This article was downloaded by: [64.187.132.166] On: 04 July 2024, At: 08:32
Publisher: Institute for Operations Research and the Management Sciences (INFORMS)
INFORMS is located in Maryland, USA

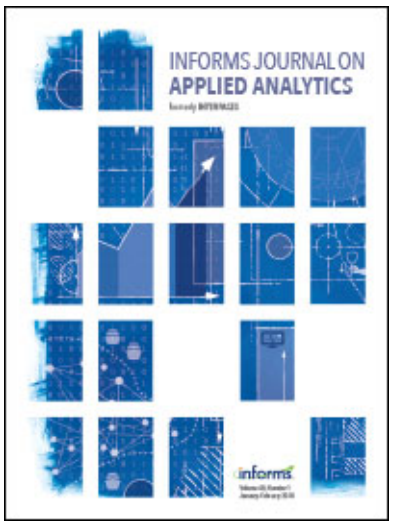

Interfaces

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Turner Optimizes the Allocation of Audience Deficiency Units

José Antonio Carbajal, Wes Chaar

To cite this article:
José Antonio Carbajal, Wes Chaar (2017) Turner Optimizes the Allocation of Audience Deficiency Units. Interfaces 47(6):518-536. https://doi.org/10.1287/inte.2017.0913

Full terms and conditions of use: https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2017, INFORMS

Please scroll down for article—it is on subsequent pages

With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.
For more information on INFORMS, its publications, membership, or meetings visit http://www.informs.org

http://pubsonline.informs.org/journal/inte/

INTERFACES
Vol. 47, No. 6, November–December 2017, pp. 518–536
ISSN 0092-2102 (print), ISSN 1526-551X (online)

Turner Optimizes the Allocation of Audience Deficiency Units

José Antonio Carbajal,[a] Wes Chaar[a]

[a] Turner Broadcasting System, Inc., Atlanta, Georgia 30318

Contact: antonio.carbajal@turner.com, http://orcid.org/0000-0003-4025-1755 (JAC); wes.chaar@turner.com, http://orcid.org/0000-0002-2743-8956 (WC)

Received: August 17, 2016
Revised: December 8, 2016
Accepted: February 15, 2017

https://doi.org/10.1287/inte.2017.0913

Copyright: © 2017 INFORMS

Abstract. Each broadcasting quarter, advertisement-based networks face the challenge of allocating make-good units that reduce the liability they have incurred because of fewer people in the audiences they have guaranteed to advertisers. Turner Broadcasting System, Inc. has developed an optimization framework that improves the process of allocating make-good units across hundreds of sales plans in a dynamic setting in which priorities change continually and customer relationship management is of paramount importance. This system, called the Liability Efficiency Optimizer (LEO), has significantly simplified the process of allocating audience deficiency units at Turner, reducing it from weeks down to minutes, and has enabled its users to make the most profitable use of the liability airtime inventory while honoring advertiser requirements in terms of selling-title mix and unit distribution across weeks. This system has been deployed to all the domestic networks operated by Turner, and its monetary benefit has been conservatively estimated to be $10 million of recurring yearly liability reduction over the manual process it replaced.

History: This paper was refereed.

Keywords: media industry • TV advertising • make goods • multicriteria optimization • integer programming

Turner Broadcasting System, Inc. creates and programs branded news, entertainment, animation and young adult, and sports media environments on television and other platforms for consumers worldwide. An important part of Turner's domestic business involves the management and operation of several major advertisement-based cable networks that include TBS, TNT, truTV, CNN, HLN, NBA TV, Cartoon Network, Boomerang, and Adult Swim. Although Turner's portfolio spans a multitude of platforms and includes a number of digital and mobile properties, this paper addresses only the linear (i.e., television) side of the business.

Most broadcasting and cable networks in the United States are advertisement based. As such, their business is the delivery of impressions or audiences to advertisers. The broadcast year in the United States typically starts in late September or early October. Networks announce their programming schedules for the new broadcast year about mid-May. These announcements are followed by an intensive sales period, known as the upfront market, in which networks sell approximately 60 to 80 percent of their commercial airtime. The remaining unsold commercial airtime is typically sold at higher rates throughout the broadcast year in what is known as the scatter market.

Independently of the market type, networks and advertisers, which are usually represented by advertiser agencies, may negotiate during several sessions in which they finalize the details of a sales proposal. The sales proposal specifies the flighting dates (i.e., the date range in which the spots will air), a total budget, cost per thousand impressions (CPM), quarterly impression distribution, selling-title mix (i.e., proportion of units or impressions on each of the selling titles), spot-length mix (e.g., proportion of 15-second or 30-second spots), a primary demographic that the advertiser is trying to reach (e.g., 25-to-54-year-old females—F25–54), and a guaranteed audience level in that target demographic. The proposal may also specify weekly distribution goals; in instances in which it does not specify weekly distribution goals, the expectation is typically equitable distribution across weeks. A selling title is an interval of programming time that a network utilizes to sell commercial airtime. It may comprise a particular show (e.g., Angie Tribeca or Legends) or a

Downloaded from informs.org by [64.187.132.166] on 04 July 2024, at 08:32 . For personal use only, all rights reserved.

Downloaded from informs.org by [64.187.132.166] on 04 July 2024, at 08:32 . For personal use only, all rights reserved.

block of time (e.g., daytime: 8:00 AM–3:00 PM, overnight: midnight to 5:00 AM). In some situations, the term daypart may be used as an alternative to selling title; in other cases, a daypart may be an aggregation of selling titles.

Once both parties have negotiated and accepted the proposal, it becomes a deal, and the network must deliver the total guaranteed impressions specified in the deal. Deal stewardship consists of monitoring the performance of the deal and managing deal liability. When the projected number of impressions throughout the entire flight of the deal is less than the guaranteed number of impressions, the network must allocate additional unpaid units to meet its liability for the deal. These units are known as audience-deficiency units, make-good units, under-delivery units, or simply UD.

Program audience levels are uncertain and can be hard to forecast (Danaher et al. 2011, Danaher and Dagger 2012). They depend on a number of factors that include time-dependent attributes, such as time of the year, day of the week, and time of the day, and program-dependent attributes, such as show genre, lead-in show (i.e., the preceding program), and competing shows on other networks. Furthermore, while networks are liable for unmet guaranteed impressions, they do not receive any additional benefits for delivering impressions beyond the guaranteed levels. Therefore, networks would like to closely manage their allocations of UD units to minimize their liability without exceeding their guaranteed impressions. They must do this in a dynamic environment in which a program's unexpected performance (which might be either higher or lower than expected) changes the UD allocation priorities.

Typically, hundreds of deals receive UD units each broadcast quarter. Until recently, Turner manually developed UD allocations near the beginning of each quarter. Planners performed allocations one deal at a time. They tried to balance business goals, such as generating the most value out of the UD inventory reserve, against customer requirements, such as selling-title mix and weekly distribution of allocated units. They prioritized deals based on their index, defined as the total projected impressions to be delivered at the end of the deal divided by the total guaranteed impressions, focusing first on deals with lower indices. This was a labor-intensive process, which often took several days, even weeks, to complete and involved a substantial rework because of the problem's scale and complexity.

We identified an opportunity to develop a system that would simplify the UD allocation process, and more importantly, would make optimal use of the available inventory by simultaneously allocating units across all deals. Based on our work, we created an application that we called the Liability Efficiency Optimizer (LEO), a multistage, mixed-integer programming engine that allocates UD units across hundreds of deals and allows planners to run multiple scenarios and choose the allocation that best meets their needs.

The main contributions of this paper are that we report on an actual implementation of operations research (OR) practice in the media industry, an area of research that is relatively unexplored and is fertile ground for applying advanced analytics. We also describe the first optimization approach that addresses the UD allocation problem. We organize the paper as follows: *The Media Revenue Management Process* section describes various planning and control problems that arise in the hierarchical media revenue-management process, provide some references to the related literature, and position the UD allocation problem within this framework. In the section *The Under-Delivery Allocation Problem*, we give an overview of the allocation of audience-deficiency units and describe its complicating factors. In *An Optimization Framework for Under-Delivery Allocation*, we detail the restrictions that govern UD allocation and explain how we use mixed-integer programming in three phases to approach this problem. In the *Implementation, Results, and Conclusions* section, we discuss implementation results, summarize conclusions, and discuss directions of future work. In Appendix A, we provide a glossary of media terms, and in Appendix B, we include the mathematical programming formulations of all our models.

