessarily improve bundle sales if other shows' values are dwarfed in comparison to the highly popular act. In that case, patrons might opt to just purchase single show tickets and season subscription sales might suffer, leading to less funds available for upfront costs and perhaps into a downward death spiral for bundle purchases and venue sustainability.

When considering potential events to schedule over a season, event planners must decide whether to include each individual event in the master schedule. This decision is very complicated; however, one factor in deciding whether to include an event is its ease in scheduling and bundling. Some events could be scheduled on any date and in any bundle, but more realistically, events are appropriate (e.g., seasonal concerts) or available (e.g., out-of-town guest artists) on a restricted number of days. Additionally, bundles often have a theme that

restricts the type of event by genre, artist, performer, etc. These two levels of flexibility (bundle and scheduling) are perhaps more critical as a venue attempts to schedule events in a certain sequence.

Recently, operations management scholars have begun to investigate the importance of the sequence of events across time in the design of a service schedule (Dixon & Verma, 2011). Given that not all parts of a service have equal value, findings suggest that the placement in time of high and low value events should be considered in an attempt to maximize the utility of a schedule. In this paper, we develop an approach to schedule and bundle events while maximizing the theoretical psychological effects of sequence. We explicitly model sequence effects and provide a mathematical representation of a potential optimization problem that includes realistic constraints and conditions. The mathematical representation of the problem yields insight about the potential complexity of maximizing sequence effects. We solve the problem using a meta-heuristic algorithm approach. Because of its ability to solve a myriad of different problems, the meta-heuristic solution procedure provides managerially relevant insights into the structure of near-optimal service bundles and schedules. Specifically we investigate the relative flexibility that specific events allow bundle membership (i.e., *what* bundles could this event be a part of) and event timing (i.e., *when* could this event be scheduled). Problems with various levels of *what* vs. *when* flexibility are solved and compared in order to gain insights on the importance of event flexibility in sequence-effect-based scheduling and bundling efforts.

To further frame the research question, we will present the scheduling complexities of a world renowned performing arts venue with which we define and describe the sequence effect optimization problem. We describe the meta-heuristic solution procedure and describe the

experiments designed to provide insight into the research question. We provide the results and discuss their managerial implications.

Initial Problem Description

We began considering this problem after receiving a unique archival dataset from a renowned performing arts venue (identification of venue held for review). The venue hosts over 300 performing arts events a season (year) coming from twelve different genres consisting of local or in-house performers and artists and touring or guest performers. Most performance events that are scheduled in the venue become a part of a subscription bundle. The venue was suffering from a decline in season subscription sales and expressed interest in investigating novel ways to improve the popularity of their bundles.

Bundles for this venue are usually theme, genre, or market oriented, e.g., American composer (theme), jazz (genre), or family matinee (market) subscriptions. The same event can be a part of multiple bundles and there are often multiple showings of the same event. Additionally, they have six event halls within the venue that can be scheduled simultaneously each with different capacity and varying performance attributes (acoustic, lighting, staging, etc).

Each event needs to be assigned a date and a time, a hall, and to one or multiple bundles. With 300 events a year, six halls, and nearly 50 bundles we were initially impressed by size of the problem and began to consider ways to improve the overall schedule of events. The maximization of sequence effects was considered after an econometric model provided evidence that repurchase of season subscription tickets were correlated with sequence of a bundle's events (Dixon & Verma, 2011). Dixon and Verma (2011) empirically model and find evidence that the order of events — more specifically, the order of the events' utilities — within a subscription

impact customer repurchases of the same subscription the following season. They provide insight on how to design subscription bundle schedules independent of one another; mainly, that the last event should have high utility or value, the trend of event utility over time within a subscription should be positive upward, and that the peak or highest utility event should be placed either at the end or near the beginning depending on the homogeneity of the event utilities within a bundle. These suggestions seem easily implemented when considering one bundle, but they suggest that a harder, more interesting problem would be to consider a number of bundles that share the same space, calendar, and perhaps even events. Implementation is further complicated by considering which events from a large set of possible events should be bundled together. In this paper, we create a model to implement the findings of Dixon and Verma (2011) across a set of interrelated service bundles in order to provide insights on the level of complexity associated with incorporating sequence effects into service design of a large scheduling problem.

Psychology and Event Scheduling Design

Service operations management is concerned with how to effectively design and deliver desired elements of a service concept (Goldstein et al., 2002; Roth & Menor, 2003). Some researchers have compared a service concept to a theater production (Grove & Fisk, 2001; Pine & Gilmore, 1999) and compared the work of operations to that of a choreographer (Voss et al., 2008). We introduce a similar metaphor by comparing the work of service design, and more specifically event scheduling design, to the work of *orchestration*. Orchestrate, as defined in the Collins English Dictionary has two similar meanings (“orchestrate,” n.d.): (1) to score or arrange (music) for an orchestra (2) to arrange, organize, or build up for a special or maximum effect. Comparing event scheduling design to orchestration implies that the design of a series of events should be intentionally arranged and organized to affect customer perception. For example,

service providers consider details such as the physical surroundings and layout (Bitner, 1992), customer contact and involvement (Chase, 1981), process transparency (Buell & Norton, 2011), employee staging (Grove & Fisk, 2001) and scripting (Victorino, 2008).

Effective music orchestration requires both a thorough knowledge of music theory and being able to harness and utilize the strengths of the multiple players all to craft a score that is playable and evokes an appropriate response from the audience. The result of orchestration, whether for a music orchestration or service design, is a purposeful creation that evokes a desired response. Service providers must orchestrate various service design elements in an attempt to optimize the intangible psychological responses of a customer's experience that may define the perception of quality and value of a given service. The psychological responses may impact customer perception to a degree greater than the actual service or facilitating good provided, i.e., in some cases, the utility that a customer derives from a service may have more to do with an intangible psychological reaction than to the actual tangible result of the service or good (e.g., the food was lousy, but the service was excellent).

One design element that can lead to a purposeful desired response is event scheduling. A recent survey (Kendall, Knust, Ribeiro, & Urrutia, 2010) found 162 articles published from 1968 to 2008 concerned with scheduling sporting events within leagues and tournaments. Most of these articles are interested in finding efficient schedules from the perspective of league or tournament organizers and not likely from the perspective of spectators. Sampson takes a more service management perspective to scheduling problems by researching implications of attendee preference based conference scheduling (Sampson, 2009; Sampson & Weiss, 1995, 1996) and student preference based classroom scheduling (Sampson, Freeland, & Weiss, 1995). The shift in scheduling with an objective to maximize customer perception is likely the key difference

between event scheduling design as a sub-field of service design and event scheduling as it more closely relates to traditional manufacturing scheduling, i.e., event scheduling design brings with it a focus of a customer journey in mind. That is not to say that other aspects important to a schedule should not be considered when developing a schedule, but this change in objective changes the purpose of a schedule from one of efficiency to one of orchestrated service delivery.

Chase and Dasu (2001) suggest that to perfect a service, managers must understand fundamental psychological effects that the service may invoke. Among other things, they suggest that service designers should consider the sequence of the levels of pain and pleasure experienced overtime during a service. They point to psychology and behavioral economics literature as evidence that humans prefer certain sequences over others. For example, they suggest positioning painful parts of a service together and far away from the end and leaving those parts that are most pleasurable for the last. Dixon and Verma (2011) provide a thorough review of the psychology and behavioral economics literature concerned with sequence effects and cite four main effects that emerge as relevant to service scheduling: the impact of the highest point, most intense, or highest utility part of an experience (Peak Effect), the impact of the last point of an experience (End Effect), the impact of the placement of the peak in time (Spreading Effect) and the overall trend of the experience over time (Trend Effect).

In the context of our problem, we suggest that the design of an event schedule can lead to an intangible effect influencing customer behaviors, and therefore the venue can expect an increase in customers' perception of bundle utility as sequence effects are designed into the schedule. Restated, we believe that by appropriately scheduling and bundling events, a venue can increase the value of its offering without changing anything about the set of events.

Interrelated Service Bundling and Scheduling Problem

The problem we address is scheduling and bundling a set of events, considering one master calendar and a set of bundles. Each event can be scheduled into only one date and time (here to referred to as datetime), one hall, and one or more bundles. The objective of the multifaceted assignment problem is to maximize event sequence effects across all bundles b (see equation (1)). We focus on the four sequence effects tested and developed by Dixon and Verma (2011) mainly, peak effect, end effect, spreading effect and trend effect. Additionally, we allow for different weighting for each effect ($w_1 - w_4$).

$$\max \sum_{\forall b} (w_1 \text{EndEffect}_b + w_2 \text{PeakEffect}_b + w_3 \text{SpreadEffect}_b + w_4 \text{TrendEffect}_b) \quad (1)$$

Each of the four effects are explicitly formulated in terms of decision variables (discussed below) and can be seen in the appendix; however, in words the end effect is the utility of the last event in a bundle, the peak effect is the utility of the highest utility event in a bundle, the spread effect is the time between the peak event and the last event, and the trend effect is the slope of the least squares regression line of event utility and days from the first event in the bundle.

Assuming that events have different utility, the temporal placement of the events within a bundle will alter the level of the bundle's sequence effects. Each event has a measure of independent utility that we assume can be determined a priori through means of past performance data (e.g., forecasting), customer surveying (e.g., choice modeling), or with expert content knowledge. We refer to this utility as independent because it is a measure of the event's worth independent of the bundle it is in or the time at which it is scheduled. Similarly, the utility

measure is independent of customer type, i.e., it is an aggregate estimate across all customers. Therefore, when we refer to event utility, we mean the utility that the event would have on average across all customers as opposed to the utility that each individual customer might get from each single event. This aggregate measure can be considered the utility of the event from the perspective of the venue as they make planning decisions such as pricing, scheduling, and bundling.

The decisions for the event planner are which bundle(s) is (are) appropriate for the event and what datetime the event should occur. One approach would be to first bundle all events appropriately and then schedule all events, or vice versa, i.e. schedule all events and then try to find appropriate bundles from within the schedule. A separate approach would be to make the decisions simultaneously. Events across different bundles share the same master schedule and so must be considered together. Similarly, since events can be a part of multiple bundles, it may be more logical to consider event schedules at the time of bundling, considering different bundle schedules and constraints. In addition, since the objective is to maximize the sum of bundle sequence effects, changing the datetime of an event may impact multiple bundles' sequences. Also, the same event in a different bundle might have a drastic impact on the bundle's sequence effects. For example, the same event might be a peak event in one bundle, but only a moderate utility event in another. For these reason we call our problem "interrelated" since the scheduling and bundling are related and ideally should be considered together.

To account for this interrelationship, we present the decision variables of the problem as a multi-indexed binary integer program, indexed by event, bundle, hall, datetime, and order within the bundle. A parallel decision variable is used to maintain relationship between what we refer to as "cluster" of events, or events that are the different showings of the same or similar

events. These event clusters have specific timing requirements across events in different bundles. These two decision variables are used in defining all aspects of the objective and constraints.

We attempt to capture all realistic constraints and in so doing vie for completeness in place of simplicity in order to fully understand the implications of applying sequence effects into a complex scheduling problem. Several constraints are implicitly controlled by subsets that are pre-defined to maintain the availability of decisions. For example, the set of bundles, the set of

Figure 3.1: Constraints Considered

Implicit Constraints (predefined by allowable sets):

- The allowable set of bundles for each event
(correct genre, theme, market, artist, etc);
- The allowable set of datetimes for each event
(correct day of the week, time of the day, day of the year, etc); and
- The allowable set of halls for each event
(correct stage, equipment, lighting, capacity, etc).

Explicit Constraints (explicitly notated in terms of decision variables):

Bundle Related Constraints:

- The minimum number of events in a bundle;
- The maximum number of events in a bundle; and
- The minimum number of days between events in a bundle.

Event Related Constraints:

- The minimum number of bundles an event must be scheduled into; and
- The maximum number of bundles and event can be scheduled into.

Cluster Related Constraints:

- The minimum number of events in a cluster actually scheduled;
- The minimum days between events in a cluster;
- The maximum days between events in a cluster;
- The number of days from the first to the last event in a cluster; and
- Events of the same cluster cannot be in the same bundle.

Datetime / Hall Constraints:

- Each datetime can only be scheduled once for each hall; and
- Events in a hall cannot overlap in time.

datetime, and the set of halls that are *allowable* for each event are defined by the problem itself and constrict the available choices for the decision variables. All other constraints are explicitly notated in terms of the decision variables and define the boundaries of feasible solutions. The explicit constraints can be separated into those relating to bundle, event, cluster, and hall requirements. See Figure 3.1 for a complete verbal listing of all constraints and their classification. Also, see the appendix for a detailed description and mathematical notation of the problem including a definition of all sets, subsets, indices, parameters, decision variables, objective measures, and constraints.

The specifics of the mathematical modeling developed for our problem is confined to the appendix primarily because it is secondary to the research question that we wish to address. However, from the process of mathematically representing a sequence effect optimization in an interrelated bundle problem we learn to appreciate the potential difficulties in applying findings from psychological and behavioral research into a complex service design problem. In our case, the difficulties lie in representing a decision variable that can appropriately capture the order of events within each bundle while still assigning events to bundles, halls, and datetimes; i.e., maintaining a five index decision variable; and, in representing a large list of realistic constraints using the decision variables.

Research Question: Event Scheduling and Bundling Flexibility

The difficulty in modeling the problem can be further magnified by realizing that the level of flexibility of events scheduling and bundling likely impacts sequence-effect-based scheduling. Flexibility as a research question has been investigated within the context of labor scheduling, for example: start time of breaks (Bechtold & Jacobs, 1990; Brusco & Jacobs, 2000),

employee cross training and resource flexibility (Daniels & Mazzola, 1994; Daniels, Hoopes, & Mazzola, 1996; Daniels, Mazzola, & Shi, 2004; Iravani, Oyen, & Sims, 2005), and using part-time workers (Mabert & Showalter, 1990). Additionally, flexibility has been researched in regard to queue design (Sheu, McHaney, & Babbar, 2003), process flow design (Chow, 1986) demand forecasts (Tsay, 1999), geographic production capabilities (Kogut & Kulatilaka, 1994), product production capabilities (Jordan & Graves, 1995; Graves & Tomlin, 2003) to name a few.

Sethi and Sethi (1990) provide a thorough, yet dated review of the flexibility literature and found 11 types of flexibility relating to manufacturing. Two types are of interest to us: machine flexibility and routing flexibility (or operation flexibility as defined by Koste and Malhotra (1999)). Machine flexibility is the ability for a machine to process many different types of jobs; routing flexibility is the ability of a system to produce a part by alternate routes through the system. As early as 1952, Diebold (1952) envisioned a machine that could perform a *bundle of functions*. We will refer to *bundle flexibility* — the ability of a bundle to accept many different types of events — as a corollary of machine flexibility. We will refer to *scheduling flexibility* as the ability of an event to be scheduled liberally across any datetime, similar to the routing flexibility in which a product can be produced with different sequence of operations.

Bundle flexibility is mostly a function of the bundle theme restrictions and hence a decision that a venue planner must make considering customer demands on subscription bundles. We refer to bundle flexibility as an attribute of an event, but in reality it is an attribute of the bundle, much like machine flexibility is likely a machine attribute not a product attribute. On the other hand, product attributes dictate the level of flexibility in operational production routing (Koste & Malhotra, 1999); similarly, event attributes more directly dictate scheduling flexibility, i.e., event flexibility is dictated by the event (e.g., performer's schedule, appropriate season) and

not the schedule itself. Thinking of bundling flexibility as a bundle design attribute and scheduling flexibility as an event design is helpful in considering managerial implications of our findings. However, for the sake of our research design, we will translate bundle requirements into event attributes, i.e., we will determine which events can meet bundle requirement and refer to both dimensions of flexibility as event attributes. Doing so will allow us to design research scenarios that are more simple and more comparable.

Problems with different levels of flexibility across these two dimensions are apparent across different industries, e.g., sporting event and conference scheduling. Professional sports teams that play many games over the season (e.g., NBA, NHL) often design and market mini-bundles of matches to sell in place of entire season packages. In this case it is not unrealistic to consider that any of the events (matches) could be in any of the bundles; i.e., they have complete bundle flexibility. Conversely, the sports team might not have any flexibility in datetime scheduling if the schedule is dictated by the league. The opposite extreme is apparent in academic conferences: when considering presentation scheduling across a theme-based track, most presentations are rigid in their specialized track, but more flexible in their time scheduling if presenters are assumed to attend the entire conference. Therefore we consider the following research question: Given a choice between flexibility in only one of the two dimensions of bundle flexibility and datetime flexibility, which would be preferable for creating effective sequence-effect-based schedules?

In order to more fully examine this question, it is helpful to consider the type of problem facing us. The challenge in positioning this research within the context of other scheduling or sequencing based research is that it little resembles historical scheduling literature, other than in name. Scheduling research stems from job shop scheduling in which jobs are assigned a series of

resources in such a way to reduce the time it takes to produce a product. Sequencing, then, is seen as a condition for the order in which jobs need to be performed to make the product. At its simplest, our problem assigns an event to a time and date (similar to assigning a job to a resource), but there is no consideration for minimizing time through a system, nor is there a predefined sequence that must be followed, i.e. the sequence is not a condition but instead directly affects the objective. We explicitly define effects that a sequence of events might have on the customers who view them and try to maximize these effects.

Events are scheduled considering a master schedule, i.e., all events must be scheduled with regard to the same space and time. In addition, bundles are also created from among all the events to create subscription packages. The sequence effects are determined within each bundle — and not across all events. Therefore, both scheduling and bundling must happen concurrently or at least before the objective statement is calculated. Looking at the problem one bundle at a time, one might think the problem is related to the traveling salesman problem (TSP) because we are trying to find an appropriate “route” through all the events within each bundle and consider all permutation of the available routes. In reality, this problem is different from the traditional TSP because we do not try to reduce the total distance traveled; there is no corollary to the value of the arcs between nodes (distance between cities) that we can minimize in our case. In fact, we do not know what the impact to the objective statement will be when choosing one arc versus another until the entire sequence is determined because our objective is dependent on the specific ordering, e.g., starting and stopping points within each bundle. In this sense our solution is not cyclical, but is ordinal specific: the end effect is a function of what event is the last event in a bundle; the trend effect is a function of all the events’ utilities over time; and the spreading effect is a function of the peak event placement in time. Even with a static predetermined group of

events for a bundle, there is no simple value that can be pre-determined on the arcs between the nodes without also considering the ordinal placement of the events. An analogy to the TSP might be that we are trying to find optimal routes, but, among other things, we are interested in knowing where to start and where to stop with no need to cycle.

To further complicate the problem, an event can be a member of several bundles. In the traditional TSP problem, the solution typically only goes to each node once. An abstract but incomplete analogy is to consider a network of nodes representing events and arcs representing the feasibility of interconnected nodes being in the same bundle. The resulting solution then would be a series of un-cyclical sub-routes representing a bundle. These sub-routes would be unconnected except when an event is shared among different bundles. The un-cyclical sub-route analogy only works if the arcs between events are always feasible; however, there are different conditions controlling the thematic nature of each bundle so an arc may be feasible in one thematic bundle but not another. For example two events might both be of similar genre and so could be in the same genre-based bundle, but the performing groups in the two events might be different enough disqualify them from being together in a performer-based bundle. In this case, we could extend the TSP network analogy to include the possibility of multiple arcs between nodes, each arc a different color representing possible bundle membership.

This analogy still does not completely capture the complexity of the problem since it only deals with the ordinal arrangement of events within a bundle, but ignores the actual assignment of a date and time for each event. This is problematic; first, because it does not consider the idea of a master schedule that all events in all bundles must share; and second, because events have limited flexibility in the assignment of a date and time. These restrictions might make certain ordinal permutations infeasible, but this infeasibility is difficult to capture in a traditional TSP

node network because the arcs are a function of the datetime availability which will change as the master schedule fills. Instead, the ordinal or sequence nature of events within a bundle must be a function of the assignment of datetime, and a TSP analogy approach is replaced by something that more closely resembles a multi-stage source and sink assignment problem in which an event is assigned a bundle, a datetime, and an ordinal placement within the bundle. The ordinal assignment is a function of datetime, bundle membership, and other events within the bundle. Datetime assignment is both dependent and independent of bundle assignment since a priori event /datetime feasibility is independent of bundle membership, but constraints dictating time between events within a bundle make datetime assignment conditional to bundle assignment. Similar to TSP, this analogy is still incomplete because not all arc values are known a priori.

Even though our problem is not easily described in terms of traditional problems, it is helpful to think of arcs as feasible moves within the problem. The answer to the research question of which dimension of flexibility would be preferred may be a matter of the number of arcs the flexibility provides. This leads to our first proposition:

Proposition 1: *As a flexibility dimension gains more possible options, it become more useful in allowing for effective sequence-effect-based schedules.*

This first proposition, though simplistic, is important to help realize that flexibility may be problem specific and different classes of problem come with different levels of implicit flexibility. For example, our data provider's problem has over 300 possible datetimes, but only 50 possible bundles. In contrast, a multi-track mega-conference may have 50 concurrent tracks (bundles) across only 30 timeslots (datetimes). Proposition 1 states that as the number of

datetimes or timeslots increase, scheduling flexibility becomes more important and as the number of bundles or tracks increases bundling flexibility becomes more important.

In its simplicity, this first proposition ignores the specific attributes that each flexibility dimension might play in improving our non-linear objective of maximizing sequence effects across all bundles. The ability to spread out events across bundles is akin to achieving a more balanced load across machines within a factory — a benefit of machine flexibility. The peak effect portion of the objective statement is simply the utility of the highest utility event within a bundle; maximizing peak effects across all bundles then acts as a force to balance high-utility events across all bundles. For example, in a simple problem of four events scheduled into two bundles, the peak effect would be maximized if the two highest-utility events were separate rather than together. Similarly, low-utility events matched with high-utility events can maximize the trend effect suggesting that an effort to balance might extend to events of both high and low utility. This leads to our next proposition, split into three parts:

Proposition 2a: *Event balancing becomes more necessary as variation of event utility increases.*

Proposition 2b: *As event balancing across bundles becomes more necessary, bundling flexibility plays a larger role in finding effective sequence-effect-based schedules.*

Proposition 2c: *Conversely, as event balancing across bundles becomes less necessary, scheduling flexibility plays a larger role in finding effective sequence-effect-based schedules.*

If event utility is relatively homogeneous across all events then there is less to be gained by switching events across bundles. However, as variation increases among event utilities, balancing the high and low-utility events across all bundles can make a bigger difference in

terms of maximizing the peak and trend effects. As variation decreases, less is gained by switching events across bundles and scheduling flexibility becomes more important.