The Media Revenue-Management Process

The media revenue-management process is highly complex, and involves the planning and control of several interrelated components. However, because of organizational structures, technical constraints, and industry practice, it lends itself to a hierarchical planning approach that involves strategic, tactical, and operational decisions. In this section, we outline the planning and control problems in this hierarchical process

Downloaded from informs.org by [64.187.132.166] on 04 July 2024, at 08:32 . For personal use only, all rights reserved.

Figure 1. The Media Revenue-Management Process Comprises Several Interrelated Planning and Control Problems That Are Addressed in Practice in a Hierarchical Framework That Involves Strategic, Tactical, and Operational Decisions

	Components of the media revenue management process	
	Planning and control problems	Forecasting
Strategic	• Programming decisions: Program acquisitions, original series development, program scheduling • Commercial airtime capacity allocation between upfront and scatter markets • Upfront sales: Proposal creation, unit allocation and pricing • Development of new business opportunities	Long-term forecasts (quarterly and yearly)
Tactical	• Inventory mix distribution: Allocation of remaining nonprogramming airtime into scatter sales, promo time, and audience deficiency units • Allocation of audience deficiency units • Brand allocation • Scatter sales: Proposal creation, unit allocation and pricing	Medium-term forecasts (weekly and quarterly)
Operational	• Commercial spot scheduling • Promo scheduling • Opportunistic sales	Short-term forecasts (daily and weekly)

Note. The positioning within the process of the problem addressed in this paper, the allocation of audience deficiency units, is highlighted in light gray.

(Figure 1), and provide references that address some of these components.

At the strategic level, networks face multiyear programming decisions, such as original series development and content acquisitions, and single-year programming schedule decisions (e.g., which shows to air and when to air them during the upcoming broadcast year), as Horen (1980), Reddy et al. (1998), and Danaher and Mawhinney (2001) discuss. The programming schedule defines an aggregate capacity of nonprogramming airtime that is available; possible uses are for commercial spots, promos, filler spots, and UD units. Commercial airtime is usually sold in 30-second spots. Although some other spot lengths are possible (e.g., 15- or 60-second spots), negotiations on rates and unit distribution across selling titles and weeks are usually performed on equivalized 30-second units, called EQ30s. When units are equivalized, units of length other than 30 seconds have their impressions adjusted upward or downward, in proportion to their

Downloaded from informs.org by [64.187.132.166] on 04 July 2024, at 08:32 . For personal use only, all rights reserved.

lengths ratio, to 30 seconds. In the United States, the majority of the commercial airtime for the upcoming broadcast year is sold in the upfront market, a short period that starts around mid-May when networks announce their programming. A significant proportion (about 60–80 percent) of the commercial airtime for the entire broadcast year is sold during the upfront season. Balancing upfront and scatter inventory allocation, as Bollapragada and Mallik (2007) and Araman and Popescu (2010) discuss, and creating proposals for the upfront sales, as Bollapragada et al. (2002) and Popescu and Seshadri (2016) discuss, are part of this strategic stage.

At the tactical level, networks face a few planning problems, including inventory-mix distribution, scatter sales, distribution of branded units (i.e., distribution of units across brands for multibrand advertisers such as Procter & Gamble or General Mills), and allocation of UD units, and address these problems quarterly. The commercial airtime that is not sold during upfront market period can be sold throughout the broadcast year in the scatter market or used for show promotion or to clear liability via allocation of UD units. The inventory-mix distribution problem consists of determining the proportions of nonprogramming airtime that should be devoted to scatter sales, promotional efforts, and UD allocation. Scatter sales involve deals that typically cover a period that can range from a few weeks up to two quarters (as opposed to an entire broadcast year), and scatter units usually sell for higher rates than their corresponding upfront units (Bollapragada et al. 2002, Popescu and Seshadri 2016). The distribution of branded units involves breaking down the units purchased by large advertisers into smaller branded deals that must meet certain budget and (or) impression distribution conditions across, for example, selling titles, weeks, or spot lengths. The allocation of audience-deficiency units involves determining what additional unpaid units should be allocated to deals that have a guaranteed audience level, but are expected to underperform, subject to operational constraints and advertiser restrictions, such as selling-title mix and weekly distribution. In this paper, we address this problem, which we explain in detail in the *Under-Delivery Allocation Problem* section.

Finally, at the operational level, networks face opportunistic sales and daily and (or) weekly scheduling problems involving commercial spots and promos. Opportunistic (i.e., filler) sales are the sales of distressed inventory just before airing. These usually involve units in selling titles whose audiences are lower than expected or that air in periods that traditionally face low sellout rates, such as overnight (i.e., midnight to 5:00 AM). The scheduling of individual commercials, which include upfront and scatter paid spots, UD units, and filler units in nonprogramming breaks, is performed daily and (or) weekly, and must honor a number of constraints that include association constraints, product conflicts, time-separation guarantees, and specific positions during a break. This operational problem has received the most attention in the literature. Examples include Bollapragada and Garbiras (2004), Bollapragada et al. (2004), Bai and Xie (2006), Zhang (2006), Kimms and Muller-Bungart (2007), Brusco (2008), Gaur et al. (2009), Popescu and Crama (2016), and Giallombardo et al. (2016). Promos are spots that advertise shows in the same or other sister networks and are also scheduled daily and (or) weekly, usually separately from commercial spots; they follow a distinct set of constraints that include distribution goals for franchises, titles, and spot lengths, and break-separation bounds, while attempting to maximize the total gross rating points that will be delivered because of showing promos.

Some constraints that are strictly enforced until the scheduling level, such as time-separation guarantees (i.e., minimum time separation between spots from the same advertiser or spots for the same brand for advertisers with multiple brands) and product conflicts (i.e., competition-avoidance requirements that prevent advertising competing products within the same commercial break), are enforced at an aggregate level in higher planning levels. For example, when a proposal is created or a brand allocation or UD allocation is performed, the number of units assigned to a particular selling-title week should not exceed the number of total planned commercial breaks in that selling-title week minus the number of spots with the same product conflict that are already assigned to that selling-title week. Furthermore, forecasting the numbers of audience members who will show up for the various shows at different times is a problem that affects the three decision levels; see Gensch and Shaman (1980), Meyer and Hyndman (2006), Danaher et al. (2011), and Danaher

Downloaded from informs.org by [64.187.132.166] on 04 July 2024, at 08:32 . For personal use only, all rights reserved.

and Dagger (2012). As decisions move from strategic to tactical to operational, the required forecasts go from long- to short-term horizons and from aggregate estimates (program quarterly estimates) to very granular (30-minute or even one-minute estimates).

The Under-Delivery Allocation Problem

Advertisers can be classified by industry category, which is commonly referred to as a product conflict code (e.g., fast-food restaurants, movies), and are usually trying to reach a specific demographic that they believe is most likely to consume their products; for example, an advertiser of male grooming products might target males age 18 and over. Most upfront and scatter deals have performance guarantees on the total audience in the advertiser's target demographic. The total audience delivered to a deal is the gross sum of impressions that fall into the target demographic regardless of duplication. Different classifications (called ratecard types) are used to indicate the components of a valid audience for a specific deal. Examples include live (i.e., average number of people watching a particular show when it airs, including both commercial airtime and programming), ACM+3 (i.e., average audience watching only commercial airtime either live or within three days of airing via DVR systems), and ACM+7 (i.e., similar to ACM+3, but with seven days of delayed viewing). Although the specific ratecard type to be applied to a deal is determined during the sales process, specific audience levels are merely forecasts at that time. The actual audience levels are provided by media rating agencies after the commercials have aired. In the United States, Nielsen Media Research is the standard for broadcast and cable ratings.