While bundling flexibility may ensure that events are spread out more appropriately across all bundles, scheduling flexibility ensures that a given set of events within a bundle can be appropriately sequenced. Again, referring to our specific sequence effects, three out of the four are concerned with the placement of events within a bundle. As mentioned before, the trend effect can be maximized if low and high-utility events are paired, but only if the sequencing leads to a positive trend. If the sequence leads to a negative trend, the pairing of the extreme-utility events could deplete the objective statement. The end effect and spreading effect compete for the placement of the peak event: end effect maximization attempts to place the highest utility event as the last event, but the spreading effect maximization attempts to pull the peak event as far away from the end as possible. Coupled with an upward trend effect maximization, the complexity and competitive nature of the various parts of the objective statement become more apparent. The resultant maximization, then, is a matter of the weighting given to each sequence effect. Under equal weighting of the individual sequence effects, one of two resultant inner-bundle event utility profiles emerge: first, a peak, trend, and end effect optimization in which the profile starts with the lowest utility event and increase over time to end on the peak; or second, spread, trend, and end effect balanced optimization in which the profile starts again with the lowest utility event, but quickly jumps to the highest utility event only to bounce back down and make it way back up to finish on the second to highest utility event. This second profile (a skewed U shape) attempts to pull the peak event away from the end while still maintaining an upward slope and high ending event utility. Which is better of the two event profiles is

dependent upon the homogeneity of the events within the bundle according to Dixon and Verma (2011). This discussion on possible optimal sequence profiles leads to our final proposition:

***Proposition 3:** Under equal sequence effect weightings, scheduling flexibility will always provide better sequence-effect-based schedules than will bundling flexibility.*

This, the boldest of our propositions, is supported by the discussion of the dimensions that influence the various elements of our objective statement. With complete bundle flexibility lacking any scheduling flexibility, only the peak effect can be certain to be positively influenced by appropriate spreading of high-utility events across bundles. However, with complete scheduling flexibility lacking any bundling flexibility, one of two sequence profiles can emerge that impact three out of four elements of the sequence effect optimization.

Meta-Heuristic Algorithm

The complexity and non-linear nature of some aspects of the objective statement and several of the constraints of the scheduling and bundling problem make solving it as normal integer programming model difficult. Instead, we propose a solution using a meta-heuristic algorithm incorporating simulated annealing (SA) concepts. The details are described in the appendix, but the basic premise of the algorithm is to find a feasible solution and then iteratively perturb it slightly and with randomness. At each iteration a feasible solution is rebuilt and compared with the previous solution; parameters control whether or not a worse solution will be kept. Keeping worse solutions allows the algorithm to investigate more solutions, reducing the chance of it being trapped in a local optimum. SA is a well known and used meta-heuristic search procedure and we find that it works well to find satisfactory solutions to our data provider's problem, and in addressing our propositions concerning flexibility dimensions.

Similar to our approach to mathematical modeling, we approached algorithm design with the full complexity of the problem in mind. Being able to solve a problem as complex as ours gives us the capability to solve problems with different attributes, and the opportunity to test the impact of scheduling efforts across different problem types.

Experiment Design and Results

To test our propositions we designed a series of problem attributes that can be used to generate problem sets. First, we altered the percentage of events with unconstrained datetime flexibility (i.e., all events are initially allowed to be scheduled in any datetime) and unconstrained bundle flexibility (i.e., all events are initially allowed to be scheduled in any bundle). For simplicity in research design, if an event is not unconstrained in a dimension, it is only allowed one option, i.e., it is completely *constricted* to only one possible option. Although, these extremes might rarely be realistic, they provide ample divergence in order to make scientific observations. We begin then with experiments 1 & 2 as shown in Figure 3.2 wherein the extremes are inverted across the two flexibility dimensions:

Figure 3.2: Experiments 1 & 2

	Experiment Number	
	1	2
Bundle Flexibility	Completely Constrained	Completely Unconstrained
Datetime Flexibility	Completely Unconstrained	Completely Constrained

Completely Constrained = only 1 allowable option
 Completely Unconstrained = all options are initially allowable

We initially run these experiments for problem sets with attributes (except those of bundle and scheduling flexibility) roughly matching our data provider's, i.e, 200 events, fifty bundles, six halls, 300 datetimes. Twenty problem sets are generated each with random assignment of event utility (exponential distributed with mean = 50; this is similar to the distribution of the event utilities of our data provider). Each problem set is used to generate both experiment 1 and experiment 2 using random assignment in datetime and bundle availability for each event according to the experiment conditions. The result is a paired sample design since the solution of a given problem set can be compared across the two experiment conditions; there is no difference in any other attribute of the problem other than the experiment condition. Each problem is solved five times and the maximum objective solution for each condition and each problem set is preserved and the average of these maximum solutions across all problem sets are compared across conditions.

Figure 3.3: Experiment 1 & 2 - Paired Sample T-Test

	Mean	N	Std. Deviation	t	df	Sig. (2-tailed)
Experiment 1: Flexible Datetime	.9550	20	.02016	35.177	19	.000
Experiment 2: Flexible Bundle	.7702	20	.02089			

For a degree of comparison, the outcomes of each experiment are scaled as percentages of the objective of the completely unconstrained version (unconstrained in both dimensions) of the same problem. We will discuss these and other results at length in the next section, but quickly the results show that the flexible datetime problems significantly outperformed the flexible bundle problems (95.5% to 77.02%).

Next, we determine the rate of objective improvement across levels of flexibility for the two dimensions. We are interested in knowing the level of flexibility that will allow a schedule to be relatively “free” to optimize across the two dimensions; i.e., the amount of flexibility in each dimension that will allow a schedule to be nearly as good as the most unconstrained model. To do this we can hold one dimension constantly unconstrained and alter the level of flexibility across the other dimension at 10% intervals, meaning that some percentage of the events will be allowed to be unconstrained in the opposing dimension. These factor levels within each experiment can answer two questions: which dimension increases the objective the quickest as flexibility increases and at what point do incremental levels of flexibility lead to meager improvements in the objective. Figures 3.4 and 3.5 describe factor levels for both experiments.

Figure 3.4: Experiment 1 Factors 1 to 9 – Unconstrained Datetime Flexibility and Varying Degrees of Bundle Flexibility

	1	2	3	4	5	6	7	8	9
Bundle Flexibility % Unconstrained	10%	20%	30%	40%	50%	60%	70%	80%	90%
Datetime Flexibility % Unconstrained	100%	100%	100%	100%	100%	100%	100%	100%	100%

Figure 3.5: Experiment 2 Factors 1 to 9 – Unconstrained Bundle Flexibility and Varying Degrees of Datetime Flexibility

	1	2	3	4	5	6	7	8	9
Bundle Flexibility % Unconstrained	100%	100%	100%	100%	100%	100%	100%	100%	100%
Datetime Flexibility % Unconstrained	10%	20%	30%	40%	50%	60%	70%	80%	90%

Each factor levels for both experiments are run across the same twenty problem sets run as before. Each problem is solved five times and the maximum objective is maintained. Statistics are estimated for each experiment across the twenty problems and averages and confidence intervals are plotted. As before, the outcomes are scaled in terms of the completely unconstrained version of the same problem and condition.

The findings of these first experiments show that, for a problem that resembles the original problem of our data provider, bundle flexibility is much less important than equivalent levels of datetime flexibility — Proposition 3 states that this will always be the case under balanced sequence effect weightings. Experiment 1 showed that with no bundle flexibility at all,

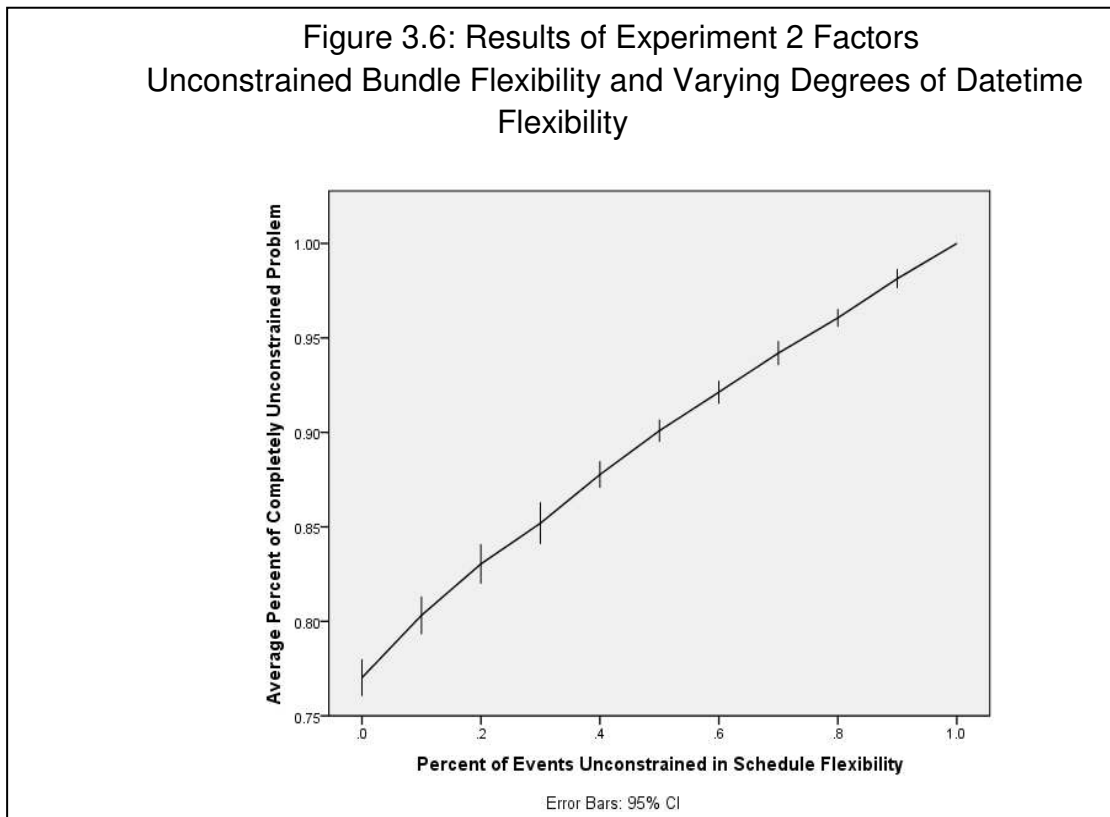


Figure 3.7: Results of Experiments 1 Factors
Unconstrained Datetime Flexibility and Varying Degrees of Bundle Flexibility

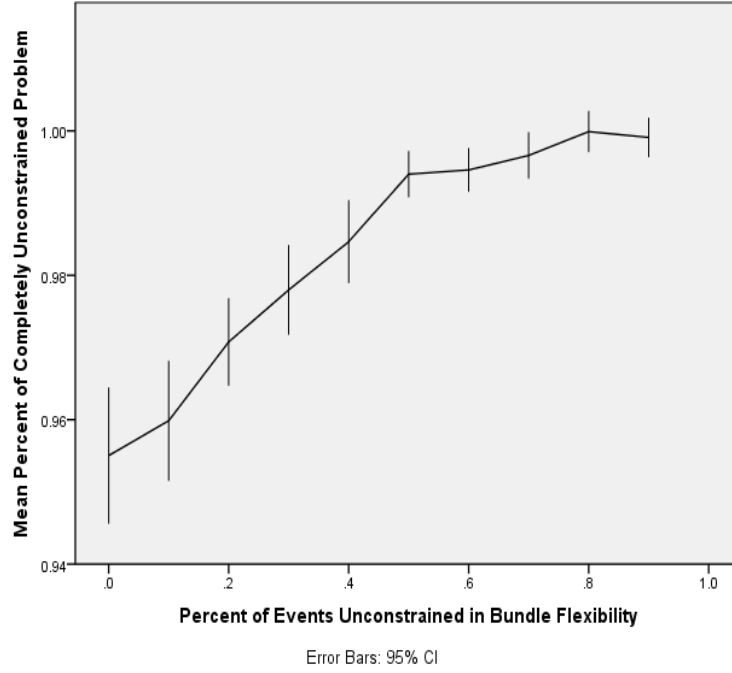
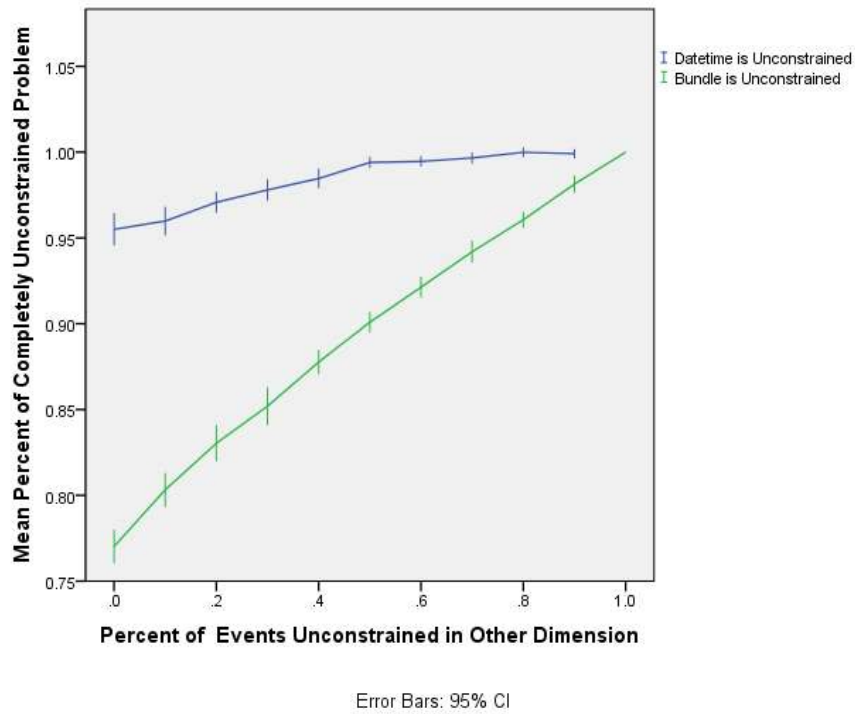


Figure 3.8: Figures 3.6 and 3.7 plotted together



the algorithm was still able to find an answer that reached 95% of objective of the unconstrained problem. Alternatively, the bundle flexibility problem was able to find up to 77% of the objective of the unconstrained problem under constrained datetime conditions. Under unconstrained datetime flexibility, increasing levels of bundle flexibility did raise the objective, but by the time 50% of events had unconstrained bundle flexibility, the objective was nearly 99% as high as the completely unconstrained problem.

After solving problems that were similar to our data provider's, we designed and ran additional experiments to determine if these results were unique to the attributes of this problem and in order to investigate our propositions. We designed problems categories with attributes varying in two dimensions each across two levels, (1) number of possible datetimes (high and low) and (2) event utility distributions (high and low coefficients of variation); specifics for the dimension levels are explained in the following paragraph. Combining the two dimensions across one another created four unique problem categories. We re-ran experiments 1 and 2 (original factors as shown in Figure 3.2) across all four of these new problem categories using all the same procedures as before (twenty problems created; each solved five times, with the maximum objective recorded).

The high datetime size condition was similar to our original problem with 300 datetimes and 50 bundles; the low datetime size condition kept the number of datetimes at 300, but restricted the maximum number of possible datetimes under the unconstrained datetime flexibility to 30. Under the original high datetime size condition all 300 datetimes were allowable, but under the low datetime size condition 30 datetimes were randomly selected from the original 300 for each event and the event was restricted to choose from only those 30 during unconstrained datetime conditions. Under the high utility variation condition the event utility

was kept at an exponential distribution with mean equal to 50. An exponential distribution's mean is equal to its standard deviation resulting in a coefficient of variation equal to 1. Under the low event utility condition, the event utility was modeled with a log-normal distribution with mean equal to 50 and standard deviation equal to 25 resulting in a coefficient of variation equal to 0.5.

Reported on Figures 3.9 and 3.10 are the results of experiment 1 (datetime flexibility) and experiment 2 (bundle flexibility) for the four problem types. The results are given as the average and 95% confidence interval of the percentage of the completely unconstrained problem for each problem, i.e., for each problem generated, a near-optimal answer was found for the condition in which both flexibility dimension were unconstrained.

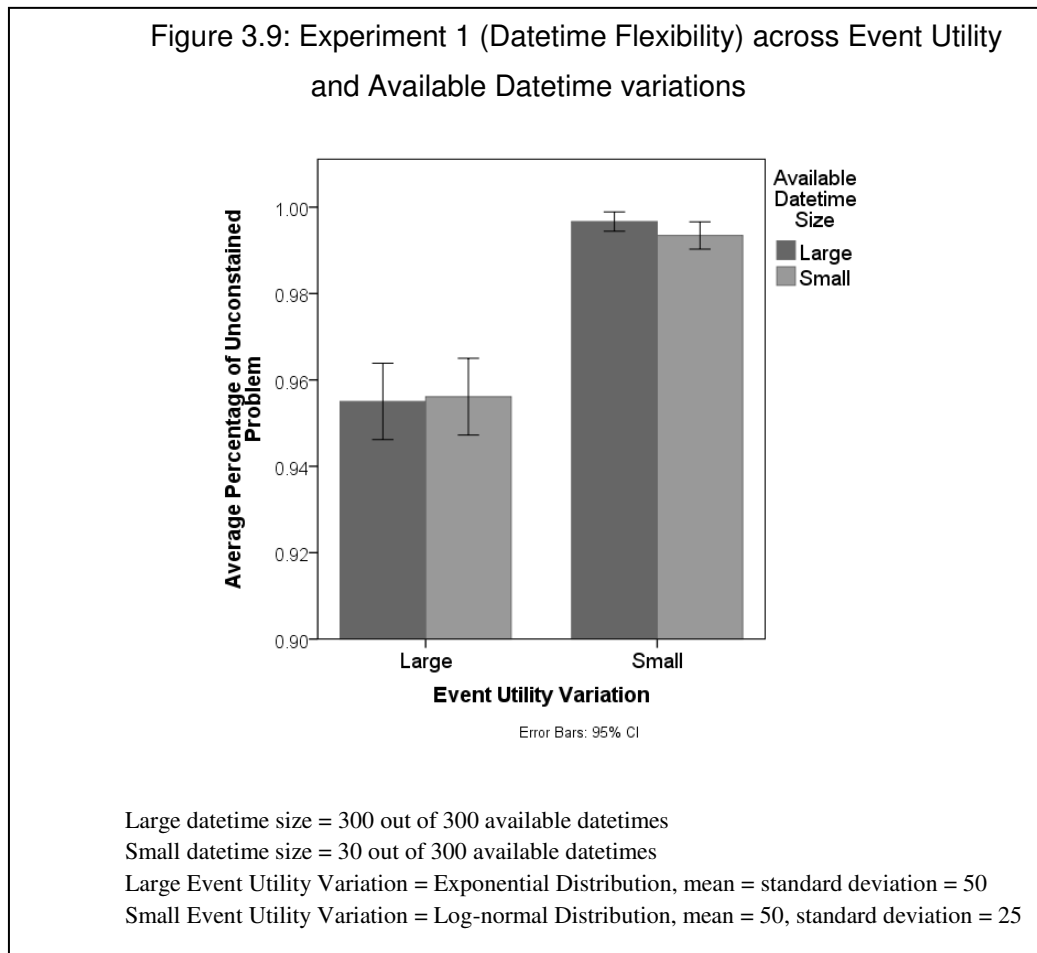
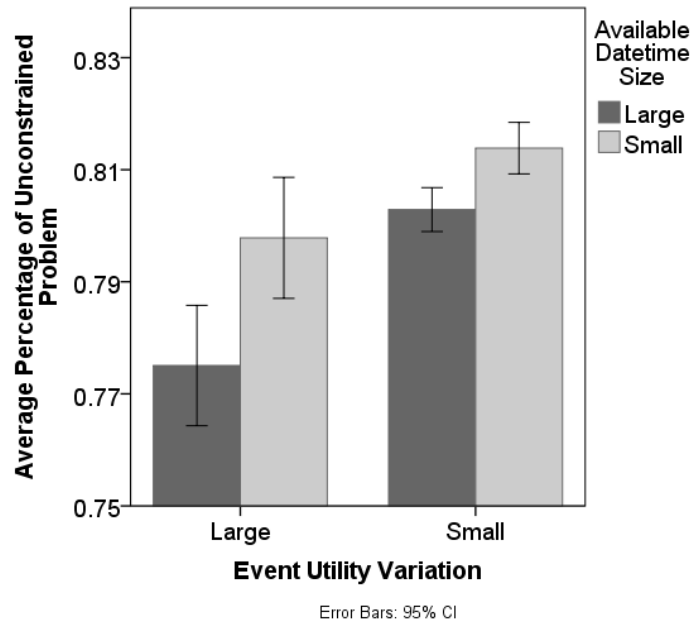


Figure 3.10: Experiment 2 (Bundle Flexibility) across Event Utility and Available Datetime variations

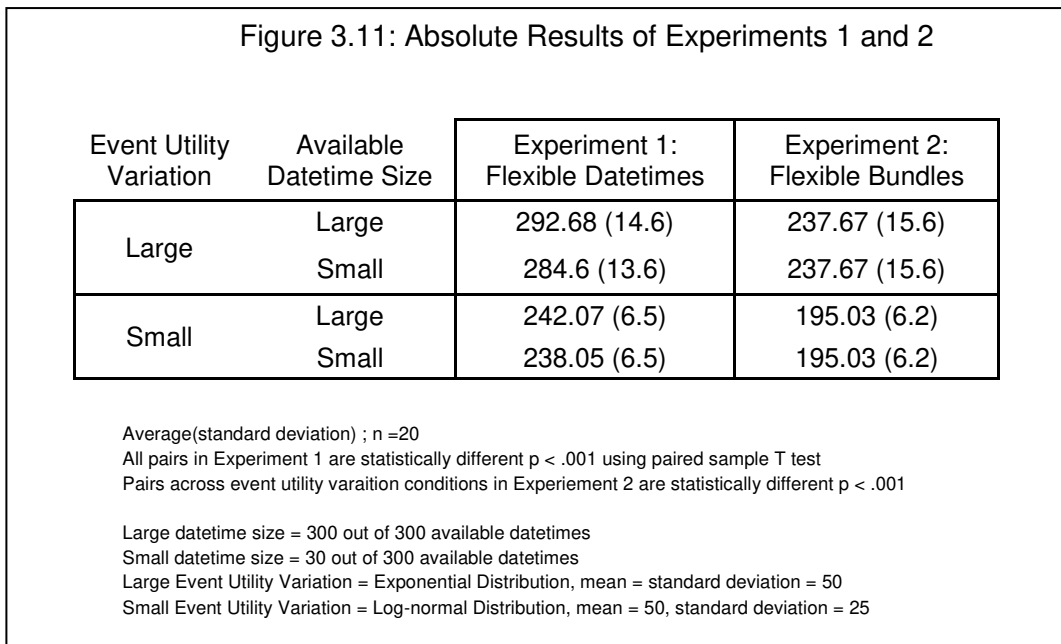


Large datetime size = 300 out of 300 available datetimes
 Small datetime size = 30 out of 300 available datetimes
 Large Event Utility Variation = Exponential Distribution, mean = standard deviation = 50
 Small Event Utility Variation = Log-normal Distribution, mean = 50, standard deviation = 25

Experiments conducted on datetime flexibility across the four categories showed that there was no significant difference relative to unconstrained problems in near-optimal solution across the size of available datetimes; this finding may be in contrast to proposition 1, which states that as a flexibility dimension gains more available options, it becomes more important. However, when considering bundle flexibility, we do find support for proposition 1 in experiment 2 as problem categories with smaller number of available datetimes provide better relative solutions than their counterparts with larger available datetimes. In our design, the number of bundle options does not increase except in proportion to the available scheduling

flexibility, but this finding suggests that even the proportion of bundle options to datetimes options can impact the effectiveness of sequence-effect-based scheduling.