A deal is considered to be underperforming if the gross sum of delivered audiences from spots that have aired plus the gross sum of estimated audiences from spots that have not aired yet is less than the audience guaranteed in the deal. The projected impression shortfall of a deal is constantly changing as new spots air and deliver additional impressions to the deal and as estimates of impressions of future spots are updated to reflect the latest viewer trends. However, the decision to allocate UD units has been typically performed once a quarter. Figure 2 illustrates this situation for a fictitious upfront deal that spans four broadcast quarters. For simplicity, we label the broadcast quarters Q1 through Q4; in practice, an upfront deal spans, for example, Q4 2015 through Q3 2016. We also assume that the deal is evaluated exactly at the end of a broadcast quarter; in practice, UD allocation decisions occur a few weeks prior to the end of a broadcast quarter and include a small lag period for receiving actual impressions from aired spots. At the end of Q1, the impressions delivered include all the impressions from spots aired during Q1, and the estimated future impressions include all spots that will air during Q2, Q3, and Q4. The difference between the total guaranteed impressions and the total projected impressions to be delivered at the end of Q4 is the projected impression shortfall for the deal. At this point, UD units that will air during Q2 are allocated with the intention of reducing

Figure 2. Deals Are Typically Evaluated Every Quarter to Be Allocated UD Units for the Next Quarter

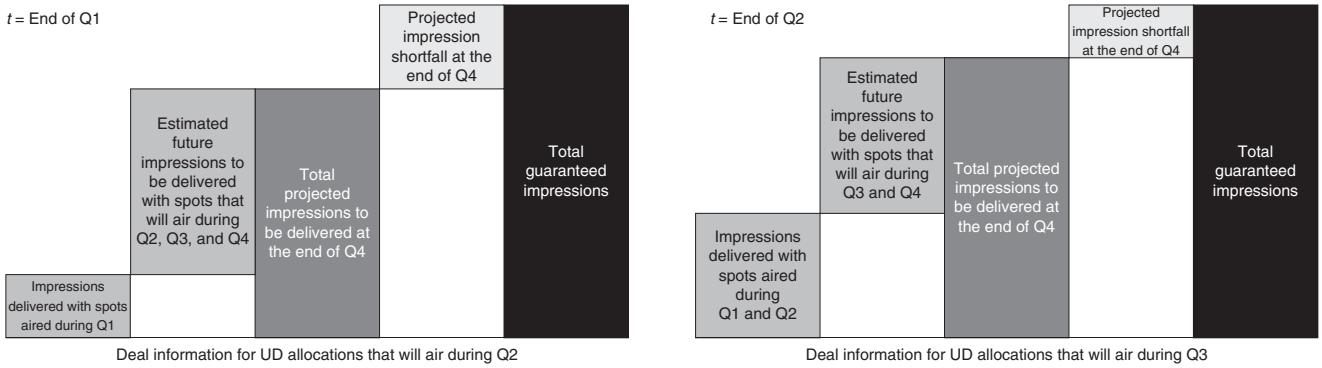

Notes. The total projected impressions to be delivered at the end of the deal include impressions of spots aired in prior quarters, which include paid spots plus UD units, and estimated future impressions of spots that will air in upcoming quarters. The difference between the number of impressions guaranteed in the deal and the total projected impressions at the end of the deal flight is the projected impression shortfall.

Downloaded from informs.org by [64.187.132.166] on 04 July 2024, at 08:32 . For personal use only, all rights reserved.

the impression shortfall. One quarter ahead, at the end of Q2, the deal is evaluated again and the impressions delivered now include all the impressions from spots aired during Q1 and Q2, including paid and UD spots. The estimated future impressions now include only any spots that will air during Q3 and Q4. The deal impression shortfall is updated and UD units that will air during Q3 are allocated. This process repeats every quarter until the end of Q4. At that point, the deal either will have no impression shortfall or will become an aged deal—a deal that has ended its flight of original paid units but has not yet met its guaranteed impressions. Aged deals continue to be allocated UD units until their impression shortfalls have been eliminated.

Turner honors any unmet guaranteed impressions and provides additional unpaid spots to make up audience deficiencies; however, it does not receive any additional revenue if it overdelivers impressions (i.e., delivers more impressions than it guaranteed in the deal). Each broadcast quarter, Turner sets up a UD inventory reserve and allocates those units across hundreds of underperforming deals. A proposal of UD allocations specifies the number of units (in EQ30s) that will be allotted to a deal for specific combinations of selling titles and weeks. The objective is to create UD proposals that maximize the value of the allocation, which is equivalent to maximizing the reduction in impression deficits across all deals. In practice, because such a maximal allocation efficiency uses the UD reserve, we often refer to this as maximizing the value of the UD reserve. Because selling titles exhibit demographic differences in their viewership, and deals have distinct primary demographics and different impression values, which are measured by their cost per thousand impressions (CPM), the value of a UD unit varies depending on its specific allocation to a deal selling-title and week combination. For example, consider a simplified situation with two liable deals: deal A with CPM = 10 and primary demo α and deal B with CPM = 50 and primary demo β. A unit available in selling title I on a given week is expected to bring 500,000 impressions in demo α and 200,000 impressions in demo β. If the unit is assigned to deal A, its value would be 500 (10) = 5,000; if it is assigned to deal B, its value would be 200 (50) = 10,000. These differences can be exploited to maximize the global value of a UD allocation; however, deals also have several characteristics that directly impact the allocation of make-good units.

Figure 3 provides more details of the upfront deal, which spans four quarters, and which we describe in Figure 2. A deal specifies the total number of paid units per selling title that will be scheduled during each quarter. Within each quarter, units are plotted across weeks, either equitably during the weeks in which the selling title airs or according to some weighted distribution that the advertiser provides. The value of a given selling title to a particular advertiser varies. Each deal specifies a selling-title mix that reflects the specific preferences of the advertiser. Therefore, advertisers want to receive a commensurable selling-title mix in their UD units. If a show is oversold, an advertiser might accept a different selling title with similar demographics, or it might choose to wait until the same selling title is available at a later time. UD allocations should also be distributed across the same weeks when paid units are in flight; specifically, UD units cannot be assigned to weeks when the original deal is on hiatus. Selling titles have additional characteristics that impact UD allocations; in particular, selling-title groups and selling-title packages are important. Selling-title groups refer to collections of selling titles with similar viewership demographics that are deemed to be of commensurate value to advertisers. Examples include world broadcast premieres of comparable blockbuster movies and original series with similar themes. Selling titles belonging to the same group can be used as surrogates for UD allocation when the original selling titles are oversold, no longer exist, or might not air for a while. In the example shown in Figure 3, ST2 and ST3 are assumed to belong to the same group, and ST5 and ST9, of which the latter is not part of the original deal, are assumed to belong to another group. Selling-title packages refer to related shows that are sold in tandem. Examples include movies with multiple telecasts and shows that air in prime time and are repeated overnight; such as a Conan premiere that airs at 11:00 PM and a Conan repeat that airs at 12:30 AM. UD allocations should honor the same packaging rules from the original deal. In the example in Figure 3, ST6, ST7, and ST8 are assumed to be a package.

Downloaded from informs.org by [64.187.132.166] on 04 July 2024, at 08:32 . For personal use only, all rights reserved.