Additional support for proposition 1 is found by investigating absolute results (shown in Figure 3.11) as opposed to results relative to unconstrained problems. For experiment 1, absolute results improve as available datetime size increases (292.68 versus 284.6 and 242.07 versus 238.05) indicating that as more datetime options are available the absolute schedule will improve. The absolute solutions for experiment 2 across large and small datetime conditions did not change since datetime availability was always restricted to only one available datetime per event. The relative difference comes, then, because the completely unconstrained answer was reduced, i.e. the denominator shrunk.



The experiment results show a significant relative difference across event utility variability under both dimensions of flexibility. For datetime flexibility, the relative results from experiment 1 show that solutions to problems with smaller variability in event utility distribution

could be nearly as good as solutions found in completely unconstrained problems. This finding supports the proposition 2b that indicate that as event utility variation decreases, scheduling flexibility will be more importance since little will be gained in switching events across bundles. The trend was repeated for bundle flexibility in experiment 2; however, these specific results are reverse of what is predicted in proposition 2a; namely, as event utility variability decreases, bundling flexibility becomes *more* rather than less important. Proposition 2a states that as event utility becomes more variable, solutions will be impacted greater by the ability to move events across bundles; however, the current results do not support this idea.

A general finding across our experimentation is that, all else equal, datetime flexibility trumps bundle flexibility, i.e., the freedom to schedule is more important than the freedom to bundle. In all cases, scheduling flexibility alone provided for a better solution than bundling flexibility alone. This was captured in our final proposition; we believe that this is largely driven by the specific sequencing requirements of our sequence-effect-based objective statement.

Conclusions

Recall that we define event schedule design as the purposeful orchestration of events across time. From the perspective of a venue concerned about creating a master schedule with the objective of optimizing sequence effects across all bundles, the value in our algorithm comes from gaining an understanding of the ramifications of considering including a specific event into the master schedule. Our experiments were designed to elicit what impact the comparative flexibility of an event might have on an ideal schedule. Our findings suggest two main conclusions concerning event selection:

Conclusion 1: *The scheduling flexibility of an event is more likely to influence sequence effect scheduling than is the bundling flexibility of an event if sequence effects are equally weighted.*

Conclusion 2: *As event utility variation decreases, the scheduling flexibility of an event becomes more important in finding properly sequenced schedules.*

We draw two other conclusions from our findings that are related to the design of bundles. As discussed earlier, bundle flexibility is more likely an attribute of the bundle than the event. Event designers can dictate the specific attributes (e.g., genre, performer, market appeal) that an event must have to be a member of a specific bundle. The final two conclusions are as follows:

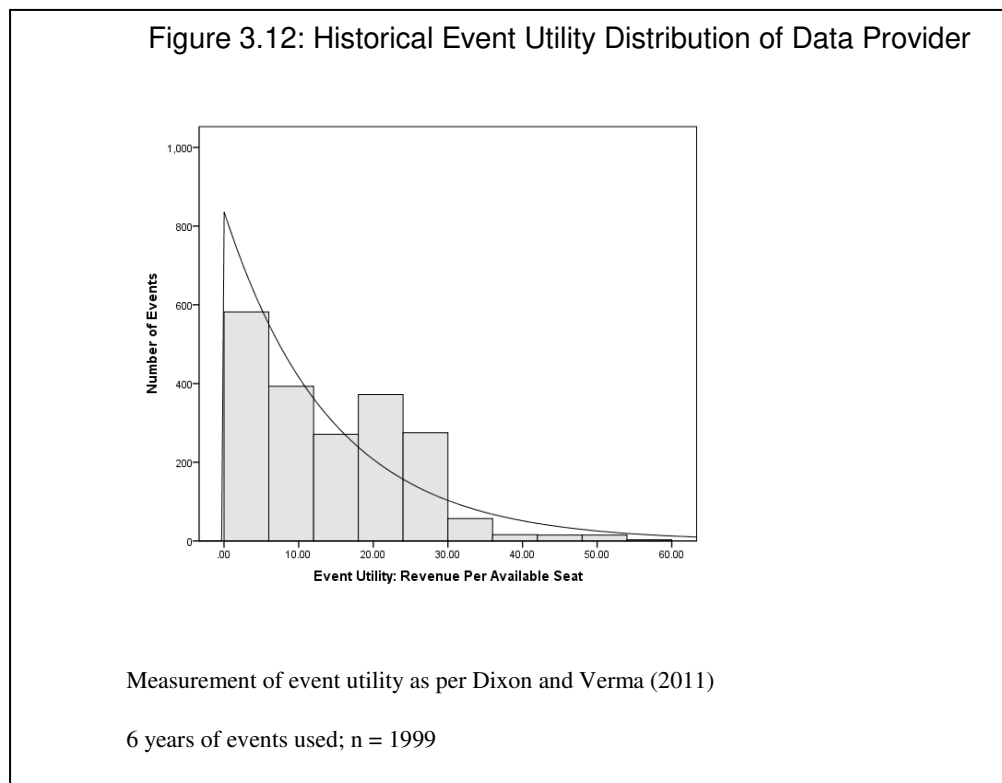
Conclusion 3: *As scheduling flexibility options decrease, bundling flexibility becomes more important in sequence scheduling.*

Conclusion 4: *Bundling flexibility alone does not adequately allow for sequence-scheduling.*

Managerial Implications

The managerial implication of our first conclusion is that a scheduler should favor events that have the ability to be widely scheduled. Even with fewer available datetimes (30, versus 300) unconstrained scheduling flexibility conditions performed in a statistically equivalent manner (relative results), indicating that scheduling flexibility need not be completely unconstrained in order to adequately sequence bundles. An appropriate follow-up study would be to investigate the limit of available datetimes that begins to impact relative solutions. With this knowledge, venue schedulers are able to classify events into those that will help or hinder sequences-based scheduling efforts.

Our second conclusion suggests that event planners should be concerned with event utility distribution; however, specific implications for event consideration decisions are more complex than simply booking events aiming to achieve a specific utility distribution. In the case of our data provider, event distribution was roughly shaped like an exponential distribution with most events with very low utilities and a few very high-utility events as shown in Figure 3.12. In order to reduce events utility variability an event planner could be asked to forego high-utility events. This seems counterintuitive and we are not ready to recommend it, but it does agree with traditional operations management advice of reducing process variability. Another recommendation would be to shift the distribution to one that might have fewer very high-utility events, but will also have fewer very low-utility events. Reducing event utility variation this way does two things: first, it allows for easier sequence effect scheduling as shown in this research; and second, it should allow for more consistency across events, which may lead to higher customer satisfaction.



Getting from the current high-variation exponential distribution of event utility to a future state of a low-variation log-normal distribution might be achieved by changing resource prioritization. Event planners typically spend a great deal of their time attracting high-utility acts. These acts fill the house, generate buzz and can lead to a one-time quick revenue burst. High-utility events are likely very rigid in their scheduling flexibility, moderate-utility events are moderately rigid and low-utility events are the least rigid. The consequence of spending time trying to book high-utility acts is that many moderate acts are passed up or unable to be fit in, leaving the remainder of the schedule to be filled in by available low-utility events. If priorities and resources are shifted to find a steadier stream of moderate-utility events, event utility distribution will shift and variation will likely reduce.

Our final conclusion is similar to our first, but we believe a further discussion of bundle design is appropriate. Even under conditions for which event spreading would seemingly improve the objective statement (high event utility variability) bundling flexibility alone only provided moderately sequenced solutions (highest was 81% of completely unconstrained near-optimal) compared to scheduling flexibility alone (highest was 99% of completely unconstrained near-optimal). It could be the case that we have not yet designed a set of problem characteristics for which bundling flexibility can shine (e.g., the preceding paragraph), but findings from our tests point to the conclusion that bundling flexibility cannot stand on its own as a sequence scheduling attribute.

We learned that added flexibility in datetime assignment may lead to a better placement of an event within a sequence of events compared to bundle flexibility. Since event content most likely influences bundle flexibility, this finding suggests that event content may be less important than event scheduling flexibility at least in the context of sequence effect optimization. This

means that rather than expending resources designing events that can be placed in multiple bundles, resources should be spent on adding flexibility to event scheduling. Additionally, bundle design does not necessarily need to be all inclusive, i.e., the temptation to allow themes to grow wider and wider in scope in order to allow for better schedules is not justified. There may be other reasons to widen scopes of bundle themes (e.g., changing audience preference, variety-seeking market segment), but improving possible schedules should carry little weight in this decision.

Research Implications

The finding that datetime flexibility is more important than bundle flexibility may indicate that there is less need to solve the full interrelated bundle-scheduling problem (i.e., bundling and scheduling taking place simultaneously) in order to achieve good schedules. Instead, the results indicate that a step-wise optimization approach might yield adequate results and that a simpler heuristics might be put in practice. A proper follow-up study would include developing simpler heuristics and testing the solutions generated against those of the full model across different problem conditions to determine when a full interrelated bundling and scheduling model would be preferred over a simpler heuristic.

We believe that the study of event utility distribution should become a focus of researchers interested in event scheduling design. Our findings that event utility distributions impact sequence-based scheduling efforts will be among the first of many indicating the importance of an event utility distribution in the success of an event sequence. Along these lines, a research question not fully addressed in this paper is how scheduling efforts will perform given events with high, moderate and low utilities across different levels of scheduling and bundle

flexibility as suggested in the previous section. This question will provide guidance to event schedulers about whether to accept an event given its expected utility and level of flexibility. This may also lead researchers to a better understanding of the importance of low-utility events in a schedule. As a contrast to high-utility events, it seems appropriate to include low-utility events, but how many and how low of utility? Empirical methods can help researchers understand the relationship between low-utility events and the importance of consistence or variety in overall events scheduling efforts.

Solutions to our problem sets indicate that as the maximum allowable datetimes in unconstrained scheduling flexibility decreased, increased bundle flexibility allows for and increase in the ability to find near-optimal sequences for the given problem. An appropriate follow-up would be to find the available datetime size that would make bundling flexibility a very high proportion ($> 95\%$) of the unconstrained solution. Further, a study of the impact of an increase and decrease in the possible number of bundles might indicate a further need for bundling flexibility. We consider 50 bundles to be a high number of bundles, but perhaps there are situations that might have more, e.g., a very large academic conference might have more than 50 concurrent tracks. Our results suggest that a larger number of bundles will increase the usefulness of bundle flexibility.

In closing, we concede that when deciding whether to add or ignore a proposed event into a schedule, an event planner considers more than the ability to create good sequence-effect-based schedules. An analysis of an event's success in isolation may be appropriate regardless of its scheduling and bundling flexibility and there are probably cases in which events that have no flexibility in either dimension still should be considered on their revenue generation alone. Still, a schedule planner must be aware of the impact that decision in isolation might have on multiple

bundles. This study highlights the importance of considering events in context of all other events as opposed to in isolation. The mathematical model provided an example of the complexity entailed in sequence effect scheduling in interrelated bundling and scheduling problems. Our algorithm provides a means to solve the mathematical model and provide solutions across different problem types. Solving different problem types has allowed us to make meaningful observations about specific problem attributes and led us to further questions. This paper marks a beginning into an operational view of event scheduling design. Dissimilar in many ways from traditional scheduling problems, event scheduling design must aim to support customer experiences. This means that there may be a trade-off between efficient design and service-oriented design, requiring modelers to creatively consider how design will impact customers. Adding an operational lens to the problem of event scheduling design will no doubt lead to better decisions and experiences in complex services.

CHAPTER 4

FUTURE DIRECTIONS

Abstract

This final chapter discusses future research opportunities presented by the main findings of this dissertation. It further addresses a number of weaknesses in the dissertation. Several short proposals for future research are discussed and ideal data sources are described. We address some remaining research opportunities for applying sequence effects to service design and delivery and present other ways behavioral-based research can be applied to a service operations management research agenda. We conclude with some key learning points not yet expressed concerning future research direction of event schedule design.

Introduction

This dissertation has applied one behavioral finding to the design and delivery of a service. Chapter 2 provides evidence that sequence effects influence performing arts patrons deciding to renew a venue's season subscription of events: the order of the event utilities within a subscription correlates with the subsequent year's repurchase decision. Chapter 3 applies these findings; creating a method to formulate a master schedule across multiple bundles, while optimizing theoretical sequence effects. In doing so, we developed a complex algorithm capable of solving many varieties of event scheduling problems. The algorithm provides the capability to then test "what-if" scenarios experimentally across many different conditions. Chapter 3 addresses one such scenario by testing the importance of event scheduling and bundling flexibility in creating optimal schedules. The results suggest that datetime scheduling flexibility is more important than bundle membership flexibility in creating near-optimal solutions across

the entire set of bundles. The findings of these two chapters collectively suggest practitioners should begin applying sequence effect scheduling principles into their service design, and that the choices making up the service package might play a role in how successful their attempts will be.

In this chapter we continue discussing opportunities for improvement and advancement presented by the two main research projects (chapter 2 and chapter 3). This chapter points out possible research directions implied by this dissertation and suggests proposals to begin such research.

Opportunities Identified — Chapter 2

The econometric estimation of chapter 2 is hampered by its use of an aggregate measure for utility, i.e., in place of a customer level measure of event utility, a measure across all customers for each event is estimated and used to determine sequence characteristics. We believe that this measure performs well as a proxy for individual utilities on average across all customers, but it limits the level of analysis that can be performed adequately. Most notably, customer segments can be identified easily using cluster analysis methodologies and once identified, specific preferences for sequence characteristics could be researched. We theorize that more loyal customer segments rely less on sequence effects in decision making, while newer customers might rely more heavily on them. The findings of the simple optimization found in the appendix 2 support this theory; however, a customer level utility measure will be able to provide more specific evidence in this regard. Similarly, it would be interesting to identify and cluster customer segments by their level of preference for specific sequence effects. We can further analyze the characteristics of those customers who use sequence effects more freely versus those

that do not. These types of questions can be more easily addressed under a pure customer-level utility measure for specific events.

Our estimation did not include a customer level measure of utility because the dataset did not include one explicitly. While we may have been able to derive individual utilities by specifying a choice model across all events for all customers, the results would likely have been very sparse in utility for most events, with little variability among those events that each customer attended, leaving little usefulness for our purposes. Instead, an explicit measure of event-level utility for each customer would have been ideal. An ideal dataset for this type of econometric estimation, then, is one that would include many specific characteristics of event, bundle, and customer attributes, and each customer's individual measures of utility for each event that they attended.

This individual utility measurement would ideally occur at three specific times of the service: before, during, and after each event. Economists refer to these different times as expected (or predicted) utility, experienced utility, and remembered utility (Kahneman & Snell, 1992; Kahneman et al., 1997). It might prove interesting to understand which utility perspective is most useful in designing sequence effects into a service. Largely our measure is a “before” or expected utility as it utilizes the number of seats sold and the revenue generated — all considered before the event begins. An analysis of event utility during an event might more akin to the pain literature that asked people for assessment of pain during the experiment. Similarly, in the pain sequence literature, the “during” assessments were later correlated to “after” assessments. These post assessments might be more heavily used by customers in determining future purchase decisions. A post assessment of an event will most likely include more than the actual core service: a post assessment might be confounded by the difficult of leaving the venue, the

availability of restrooms, parking and driving congestion, or other non-core attributes of the service. This comparison of during and post assessment is akin to a gap model of service quality (Parasuraman, Zeithaml, et al., 1985) in that we are interested in knowing if there is a gap between how customer felt *during* the delivery of the core service and how they felt about the entire service package *after* the service completed. If there is a difference between during event assessment and post event assessment, venues might devote more resources to designing a more complete service package that can seamlessly maintain or elevate assessment of the core service.

The during–post assessment gap implies that the portion of event utility considered when sequencing events may impact optimal sequence design. The danger of using a during or strictly core service assessment as a measure of utility to sequence lies in potential periphery service failures that can reverse the utility direction for an event; e.g., at peak mega-events high attendance numbers may stretch thin venue and service capacity and service quality capabilities. If an expected peak event is scheduled appropriately at the end of a sequence, but because of the high demand of the event the venue is unable to support periphery service adequately, the event might transform into a low utility event for some customers. This event’s placement in the sequence would then be detrimental to future repurchase assessments. Considering core event utility sequencing might then need to account for the capabilities of auxiliary service, and service providers need to better understand how auxiliary services impact overall impressions.

Expected or predicted utility must also play a role in creating future sequences. Our study only focused on the current sequence as it correlated to future decisions, and did not include how the current sequence fit into the larger sequence of events stretching into future seasons. Another question is to determine how important expectations of future sequences are as they relate to purchase decisions, experienced utility, and repurchase decisions (in subsequent seasons). This

line of work will help researchers begin to understand the role that anticipation can play as a sequence effect (e.g., Loewenstein, 1987). Predicted utility may increase experienced evaluations of a service through anticipation and savoring expected pleasure, even when it was worse than expected (Chun, 2009). This suggests that an appropriate way to increase aggregate average evaluations may be to design a sequence in which savoring the future peak is maximized.

A measurement of utility across these three different stages could certainly be accomplished with survey instruments designed to capture expected utility, experienced utility, and remembered utility. The difficulty in capturing individual utility at a scale as large as was used in chapter 2 is finding participants willing to measure their utility at these levels across several events and maybe even across several seasons of events. This realistic restriction makes a level of analysis similar to chapter 2 unlikely, but a much smaller scale design could provide results to help us better understand the implications of sequence and time related satisfaction. A simple research design would be to determine if customers can perceive an anticipated sequence before experiencing it: after having seen a schedule, but before experiencing it, can patrons identify which events they are more likely to enjoy? If so, then does the pre-experienced schedule *alone* change an expected utility of the entire schedule? A simple survey experiment could be conducted online, with the collaboration of the data provider for this dissertation, to identify if customers are aware of sequences and if they use them to develop anticipated utility across an entire bundle of events. We could follow up with the customers who actually participate in the events, and test to see if there is a significant difference between expected utility and remembered utility across different event utilities and time spans.

Similarly, experienced utility could be captured with electronic scale devices that can be used by participants to scale their utility in real-time during a performance. These devices were

in use during the 2008 presidential election debates by panels of listeners hired by television news broadcasters to determine when certain contingencies were more or less satisfied with the candidate's responses. A device called a Perception Analyzer© sold by DialSmith was given to participants, and functions as follows:

The software detects movement of the dials as participants make their selections. When the dials stop, the software captures the value from each dial, and then processes the data for analysis and display...

Each participant is associated with a unique dial ID... the software creates a profile of each person. This allows you to organize results into subsets to more deeply understand why people feel the way they do, and what it means to you. ("Dialsmith and the Perception Analyzer - Measure Responses," n.d.)

A similar real-time capture of utility or satisfaction would provide a more detailed sequence within each individual event. Whereas in chapter 2 we identified a sequence *across* a series of events, the sequence of utility *within* the event could yield interesting research questions related to the specific design of each event. More interesting is trying to understand the relationship within event sequences and between event sequences. Are there relationships between several nested levels of sequences and should different levels be considered jointly? What level is most noticeably important to customers?

Similarly, it is interesting to understand the cutoff points, when one sequence ends and another begins in the mind of a customer. In this dissertation we assumed that a sequence of events began with the first event in a bundle and ended with the last event in a bundle; however, we have little to justify this assumption except the natural design and sale of the bundles. If, though, customers are participating in multiple consecutive bundles, do sequences have such

strict cut-off points, or should sequences instead be considered as a modular arithmetic problem; i.e., should the problem of sequencing be considered cyclical instead of linear? If the sequence is cyclical, the theoretical design implication might be altered to allow peak events to be more spread out across time; instead of peaks at beginning and end, perhaps it would be more appropriate to just keep them separated from one another. Similarly interesting would be considering ways to signal the beginning and end of a sequence to customers in order to maintain a simpler design. Cognitive psychologists have devoted a stream of research to better understand the concepts of an internal clock (Treisman, Faulkner, Naish, & Brogan, 1990) and human perception over time (Dehaene, 1993) and continue to question if perception is continuous or discrete (VanRullen & Koch, 2003). Certainly, these ideas should be considered in light of service design and sequence cycles if service providers hope that customers will experience their service repeatedly.

One aspect of behavioral based research in operations, economics, accounting, and others fields is the consideration of the rational behavior of decision makers. In particular, many studies show that the behavior of decision makers is not rational and that irrational behavior can be predicted to some degree across certain groups. Systems can be devolved to either take advantage of the irrational decisions or to help prevent them from happening. Along these lines, sequences effect research to date has focused of understanding how sequences might influence behavior — rational or not. Of interest to service designers is to know what degree the sequence-induced behavior is predictable and under what conditions would sequence-induced behavior be rational and irrational from an economic utility perspective. Doing so can tie our research more closely to that of behavioral operations and behavioral economics. The current dataset provided to us for this dissertation can be used to further explore these aspects.

Opportunities Identified — Chapter 3

The algorithm developed in chapter 3 will continue to be a resource to address important research questions in two important ways. First, improving the model and algorithm itself, and second, by developing additional experiments to test against the algorithm and model.