Figure 3. A Deal Specifies the Total Number of Paid Units per Selling Title That Will Be Scheduled During Each Quarter

Selling titles	Q1		Q2		Q3		Q4		Total		ST unit percent. (%)
	EQ30s	Imps (000)	EQ30s	Imps (000)	EQ30s	Imps (000)	EQ30s	Imps (000)	EQ30s	Imps (000)	
ST 1	50	1,951	59	2,512	85	3,773	100	4,489	294	12,725	30
ST 2	60	1,037	40	861		0		0	100	1,899	10
ST 3		0	25	519	50	1,162	75	1,468	150	3,149	15
ST 4	50	600	50	831	50	836	50	818	200	3,085	20
ST 5	15	75	35	173	25	132	25	125	100	505	10
ST 6	5	161	15	416	15	402	15	444	50	1,423	5
ST 7	5	86	15	214	15	249	15	241	50	790	5
ST 8	5	51	15	145	15	160	15	174	50	530	5
Totals	190	3,962	254	5,671	255	6,714	295	7,759	994	24,105	100

Deal impression pacing (%)		
Q1	3,962	16
Q2	9,633	40
Q3	16,346	68
Q4	24,105	100

Selling titles	W1	W2	W3	W4	W5	W6	W7	W8	W9	W10	W11	W12	W13	Qtr total
	EQ30s	EQ30s	EQ30s	EQ30s	EQ30s	EQ30s	EQ30s	EQ30s	EQ30s	EQ30s	EQ30s	EQ30s	EQ30s	EQ30s
ST 1		7	7		6	7	6		6	7	7	6		59
ST 2	8	8	8		8	8								40
ST 3							5		5	5	5	5		25
ST 4	5	5	5		5	5	5		5	5	5	5		50
ST 5	5		5		5	5			5		5	5		35
ST 6		5				5				5				15
ST 7		5				5				5				15
ST 8		5				5				5				15
Totals	18	35	25	0	24	40	16	0	21	32	22	21	0	254

Notes. These units are plotted across weeks within each quarter, except during hiatus weeks when the advertiser does not want to air any spots. The total impressions across the entire deal are the guaranteed impressions the network must meet.

Figure 4 shows a potential UD allocation for Q2 for the deal shown in Figure 3. Notice that although ST2 and ST5 do not appear in the UD allocation, the UD units closely match the selling-title mix from the original deal. Because ST2 and ST3 belong to the same group, their overall percentage of 25 percent is the group target in the UD allocation. Similarly, the target percentage for ST9 is the percentage from ST5 in the original deal. Observe also that the same number of units per week are assigned to the packaged selling titles ST6, ST7, and ST8, and that UD units are not assigned during W4, W8, and W13 because they are hiatus weeks for the original deal.

Planners manage UD allocations using indices that summarize the deal impression performance. In particular, the deal index is defined as the total projected impressions to be delivered at the end of the deal, divided by the total guaranteed impressions; however, deal indices can also be referred as percentages. Deal indices of less than one (i.e., less than 100 percent) imply a projected impression shortfall. Additionally, the season-to-date index is defined as the total

Figure 4. The Chart Illustrates UD Allocation Requirements

Selling titles	W1	W2	W3	W4	W5	W6	W7	W8	W9	W10	W11	W12	W13	Qtr total	ST unit percent. (%)
	EQ30s	EQ30s	EQ30s	EQ30s	EQ30s	EQ30s	EQ30s	EQ30s	EQ30s	EQ30s	EQ30s	EQ30s	EQ30s	EQ30s	
ST 1	2	1	2						2	1	2	1		11	26
ST 3							2		2	2	2	2		10	24
ST 4	1	1	1		1	1	1		1	1	1	1		10	24
ST 9			1		1	1			1			1		5	12
ST 6		1										1		2	5
ST 7		1										1		2	5
ST 8		1										1		2	5
Totals	3	5	4	0	2	2	3	0	6	4	5	8	0	42	100

Notes. UD allocations should attain a selling-title mix comparable to their corresponding deal paid units, should be distributed across weeks when paid units are in flight, should honor packaging rules from the original deal, and could use surrogate units for selling titles that are oversold, no longer exist, or will not air for a given period.

Downloaded from informs.org by [64.187.132.166] on 04 July 2024, at 08:32 . For personal use only, all rights reserved.

projected impressions to be delivered at the end of the current quarter divided by the total guaranteed impressions for the entire deal. In the last quarter that the deal is in flight, both indices are the same. Planners typically use the next quarter season-to-date indices to plan their UD allocations for the upcoming quarter, and together with the expected impression pacing from the original deal (i.e., the percentage of impressions that are expected to be delivered based on forecasts), they define lower and upper targets on these indices. Fundamentally, they would like to set up lower and upper targets that reduce the likelihood of under- or overdelivery, respectively. For example, a planner performing a UD allocation for Q2 for the deal illustrated in Figure 3 would look at the deal impression pacing that is expected at the end of Q2 from the original deal (i.e., 40 percent), and would allocate units to bring its season-to-date index for Q2 to a value close to it. The expertise of the planner regarding the deal, its industry category, and the specific knowledge of the client and the client's relationship with Turner are important factors in determining this target. Based on the deal impressions alone, the target for the season-to-date index can simply be set to 40 percent. However, knowledge of the advertiser might indicate the need to set this target to a higher value because the client might not advertise as much during Q3 (i.e., the client has many hiatus weeks) or might advertise too heavily during Q4; in the latter case, fitting many other additional UD units in that quarter might not be feasible. Some advertisers also exhibit preferences for the quarters when they like to receive their UDs. Some have time sensitivity in airing their spots and (or) they might be interested in advertising only during specific times of the year; for example, a tax preparation service wants to advertise during tax season and a retailer wants to advertise toys during the holiday season. This puts additional pressure on the network to determine the priorities for UD allocations.

Another UD allocation challenge is the handling of aged deals—the deals that are out of flight but have not yet met their guaranteed audience. The UD allocations for aged deals should follow a selling-title distribution similar to that in the original deal, using surrogate selling titles from the same group if necessary, and are typically distributed in weeks in which other current deals from the same advertiser are in flight. On rare occasions, some advertisers with aged deals might decide together with the network to no longer receive any additional spots on their plan and instead receive a write-off equivalent to the value of the unmet impressions. In some instances, this write-off might also be accompanied by a penalty depending on the conditions of the deal.

Traditionally, the allocation of audience-deficiency units at Turner was a manual process that required a myopic allocation of units of one, or at most, a few deals at a time. It was a long, time-consuming process that took weeks to complete and involved significant rework because of the problem's scale and complexity. UD allocation planners usually know the appropriate range of the lower and upper targets for a season-to-date index to ensure that the deal is good and will most likely not under- or overdeliver at the end of its flight. These lower and upper targets change depending on factors such as the advertiser category, the size of the deal, and the time sensitivity. However, because of the sequential nature of the manual process, planners were unable to enforce these targets for all deals and frequently had to rework allocations to release or swap units needed for other deals. Although the UD allocation problem cannot be fully automated because of the nature of the customer relationship management that any business-to-business operation entails, it was evident that approaching this problem via an optimization framework could yield significant improvements in time reduction and value derived from the UD inventory reserves. After several discussions with the UD allocation planners from all Turner domestic networks, we realized that they wanted to retain control of determining the lower and upper targets to manage UD allocations, but were open to letting a mathematical model determine the specific UD allocations across all deals. The course of action we decided to take was to develop an optimization engine that would allow planners to run multiple scenarios in which they could change the values of the lower and upper targets for various sets of deals and obtain an optimal UD allocation for those parameters. In the next section, we describe the multistage, mixed-integer programming approach that we developed to support the UD allocations at Turner.

An Optimization Framework for Underdelivery Allocation

UD allocations impact several levels of the media revenue management process. Araman and Popescu (2010)

Downloaded from informs.org by [64.187.132.166] on 04 July 2024, at 08:32 . For personal use only, all rights reserved.

investigate the problem of media capacity planning across upfront and scatter markets and conclude that such decisions must include the effect of audience uncertainty and make-good allocations. Our focus is at the tactical level when upfront versus scatter distribution decisions have already been made. The former work also reports that no mathematical model has been developed for capacity planning or make-good allocation. To the best of our knowledge, this is still the case; Turner is the only media company that reports using optimization technology to approach this problem. Although the decision of optimally distributing the nonprogramming airtime that remains after upfront sales into commercial airtime, promo time, and UD time is a complex problem in its own right, the current state of the practice is that the various departments responsible for scatter sales, promotional inventory, and UD allocation negotiate this unit distribution depending on market trends. For the purposes of this paper, we assume that this distribution has already been decided and that UD inventory reserves are fixed; therefore, we concentrate on allocating the UD reserve across all underperforming deals.

The UD allocation problem can be modeled as a mathematical programming problem in which the main decision variables are the allocations of UD units (in integer EQ30s) to allowable combinations of deals, selling titles, and weeks. The objective function is to maximize the overall value of the UD allocated across all deals, while honoring advertiser requirements and operational constraints. Next, we list the types of constraints that limit UD allocations.

Airtime Capacity Constraints

- For every selling title and week, the total number of UD units allocated across all deals cannot exceed the available units on the reserve.
- For every conflict code, selling title, and week, the total number of UD units associated with that conflict code allocated across all deals cannot exceed the maximum number of units that can still be scheduled. This restriction prevents allocating an excessive number of units with the same conflict code (e.g., insurance, soft drinks) that could not be feasibly scheduled at the spot-scheduling stage.
- The number of units allocated to a specific deal, selling title, and week, cannot exceed a threshold that takes into account the numbers of paid units and additional UD units for that advertiser that are already assigned. This restriction prevents scheduling allocations that could later be infeasible because of spot time-separation restrictions.

Season-to-Date Index Upper and Lower Targets

- For every deal, the projected season-to-date index must not exceed its upper target. These constraints are necessary to prevent deal over delivery.
- For every deal, the projected season-to-date index should exceed its lower target. These are soft constraints because the reserve might not contain enough units to guarantee that all deal indices can be brought within their desired range.
- Feasible allocations in which a deal meets its lower target must dominate feasible allocations in which the deal does not meet its lower target. Enforcing this constraint in a simple way is difficult because the lower-target constraint is a soft constraint.

Usage of Surrogate Selling Titles

- A deal can be assigned units on a selling title that belongs to the same group of an original selling title in the deal if and only if all available units of the original selling title have been exhausted. These constraints are particularly important for aged deals for which the original selling titles may no longer exist, but are also relevant to substitute units across selling titles that are deemed commensurate in value to advertisers.

Allocation of Packaged Selling Titles

- UD allocations should honor the packaging constraints in the original deal; that is, the number of units assigned to a given deal, selling title, and week must be equal for all the selling titles that were sold as a package to that advertiser.

Quantification of UD Value

- The value of all UD units assigned to a given deal is measured by the corresponding total impressions multiplied by the deal CPM. However, UD value for a given deal should only be quantified when it does not exceed the current dollar deficit for the corresponding deal. The reasons behind this business rule for the quantification of UD value are slightly esoteric and ingrained in accounting principles; however, for our purposes in this paper, it is sufficient to understand

Downloaded from informs.org by [64.187.132.166] on 04 July 2024, at 08:32 . For personal use only, all rights reserved.

that the value of the UD units assigned to a deal is capped by a deal dollar deficit; that is, the value of UD units assigned to a deal equals the minimum of the deal CPM times the UD allocation impressions (in thousands) and the deal dollar deficit.

Selling-Title Mix Distribution and Weekly Distribution

- The distribution of UD units allocated to a given deal should closely match the selling-title distribution of paid units in the original deal.
- UD units should be allocated only in weeks in which original paid units were purchased (i.e., no UD should be allocated for a given advertiser during that advertiser's hiatus weeks).
- The distribution of UD units allocated to a given deal should closely match the deal weekly weights. If no weights are specified in the deal, the units should be distributed equitably across all weeks in which allocations are made.

Incorporating some of these constraints into a mixed-integer program is straightforward; however, some, as we list, add to the complexity of the problem. The proper quantification of UD value is a nonlinear expression; the restrictions on the usage of surrogate selling titles lead to disjunctive constraints; the lower target limits on the season-to-date indices are hard constraints only if all deals have enough units to meet that condition; and the conflicting objectives between the business goals (e.g., maximizing UD value) and the advertiser requirements and operational constraints have vastly different metrics that are difficult to balance.

The nonlinear expression required to enforce the proper quantification of UD value states that a deal UD value equals the minimum between the gross value of the units allocated to the deal and the deal dollar deficit. This expression is a piecewise linear function (i.e., the minimum of a set of linear functions) that can be linearized by reformulating it into a set of linear inequalities: the UD value is less than or equal to the gross value of the units allocated to the deal, and the UD value is also less than or equal to the dollar deficit. Although this reformulation expands the feasible region, when combined with a maximization objective, it ensures that the optimal solution to the model will appropriately quantify the value of the allocated UD. Another reformulation we perform is to model the disjunction, which states that a deal can be assigned units on a selling title from the same group of a selling title in the original deal if and only if all available units of the original selling title have been exhausted. Additional auxiliary binary variables are defined for each selling title that reflect whether all its units in the reserve have been exhausted and additional constraints are added to relate these variables to the decisions that allocate UD units on surrogate selling titles.

Note that modeling the requirement that all deals should be allocated enough UD units so that their season-to-date indices meet their lower targets would imply a complex optimization model, because this is a hard requirement only if all deals have enough units to meet that condition. This requirement states that priority should be given to allocating units to deals that have not yet met their lower targets; however, if no units are available to allocate to a deal that has not met its lower target, then it is permissible to give additional units to deals that have already met their lower targets. Instead of attempting to incorporate all these requirements into a single model, this lends itself to a hierarchical two-phase optimization approach. Phase 1 allocates units to bring all deals beyond their lower targets if possible (or determines the maximum index that can be attained for each deal); Phase 2 maximizes the resulting UD value subject to the (possibly) new lower targets, which become hard constraints.

The mathematical formulation for Phase 1 determines whether the desired lower targets for the season-to-date indices across all deals are attainable, and UD value is used to prioritize allocations in which lower targets are not attainable. During this phase, impression shortfall reduction for a deal is valuable only while the new season-to-date index is less than or equal to the desired lower target for that deal; that is, the goal is to get each deal past its lower target. However, any additional impressions should occur only because of the discrete nature of the allocations. We can enforce this logic by adding one more element to the quantification of UD value, making it equal to the minimum of the gross value of the units allocated to the deal, the deal dollar deficit, or the value of the impression deficit below the lower target (i.e., the deal CPM times the impressions required to bring the deal to its lower target). The same reformulation logic described

Downloaded from informs.org by [64.187.132.166] on 04 July 2024, at 08:32 . For personal use only, all rights reserved.

above for the minimum of a set of linear functions still applies. Thus, Phase 1 maximizes this UD value subject to all pertinent constraints with the exception of the selling-title mix distribution and weekly distribution constraints. The allocations from Phase 1 determine whether original lower targets should be lowered to be attainable. The (possibly) updated lower targets that result from Phase 1 become hard constraints in Phase 2.

Note that the selling-title ratio mix and weekly distribution constraints are essentially soft constraints; therefore, they can be modeled as goals, and the deviation from these goals can be added as penalties in the objective function. However, this approach presents a challenge—determining the correct penalties that trade off dollar value against deviation from distribution goals (measured in EQ30s) is complex and can vary significantly. In addition, planners would like the ability to run different scenarios in which they could see the effect of varying the emphasis placed on the selling-title mix and weekly distribution, and thus being able to understand the impact that this would have on the reduction of UD value that is derived from the reserve. Therefore, we further subdivide Phase 2 into two more stages: Phase 2A maximizes the UD value that meets all constraints and the attainable lower targets computed in Phase 1; and Phase 2B minimizes the total penalty from deviations from selling-title mix distributions and weekly distributions subject to an additional lower bound on the overall UD value, which is a proportion of the maximum UD value determined in Phase 2A.

In addition to the constraints that appeared in Phase 2A, Phase 2B explicitly models the deviations from the desired selling-title group mix and weekly distributions. For both cases, the deviations in EQ30s are defined as the absolute value of the difference between the attained value and the goal value; for example, if 10 units are allocated to deal A, three of which are assigned to ST1, but the desired deal goal for that selling title is 20 percent, the difference between the attained value of 3 and the goal value of 2 is 1. These absolute values can be linearized by modeling the deviations as the difference between two nonnegative variables and defining the absolute value as the sum of those two nonnegative components. Different options can be used to minimize these deviations; in particular, one could use minisum or minimax objectives. For this application, the minimax objective yields better results because it spreads deviations more equitably. We quantify the maximum deviations across selling-title group unit goals by deal, and the maximum deviations across weekly unit goals by deal-selling title. Mirroring the logic to linearize the UD value quantification, the maximum across a set of linear functions can be linearized noting that the maximum must be greater than or equal to each of the functions and exactly equal to one of them. The minimization objective ensures the correctness of this reformulation. Bertsimas and Tsitsiklis (1997) provide additional details on all the linearizations we use in our models. Finally, Phase 2B has an additional constraint that restricts the minimum UD value that must be attained with the resulting allocation, which is essentially a proportion of the maximum UD value that was identified in Phase 2A. This proportion can be changed with a parameter (a dial) that represents the total percentage of UD value that the planners are willing to forgo to obtain a unit distribution that conforms more closely to the advertiser requirements. This parameter is intuitive to the planners, and they can run scenarios with different percentages until they find an allocation that they are confident will be amenable to advertisers without sacrificing too much UD value. We summarize this multistage optimization process in Figure 5 and provide the explicit mathematical programming formulations for each of these phases in Appendix B.