While much of the effort of chapter 3 was spent developing and considering the model and algorithm used to conduct our research, there is still room to improve the model itself. The primary weakness is in developing the objective statement; currently the objective statement only addresses an optimization of sequence effects, but does not address capacity, demand, or revenue. During early phases of the mathematical representation of the problem we began to address capacity in particular; however, it was not fully developed and was left for future work. Specifically, the worry is that if a popular event is placed into a bundle, it will impact the sale of the bundle and require higher capacity for all other events in the bundle. Similarly, events assigned to several bundles might increase the capacity requirements for the event, i.e. bundle demand may in turn influence event demand. This endogenous relationship between event demand and bundle demand then influences hall assignment constraints as an event should be scheduled in a hall that will fulfill as much demand as possible.

A further difficulty lies in determining what portion of a bundle's demand is attributable to an event's demand if the event is shared across multiple bundles. The proportion of the event's demand attributed to each bundle may be a function of the correlation between events within the bundle. For example, suppose an event is placed into two bundles; the first bundle consist of events that are similar to the event in utility; the second bundle consist of events that are all of

lower utility than the event. The increase in demand with the inclusion of the event then would be lower for the first bundle, but we could say the event is being used as leverage in the second bundle because its higher utility can attract customers to the bundle. The proportion of the event's demand might be more highly attributed to the second bundle.

The first step in modeling these complexities is to address two underlying research questions: (1) what impact does adding an event have on to bundle demand, and (2) what impact does adding an event to a bundle have on event demand? The endogenous relationship between event demand and bundle demand makes these research questions difficult to approach using our original dataset, but an appropriate econometric method may be applicable.

To finish discussing event demand, it is difficult to distinguish between event utility, as we use the term, and event demand. In chapter 3 we state that our definition of event utility is an aggregate measure of value that customers derive from the event independent of bundle membership or sequence. We assume that bundle demand — more specifically bundle repurchase — can be influenced by the schedule of event utilities within a bundle. In this sense we implicitly try to maximize bundle demand using only sequence effects. While event utility is a measure of value, closely related to a price, event demand is the number of customers who will purchase at a price. Event demand then is a function of event utility, bundle membership, and price. Bundle demand is a function of individual event demands, event schedule design, and price.

Therefore, the objective could also include some sort of expected revenue function by including pricing for each bundle. It would be interesting to determine if an appropriately sequenced bundle could support or justify a higher price than a poorly sequenced bundle. To test

this, a choice experiment could be developed and delivered via survey to performing arts patrons. So far in our research, we have tied sequencing to revenue only through repurchase (Chapter 2), and it would be interesting to see if sequence effects can be related to first time sales.

In addition to altering the objective statement, the operationalization of individual sequence effects could be reconsidered experimentally. In particular, the model of the trend effect is the linear slope of the line through the event utility and days from first event. While this is a simple way to represent a trend, it is not linear in the sense necessary to simplify optimization. Furthermore, it may not fairly represent the way that a trend is felt; e.g., a linear trend assumes a continual and steady change from one event to the next, but perhaps trends are experienced more like a step-wise function for which a level of utility is achieved after having experienced an event, and the utility level then stays constant until the next event is experienced. Or perhaps the utility profile is curvilinear in nature as anticipation for future events impact sequence trends. Further study comparing how the trend of an experience is actually felt to the intended design can give insights on how to appropriately model the trend effect. One way to determine the nature of how customers interpret an intended trend would be to design a trend into an experience and ask research participants to draw a utility profile. In addition to getting a better understanding of how customers might interpret trends, this will help in understanding if customers can perceive trends that are purposefully designed, giving insight on how to design a trend.

An additional way to improve the objective statement is to recognize that different customers will find different utility in the same event. Our current models assume that utility is a static measurement that can be determined at an aggregate level across all customers. Instead, we can determine that different customers have different utilities (and preferences) for different

events and build bundles and schedules that will maximize schedule utility across all customer segments. This approach would require us to identify customer segments and appropriately define event utility for each event for each customer segment. Suppose that there are two customer segments: existing customer and new customers. Suppose further that these two segments have very different event preferences and perhaps event different sequence preferences. These two segments can then be thought of as competing for event schedule design decisions to favor their preferences. Such research can give insight on how to manage an existing loyal customer base, while inviting a new customer base to become loyal—a situation very relevant to performing arts venues. Additionally, it would give managerial insight on the wisdom of trying to attract diverse customer segments from a service design standpoint. How to achieve a schedule that can accommodate this balance is likely complex.

Another approach that considers event utility differently across customers is modeling each event utility as a stochastic measure for each customer. Even within customer segments, events are perceived differently for each individual customer. A stochastic representation of event utility would allow us to determine how design profiles might be similar to or different from those found using static aggregate measures. Furthermore, we could determine the impact of highly variable events on schedule design. While in chapter 3 we discussed the variability of event utility across events, this question would consider event utility across customers.

In addition to changes in the model itself, the algorithm can be altered to investigate efficiencies in exploring the solution space. While much has already been done in this regard across one general size of problems, a future step would be to determine how the algorithm responds to problems of different size and how the parameters of the algorithm should best be determined. Currently, many parameters are problem specific and are determined with sacrificial

iterations at the beginning of each solution (e.g., T, alpha, T-stop); however, a better understanding of the dynamic evolution of the solution could be helpful in order to efficiently move through phases of the solution space. For example, as the algorithm progresses, the solution improvement profile is often shaped with immediate improvement during the early iterations of the algorithm, little or slow improvement in the middle iterations and more quick improvement during the final iteration during greedier phases. The solution then is largely a function of where the solution is in the solution space near the end of the beginning phases, and so a more efficient search might include more early searches. One test that has not been performed then is determining the point at which the temperature parameters should be drastically altered in order to speed through the slower improvement phases of the algorithm. If this approach yields sufficiently high results, it could mean more solutions could be run with less iterations in the same amount of time raising the likelihood of finding a higher objective.

In creating problem sets to solve for our experiment, we assumed that all problem parameters were independent of one another, e.g., event utility was drawn independently from the bundle or datetime flexibility. This source of variation certainly seems acceptable given the purpose of our experiment, but it is more likely in practice that most problem conditions are correlated. Understanding the nature of this correlation will require either a survey of event schedulers or a dataset that includes problem parameters. The existing archival data includes the actual schedule and bundling of events, but does not give any indication of *possible* dates, bundles, or halls that could have been used in creating the schedule. In this sense, our algorithm has yet to be tested with real data, but we look forward to sharing the results with our data supplier and getting feedback on the practicality of our findings.

The algorithm that was built for chapter 3 was no small undertaking. The code is over 4,000 lines long and if printed would span nearly 80 pages. Pain was taken to make the algorithm robust and capable of handling many different scheduling and bundling problems in order to make a more complete examination of the ramification of sequence effect scheduling efforts. In chapter 3, we began these investigations, but there remain many more research questions that can be answered now that the algorithm is functional.

A thorough understanding of what an optimal sequence looks like in terms of predominant sequence profiles remains to be fully investigated. From chapter 2, we found that two possible sequence profiles were likely: start low and go high and start high and end high (skewed U shaped). An understanding of what profiles are more likely to emerge across a myriad of different problem characteristics would be helpful in providing managerial guidance or heuristics to follow under certain conditions. These characteristics will largely be a function of the weights of the sequence effects and may therefore be venue specific. In this regard it would be interesting to continue to investigate the importance of each sequence effect individually on its impact to sequence profiles. Currently the weights of all four effects are roughly equal; an experiment could be designed to change the relative weighting of all four effects, keeping all else equal. Of interest to event planner who may believe that different aspects of sequence effects are more important than others would be to see what sequence weighting impact large changes in solution profiles. For example, if the trend effect is more highly weighted we would expect to see more low-to-high sequences, but if the spreading effect is more highly weighted we should see more U shaped profiles.

The event utility distribution was modeled first as an exponential function with mean equal to 50 and then later a log-normal distribution with mean equal to 50 and standard deviation

equal to 25. The exponential distribution calls for most events to be small in utility, with a few larger ones. This distribution roughly matches that of the actual event utilities of the events in the archival dataset as estimated in chapter 2. However, chapter 3 results call into question whether the exponential distribution is the right distribution in terms of product mix and in terms of scheduling and bundling. The problem of event utility mix should be addressed both analytically and empirically. Analytically we can test the implications of different utility distributions on final solutions by changing the shape of the distribution curves while maintaining constant the area under the curve, i.e., keeping constant the sum of all event utilities. Is it more appropriate, from a sequence effect optimization standpoint, to have a small number of mega events with the remaining events being low utility events, or should the event utility be more normally distributed around a moderate utility mean? Addressing this empirically—via survey or archival data analysis — will help in understanding the role of mega-events in event bundling. The results from chapter 2 suggest that those bundles with true peaks were less likely to be repurchased than those bundles made up with homogenous event utilities. One interpretation of this result is that customers bought the bundle package to attend the high utility event and were not impressed with the remaining, relatively lower utility events and so decided not to repurchase. Venues use mega-events as leverage to get customers to buy packages, but the results suggest that this is not a sustainable, loyalty inducing practice. Still, more directed research can be done in this regard to make clear the suggested findings.

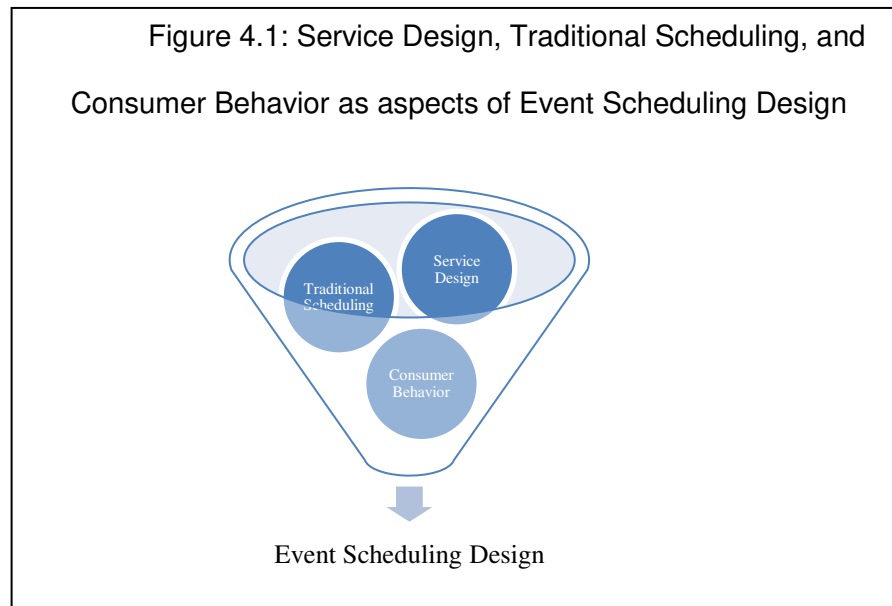
Event Schedule Design

We began this dissertation intending to better understand the elements of time in service design. Behavioral research was investigated and theory concerning the appropriate event sequencing was developed, tested, and applied. Although our intent was to improve our

understanding of service design in general, we noticed a gradual shift in considering specific aspects of the design of an event schedule. As discussed in chapter 3, events scheduling design is akin to orchestration in that its purpose is to combine different parts of a service to create a framework that can elicit expected responses from customers. In this way, we propose that event scheduling design can be considered a sub-discipline of service design, that considers how to appropriately design the placement of discrete events.

Scheduling as a research topic has been studied from the turn of the last century, beginning when Henry Gantt (1903) published alongside Frederick Taylor (Wilson, 2003). In discussing the role of traditional scheduling, one prominent author states, “Scheduling concerns the allocation of limited resources to tasks over time” (Pinedo, 2008). Pinedo continues to describe that scheduling efforts are developed to meet some objective, usually reducing a cycle time or increasing production rates. While this perspective of scheduling has proven important through the years, it does not adequately capture some of the fundamental differences found in service operations discussed in chapter 1; mainly, that customers are often part of the production process. For this reason, event scheduling design is more than just an effort to schedule events, event scheduling design is concerned with how the schedule can impact customer perceptions given that customers experience the schedule. This is somewhat different from a production scheduling perspective in which customer’s perceptions might be influenced by the schedule only because the promised good or service is delivered to them in a timely manner. For this reason, even scheduling efforts applied to the service industry do not always meet the criterion that we defined above for event scheduling design.

Researchers interested in event scheduling design must keep in mind “the allocation of limited resources,” while considering objectives that might be very different from the reducing



cycle time perspective found in production scheduling. Still, many aspects of production scheduling must be considered when attempting event scheduling design; first and foremost are constraint considerations that might limit the ability to create a desired, designed schedule. Underlying the purpose of event schedule design is the scheduling effort itself, which with even a moderate degree of complexity, can quickly become a difficult problem as we found in chapter 3. For these reasons, we submit that: (1) event scheduling design most likely will build off traditional scheduling design principles and methodology; and (2) operations management researchers bring the ideal background and disciplinary lenses to make headway in developing complex models that can advance managerial understanding of events scheduling design.

The objectives of event scheduling design differ from the system efficiency objectives of production scheduling. Instead, they focus on a more service oriented experience orchestration objective. This new type of objective requires that researchers be able to find ways to quantify

and codify the behavioral and cognitive responses of customers to design parameters. Importantly, we repeat our challenge in chapter 1 and call for operations management researchers to become more familiar with behavioral research in order to appropriately model these new objectives. Our focus in this dissertation was with event placement over time, but the proposals in this chapter call for addressing the issues of the appropriate number of events, event utility distribution, event length, and event sequence length, among others. In order to develop principles and theory about these issues, researchers can seek out a better understanding of how humans react under certain conditions. This search will largely be in behavioral and cognitive research.

A first step in increasing research interest in event scheduling design is to perform a comprehensive literature review in order to elicit best practices, provide guiding principles, and develop a research agenda. This review would assuredly cross disciplines and industries, but we propose that an appropriate context for event scheduling design is the field of experience management. A recent study found that firms that are openly designing services with experience in mind often use terms such as customer “journey” to describe how their customers interact with their service overtime. The idea of a “journey” is useful in terms of considering the types of research that addresses how customers traverse through a service encounter or series of encounters and how service providers plan and design this journey.

The End

As we conclude this chapter and this dissertation, we consider how we can apply some of our own results; mainly, how we can end on a peak. While the “journey” of this dissertation has brought us to think about the elements of service design with more depth, we certainly don’t

consider this treatise the end of the discussion. Instead, it is evidence that the investigation of our initial research questions has led us to consider even more questions. For this reason, we are excited for future investigations and discoveries and anticipate “peaks” yet to come. With the writing of this chapter, we hope to leave the reader with similar anticipation and in so doing hope that we will end on a high note. Perhaps Winston Churchill said it well after the British had defeated Germany in WWII: “This is not the end. It is not even the beginning of the end. But it is, perhaps, the end of the beginning” (Churchill, 1942).

APPENDIX 1:

LOGISTIC REGRESSION MODEL SPECIFICATIONS

It is prudent to discuss the appropriateness of the econometric method in estimation. We are using a logistic regression run with a Generalized Estimating Equation (GEE) estimator since we are accounting for within-group correlation by clustering on customer ID. To test for severe multicollinearity, we calculated variance inflation factors (VIF) for the final model showing that none of the variables have a VIF of higher than 10 with the highest of 8.3 and only three variables with VIF greater than 3. Similarly, by nesting the models and observing very little changes in previously estimated variables, we can conclude that multicollinearity is not severe in the models. Certainly there is correlation between observations as we often have the same customer over several seasons and across different cycles. By plotting the Deviance difference and the Pearson Chi squared difference against the predicted probability, we were able to identify 1 observation that had a high level of influence on the model. Upon investigation, the observation was from a customer who had purchased 40 bundles for the same season subscription. This outlier proved to be a significant influence on the model and was removed from the final results since a purchase of 40 subscriptions for the same bundle was not typical (mean = 1.7) and did not represent a normal customer. No other single observations were left as significant influencers.

The overall models are significant shown by the Likelihood ratio, Score, and Wald test statistics indicating that at least one of the predictors has a beta not equal to zero for all three models. The R squared values are increasing across the three models. Predictive accuracy of the models was determined by calculating the probabilities of repurchase for the excluded 10% and calculating a Brier score (the average of the squared difference between the prediction and the


outcome). Brier Scores range from 0 for a perfect prediction to 1 for a perfectly incorrect prediction, so a smaller score indicates an improved prediction. The scores for the 3 models improve across models (REVPAS: 0.1649, 0.1617, and 0.1616). By excluding a random set of observations in estimation, we were able to avoid bias that would result in using the same data to test the model as was used to fit the model.

APPENDIX 2:

AN ILLUSTRATION OF RECOMMENDATIONS FROM CHAPTER 2

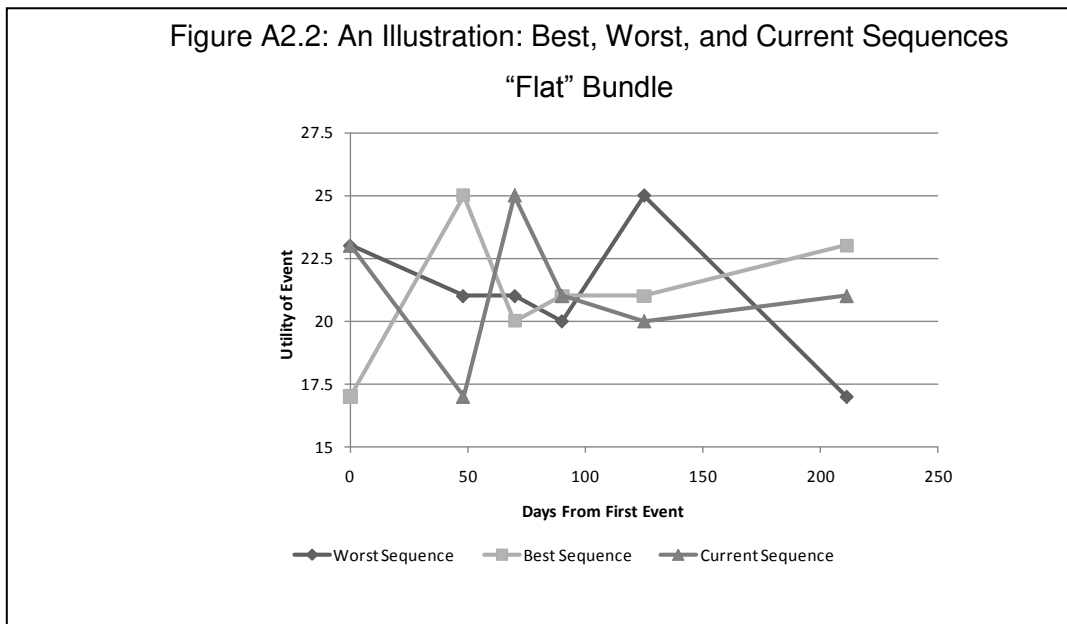
In this section we illustrate the impact that the sequence has on the probability of repurchase. Using the estimates from the REVPAS model, we show what the probability of repurchase would be under different event sequences of specific bundles for one customer. Next, we show average repurchase probability changes for different events sequences across a large population of customers and bundles.

A subscription bundle found in the dataset consists of the following events with utilities on the appointed day: on day 0 utility 23, day 48 utility 11, day 70 utility 41, day 90 utility 21, day 125 utility 20, and day 211 utility 20. If we keep the day of the event constant we can optimize the impact that the coefficients of the sequence variables will have on the overall probability of repurchase and identify the best and the worst sequence by using exhaustive search optimization, i.e., we solved for every permutation and found the sequences that maximized and minimized the effects of the sequence variable coefficients found in our estimation of Model 3. Figure A2.1 shows the current, best and worst sequence plotted. We notice that in this example

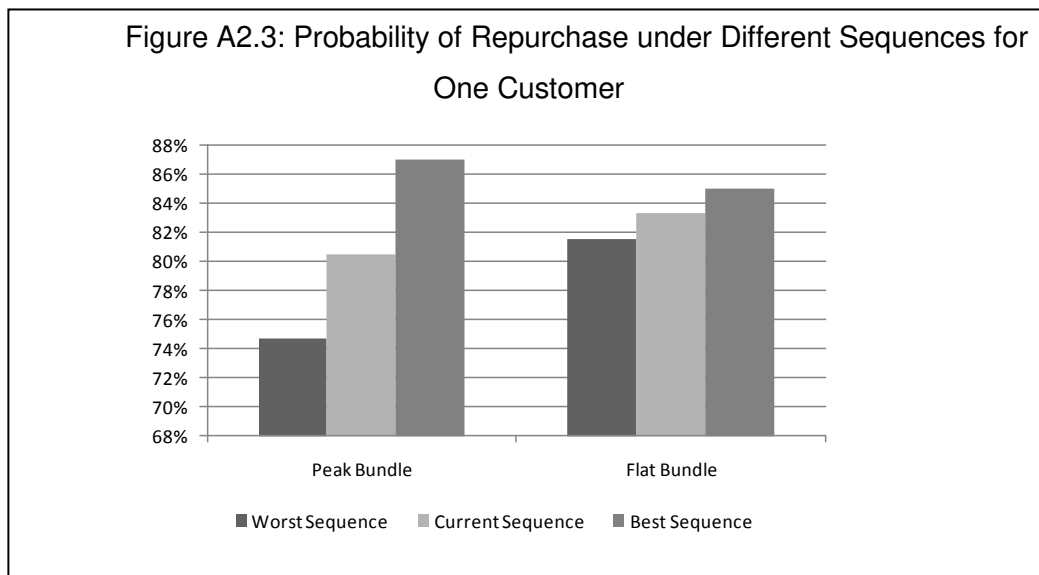


there is a clear peak (utility= 41) and the peak is placed at the end under the best sequence and at the beginning for the worst sequence.

For comparison, we can imagine a separate bundle that has events on the same days, but with different event utilities: 23, 17, 25, 21, 20, and 21. The event utilities are more homogenous and the bundle would be classified as “flat”. When we find the optimal sequence for this set of event utilities we see a different story. Figure A2.2 shows that rather than the peak event being placed at the end, the best sequence places the peak event (utility = 25) as the second event. These two examples illustrate that different solutions can be reached based on the different mix of the events within the bundles. In the first example, a clear peak was placed at the end of the sequence magnifying the end effect and the trend effect. With the second example, there was no clear peak, but in fact the two top events with relatively close utilities (25 and 23) get spread out across the sequence magnifying the days from peak to end impact. In the first example we can see an example of the peak and end effects, while in the second we see the spreading effect.



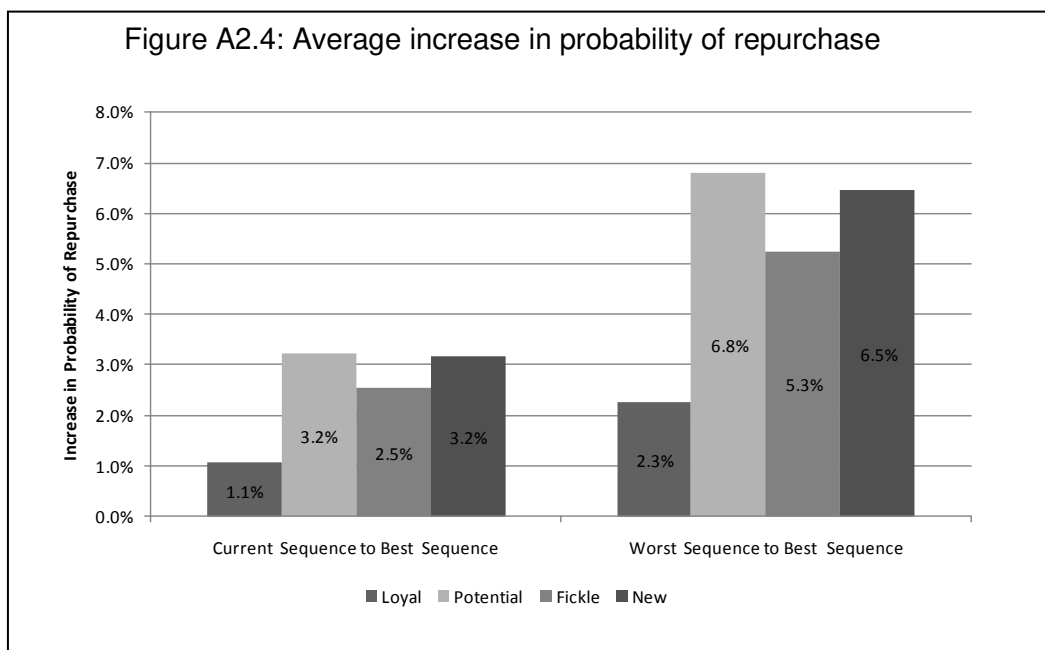
Within the dataset we find an individual customer with unknown gender who has purchased 3 bundles from only this one cycle 72 days before the first event all in the price category three, who has not purchased a membership, but is a loyal customer who has purchased the same cycle the past 3 years. For this customer we can show the probability of repurchase for the worst, current and best sequence using the coefficients estimated earlier. Figure A2.3 shows the probabilities of repurchase under the three sequences for both the examples used above, the “peak” bundle from Figure A2.1 and the “flat” bundle of Figure A2.2. The increase in probability of repurchase from the current to the best sequence is 7% for the “peak” bundle, but only 2% for the “flat” bundle. It appears from this example that an improvement in the sequence for a “peak” bundle is much more impactful than an improvement in a “flat” bundle.



Following the same procedures, we have found optimal sequences for all the bundles with less than 8 events. Since solving for the optimal sequence was not the objective of this paper, we stopped at bundles with 7 events leaving us with a total of 19,606 observations from

98 bundles for which we had found the probability of repurchase under the current sequence, the worst sequence, and the best sequence.

Across this sample, we experience an average increase of 2% of repurchase probability from the current sequence to the best (68% to 70%) and 4% from the worst sequence to the best (66% to 70%). Among *loyal* customers, the increase is smaller, 1% and 2%, but for the remaining segments we see a much higher increase (see Figure A2.4).



To further the illustration, we arbitrarily choose a cutoff point of probability for which we believe that a customer will repurchase. For illustration we choose 50% as our cut off, i.e., for any customer whose predicted probability of repurchase is greater than 50%, we believe that they will repurchase. Figure A2.5 shows the results in the percentage of repurchases given the 50% cutoff for the worst, current, and best sequences. On average, 3.7% representing 725 total customers move from not repurchasing to repurchasing by moving from the current sequence to the best sequence. This number is purely illustrative, since the cutoff that we chose may not be appropriate; however, it illustrates the impact that the sequence has on the probability of

repurchase in our model. *Loyal* customers don't show any increase, i.e., there are no *loyal* customers who have a probability of repurchase lower than 50% even with the worst possible sequence. However, among *potentially loyal* customers, only 55% will purchase under the current sequence while nearly 71% will purchase under the best sequence. Customers in *flat* bundles show an increase of 1.9%, but *peak* and *valley* bundles show a much higher increase (9.3% and 8.1%) illustrating again that homogenous bundles do not benefit from an improved sequence as much those with more variability in event utility.

Figure A2.5: Percentage of repurchases given a cutoff value of 50% probability

		Worst	Current	Best	Change from Current to Best Sequence
Customer Type	Loyal	100.0%	100.0%	100.0%	0.0%
	Potential	41.1%	55.2%	70.9%	15.8%
	Fickle	86.0%	92.0%	94.5%	2.5%
	New	13.5%	19.1%	28.2%	9.1%
Sequence Type	Flat	74.6%	76.3%	78.2%	1.9%
	Peak	45.8%	51.1%	60.3%	9.3%
	Valley	68.2%	79.8%	88.0%	8.1%
	Grand Total	68.8%	71.8%	75.5%	3.7%

APPENDIX 3:

MATHEMATICAL REPRESENTATION OF INTERRELATED BUNDLE AND SCHEDULING PROBLEM

In this section we present the mathematical representation of the problem of scheduling events into bundles, datetimes, and halls with an objective of maximizing explicit sequence effects across all bundles. The problem becomes one of not only timetabling, but also event bundling. Each event can only be scheduled on one datetime and in one hall, but may be able to be scheduled across multiple bundles. We take an integer programming approach and strive to express the problem linearly when possible. We will explain different aspects of the notation starting with indices, constants, and sets used, followed by a definition of the decision variables. Next we discuss the objective statement and explicitly define sequences effects. Finally, we present the constraints of the model.

Indices

The following indices are used to index events, bundles, datetimes, halls, clusters, and event orders. We will define a cluster shortly. The event order in a bundle and cluster are explicitly defined in order to define first and last events used in both the objective statement and a number of constraints. The indices are:

e, e' - events;

b - bundles;

d, d' - date times;

h - event halls;

c - event clusters;

o - event order in a bundle; and

p - event order in a cluster.

Constants

We assume that event planners will provide a number of constants used to define constraints in the model. As stated in the body of this paper, we define event utility as an aggregate measure of event value or popularity in comparison to other all other events regardless of when, where, or with what other events the event is scheduled and bundled. We assume that event planners can derive this utility for each event using inside industry knowledge, forecasting past purchase data, or with a survey designed to elicit a utility measure for each event from customers. While the current model will not consider it, we hope to consider the impact of schedules with different utilities across customer segments in future iterations. In this paper, we assume that there is one utility measure for each event, which is known a priori:

$$utility_e = \text{utility of event } e.$$

Also predefined are explicit descriptions of bundle requirements — the number of days that is required between each event in a bundle and the maximum and minimum number of events required in a bundle:

$$separate_b = \text{minimum allowable time between events in bundle } b;$$

$$n_b^+ = \text{maximum number of events in bundle } b; \text{ and}$$

$$n_b^- = \text{minimum number of events in bundle } b.$$

Similarly, we define the minimum and maximum number of events in a cluster and the number of days between events in the same cluster. A cluster is a set of events that have to be performed close to one another, primarily because the events in a cluster are actually several showing of the same performance. Therefore, not only are we concerned about the minimum number of days between events in a cluster, we also need to make sure that maximum number of

days between events is not exceeded and the number of days from the first to the last event in a cluster is controlled:

n_c^+ = maximum number of shows in cluster c ;

n_c^- = minimum number of shows in cluster c ;

$days_c^-$ = minimum number of days between events in cluster c ;

$days_c^+$ = maximum number of days between events in cluster c ; and

$SpreadDays_c$ = total allowable number of days from the first to the last event in cluster c .

The weights of the different sequence effects can be derived from econometric modeling similar to what has been done in Dixon and Verma (2011), mainly by estimating coefficients for the separate sequence effects in an econometric methodology. When we begin to solve the problem, we will set the weights according to the coefficients that Dixon and Verma (2011) estimated. Left for future research is different weights for different customer segments. The appendix 4 discusses the details of solving the problem describes the method we used in solving for the individual effect weights by essentially setting them to be equal to one another at the maximum feasible value of each effect. So:

w_1 = weight of the End Effect portion of the Sequence Effects;

w_2 = weight of the Peak Effect portion of the Sequence Effects;

w_3 = weight of the Spreading Effect portion of the Sequence Effects; and

w_4 = weight of the Trend Effect portion of the Sequence Effects.

For purposes used in constraint building, we define the last possible date that an event could be scheduled as well as the date part for each datetime. Finally, each event is given specific requirements on the number of bundles that it can be a part of:

LastSeasonDate = the last possible date an event could be scheduled;
Date_d = the date part (month, day, year) of datetime *d*;
EinB_e⁺ = the maximum number of bundles that event *e* can be a part of; and
EinB_e⁻ = the minimum number of bundles that event *e* can be a part of.

Sets

We begin by identifying the following sets:

E = set of all events;
B = set of all bundles;
D = set of all datetimes; and
H = set of all event halls.

Of interest among the sets of the model are the “Possible” sets. They are a subset of larger sets restricted in order to implicitly maintain constraint parameters. For example, there may exist a series of event level requirements for membership into a certain bundle: the event may need to be the appropriate genre or theme to be considered appropriate for the bundle. Events may require specific space or equipment only available in a subset of halls or events may only be allowed to be scheduled on weekends, matinees, or certain times of the year (holiday concerts). All of these constraints can implicitly be maintained by predefining possible or allowable event sets for which constraints will sum over. All possible sets are a subset of the correlating larger set:

PossibleE_b = events that could be scheduled in bundle *b*
 (correct genre, theme, artist, etc);
PossibleD_e = date times that could be scheduled for event *e*
 (correct day of the week, time of the day, specific dates, etc); and
PossibleH_e = event halls that could be scheduled for event *e*
 (proper stage, equipment, performers preference etc).

Cluster membership is determined a priori. All events that do not have cluster membership are placed into one cluster which has non-restricting constraints:

$Cluster_c$ = set of events that are possible in cluster c .

In determining appropriate schedules for halls, the following set provides the set of events that are allowed to be scheduled on a datetime in a hall given an event is scheduled on a different datetime in the same hall. This set is determined a priori because it considers the length of each event and the amount of time it takes to prepare the hall for another event. For example, if a lengthy show is scheduled at a time early in the afternoon, an evening event is not feasible; however, if a short show is scheduled early, perhaps an evening show is allowable:

$NotAvailable_{dd'h}$ = set of events that cannot be scheduled on datetime d' in hall h given event e is scheduled on datetime d in hall h .

Variables

The two primary variables that determine event bundling and scheduling are binomial integer variables indexed across five indices. The first, $BundleOrder$ is indexed across all events, bundles, orders, datetimes, and halls and the second, $ClusterOrder$, is similar except it indexes across clusters and cluster order instead of bundle and bundle order. A series of constraints built around these two variables maintain their singularity in event datetimes and halls:

$$BundleOrder_{ebodh} = \begin{cases} 1 & \text{if event } e \text{ is in bundle } b \text{ and in the } o^{th} \text{ order, scheduled on datetime } d \text{ in hall } h, \\ 0 & \text{otherwise; and} \end{cases}$$

$$ClusterOrder_{ecpdh} = \begin{cases} 1 & \text{if event } e \text{ is in cluster } c \text{ and in the } p^{th} \text{ order, scheduled on datetime } d \text{ in hall } h, \\ 0 & \text{otherwise.} \end{cases}$$

Two other integer variables indicate if an event has the highest utility (*Peak*) and if it has the latest datetime (*Last*) among all the events of the same bundle. As you will see, these variables are defined by the bundle order variables and as such are not explicit decision variables:

$$Peak_{eb} = \begin{cases} 1 & \text{if event } e \text{ is the peak event in bundle } b \\ 0 & \text{otherwise; and} \end{cases}$$

$$Last_{eb} = \begin{cases} 1 & \text{if event } e \text{ is the last event in bundle } b \\ 0 & \text{otherwise.} \end{cases}$$

The remaining variables are derived from the above variables, but are useful in notation.

Their definitions will be explained in the constraint section:

N_b = count of events in bundle b ;

B_e = count of bundles that event e is scheduled in;

$EndEffect_b$ = the utility of the last event in bundle b ;

$PeakEffect_b$ = the utility of the peak event in bundle b ;

$SpreadEffect_b$ = the number of days from the peak event to the last event;

$AvgUtility_b$ = average event utility of events in bundle b ;

$DaysFromFirst_{eb}$ = days from the first event in bundle b to event e ;

$AvgDaysFromFirst_b$ = average days from the first event in bundle b ; and

$TrendEffect_b$ = slope of the utility and days from the first event.

Objective

The object of the model is to assign events to halls, datetimes, and bundles in such a way to maximize the sequence effects within each bundle and across all bundles. The four sequence effects are the end effect, the peak effect, the spread effect and the trend effect. Each effect will be explicitly defined below. The model allows for different weights for each effect:

$$\max \sum_{\forall b} (w_1 EndEffect_b + w_2 PeakEffect_b + w_3 SpreadEffect_b + w_4 TrendEffect_b) \quad (2)$$

Sequence Effect Definition Constraints

There can only be one peak event per bundle and there can only be one last event per bundle:

$$\sum_{e \in PossibleE_b} Peak_{eb} = 1, \quad \forall b; \quad (3)$$

$$\sum_{e \in PossibleE_b} Last_{eb} = 1, \quad \forall b. \quad (4)$$

The peak event utility is the largest utility among events scheduled in bundle b . The last event date is the largest event date among events scheduled in bundle b . These constraints check that each event's utility and date is less than or equal to the peak and last for a given bundle. Note that because decision variables (Bundle Order) are on both sides of the equation, these definition are non-linear.

$$\sum_{e' \in PossibleE_b} \left((Peak_{e'b}) \left(\sum_{o=1}^{n_b^+} \sum_{d \in D} \sum_{h \in H} (BundleOrder_{e'bodh}) (Utility_{e'}) \right) \right) \geq \sum_{o=1}^{n_b^+} \sum_{d \in D} \sum_{h \in H} (BundleOrder_{ebodh}) (Utility_e),$$

$\forall b, \forall e \in PossibleE_b;$

(5)

$$\sum_{e' \in PossibleE_b} \left((Last_{e'b}) \left(\sum_{o=1}^{n_b^+} \sum_{d \in D} \sum_{h \in H} (BundleOrder_{e'bodh}) (Date_d) \right) \right) \geq \sum_{o=1}^{n_b^+} \sum_{d \in D} \sum_{h \in H} (BundleOrder_{ebodh}) (Date_d),$$

$\forall b, \forall e \in PossibleE_b.$

(6)

Therefore, the peak effect and end effect are defined as the utility of the peak event and the last event:

$$\sum_{e \in \text{Possible}E_b} \left((Peak_{eb}) \left(\sum_{o=1}^{n_b^+} \sum_{d \in D} \sum_{h \in H} (BundleOrder_{ebodh})(Utility_e) \right) \right) = PeakEffect_b, \forall b; \quad (7)$$

$$\sum_{e \in \text{Possible}E_b} \left((Last_{eb}) \left(\sum_{o=1}^{n_b^+} \sum_{d \in D} \sum_{h \in H} (BundleOrder_{ebodh})(Utility_e) \right) \right) = EndEffect_b, \forall b. \quad (8)$$

The spread effect is defined as the number of days between the date of the last event and the date of the peak event:

$$\sum_{e \in \text{Possible}E_b} \left((Last_{eb}) \left(\sum_{o=1}^{n_b^+} \sum_{d \in D} \sum_{h \in H} (BundleOrder_{ebodh})(Date_d) \right) \right) - \sum_{e \in \text{Possible}E_b} \left((Peak_{eb}) \left(\sum_{o=1}^{n_b^+} \sum_{d \in D} \sum_{h \in H} (BundleOrder_{ebodh})(Date_d) \right) \right) = SpreadEffect_b, \forall b. \quad (9)$$

N_b is the sum of all events scheduled in bundle b :

$$\sum_{e \in \text{Possible}E_b} \sum_{o=1}^{n_b^+} \sum_{\forall d} \sum_{\forall h} BundleOrder_{ebodh} = N_b, \forall b.$$

The average bundle utility is calculated by summing the utility of all events in a bundle and dividing by the number of events in the bundle:

$$\frac{1}{N_b} \sum_{e \in \text{Possible}E_b} \sum_{o=1}^{n_b^+} \sum_{d \in D} \sum_{h \in H} (BundleOrder_{ebodh})(Utility_e) = AvgUtility_b, \forall b. \quad (10)$$

The number of days from the first event can be calculated by subtracting each event's date from the event date of the event scheduled in the order number 1:

$$\sum_{o=1}^{n_b^+} \sum_{d \in D} \sum_{h \in H} (BundleOrder_{ebodh})(Date_d) - \sum_{d \in D} \sum_{h \in H} (BundleOrder_{eb,1,dh})(Date_d) = DaysFromFirst_{eb}, \forall e, b. \quad (11)$$

The average number of days from the first event in the bundle is calculated by summing all the number of days from the first event and dividing by the number of events in the bundle:

$$\frac{1}{N_b} \sum_{e \in PossibleE_b} \sum_{o=1}^{n_b^+} \sum_{d \in D} \sum_{h \in H} (BundleOrder_{ebodh})(DaysFromFirst_{eb}) = AvgDaysFromFirst_b, \forall b. \quad (12)$$

Finally, the trend effect is calculated as the linear slope of a line that best fits the points of utility and days from the first event. The line is fit under ordinary least squares and the equation for the slope of the line is as follows:

$$\frac{\sum_{e \in PossibleE_b} \sum_{o=1}^{n_b^+} \sum_{d \in D} \sum_{h \in H} (BundleOrder_{ebodh})(utility_e - AvgUtility_b)(DaysFromFirst_{eb} - AvgDaysFromFirst_b)}{\sum_{e \in PossibleE_b} \sum_{o=1}^{n_b^+} \sum_{d \in D} \sum_{h \in H} (BundleOrder_{ebodh})(utility_e - AvgUtility_b)^2} = TrendEffect_b, \forall b; \quad (13)$$

or

$$\frac{\sum (x - \bar{x})(y - \bar{y})}{\sum (x - \bar{x})^2}. \quad (14)$$

These equations explicitly define what we mean by peak, end, spreading, and trend effect. As noted earlier, it is difficult to describe some of the effects in a purely linear fashion. In the following sections we will define the constraints of the problem, some of which continue to be difficult to express linearly. While a linear approximation would aid in solving the problem, we are more concerned about expressing the actual complexity of the problem than with solvability.

Constraints

As a reminder, N_b is the sum of all events scheduled in bundle b \therefore

$$\sum_{e \in \text{Possible}E_b} \sum_{o=1}^{n_b^+} \sum_{\forall d} \sum_{\forall h} \text{BundleOrder}_{ebodh} = N_b, \quad \forall b. \quad (15)$$

The number of events in bundle b has to be between the minimum allowable number of events and the maximum number of events:

$$n_b^- \leq N_b \leq n_b^+, \quad \forall b. \quad (16)$$

Similarly, the number of shows scheduled in cluster c has to be at least the minimum number allowable:

$$\sum_{e \in \text{Cluster}_c} \sum_{p=1}^{n_c^+} \sum_{\forall d} \sum_{\forall h} \text{ClusterOrder}_{ecpdh} \geq n_c^-, \quad \forall c. \quad (17)$$

Each event can only have one order in the same bundle. Notice this constraint is not a restriction on the number of bundles in which an event can be a member:

$$\sum_{o=1}^{n_b^+} \sum_{\forall d} \sum_{\forall h} BundleOrder_{ebodh} \leq 1, \forall b, \forall e \in PossibleE_b. \quad (18)$$

Similarly, each event can only have one order in the same cluster:

$$\sum_{p=1}^{n_c^+} \sum_{\forall d} \sum_{\forall h} ClusterOrder_{ecpdh} \leq 1, \forall c, \forall e \in Cluster_c. \quad (19)$$

Event order is determined by the event date, i.e., earlier event dates have earlier event orders. Event order $(o+1)$ must have a larger date than order o . The number of days between events in the same bundle has to be greater than or equal to a separator constant for that bundle. If $(o+1)$ is not scheduled then the constraint is satisfied taking the difference between the Event date of order o and the last known season date plus the separation amount.

$$\begin{aligned} & (LastSeasonDate + separate_b) \left(1 - \sum_{e' \in PossibleE_b | e' \neq e} \sum_{\forall d} \sum_{\forall h} BundleOrder_{e'b,o+1,dh} \right) \\ & - \sum_{e' \in PossibleE_b | e' \neq e} \sum_{\forall d} \sum_{\forall h} (BundleOrder_{e'b,o+1,dh})(Date_d) \\ & - \sum_{e \in PossibleE_b} \sum_{\forall d} \sum_{\forall h} (BundleOrder_{ebodh})(Date_d) \geq separate_b, \forall b, o = 1, \dots, n_b^+ - 1. \end{aligned} \quad (20)$$

A similar constraint exists for cluster orders:

$$\begin{aligned} & (LastSeasonDate + days_c^-) \left(1 - \sum_{e' \in Cluster_c | e' \neq e} \sum_{\forall d} \sum_{\forall h} ClusterOrder_{e'c,p+1,dh} \right) \\ & - \sum_{e' \in Cluster_c | e' \neq e} \sum_{\forall d} \sum_{\forall h} (ClusterOrder_{e'c,p+1,dh})(Date_d) \\ & - \sum_{e \in Cluster_c} \sum_{\forall d} \sum_{\forall h} (ClusterOrder_{ecpdh})(Date_d) \geq days_c^-, \forall c, p = 1, \dots, n_c^+ - 1. \end{aligned} \quad (21)$$

An additional constraint for cluster orders is included to ensure that the maximum number of days between events is not violated. If $(p+1)$ is not scheduled the constraint is

satisfied by subtracting the event date from itself, resulting in 0 which will be less than the maximum number of days between events:

$$\begin{aligned}
& \sum_{e' \in \text{Cluster}_c | e' \neq e} \sum_{\forall d} \sum_{\forall h} (\text{ClusterOrder}_{ecpdh})(\text{Date}_d) \left(1 - \sum_{e' \in \text{Cluster}_c | e' \neq e} \sum_{\forall d} \sum_{\forall h} \text{ClusterOrder}_{e'c,p+1,dh} \right) \\
& - \sum_{e' \in \text{Cluster}_c | e' \neq e} \sum_{\forall d} \sum_{\forall h} (\text{ClusterOrder}_{e'c,p+1,dh})(\text{Date}_d) \\
& - \sum_{e \in \text{Cluster}_c} \sum_{\forall d} \sum_{\forall h} (\text{ClusterOrder}_{ecpdh})(\text{Date}_d) \leq \text{days}_c^+, \quad \forall c, p = 1, \dots, n_c^+ - 1.
\end{aligned} \tag{22}$$

Bundle event order o must be scheduled if order $o+1$ is scheduled. Similarly, if order o is not scheduled, order $o+1$ cannot be scheduled:

$$\sum_{e \in \text{Possible}E_b} \sum_{\forall d} \sum_{\forall h} \text{BundleOrder}_{ebodh} \geq \sum_{e' \in \text{Possible}E_b} \sum_{\forall d} \sum_{\forall h} \text{BundleOrder}_{e'b,o+1,dh}, \quad \forall b, o = 1, \dots, n_b^+ - 1. \tag{23}$$

Similarly, cluster event order p must be scheduled if order $p+1$ is scheduled:

$$\sum_{e \in \text{Cluster}_c} \sum_{\forall d} \sum_{\forall h} \text{ClusterOrder}_{ecpdh} \geq \sum_{e' \in \text{Cluster}_c} \sum_{\forall d} \sum_{\forall h} \text{ClusterOrder}_{e'c,p+1,dh}, \quad \forall c, p = 1, \dots, n_c^+ - 1. \tag{24}$$

An example of how constraints 19 to 23 ensure the ordering of events within bundles and clusters is included in the final section of this appendix.