Implementation, Results, and Conclusions

We implemented the three-phase optimization framework described in the previous section in IBM ILOG CPLEX Optimization Studio 12.5 using the Optimization Programming Language (OPL) as the modeling language for the mixed-integer programs (MIPs) and IBM ILOG Script for OPL to execute the flow control across the models. The optimal solution to the MIP program solved in each phase is entered as a warm start to the MIP in the next phase, which reduces the solution time. This is particularly helpful to the model in Phase 2B because the solver by itself has difficulty in finding an initial feasible solution. Solution times are less than one minute for the first two phases, even for the largest instances involving hundreds of deals, dozens of selling titles, and entire broadcast quarters (13 weeks). In the third phase, as we might expect,

Downloaded from informs.org by [64.187.132.166] on 04 July 2024, at 08:32 . For personal use only, all rights reserved.

Figure 5. Three-Phase Optimization Framework

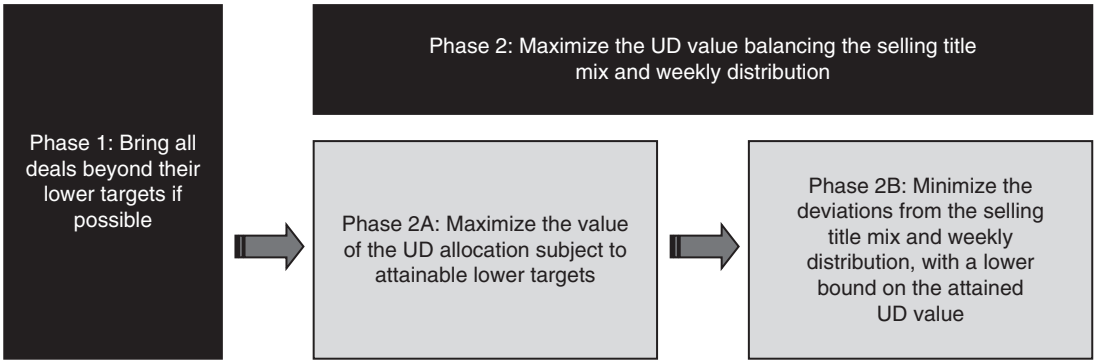

Notes. The three phases determine the minimum index that can be attained for each deal, calculate the maximum possible UD value that can be produced with the given inventory reserve, and then trade off a portion of that value to produce an allocation that has an improved selling-title mix distribution and weekly distribution.

the model takes longer to solve because of the weak linear programming (LP) relaxation that arises in minimax problems. However, the model generally produces results that are within 10 percent of optimality in 15 minutes of solution time. Given that this is a tactical problem that does not require solving in real time and Phase 3 minimizes penalties that are used as proxies to measure unit spread across selling titles and weeks, we cap the Phase 3 solution time at 15 minutes.

We built an application, the Liability Efficiency Optimizer (LEO), around this three-phase optimization engine. Figure 6 depicts the UD allocation process, which comprises creating UD reserves in the inventory system, generating optimal UD allocations across all deals in LEO whose functions include producing deal-level reserves in the inventory-management system, booking UD orders in the order-management system once the advertisers have agreed with their UD allocations, scheduling the UD spots on their corresponding selling-title weeks, and updating the pacing and (or) indexing of the deals in the deal-stewardship system once the UD spots have aired and their impressions are known. LEO is at the core of the UD allocation process. Planners use LEO to create a UD plan; depending on the timing of their allocations, they can create several plans throughout the broadcast quarter. They search for liable deals of interest to add to the plan; examples include deals whose current index is below a certain threshold and deals that belong to a specific conflict code. They set the lower and upper targets for these deals based on various scenarios that they would like to evaluate. For example, in a scenario in which they want to prioritize aged deals, they might set higher lower and upper targets for aged deals and set them lower for the other deals. In other scenarios, they might focus on deals that will soon be out of flight or on deals that belong to a specific conflict code, such as department stores near Christmas time or tax preparation services near the tax filing deadline. In yet a different scenario, they might simply want to be more aggressive or more conservative with the dial that trades off liability value for an improved unit spread. The interface enables planners to run as many scenarios as desired and provides them with tools to compare solutions across the scenarios. It also allows planners to make manual changes to the allocations so that they can enforce rules that cannot be generalized in a mathematical model, such as advertiser-specific requests. Once a planner selects a solution across scenarios, that planner can commit the solution to the system, which automatically breaks the liability inventory reserve into individual deal reserves with their corresponding audience-deficiency units. Once the advertiser has found the UD allocation acceptable, UD orders are created so that the explicit audience-deficiency spots will be scheduled on their specific weeks.

The development of LEO underwent the same rigorous software development process as all applications developed at Turner to ensure compatibility with all company applications and systems. The development team included software architects, software developers, business analysts, user interface designers, quality assurance testers, and an operations research analyst. The first version of LEO was put into production in

Downloaded from informs.org by [64.187.132.166] on 04 July 2024, at 08:32 . For personal use only, all rights reserved.

Figure 6. The Liability Efficiency Optimizer (LEO) Is at the Core of the UD Allocation Process and Integrates with Other Advertisement-Management Systems, Including Traffic, Order Management, Inventory Management, and Deal Stewardship

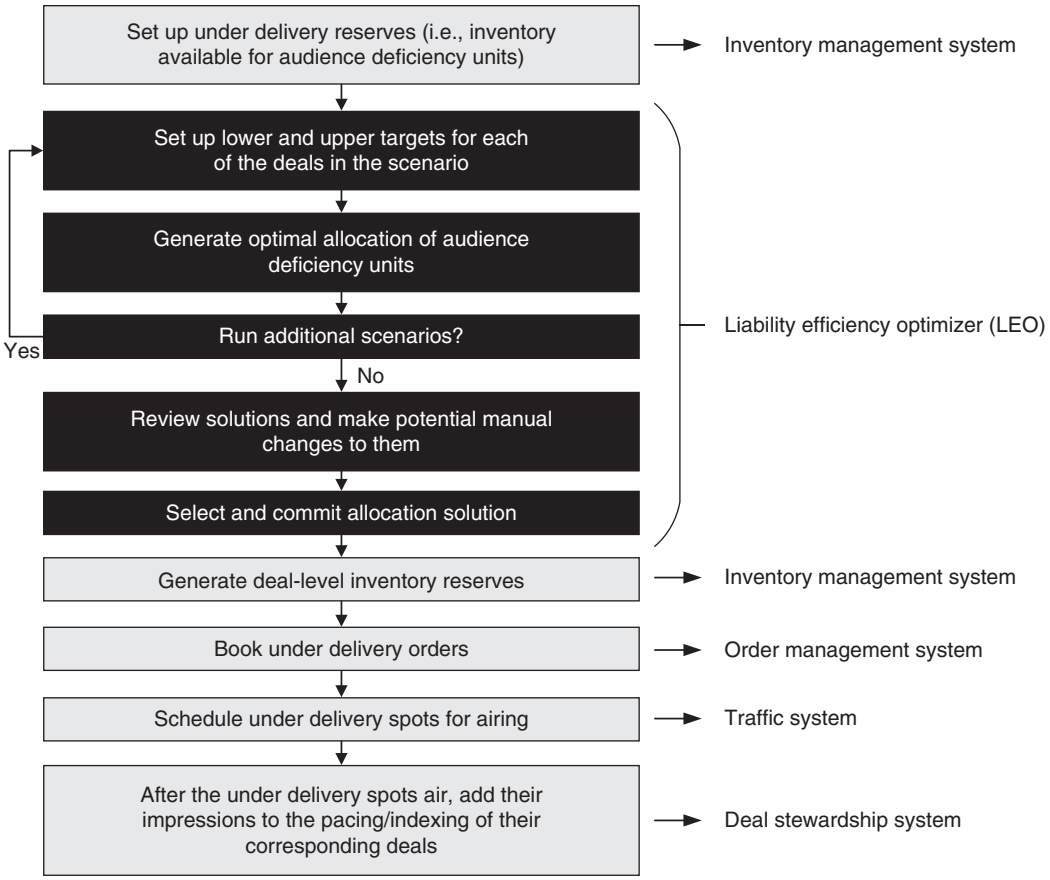

the fourth quarter of 2013 for the major entertainment networks, TBS and TNT. The development team has continued to integrate functionality to accommodate the other Turner networks and user-interface features requested by the users. To date, all domestic networks use this system each quarter to perform their allocations of audience-deficiency units. Table 1 summarizes some statistics that showcase the volume of plans and size of the allocations in terms of deals, selling titles, weeks, and units (in EQ30s).