The days between the first event in a cluster and all other events in a cluster cannot be more than the cluster spread:

$$\begin{aligned}
& \sum_{e \in \text{Cluster}_c} \sum_{\forall d} \sum_{\forall h} (\text{ClusterOrder}_{ecpdh})(\text{Date}_d) - \sum_{e \in \text{Cluster}_c} \sum_{\forall d} \sum_{\forall h} (\text{ClusterOrder}_{ec,1,dh})(\text{Date}_d) \leq \text{SpreadDays}_c, \\
& \forall c, p = 2, \dots, n_c^+.
\end{aligned} \tag{25}$$

Events in the same cluster cannot appear in the same bundle:

$$\sum_{e \in Cluster_c} \sum_{o=1}^{n_b^+} \sum_{\forall d} \sum_{\forall h} BundleOrder_{ebodh} \leq 1, \quad \forall b, \forall c. \quad (26)$$

B_e is the sum of all bundles that have event e scheduled:

$$\sum_{b \in B} \sum_{o=1}^{n_b^+} \sum_{d \in PossibleD_e} \sum_{h \in PossibleH_e} BundleOrder_{ebodh} = B_e, \quad \forall e. \quad (27)$$

If an event is scheduled in a cluster on a datetime and in a hall, it must be scheduled on the same datetime and hall in a bundle. Similarly, each event can be scheduled for only one allowable datetime and in only one allowable hall. An event can be placed in multiple bundles, but must be scheduled for the same datetime and hall in each bundle. While an event can be in multiple bundles, it is pre-assigned into only one cluster:

$$\sum_{b \in B} \sum_{o=1}^{n_b^+} BundleOrder_{ebodh} \geq \sum_{p=1}^{n_c^+} ClusterOrder_{ecpdh}, \quad \forall c, d, h, e \in Cluster_c; \quad (28)$$

$$\sum_{b \in B} \sum_{o=1}^{n_b^+} BundleOrder_{ebodh} \leq B_e \sum_{p=1}^{n_c^+} ClusterOrder_{ecpdh}, \quad \forall c, d, h, e \in Cluster_c. \quad (29)$$

Events have to be scheduled between the minimum and the maximum number of bundles. This constraint can force an event to be scheduled if the minimum constraint is greater than 0. It also ensures that an event doesn't get scheduled into too many bundles:

$$EinB_e^- \leq B_e \leq EinB_e^+, \quad \forall e. \quad (30)$$

Each datetime can have no more than one event scheduled per hall:

$$\sum_{e \in E} BundleOrder_{ebodh} \leq 1, \quad \forall b, o, d, h; \text{ and} \quad (31)$$

$$\sum_{e \in E} ClusterOrder_{ecpdh} \leq 1, \quad \forall c, p, d, h. \quad (32)$$

Events scheduled in the same hall cannot overlap in time. Recall the set $NotAvailable_{dd'h}$ consists of all events that cannot be scheduled if event e is scheduled in hall h , datetime d . If event e is not scheduled then $B_e = 0$ and the constraints are satisfied:

$$(B_e) \left[\sum_{e' \in NotAvailable_{dd'h}} \sum_{b \in B} \sum_{o=1}^{n_b^+} \sum_{d' \in D} (BundleOrder_{e'bod'h}) \right] = 0, \quad \forall e, d, h; \quad (33)$$

$$(B_e) \left[\sum_{e' \in NotAvailable_{dd'h}} \sum_{c \in C} \sum_{p=1}^{n_c^+} \sum_{d' \in D} (ClusterOrder_{e'cpd'h}) \right] = 0, \quad \forall e, d, h. \quad (34)$$

Event Ordering Example

Since the event ordering is key to our problem, we provide an example of how (19) thru (23) maintain proper orders. For simplicity, assume we have a bundle with three events, event 1, 2, and 3. Event 1 is scheduled on *Date* 100, event 2 on *Date* 50 and event 3 on *Date* 75 as shown on Table A3.1. Also assume that the *LastSeasonDate* is 360 and the events must be separated by 25 days, i.e., $separate_b = 25$. Further assume that these three events are the only three that are

Table A3.1: Events and their scheduled date	
Event	Scheduled on
1	100
2	50
3	75

possible to be scheduled in this bundle and that $n^+ = 3$.

The constraint (19) for $o = 1$. (19) reads as follows:

$$\begin{aligned}
 & (LastSeasonDate + separate_b) \left(1 - \sum_{e' \in PossibleE_b | e' \neq e} \sum_{\forall d} \sum_{\forall h} BundleOrder_{e', b, o+1, dh} \right) \\
 & - \sum_{e' \in PossibleE_b | e' \neq e} \sum_{\forall d} \sum_{\forall h} (BundleOrder_{e', b, o+1, dh})(Date_d) \\
 & - \sum_{e \in PossibleE_b} \sum_{\forall d} \sum_{\forall h} (BundleOrder_{e, bodh})(Date_d) \geq separate_b, \quad \forall b, o = 1, \dots, n_b^+ - 1.
 \end{aligned}$$

The first section will resolve to 0 since $o = 2$ is scheduled: $(360 + 25)(1-1) = 0$. We will return later to discuss what happens if the $(o+1)$ is not scheduled. Next, the *Date* for the event of $o=2$ is subtracted from the *Date* for the event of $o = 1$ and it must be greater than or equal to 25.

Consider all combinations:

Table A3.2: All possible combinations for $o=2$ and $o=1$

event $o=2$	event $o=1$	$ \sum_{e' \in PossibleE_b e' \neq e} \sum_{\forall d} \sum_{\forall h} (BundleOrder_{e', b, o+1, dh})(Date_d) \\ - \sum_{e \in PossibleE_b} \sum_{\forall d} \sum_{\forall h} (BundleOrder_{e, bodh})(Date_d) $
1	2	$100 - 50 = 50$
1	3	$100 - 75 = 25$
2	1	$50 - 100 = -50$
2	3	$50 - 75 = -25$
3	1	$75 - 100 = -25$
3	2	$75 - 50 = 25$

Only event 2 or event 3 as $o=1$ will satisfy this constraint. Now solve for $o=2$. Since $o=3$ must be scheduled, the first part resolves to 0 again. Again, considering all combinations of $o=2$ and $o=3$ we get the following table:

Table A3.3: All possible combinations for $o=3$ and $o=2$

event $o=3$	event $o=2$	$\sum_{e \in \text{Possible}E_t, e \neq e} \sum_{\forall d} \sum_{\forall h} (\text{BundleOrder}_{e', b, o+1, dh})(\text{Date}_d)$ $- \sum_{e \in \text{Possible}E_t} \sum_{\forall d} \sum_{\forall h} (\text{BundleOrder}_{e, bodh})(\text{Date}_d)$
1	2	$100 - 50 = 50$
1	3	$100 - 75 = 25$
2	1	$50 - 100 = -50$
2	3	$50 - 75 = -25$
3	1	$75 - 100 = -25$
3	2	$75 - 50 = 25$

Event 2 or 3 could be $o=2$. Since event 1 cannot be $o=1$ or $o=2$ it must be $o=3$. If $o=2$ was event 2 then there is not a solution that satisfies $o=1$ constraint (see above), therefore $o=2$ must be event 3.

The constraint is only created for $o=1..n_b^+ - 1$, so for this example we stop. However, now assume that there are 4 possible events that could be scheduled and $n^+=4$, but that event 4 is not scheduled. Create the constraint for $o=3$. The first part now resolves to $385 = (360+25)(1-0)$ since $o=4$ is not scheduled. $o=4$ can only be scheduled if $o=3$ is scheduled per (6) so we will never have a case where $o+1$ is scheduled but o is not. The remainder of the constraint will resolve to the *Date* of $o=3$ which could be any of events, even if the event is scheduled on the last day of the season. However, the other constraints ensure that it will be event 3. If we assume further that we have two unscheduled events and $n^+=5$, then the constraint for $o=4$ resolves to $385 \geq 25$ since neither $o=4$ or $o=5$ is scheduled.

APPENDIX 4:

DISCUSSION OF SIMULATED ANNEALING HEURISTIC ALGORITHM USED TO SOLVE THE INTERRELATED BUNDLING AND SCHEDULING PROBLEM

In this section we discuss the simulated annealing (SA) heuristic algorithm used to solve our problem. Simulated annealing was chosen because of the discrete nature of the solution (as opposed to continuous) and complexity of a solution. Other popular search heuristics, e.g. tabu search, genetic algorithm, require memory for multiple solutions. Our solution includes the bundle, date, and hall assignment for each event as well as the sequence characteristics of each bundle; because of the complexity in the solution, maintaining a large number of solutions may quickly exceed memory capacities. However, SA maintains only a current solution, the last solution, and a copy of the best solution.

Simulated annealing is based on the annealing process from metallurgy that allows a metal to cool at a controlled rate to ensure a more solid crystallization in its final structure. In simulated annealing, discrete non-linear solution spaces can be explored by allowing solution evolution to weaken the objective in hopes to break free from local optimum. SA includes a “cooling” parameter that controls the number of worse solutions that are maintained at any point; typically, this cooling parameter ensures that some final percentage of solution accepted are purely greedy, meaning only improved solutions are accepted.

Problem Generation

In the *Problem Generation* stage the problem is defined by specifying the number of events, bundles, datetimes, halls and clusters. At this stage we define the distribution of event utilities, the max number of events in bundles, the number of bundles each event can be

scheduled into, the number of events in each cluster, and define the events that have specific cluster membership. Additionally, the set of possible bundles, halls, datetimes, and cluster for each event are defined. Of particular interest is the form of the datetime arrays; a two dimensional Boolean (0 or 1, true or false) array is created for each event, bundle, hall, and cluster by all possible datetimes. If an event, bundle, hall, or cluster is allowed to be scheduled on a given datetime, the variable reads “true”, otherwise it reads “false”. These arrays will be changed as the algorithm proceeds and a second set of arrays keeps track of which events make a possible datetime infeasible. The bundle, hall, and cluster “spreads” are defined at this point which make clear how near other events are allowed in bundles, halls, and clusters.

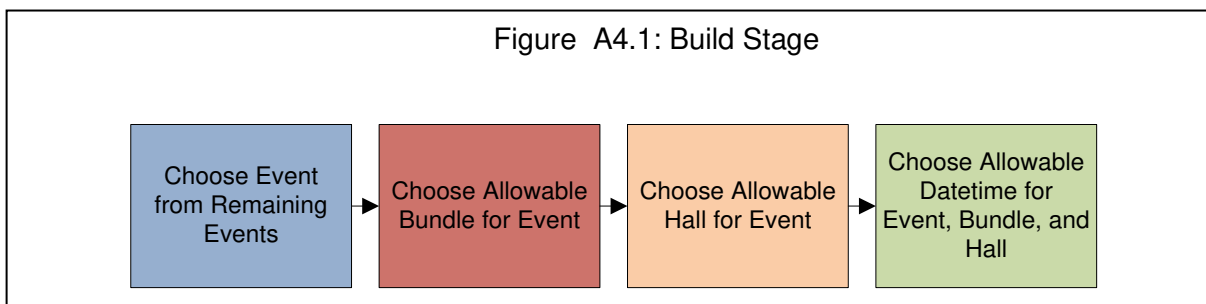
After a problem is generated, our algorithm has four stages: (1) a feasible solution is generated, (2) evaluated, (3) perturbed and (4) subsequently rebuilt. Steps 2, 3, and 4 repeat a predetermined number of iterations.

Build stage

The initial build stage creates a random feasible solution in order to begin the annealing algorithm. This is done in four stages, (1) event selection (2) bundle selection (3) hall selection and (4) datetime selection. An event is chosen from a list of *available events*; i.e., events that have not yet reached their maximum bundle membership quota. Event selection is done randomly with a bias towards those events that have more bundle and datetime restrictions. In this manner, events that are not flexible will likely be scheduled early in the build stage. A bundle and hall are randomly selected from the set of bundles and halls that are available for the specific event. Next, a datetime that is available for the selected event, bundle, and hall is

randomly selected. Finally, all availabilities are updated for future bundle, hall, and datetime selections to ensure constant feasibility during future iterations of the build stage.

The challenge of the build stage lies when there is no available bundle, hall, or datetime that is appropriate for the selected event. This becomes more likely as the schedule fills out. The algorithm can take several routes in order to fit an event into the schedule. Details of the different methods for the each stage within the build stage to find an appropriate assignment are detailed in later in this appendix in the section “Details on the Build Stage”; however the general principle is to find the reason why an assignment cannot be made and make a change to previous events that will allow both events (or perhaps several events deep) to maintain feasibility and still be assigned. After looking for an appropriate re-assignment schedule for some time, this re-scheduling effort will stop and a different event will be selected to proceed. In this manner, the build stage is not required to schedule every single event or fill all event bundle membership quotas. The build stage stops after the *remaining events* set is empty or after a certain number of attempts to schedule is reached.



Evaluation stage

Simulated annealing is based on the premise that solutions to np-hard problems are often not linearly improving as an algorithm progresses, i.e. there are both local optimal solutions and

global optimal solution and a strictly greedy evaluation could lead to a local optimum. Simulated annealing allows for both improvements and deterioration of the objective solution in order to search for a global optimum. This is accomplished by set a temperature parameter known as T that controls how and when a worse solution might be preferred.

Problem solutions consist of the all individual parts of the objective statement, as well as event datetime, bundle, and hall assignments. We track three separate solutions:

***The last solution** – the solution prior to the most recent perturb state;*

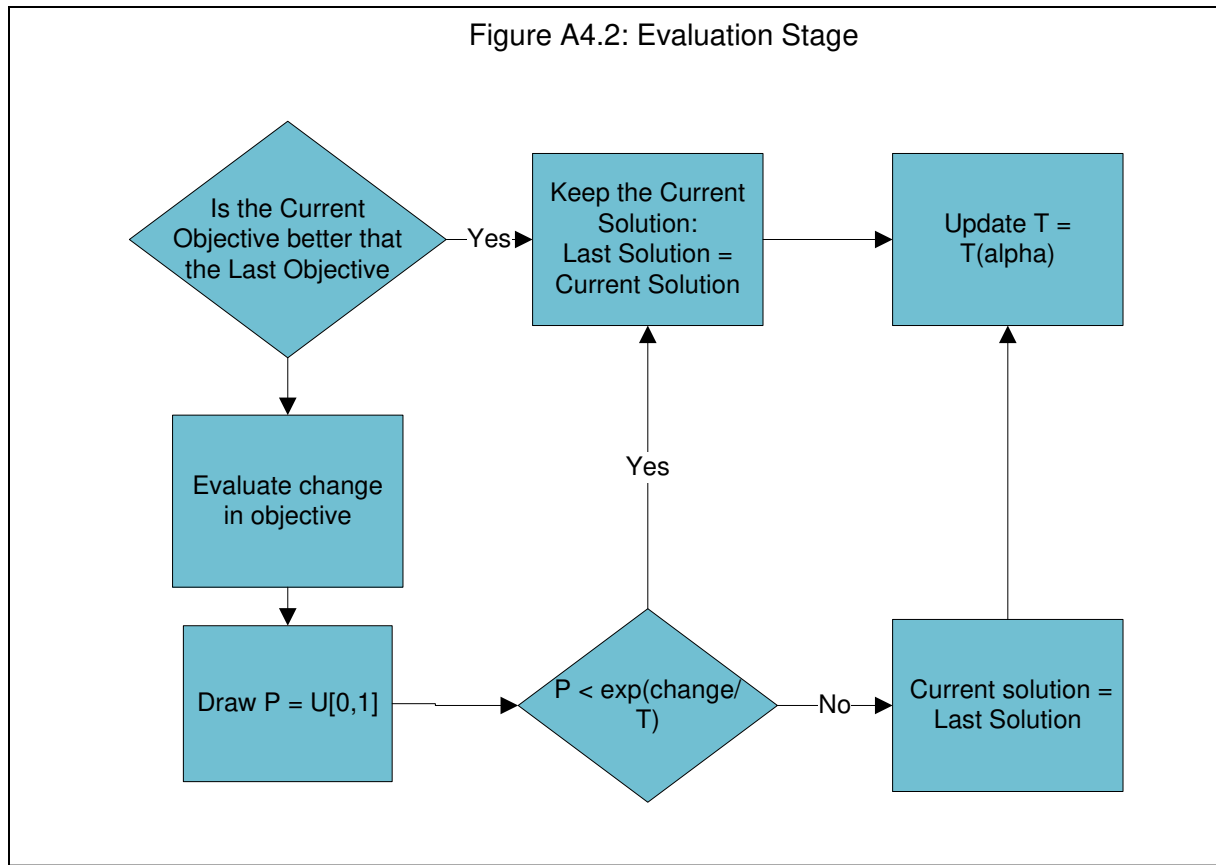
***The current solution** – the solution most recently produced by the re-build stage; and*

***The best solution** – the solution that has the highest objective statement.*

After the re-build stage (described below) completes, the current solution objective is compared with the last solution objective. If the current objective is better, then the solution is kept and the last solution becomes the current solution. If the current solution is better than the best solution, then the best solution becomes the current solution. However, if the current solution is worse than the last solution, then a random uniform number $U[0,1]$ P is drawn and the current solution is kept with probability $P < e^{(current\ objective - last\ objective)/T}$. T is controlled by a cooling factor $\alpha < 1$. At each iteration T is updated such that $T = T \cdot (\alpha)$. If T is large, then the current solution is nearly always kept, but as T gets smaller, it is less likely that large, worsening changes in the objective statement will be kept. Initial values of T are chosen to allow liberal exploration of the solution set in early stages of the algorithm and α is set to ensure a certain “cooling schedule”; i.e., to allow for a more greedy search to begin after a certain number of iterations. T and α are problem specific and a procedure for how they are determined is described later in this appendix.

If the current solution is chosen to remain, the last solution becomes the current.

Otherwise, the current solution reverts back to the last solution. Finally, the current solution is sent to the perturb stage.



Perturb stage

At this point, a random number between .005 and .01 is drawn to represent the percentage of events that will be unscheduled. The appropriate number of events are randomly drawn from scheduled events and unscheduled. The event is unscheduled from all bundles that it is a part of, the assigned hall, and datetime and is placed back into the set of remaining events. Additionally, the objective statement is updated as events are unscheduled. What remains is a partial solution with between .5% and 1% of events yet to be scheduled.

Re-Build Stage

The re-build stage is an attempt to consider a “neighboring” solution and can be accomplished simply by putting the partial solution generated by the perturb stage back in the build stage and rerun the build stage until it stops under the conditions stated above. This is not a complete rebuild, but just a partial rebuild only considering those events that were identified randomly in the perturb stage. However, because of the nature of the algorithm, other events may be unscheduled and rescheduled to make a place for the resultant perturbed events. Because of the probabilistic nature of the algorithm, the new solution will almost assuredly be different from the last solution and the current solution is looped back into the evaluation stage. To be sure the rebuilt solution is different than the last solution, the rebuild method will not allow the new solution to be identical to the last solution and will resort to leaving the perturbed events unscheduled if there are no other alternatives.

Algorithm Results

The entire evaluate-perturb-re-build loop stops once the parameter T reaches a predefined small number. In our tests of the algorithm, we set T and α so that the algorithm will cool at a rate to allow for it to conclude after 2 million iterations. To evaluate the algorithm’s effectiveness, we consider the solution of a problem with a size comparable to that of the concert venue: 200 events, 50 bundles, 6 halls, and 300 datetimes. We allow the problem to be relatively unconstrained in that all events can be scheduled in any datetimes and all events can be scheduled into any bundles. Next we run the problem through the *Build* stage from an empty solution 1000 times and capture the resulting objective providing us with a sample of random feasible solutions. Next we solve the problem 30 times with the entire algorithm and compare the

results. If we scale the *maximum* objective found in the first 1000 random solution to equal to 1, we find the average of the random solutions equal to .84 with a standard deviation of 0.05. Using the same scaling, the maximum solution found by the algorithm was equal to 1.32, the average equal to 1.28, and the standard deviation equal to 0.02. This means that the max objective found in 30 iterations of the algorithm was 132% larger than the max objective found by sampling 1000 feasible solutions. Similarly, the optimized mean was 128% larger than the sampled max, and the standard deviation has decreased. As tested, the algorithm can consistently find objectives that are significantly higher than a solution that can be found by extensive sampling. In short, the algorithm appears to be converging toward an optimal solution.