The LEO implementation has had a threefold impact at Turner. First, it has provided significant productivity improvements, because the process of allocating audience-deficiency units across hundreds of deals has decreased from weeks to minutes. The frequency of UD allocations has increased and the process is now performed two, and sometimes three, times per quarter. Second, it has enabled planners to test and evaluate multiple scenarios, a benefit they did have not have with the prior manual process. Third, it has provided significant monetary benefits, which Turner has conservatively calculated to be approximately $10 million of additional liability reduction per year over the manual

Table 1. Since its Deployment Into Production, LEO Has Been Used to Allocate Hundreds of Thousands of Audience-Deficiency Units Across All of Turner's Domestic Networks

Network	Plan count	Average deals	Average selling titles	Average weeks	Average scenarios	Unit count
TBS	57	639.9	176.0	6.2	2.8	53,931
TNT	27	283.34	103.7	6.6	2.0	20,773
TRU	9	120.1	33.6	1.9	1.0	1,008
CNN	73	330.7	110.8	6.9	2.2	94,926
HLN	34	166.8	37.2	8.2	2.5	44,190
NBA TV	15	87.4	194.5	3.2	1.13	17,154

Downloaded from informs.org by [64.187.132.166] on 04 July 2024, at 08:32 . For personal use only, all rights reserved.

process that it replaced. The monetary impact was measured during the quarter prior to deploying LEO into production. The allocations were run off-line in parallel with running the manual process. As input, we used the same inventory that had been used for the manual allocations, and we set the desired lower target for each deal index as the index that had resulted from the manual allocations. The difference between the UD value from the allocation that LEO generated and the UD value of the manual allocation was considered to be the optimization benefit. We expressed this benefit as a percentage of the total quarter liability; to be conservative, we applied this same percentage to the quarter that had the lowest liability in that year. This resulted in an estimated benefit of $2.5 million for the quarter, which we annualized to $10 million. In summary, the use of an optimization framework to solve the audience-deficiency-unit allocation problem has proven successful and represents a state-of-the-art contribution to the media industry revenue-management process, which has been stagnant in adopting OR techniques.

We envision several directions for future work, and categorize them as methods, business processes, and problem evolution as follows.

- Methods: (1) Explore strengthening the LP relaxation of the MIP we use in Phase 3 of our framework; we believe this could yield solution improvements; (2) explore more sophisticated methods to determine appropriate lower and upper targets for different sets of deals; in particular, analyze the allocations LEO has generated to extract knowledge about which targets are considered good (i.e., the ones from the selected allocations). Doing so could help test newsvendor-like models to possibly reduce the dependence on the expertise of the planners to set these targets; and (3) explore the explicit incorporation of the uncertainty of the estimates used in our models via stochastic or robust optimization; this could lead to further value extraction of the UD allocation process.

- Business processes: (1) Beyond the current problem, explore the "dual" version of UD capacity planning that, given a set of liable deals, entails determining the appropriate number of inventory units needed in specific selling-title weeks to bring the deals within a desired delivery range; and (2) integrate this process into the inventory-mix distribution problem, which determines the proportions of nonprogramming airtime that should be devoted to scatter sales, promotional efforts, and UD allocation. Turner is currently developing proof-of-concept studies to show the value of such an integrated approach.

- Problem evolution: Two emerging trends have disrupted the paradigms in the media industry: (1) The use of data fusion has allowed new consumer data sources, such as frequent-shopper data and credit card data, to be merged with viewership data to attain novel targeted audiences, which have challenged the traditional methods of forecasting, scheduling, and inventory management. Within this context, proposal units cross network boundaries (as opposed to classic, demographic-guaranteed deals, which typically involve a single network) and deal guarantees are expressed in terms of granular audiences instead of broad demographics. These deal guarantees can be given in absolute terms (i.e., a total number of targeted audience impressions) or they can be combined as secondary guarantees on top of demographic impression guarantees. In the latter case, a benchmark, usually a median or average schedule, is used to measure the relative lift (i.e., additional delivery) in targeted segment viewership. (2) Traditionally, linear-content distributors that also have digital and mobile properties have negotiated their advertising deals separately within each platform. In this context, digital and mobile video advertisements have exhibited lower sellout rates with deals that rarely require make-good allocations, which opens the possibility of using digital impressions as UD units for linear deals. Furthermore, content providers are quickly moving into cross-platform advertising solutions that offer advertisers holistic deals across their entire media portfolios; these portfolios can include cable and (or) broadcast networks and digital and mobile properties, which include a cross-platform audience-deficiency component. For both of these emerging trends, the media industry is still developing the business rules that will direct the allocation of audience-deficiency units. As these rules become paradigms, the processes and models used to allocate UD will have to be extended to accommodate them.

Appendix A. Glossary of Media Terms

In this appendix, we provide a brief glossary of terms used in the media industry.

Downloaded from informs.org by [64.187.132.166] on 04 July 2024, at 08:32 . For personal use only, all rights reserved.

- *Audience-deficiency units*: Unpaid units that are given to an advertiser to reduce the expected impression shortfall for a given deal. They are also called make-good, under-delivery units, or simply UD.
- *CPM*: Cost-per-thousand impressions (deal budget divided by the number of guaranteed impressions in thousands). The CPM can be computed at different levels, including deal, network, quarter, and selling title.
- *Conflict code or category*: The industry category in which the advertiser falls; examples of conflict codes include telecommunications, restaurants, retailers, and automotive.
- *Deal*: A media deal between an advertiser and a broadcast or cable network, which specifies flighting dates, total budget, CPM, quarterly impression distribution, selling-title mix, spot-length mix, weekly distribution goals, primary demographic, ratecard type, and a guaranteed audience level in the target demographic. A deal starts as a proposal and becomes a deal only when both parties agree to all terms.
- *Deal index*: Equals the total projected impressions to be delivered at the end of the deal (which include actual impressions from aired spots and estimated impressions for future spots) divided by the total guaranteed impressions. Deal indices that are less than one imply a projected impression shortfall.
- *Liability*: The difference between the guaranteed impressions and the total delivered impressions. Negative liability provides no additional revenue to the network; positive liability represents a financial commitment that must be met. Liability is usually projected during the flighting of a deal, and additional unpaid units (UD) can be given to the advertiser to reduce the expected liability.
- *Primary demographic* (*demo*): The main segment within the overall population that the advertiser is trying to reach with its commercial spots. This segment is considered the one most likely to consume the advertiser products. A primary demo is specified with a gender (i.e., male, female, or people) and an age range (e.g., 18–49, 25–54, 60+).
- *Ratecard*: The metrics of audience in different target demographics. They represent the expected audience of a specific selling title. Different types of ratecards include live (average number of people watching a particular show when it airs, including both commercial airtime and programming), ACM+3 (average audience watching only commercial airtime either live or within three days of airing via DVR systems), and ACM+7 (same as ACM+3 but with seven days of delayed viewing).
- *Scatter*: The selling of commercial airtime throughout the broadcast year after the upfront market has concluded.
- *Season-to-date index*: The total projected impressions to be delivered at the end of the current quarter, which include actual impressions from aired spots and estimated impressions for future spots within the current quarter, divided by the total guaranteed impressions for the entire deal.
- *Selling title*: An interval of programming time that a network uses to sell commercial airtime, it can comprise a particular show (e.g., Dallas, Cougar Town) or a block or time (e.g., Daytime—8:00 AM to 3:00 PM, Overnight—12:00 midnight to 6:00 AM). It can also be referred to as a daypart.
- *Selling-title group*: A set of selling titles that are comparable in audience composition and value (i.e., their price per spot is approximately the same). Selling titles can be used as surrogates for other selling titles in the same group when the original selling titles are no longer in flight or have no additional units available.
- *Selling-title package*: Selling titles of related shows that are sold as a package to advertisers. Examples include shows that have several telecasts per week (e.g., premiere, repeat, encore), and movies with multiple telecasts.
- *Underdelivery units*: See audience-deficiency units.
- *Upfront*: The first selling wave of units in a new broadcast year for broadcast or cable networks. It usually occurs around the middle of May after the new fall schedules have been announced and presented to major advertisers.