To determine a relative effectiveness of the solution generated by the algorithm, we can estimate a hypothetical high objective for the unconstrained problem. We do this in two ways, first assuming that the peak event is the last event for every bundle, we sort the event utilities and place the top 50 utility events into separate bundles. We then assume that each bundle has an event that is approaching 0 in event utility. Using the minimum number of days in a bundle calculated when determining the weight of the slope effect (see section titled “Starting Points for weights” later in this appendix), we assume that the peak/end event is as close as possible and calculate a slope which simplifies to the peak/end event utility divided by the minimum number of days in a bundle. Finally, we calculate the peak, end, and trend effect for each bundle and sum across all bundles. This approach sets the spread effect equal to 0 across all bundles since the peak event is the end event.

The second approach in estimating a high objective incorporates the spread effect. First we assign the two highest utility events into the first bundle, the next two highest into the next bundle, and so on until all 50 bundles have two events that are very close in utility and as high as

possible. We assume that the higher utility event is scheduled on day 0 while the other is scheduled on the last day of the season, i.e., we assume the spread for each bundle is the maximum number of days in the problem. Although all 50 events cannot be scheduled on these two days, we make this assumption for the sake of simplicity in calculation. Since the two events are close to the same in utility, we assume the slope is near zero. Finally, we can calculate our bundle objectives and sum across all bundles.

The two approaches produce similar scores: for our unconstrained test problem the slope-only approach generated a score of 262.07 and the spread approach generated 278.29. For the same problem, the algorithm consistently generates a solution near 270 giving us a degree of confidence in the algorithm's ability to resolve to high solutions.

While this analysis has been performed for one problem type, a similar automated analysis should be performed in order to ensure that all solutions to problems addressed in Chapter 3 are approaching near-optimal answers. Doing so will ensure that the solutions truly are a near-optimal answer to the problem and not just an artifact of the algorithm.

Details on the Build Stage

The *Build* stage begins by selecting an event from the set of Remaining Events. The Remaining Events set contracts as events are scheduled and can no longer be placed into additional bundles. The probability of selecting a given event is conditional on the event level requirements. For example, if an event can be scheduled on nearly any datetime, any bundle, and in any hall, it is less likely to be selected early in the algorithm. If an event has strict requirements it is given a higher probability. In this manner those events that have restrictive constraints have a higher chance of being scheduled before the schedule itself becomes

restrictive. The probability weights are updated as events leave (and re-enter) the Remaining Events set.

Next, a bundle is selected from the set of bundles that are allowable for the selected event. Again, the probability of bundle selection is based on the requirements of the bundles, mainly the constraint on the minimum number of events that need to be in each bundle. Those bundles with a higher minimum number of events have a higher probability of being selected. Next, a hall is randomly chosen among the allowable halls for the event. In this case there are no “hall requirements”, only event requirements on halls, so weighted probabilities are not used.

Finally, a datetime is selected randomly from those that are allowable for the specific event without regard to the events hall, bundle, or cluster schedules. The datetime is then tested for feasibility across all the hall, bundle, and cluster schedules. If the datetime is found to be infeasible in any of the other schedules, a new datetime is randomly selected from the remaining allowable datetimes. This random selection and testing is repeated until either a feasible datetime is found that fits across all dimensions or all available datetimes are ruled as infeasible. During the iterative attempts, the algorithm tracks the reason for infeasibility, be it hall, bundle, or cluster (or several) restrictions. If no dates are found to be available, the algorithm considers the leading cause of infeasibility and takes corrective action accordingly. For example, if the leading cause seems to be that the halls schedule is too constrictive, the algorithm will attempt to select a different hall for the event or select an event already scheduled in the hall and reschedule it into a different hall. Similarly, if the bundle schedule is too restrictive a different bundle can be considered or an event within the same bundle can be rescheduled into different bundles. If the cluster schedule is too restrictive, a different event in the cluster can be rescheduled.

Additionally, the algorithm can identify the specific events that have already been scheduled that

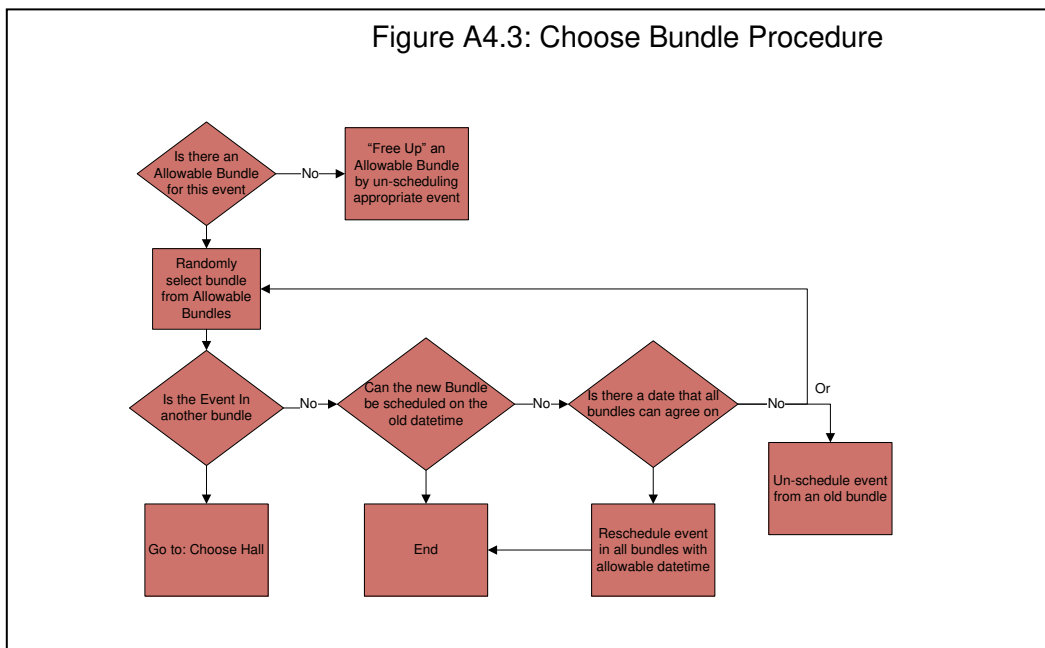
make a certain datetime infeasible for an event in a hall and bundle. The algorithm can then reschedule these events and attempt again to allow the new event to be scheduled on the certain datetime.

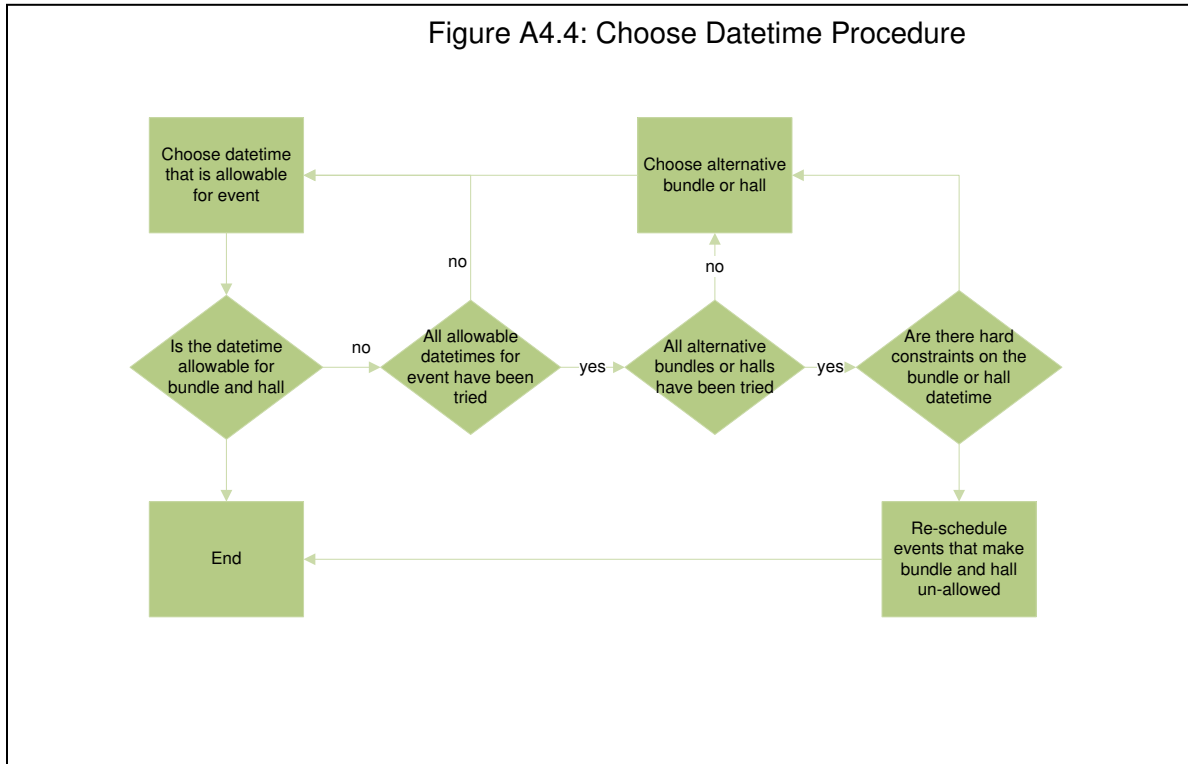
The algorithm iteration concludes by updating the objective statement for the selected bundle given the new event and its datetime. Finally, all the “allowable” sets are updated with the new information. The event is removed from the Remaining Events set if it can no longer be scheduled into another bundle. If the selected bundle has reached its maximum allowable event capacity it is removed from the set of possible bundles for events that allow it. The allowable datetimes are updated to conform to the constraints that restrict the number of days between event in bundles, hall turnover times, and cluster spreads. This iterative update of possible selections means that the algorithm can only choose feasible moves in developing the solution.

As the algorithm proceeds, an event that is allowed to be scheduled into multiple bundles can be drawn from the remaining events set; i.e., an event that has already been through the algorithm at least once. This event already has a datetime and hall assigned to it, but only needs another bundle assigned to it. In this case a new bundle is selected and the prior datetime is evaluated for feasibility across all the bundles. If the datetime is not feasible, then either another bundle is selected and re-evaluated or a new datetime is selected by creating and selecting from the set of datetimes that are feasible for the event, across all bundles, and in the hall. If a datetime is still not found, one can be forced through the process described above or the event is unscheduled from one of its previously scheduled bundles and datetimes are re-evaluated again.

The build phase of the algorithm concludes when there are no longer events in the remaining events set, or if a high number of attempts have been made on the same small number

of events. The solution is then run through a constraint checker to ensure that the solution is feasible. The constraint checking is really a check on the code to ensure that the build stage code is working as hoped, as most constraints are implicitly being kept by the restriction of selection sets during the algorithm. There are however, portions of the solution that are not implicitly handled, such as the minimum number of events in a bundle and cluster, and the minimum number of bundles for an event. We attempt to mitigate these constraints by incorporating the probabilistic selection criteria, but if these constraints are still not met, the algorithm will un-schedule a specific portion of the solution and rerun. The un-scheduling is handled much like the following *perturb* stage, but is more specific in the selection of which events get un-scheduled.





Starting points for weights for Peak, End, Spread, and Trend effects

The weights of the peak, end, spread, and trend effects were guided by the findings in Dixon and Verma (2011); however, changes have been made in order to determine what an optimal solution would be for specific interaction of the effects. The main findings in Dixon and Verma (2011) are that as the last event and peak event increase in utility, as the slope increases, and as the number of days from the peak to the end increase, probability of subscription repurchase increases. The positive nature of the spread effect (number of days from peak to end) could contradict or bound the effect of the trend and the end effect as it may place a peak event near the beginning of a schedule. We are interested in knowing when an optimal solution will place a peak event near the end as opposed to near the beginning, given near equal weights across all effects. For this reason, we set the weights of the four effects such that at an expected maximum value of their respective variables, the weights will produce nearly equal contribution to the objective statement. If we set the highest value of an event utility at 200, the maximum

value that the peak and end effect variables could take is 200; therefore, the weight associated with these two variables are set equal to one another. Similarly, the maximum number of days between a peak and the end is simplified to maximum number of days in the problem. The weight for this variable is then calculated so that when it is multiplied by the maximum number of days, it equals the weights of peak and end multiplied by the maximum utility event.

Estimating a maximum slope is a bit more difficult; first, a maximum slope would have events with the lowest and highest utilities, approaching 0 and 200. A bundle with a maximum slope would be one in which these are the only two events as close as possible. In reality, bundles have a minimum number of events and there are a minimum number of days required between events in a bundle. Assuming that an optimal slope would then be one in which the number of days between the first event and the last event is the shortest possible, we multiply the minimum number of events per bundle by the minimum number of days between events within a bundle. We can then calculate the slope of a line between the minimum utility event on day 0 and the maximum utility event on this shortest number of days calculated. Assuming the minimum event utility is approaching 0, this calculation can be simplified as the maximum utility divided by the shortest number of days possible. Finally, the weight for the slope variable is calculated so that when it is multiplied by the maximum possible slope the outcome will be equal to the weights of the other three weights multiplied by their respective maximums. As such, here are the weights and maximums used in the objective statement:

Table A4.1: Peak, End, Spread, and Slope Effect weights

	Weight	Maximum Possible Value	Product
Peak effect	0.015	200	3
End effect	0.015	200	3
Spread effect	0.01	300	3
Slope effect	2.25	1.333	3

The value of a product of three was arbitrarily chosen and is meaningless, except to say that the weight of the effects are in a similar scale as what was found in Dixon and Verma (2011). By normalizing the weights we hope to understand what conditions will lead to one effect dominating over another in an optimal solution.

Method for Determining T , α , and T_{stop}

The parameters T and α wholly determine the rate at which the algorithm will “cool” given that for each iteration k , $T_k = (T_{k-1})\alpha$ and $T_k = \alpha^k T_0$ when T is updated after every iteration. Additionally the algorithm will stop once $T < T_{stop}$. The starting T_0 should be set high enough in order to allow for a high percentage of worse solutions to be accepted in order to fully explore solution space; however, a T_0 too high will result in unnecessary iterations of the algorithm. An appropriate T_0 can be determined by estimating the expected average change in the objective statement and setting T_0 such that some high $P_0\%$ of worse solutions will be accepted. The probability of accepting a solution at iteration k is expressed by:

$$P_k = e^{(objective_k - objective_{k-1}) / T_k};$$

or

$$P_k = e^{(objective_k - objective_{k-1}) / T_0(\alpha^k)}.$$

If we can estimate an expected change in the objective statement the probability of accepting a worse solution at iteration 0 can be expressed by:

$$P_0 = e^{-Avg\Delta Objective / T_0}.$$

And finally, if P_0 is provided, the value of T_0 can be calculated by:

$$T_0 = \frac{-Avg\Delta Objective}{\ln(P_0)}.$$

We estimate an average change in objective by initially running the algorithm with an extremely high T for 2000 iterations (about one second in computing time on the unconstrained problem) and capture the change in objective. The average is then calculated and T_0 is determined with a probability $P_0 = 95\%$. Through trial and error, we have determined that a million iterations are satisfactory for finding consistent, near-optimal solutions. With this determination we can set α such that after some G iterations, the algorithm is mostly greedy, i.e., accepts some small P_G number of worse solutions:

$$P_G = e^{-Avg\Delta Objective / T_0(\alpha^G)}; \text{ and}$$

$$\alpha = \left[\frac{-Avg\Delta Objective}{(T_0)\ln(P_G)} \right]^{1/G}.$$

Finally, we define some number of iterations (*End*) to run after *G* and can calculate when the algorithm should stop:

$$T_{stop} = (\alpha^{G+End})T_0.$$

Algorithm stops once $T < T_{stop}$.

In solving our problems, we set $End = 100,000$, $P_0 = 95\%$, $P_G = 5\%$. This means that the algorithm will start with a T_0 that will allow 95% of worse solutions to be accepted. After 900,000 iterations, only 5% of worse solutions will be kept and an additional 100,000 iterations will run and will be mostly greedy in accepting solutions; i.e., it will mostly only accept an improvement in the objective statement.

Method For adjusting T , α , and T_{stop}

The previous method assumes that the average change in solutions is uniform across all areas of the solution space. For a problem as complex as ours, this assumption may not hold, which can lead to a disrupted cooling schedule. To be more exact in a cooling schedule, we iteratively check the probability of a worse solution being selected and compare it to what is expected at that iteration. From above we know that the probability of a worse solution being selected at iteration k can be estimated as follows:

$$P_k = e^{-Avg\Delta Objective/T_k}.$$

We capture the outcome of x number of previous worse solution decisions and determine the actual percentage of accepted solution:

$$Outcome_i = \begin{cases} 1 & \text{if } U[0,1] < e^{(objective_k - objective_{k-1})/T_k} \\ 0 & \text{otherwise} \end{cases} ; \text{ and}$$

$$P_{actual @ k} = \frac{1}{x} \sum_{i=0}^x Outcome_i.$$

Next, we reevaluate the average change in the objective for the current neighborhood. This gives us an estimate of what the average change is given the percentage of actual accepted solutions for a given T_k :

$$Avg\Delta Objective_k = (T_k) \ln(P_{actual @ k}).$$

Then we can reevaluate T_k with the new average change information. If the percentage of actual accepted solutions is too low (high), this will raise (lower) T_k :

$$T_k = \frac{-Avg\Delta Objective_k}{\ln(P_k)}.$$

And finally, we can update α_k and T_{stop} to ensure that the algorithm will stay on schedule to run for a total of $G + E$ iterations even with the change in T_k :

$$\alpha_k = \left[\frac{-Avg\Delta Objective_k}{(T_k) \ln(P_G)} \right]^{1/(G-k)} ; \text{ and}$$

$$T_{stop} = (\alpha_k^{G+E-k}) T_k.$$

For our problems, we adjust T_k 100 times — every 10,000 iterations — over the entire algorithm. We set the number of previous worse solution decisions, x to 500. This means that every 10,000 iterations, the previous 500 worse solution decisions are used to reevaluate the average change in objective and adjust T_k accordingly in order to maintain an expected cooling schedule that will result in a predefined probability of acceptance (P_G) at time G .

BIBLIOGRAPHY:

- Allon, G., & Bassamboo, A. (2009). Cheap talk in Operations: Role of intentional vagueness. *Consumer-Driven Demand and Operations Management Models*, 3–36.
- Anand, K. S., Paç, M. F., & Veeraraghavan, S. K. (2011). Quality-speed conundrum: Tradeoffs in customer-intensive services. *Management Science*, 57(1), 40–56.
- Anderson, N. H., & Barrios, A. A. (1961). Primacy effects in personality impression formation. *The Journal of Abnormal and Social Psychology*, 63(2), 346-350. doi:10.1037/h0046719
- Ariely, D. (1998). Combining experiences over time: The effects of duration, intensity changes and on-line measurements on retrospective pain evaluations. *Journal of Behavioral Decision Making*, 11(1), 19–45.
- Ariely, D., & Carmon, Z. (2000). Gestalt characteristics of experiences: the defining features of summarized events. *Journal of Behavioral Decision Making*, 13(2), 191-201.
doi:10.1002/(SICI)1099-0771(200004/06)13:2<191::AID-BDM330>3.0.CO;2-A
- Ariely, D., Loewenstein, G. F., & Prelec, D. (2003). “Coherent Arbitrariness”: Stable Demand Curves Without Stable Preferences. *Quarterly Journal of Economics*, 118(1), 73-105.
doi:10.1162/00335530360535153
- Asch, S. E. (1946). Forming impressions of personality. *Journal of Abnormal and Social Psychology*, 41(3), 258–290.
- Aydin, G., & Ziya, S. (2008). Pricing Promotional Products Under Upselling. *Manufacturing and Service Operations Management*, 10(3), 360.
- Bechtold, S. E., & Jacobs, L. W. (1990). Implicit Modeling of Flexible Break Assignments in Optimal Shift Scheduling. *Management Science*, 36(11), 1339-1351.
- Ben-Akiva, M., & Lerman, S. (1985). *Discrete Choice Analysis: Theory and Application to Travel Demand* (1st ed.). Cambridge, MA: The MIT Press.

- Bitner, M. J. (1992). Servicescapes: the impact of physical surroundings on customers and employees. *The Journal of Marketing*, 56(2), 57–71.
- Bitran, G. R., & Caldentey, R. (2003). An overview of pricing models for revenue management. *Manufacturing and Service Operations Management*, 5(3), 203-229.
- Bitran, G. R., & Ferrer, J.-C. (2007). On pricing and composition of bundles. *Production and Operations Management*, 16(1), 93-108.
- Bitran, G. R., Ferrer, J.-C., & Rocha e Oliveira, P. (2008). OM Forum--Managing Customer Experiences: Perspectives on the Temporal Aspects of Service Encounters. *Manufacturing and Service Operations Management*, 10(1), 61-83.
- Bolton, R. N., Lemon, K. N., & Bramlett, M. D. (2006). The Effect of Service Experiences over Time on a Supplier's Retention of Business Customers. *Management Science*, 52(12), 1811–1823.
- Brusco, M. J., & Jacobs, L. W. (2000). Optimal Models for Meal-Break and Start-Time Flexibility in Continuous Tour Scheduling. *Management Science*, 46(12), 1630-1641.
- Buell, R. W., & Norton, M. I. (2011). The Labor Illusion: How Operational Transparency Increases Perceived Value. *Management Science*, *Forthcoming*. Retrieved from http://scholar.google.com/scholar?hl=en&q=labor+illusion+buell&btnG=Search&as_sdt=0%2C33&as_ylo=&as_vis=0
- Burt, C. D. B., Mitchell, D. A., Raggatt, P. T. F., Jones, C. A., & Cowan, T. M. (1995). A snapshot of autobiographical memory retrieval characteristics. *Applied Cognitive Psychology*, 9(1), 61-74. doi:10.1002/acp.2350090105

- Caro, F., & Martínez-de-Albéniz, V. (2009). The Effect of Assortment Rotation on Consumer Choice and Its Impact on Competition. *Consumer-Driven Demand and Operations Management Models*, 63–79.
- Chase, R. B. (1978). Where does the customer fit in a service operation? *Harvard Business Review*, 56(6), 137.
- Chase, R. B. (1981). The Customer Contact Approach to Services: Theoretical Bases and Practical Extensions. *Operations Research*, 29(4), 698-706.
- Chase, R. B. (1996). The mall is my factory: reflections of a service junkie. *Production and Operations Management*, 5(4), 298–308.
- Chase, R. B. (2004). It's time to get to first principles in service design. *Managing Service Quality*, 14(2/3), 126–128.
- Chase, R. B., & Apte, U. M. (2007). A history of research in service operations: What's the big idea? *Journal of Operations Management*, 25(2), 375-386.
doi:10.1016/j.jom.2006.11.002
- Chase, R. B., & Dasu, S. (2001). Want to perfect your company's service? Use behavioral science. *Harvard Business Review*, 79(6), 78-84, 147. doi:11408979
- Chase, R. B., & Dasu, S. (2008). Psychology of the Experience: The Missing Link in Service Science. *Service Science, Management and Engineering Education for the 21st Century*, Service Science: Research and Innovations in the Service Economy (pp. 35-40). New York, NY: Springer.
- Chow, W.-M. (1986). Design for Line Flexibility. *IIE - Transactions*, 18(1), 95.
- Chun, H. E. (2009). *Savoring Future Experiences: Antecedents And Effects On Evaluations Of Consumption Experiences*. University of Southern California.

- Churchill, W. (1942, 10). The End of the Beginning. Retrieved from <http://www.churchill-society-london.org.uk/EndoBegn.html>
- Cook, L. S., Bowen, D. E., Chase, R. B., Dasu, S., Stewart, D. M., & Tansik, D. A. (2002). Human issues in service design. *Journal of Operations Management*, 20(2), 159-174.
- Daniels, R. L., & Mazzola, J. B. (1994). Flow Shop Scheduling with Resource Flexibility. *Operations Research*, 42(3), 504-522.
- Daniels, R. L., Hoopes, B. J., & Mazzola, J. B. (1996). Scheduling Parallel Manufacturing Cells with Resource Flexibility. *Management Science*, 42(9), 1260-1276.
- Daniels, R. L., Mazzola, J. B., & Shi, D. (2004). Flow Shop Scheduling with Partial Resource Flexibility. *Management Science*, 50(5), 658-669.
- Davidson, M. (2003). An Integrated Approach to Service Quality in Hotels. *Journal of Quality Assurance in Hospitality & Tourism*, 4(1), 71-85.
- Davidson, M., Manning, M. L., Brosnan, P., & Timo, N. (2001). Organizational Climate, Perceived Customer Satisfaction, and Revenue per Available Room in Four- and Five-Star Australian Hotels. *Tourism Analysis*, 6(2), 123-137.
- Debo, L., & Veeraraghavan, S. K. (2009). Models of Herding Behavior in Operations Management. *Consumer-Driven Demand and Operations Management Models*, 81-112.
- Dehaene, S. (1993). Temporal Oscillations in Human Perception. *Psychological Science*, 4(4), 264-270. doi:10.1111/j.1467-9280.1993.tb00273.x
- Dialsmith and the Perception Analyzer - Measure Responses. (n.d.). Retrieved April 1, 2011, from <http://www.perceptionanalyzer.com/pa/measure-response.html>
- Diebold, J. (1952). *Automation: the advent of the automatic factory*. Van Nostrand.

- Ding, X., Hu, P. J.-H., Verma, R., & Wardell, D. G. (2010). The Impact of Service System Design and Flow Experience on Customer Satisfaction in Online Financial Services. *Journal of Service Research*, 13(1), 96 -110. doi:10.1177/1094670509350674
- Ding, X., Verma, R., & Iqbal, Z. (2007). Self-service technology and online financial service choice. *International Journal of Service Industry Management*, 18(3), 246–268.
- Dixon, M. J., & Verma, R. (2011). *Sequence Effects in Service Bundles*. Working Paper, Cornell University, Ithaca NY.
- Ebbinghaus, H. (1902). *Grundzuge Der Psychologie*. Leipzig, Germany: Viet.
- Fynes, B., & Lally, A. M. (2008). Innovation in Services: From Service Concepts to Service Experiences. *Service Science, Management and Engineering Education for the 21st Century*, 329–333.
- Gaeth, G. J., Levin, I. P., Chakraborty, G., & Levin, A. M. (1991). Consumer evaluation of multi-product bundles: An information integration analysis. *Marketing Letters*, 2(1), 47-57.
- Gantt, H. L. (1903). *A graphical daily balance in manufacture*.
- Gensch, D. H., & Recker, W. W. (1979). The Multinomial, Multiattribute Logit Choice Model. *Journal of Marketing Research*, 16(1), 124-132.
- Goldstein, S. M., Johnston, R., Duffy, J. A., & Rao, J. (2002). The service concept: the missing link in service design research? *Journal of Operations Management*, 20(2), 121-134.
- Graves, S. C., & Tomlin, B. T. (2003). Process Flexibility in Supply Chains. *Management Science*, 49(7), 907-919.
- Green, P. E., & Krieger, A. M. (1989). Recent contributions to optimal product positioning and buyer segmentation. *European Journal of Operational Research*, 41(2), 127–141.

- Green, P. E., & Krieger, A. M. (1991). Product design strategies for target-market positioning. *Journal of Product Innovation Management*, 8(3), 189–202.
- Grove, S. J., & Fisk, R. P. (2001). Service Theater: An Analytical Framework for Services Marketing. *Services Marketing*, 83–92.
- Guiltinan, J. P. (1987). The Price Bundling of Services: A Normative Framework. *Journal of Marketing*, 51(2), 74-85.
- Hansen, D. E., & Danaher, P. J. (1999). Inconsistent Performance during the Service Encounter: What's a Good Start Worth? *Journal of Service Research*, 1(3), 227-235.
doi:10.1177/109467059913004
- Harlam, B. A., Krishna, A., Lehmann, D. R., & Mela, C. (1995). Impact of bundle type, price framing and familiarity on purchase intention for the bundle. *Journal of Business Research*, 33(1), 57-66.
- Heskett, J. L., & Sasser, W. E. (2010). The service profit chain. *Handbook of Service Science*, 19–29.
- Heskett, J. L., Sasser, W. E., & Hart, C. W. L. (1990). *Service breakthroughs: Changing the rules of the game*. New York: Free Press.
- Ho, T. H., & Zheng, Y.-S. (2004). Setting Customer Expectation in Service Delivery: An Integrated Marketing-Operations Perspective. *Management Science*, 50(4), 479-488.
- Iravani, S. M., Oyen, M. P. V., & Sims, K. T. (2005). Structural Flexibility: A New Perspective on the Design of Manufacturing and Service Operations. *Management Science*, 51(2), 151-166.

- Jerath, K., Netessine, S., & Veeraraghavan, S. K. (2009). Selling to strategic customers: Opaque selling strategies. *Consumer-Driven Demand and Operations Management Models*, 253–300.
- Jerath, K., Netessine, S., & Veeraraghavan, S. K. (2010). Revenue Management with Strategic Customers: Last-Minute Selling and Opaque Selling. *Management Science*, 56, 430–448.
doi:<http://dx.doi.org/10.1287/mnsc.1090.1125>
- Johnston, R. (1995). The Zone of Tolerance: Exploring the relationship between service transactions and satisfaction with the overall service. *International Journal of Service Industry Management*, 6(2), 46–61.
- Jordan, W. C., & Graves, S. C. (1995). Principles on the Benefits of Manufacturing Process Flexibility. *Management Science*, 41(4), 577-594.
- Kahneman, D., & Snell, J. (1992). Predicting a changing taste: Do people know what they will like? *Journal of Behavioral Decision Making*, 5(3), 187-200.
doi:[10.1002/bdm.3960050304](https://doi.org/10.1002/bdm.3960050304)
- Kahneman, D., & Tversky, A. (1979). Prospect Theory: An Analysis of Decision under Risk. *Econometrica*, 47(2), 263-291.
- Kahneman, D., Fredrickson, B. L., Schreiber, C. A., & Redelmeier, D. A. (1993). When more pain is preferred to less: Adding a better end. *Psychological Science*, 4(6), 401-405.
doi:[doi:10.1111/j.1467-9280.1993.tb00589.x](https://doi.org/10.1111/j.1467-9280.1993.tb00589.x)
- Kahneman, D., Wakker, P. P., & Sarin, R. (1997). Back to Bentham? Explorations of Experienced Utility. *Quarterly Journal of Economics*, 112(2), 375-405.
doi:[10.1162/003355397555235](https://doi.org/10.1162/003355397555235)

- Karmarkar, U. S. (1996). Integrative Research in Marketing and Operations Management. *Journal of Marketing Research*, 33(2), 125-133. doi:10.2307/3152141
- Kendall, G., Knust, S., Ribeiro, C. C., & Urrutia, S. (2010). Scheduling in sports: An annotated bibliography. *Computers & Operations Research*, 37(1), 1-19.
doi:10.1016/j.cor.2009.05.013
- Kim, H. B., & Kim, W. G. (2005). The relationship between brand equity and firms' performance in luxury hotels and chain restaurants. *Tourism management*, 26(4), 549–560.
- Kim, H. B., Kim, W. G., & An, J. A. (2003). The effect of consumer-based brand equity on firms' financial performance. *Journal of consumer marketing*, 20(4), 335–351.
- Kimes, S. E. (1999). The Relationship between Product Quality and Revenue Per Available Room at Holiday Inn. *Journal of Service Research*, 2(2), 138-144.
doi:10.1177/109467059922002
- Kimes, S. E. (2001). How product quality drives profitability. *The Cornell Hotel and Restaurant & Administration Quarterly*, 42(3), 25–8.
- Kogut, B., & Kulatilaka, N. (1994). Operating Flexibility, Global Manufacturing, and the Option Value of a Multinational Network. *Management Science*, 40(1), 123-139.
- Koste, L. L., & Malhotra, M. K. (1999). A theoretical framework for analyzing the dimensions of manufacturing flexibility. *Journal of Operations Management*, 18(1), 75-93.
doi:10.1016/S0272-6963(99)00010-8
- Kotler, P. (2002). *Marketing Management* (11th ed.). Upper Saddle River, NJ: Prentice Hall.

- LaBarbera, P. A., & Mazursky, D. (1983). A Longitudinal Assessment of Consumer Satisfaction/Dissatisfaction: The Dynamic Aspect of the Cognitive Process. *Journal of Marketing Research*, 20(4), 393-404.
- Liu, Q., & van Ryzin, G. J. (2008). Strategic Capacity Rationing to Induce Early Purchases. *Management Science*, 54(6), 1115-1131.
- Loewenstein, G. F. (1987). Anticipation and the Valuation of Delayed Consumption. *The Economic Journal*, 97(387), 666-684.
- Loewenstein, G. F., & Prelec, D. (1993). Preferences for Sequences of Outcomes. *Psychological Review*, 100(1), 91-108.
- Loewenstein, G. F., & Sicherman, N. (1991). Do Workers Prefer Increasing Wage Profiles? *Journal of Labor Economics*, 9(1), 67-84.
- Louviere, J. J. (1988). *Analyzing decision making : metric conjoint analysis*. Newbury Park: Sage Publications.
- Lovelock, C. H., & Wright, L. (1999). *Principles of service marketing and management*. Prentice Hall.
- Luce, R. D. (1959). *Individual choice behavior*. New York, NY: Wiley.
- Mabert, V. A., & Showalter, M. J. (1990). Measuring the impact of part-time workers in service organizations. *Journal of Operations Management*, 9(2), 209-229. doi:10.1016/0272-6963(90)90096-V
- MacDonald, L., Anderson, C. K., & Verma, R. (2010). Using revealed- and stated-preference customer choice models for making pricing decisions in services: An illustration from the hospitality industry. *Journal of Revenue and Pricing Management*. doi:10.1057/rpm.2010.21

- Markie, P. (2008). Rationalism vs. Empiricism. (E. N. Zalta,, Ed.)*Stanford Encyclopedia of Philosophy*. Retrieved from <http://plato.stanford.edu/entries/rationalism-empiricism/>
- McCarthy, J. E. (1972). *Basic Marketing A Managerial Approach*. Richard D Irwin, Inc.
- McFadden, D. (1980). Econometric Models for Probabilistic Choice Among Products. *The Journal of Business*, 53(3), S13-S29.
- McFadden, D. (1986). The Choice Theory Approach to Market Research. *Marketing Science*, 5(4), 275-297.
- Mcguire, K. A., & Kimes, S. E. (2006). The perceived fairness of waitlist-management techniques for restaurants. *Cornell Hotel and Restaurant Administration Quarterly*, 47(2), 121.
- Menor, L. J., Tatikonda, M. V., & Sampson, S. E. (2002). New service development: areas for exploitation and exploration. *Journal of Operations Management*, 20(2), 135-157.
- Metters, R. D. (2010). The Neglect of Service Science in the Operations Management Field. *Handbook of Service Science*, 309–319.
- Metters, R. D., & Maruchek, A. (2007). Service management-academic issues and scholarly reflections from operations management researchers. *Decision Sciences*, 38(2), 195.
- Metters, R. D., King-Metters, K. H., Pullman, M. E., & Walton, S. (2008). *Successful service operations management*. South-Western College Publishing.
- Moore, W. L., Louviere, J. J., & Verma, R. (1999). Using conjoint analysis to help design product platforms. *Journal of Product Innovation Management*, 16(1), 27–39.
- Morgan, L. O., Daniels, R. L., & Kouvelis, P. (2001). Marketing/manufacturing trade-offs in product line management. *Iie Transactions*, 33(11), 949–962.

- Nguyen, T.-D. T., & Belk, R. W. (2007). This We Remember: Consuming Representation via the Web Posting of War Photographs. *Consumption Markets & Culture*, 10(3), 251-291.
- Oliva, R., & Sterman, J. D. (2001). Cutting Corners and Working Overtime: Quality Erosion in the Service Industry. *Management Science*, 47(7), 894-914.
- Oliver, R. L. (1980). A Cognitive Model of the Antecedents and Consequences of Satisfaction Decisions. *Journal of Marketing Research*, 17(4), 460-469. doi:10.2307/3150499
- orchestrate. (n.d.). *Collins English Dictionary*. HarperCollins Publishers.
- Parasuraman, A., Zeithaml, V. A., & Berry, L. L. (1985). A Conceptual Model of Service Quality and Its Implications for Future Research. *The Journal of Marketing*, 49(4), 41-50. doi:10.2307/1251430
- Pine, B. J., & Gilmore, J. H. (1999). *The Experience Economy: Work is Theatre & Every Business a Stage*. Harvard Business School Press.
- Pinedo, M. (2008). *Scheduling: theory, algorithms, and systems*. Springer Verlag.
- Pogrebin, R. (2002, October 16). Uncertain Times: Impulse Buyers Replace Ticket Subscribers. *New York Times*.
- Pullman, M. E., & Gross, M. A. (2004). Ability of Experience Design Elements to Elicit Emotions and Loyalty Behaviors. *Decision Sciences*, 35(3), 551-578. doi:10.1111/j.0011-7315.2004.02611.x
- Pullman, M. E., & Moore, W. L. (1999). Optimal service design: integrating marketing and operations perspectives. *International Journal of Service Industry Management*, 10(2), 239-261.

- Pullman, M. E., Goodale, J. C., & Verma, R. (2000). Service capacity design with an integrated market utility-based method. *New Service Development: Creating Memorable Experiences* (pp. 111–37).
- Pullman, M. E., Verma, R., & Goodale, J. C. (2001). Service design and operations strategy formulation in multicultural markets. *Journal of Operations Management*, *19*(2), 239-254. doi:10.1016/S0272-6963(00)00059-0
- Raman, N., & Chhajed, D. (1995). Simultaneous determination of product attributes and prices, and production processes in product-line design. *Journal of Operations Management*, *12*(3-4), 187–204.
- Redelmeier, D. A., & Kahneman, D. (1996). Patients' memories of painful medical treatments: real-time and retrospective evaluations of two minimally invasive procedures. *Pain*, *66*(1), 3-8. doi:8857625
- Redelmeier, D. A., Katz, J., & Kahneman, D. (2003). Memories of colonoscopy: a randomized trial. *Pain*, *104*(1-2), 187-94.
- Rosenthal, E. C., Zydiak, J. L., & Chaudhry, S. S. (1995). Vendor Selection with Bundling. *Decision Sciences*, *26*(1), 35-48. doi:10.1111/j.1540-5915.1995.tb00836.x
- Ross, W. T., & Simonson, I. (1991). Evaluations of pairs of experiences: A preference for happy endings. *Journal of Behavioral Decision Making*, *4*(4), 273-282.
doi:10.1002/bdm.3960040405
- Roth, A. V., & Menor, L. J. (2003). Insights into service operations management: A research agenda. *Production and Operations Management*, *12*(2), 145–164.
- Rust, R. T., & Bhalla, G. (2010). Customer Equity. *Handbook of Service Science*, 61–78.

- Rust, R. T., & Zahorik, A. (1993). Customer satisfaction, customer retention, and market share. *Journal of Retailing*, 69(2), 193–215.
- Sampson, S. E. (2009). Practical Implications of Preference-Based Conference Scheduling. *Production and Operations Management*, 13(3), 205-215. doi:10.1111/j.1937-5956.2004.tb00506.x
- Sampson, S. E., & Froehle, C. M. (2006). Foundations and implications of a proposed unified services theory. *Production and Operations Management*, 15(2), 329.
- Sampson, S. E., & Weiss, E. N. (1995). Increasing Service Levels in Conference and Educational Scheduling: A Heuristic Approach. *Management Science*, 41(11), 1816-1825.
- Sampson, S. E., & Weiss, E. N. (1996). Designing Conferences to Improve Resource Utilization and Participant Satisfaction. *The Journal of the Operational Research Society*, 47(2), 297-314. doi:10.2307/2584349
- Sampson, S. E., Freeland, J. R., & Weiss, E. N. (1995). Class Scheduling to Maximize Participant Satisfaction. *Interfaces*, 25(3), 30-41. doi:Article
- Sasser, W. E., Olsen, R. P., & Wyckoff, D. D. (1978). *Management of service operations: Text, cases, and readings*. Boston, MA: Allyn and Bacon.
- Schneider, B., & Bowen, D. E. (2010). Winning the service game. *Handbook of Service Science*, 31–59.
- Schoenherr, T., & Mabert, V. A. (2008). The use of bundling in B2B online reverse auctions. *Journal of Operations Management*, 26(1), 81-95. doi:10.1016/j.jom.2007.05.001
- Sethi, A. K., & Sethi, S. P. (1990). Flexibility in manufacturing: a survey. *International Journal of Flexible Manufacturing Systems*, 2(4), 289–328.

- Shen, Z. M., & Su, X. (2007). Customer Behavior Modeling in Revenue Management and Auctions: A Review and New Research Opportunities. *Production and Operations Management, 16*(6), 713-728. doi:10.1111/j.1937-5956.2007.tb00291.x
- Sheth, J. N., Newman, B. I., & Gross, B. L. (1991). Why we buy what we buy: A theory of consumption values. *Journal of Business Research, 22*(2), 159-170. doi:10.1016/0148-2963(91)90050-8
- Sheu, C., McHaney, R., & Babbar, S. (2003). Service process design flexibility and customer waiting time. *International Journal of Operations & Production Management, 23*(8), 901-917. doi:10.1108/01443570310486347
- Shocker, A. D., & Srinivasan, V. (1979). Multiattribute approaches for product concept evaluation and generation: a critical review. *Journal of Marketing Research, 15*9–180.
- Stremersch, S., & Tellis, G. J. (2002). Strategic bundling of products and prices: a new synthesis for marketing. *The Journal of Marketing, 55*-72.
- Thaler, R. H., & Johnson, E. J. (1990). Gambling with the House Money and Trying to Break Even: The Effects of Prior Outcomes on Risky Choice. *Management Science, 36*(6), 643-660.
- Treisman, M., Faulkner, A., Naish, P. L. N., & Brogan, D. (1990). The internal clock: Evidence for a temporal oscillator underlying time perception with some estimates of its characteristic frequency. *Perception, 19*(6), 705–743.
- Tsay, A. A. (1999). The Quantity Flexibility Contract and Supplier-Customer Incentives. *Management Science, 45*(10), 1339-1358.
- Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science, 185*(4157), 1124–1131.

- VanRullen, R., & Koch, C. (2003). Is perception discrete or continuous? *Trends in Cognitive Sciences*, 7(5), 207–213.
- Vargo, S. L., & Lusch, R. F. (2004). Evolving to a new dominant logic for marketing. *Journal of marketing*, 68(1), 1–17.
- Veeraraghavan, S. K., & Debo, L. (2008). Joining Longer Queues: Information Externalities in Queue Choice. *Manufacturing and Service Operations Management*.
doi:<p>10.1287/msom.1080.0239</p>
- Verhoef, P. C., Antonides, G., & de Hoog, A. N. (2004). Service Encounters as a Sequence of Events: The Importance of Peak Experiences. *Journal of Service Research*, 7(1), 53-64.
doi:10.1177/1094670504266137
- Verma, R., & Plaschka, G. (2003). The Art and Science of Customer Choice Modeling. *Cornell Hotel and Restaurant Administration Quarterly*, 43(6), 15–24.
- Verma, R., & Plaschka, G. (2005). Predicting customer choices. *MIT Sloan Management Review*, 47(1), 7–10.
- Verma, R., & Thompson, G. M. (1999). Managing service operations based on customer preferences. *International Journal of Operations & Production Management*, 19(9), 89-90.
- Verma, R., Thompson, G. M., & Louviere, J. J. (1999). Configuring Service Operations in Accordance with Customer Needs and Preferences. *Journal of Service Research*, 1(3), 262-274. doi:10.1177/109467059913007
- Verma, R., Thompson, G. M., Moore, W. L., & Louviere, J. J. (2001). Effective Design of Products/Services: An Approach Based on Integration of Marketing and Operations

- Management Decisions. *Decision Sciences*, 32(1), 165-194. doi:10.1111/j.1540-5915.2001.tb00957.x
- Victorino, L. (2008). *Scripting the service encounter: An empirical analysis*. University of Utah.
- Voss, C. A., & Zomerdijk, L. G. (2007). *Innovation in experiential services: An empirical view*. Citeseer.
- Voss, C. A., Roth, A. V., & Chase, R. B. (2008). Experience, Service Operations Strategy, and Services as Destinations: Foundations and Exploratory Investigation. *Production and Operations Management*, 17(3), 247-266. doi:10.3401/poms.1080.0030
- Walker, L., Thibaut, J., & Andreoli, V. (1972). Order of Presentation at Trial. *Yale Law Journal*, 82, 216.
- Wilson, J. M. (2003). Gantt charts: A centenary appreciation. *European Journal of Operational Research*, 149(2), 430–437.
- Zeithaml, V. A., Bitner, M. J., & Gremler, D. D. (1996). *Services marketing*. McGraw-Hill New York.
- Zomerdijk, L. G., & Voss, C. A. (2010). Service Design for Experience-Centric Services. *Journal of Service Research*, 13(1), 67 -82. doi:10.1177/1094670509351960