The definitions of additional media terms can be found in Nielsen Media Research's Glossary at http://www.nielsenmedia.com/glossary/.

Appendix B. Mathematical Formulations

This appendix contains the formulations of the models described in the *An Optimization Framework for Under-Delivery Allocation* section. First, we list all the sets, parameters, and decision variables that are used across all formulations; we then describe the specific MIP formulations used in each phase.

Sets

$\mathcal{D}$ = the set of all deals, indexed by d.
$\mathcal{S}$ = the set of all selling titles, indexed by s.
$\mathcal{S}_d$ = the subset of selling titles on deal d, indexed by s.
$\mathcal{W}$ = the set of all weeks in the allocation period, indexed by w.
$\mathcal{C}$ = the set of all product conflict categories, indexed by c.
$\mathcal{G}$ = the set of all selling-title groups, indexed by g. $\bigcup_{g\in\mathcal{G}} = \mathcal{S}$ and $g1 \cap g2 = \phi$.
$\mathcal{P}$ = the set of all selling-title packages, indexed by p.

Parameters

T_d = Target demographic for deal d.
V_d = Dollar value of UD impressions for deal d (current quarter CPM).
B_d^M = Dollar deficit for deal d.
J_d = Guaranteed impressions for deal d.
B_d^E = Season-to-date impression balance for deal d.
I_d^C = Current season-to-date index of deal d.
I_d^L, I_d^U = Lower and upper season-to-date index targets for deal d, $I_d^L \geq I_d^C, I_d^U \geq I_d^L$.
P_d = Product conflict code associated with deal d.
R_d = Ratecard type for deal d.
F_{gd}^G = Target proportion of units in deal d to be allocated to selling titles in group g.
F_{dw}^W = Target proportion of units in deal d to be allocated in week w for a given selling title.

Downloaded from informs.org by [64.187.132.166] on 04 July 2024, at 08:32 . For personal use only, all rights reserved.

A_{sw} = Number of UD units available on selling title s in week w.
Q_{swc} = Number of c-category units that can still be scheduled on selling title s in week w without violating product-conflict constraints.
O_{dsw} = Number of units that can still air on selling title s in week w for deal d without violating time-separation constraints.
E^{TR}_{sw} = Number of people in target demographic T who are expected to watch selling title s in week w according to ratecard type R.
β = Proportion of liability reduction that can be given up to attain a better unit spread across selling titles and weeks.
π^G, π^W = Penalties associated with the deviations from unit targets across selling-title groups and across weeks, respectively.
M = A large positive constant (commonly known as big-M).

Decision Variables

x_{dsw} = Number of UD units allocated to deal d on selling title s in week w.
Θ_d = Total number of UD units allocated to deal d.
y_d = Total UD impressions allocated to deal d.
z_{dpw} = Number of units from package p allocated to deal d in week w.
i^N_d = New season-to-date index for deal d.
v_d = Value of the UD units assigned to deal d.
λ = Overall liability reduction across all deals.
$\delta^{G+}_{gd}, \delta^{G-}_{gd}$ = Positive and negative deviations from the goal on the number of UD units allocated to deal d on selling titles in group g.
Δ^G_d = Maximum deviation across the selling-title group unit goals for deal d.
$\delta^{W+}_{dsw}, \delta^{W-}_{dsw}$ = Positive and negative deviations from the goal on the number of UD units allocated to deal d on selling ti; le s in week w.
Δ^W_{ds} = Maximum deviation across the weekly goal unit goals for deal d in selling title s.
α_s = Remaining units on selling title s after allocation across all weeks.
ψ_{ds} = Boolean variable that equals 1 if selling title s can be substituted by other selling titles in its group; 0 otherwise.

Phase 1 Formulation

In Phase 1, audience-deficiency units are allocated to bring all deals as close as possible to their desired lowered target and UD value is used to prioritize allocations in which lower targets are not attainable.

$$\text{Maximize } \lambda = \sum_{d\in\mathcal{D}} v_d \tag{O1}$$

Subject to:

$$v_d \le V_d\, y_d, \quad \forall d \in \mathcal{D} \tag{B.1}$$
$$v_d \le B^M_d, \quad \forall d \in \mathcal{D} \tag{B.2}$$
$$v_d \le (I^L_d - I^C_d) V_d J_d, \quad \forall d \in \mathcal{D} \tag{B.3}$$
$$\sum_{d\in\mathcal{D}} x_{dsw} \le A_{sw}, \quad \forall s \in \mathcal{S}, \forall w \in \mathcal{W} \tag{B.4}$$
$$\sum_{\{d\in\mathcal{D}:P_d=c\}} x_{dsw} \le Q_{swc}, \quad \forall c \in \mathcal{C}, \forall s \in \mathcal{S}, \forall w \in \mathcal{W} \tag{B.5}$$
$$x_{dsw} \le O_{dsw}, \quad \forall d \in \mathcal{D}, \forall s \in \mathcal{S}, \forall w \in \mathcal{W} \tag{B.6}$$
$$y_d = \sum_{s\in\mathcal{S}} \sum_{w\in\mathcal{W}} E^{T_d R_d}_{sw} x_{dsw}, \quad \forall d \in \mathcal{D} \tag{B.7}$$
$$i^N_d = I^C_d + y_d / J_d, \quad \forall d \in \mathcal{D} \tag{B.8}$$
$$i^N_d \le I^U_d, \quad \forall d \in \mathcal{D} \tag{B.9}$$
$$\alpha_s = \sum_{w\in\mathcal{W}} A_{sw} - \sum_{d\in\mathcal{D}} \sum_{w\in\mathcal{W}} x_{dsw}, \quad \forall s \in \mathcal{S} \tag{B.10}$$
$$\sum_{\{s'\in g: s\in g, s'\neq s\}} x_{ds'w} \le \mathrm{M}\, \psi_{ds}, \forall d \in \mathcal{D}, \forall s \in \mathcal{S}_d \tag{B.11}$$
$$\alpha_s \le \mathrm{M}\,(1-\psi_{ds}), \quad \forall d \in \mathcal{D}, \forall s \in \mathcal{S}_d \tag{B.12}$$
$$z_{dpw} = x_{dsw}, \quad \forall d \in \mathcal{D}, \forall p \in \mathcal{P}, \forall s \in p \cap \mathcal{S}_d, \forall w \in \mathcal{W} \tag{B.13}$$
$$x_{dsw} \in \mathbb{Z}^+$$
$$y_d, z_{dpw}, i^N_d, v_d, \alpha_s \in \mathbb{R}^+$$
$$\psi_{ds} \in \{0,1\}$$

The objective function (O1) maximizes the total value of the UD allocation across all deals. The effective UD value of a given deal is the minimum of the gross value of the units allocated to the deal (current quarter CPM times the number of impressions resulting from the allotted units) or the deal dollar deficit. In addition, for Phase 1 this value is also capped once the deal reaches its desired lower target, which ensures that deals will be brought as close as possible to their lower targets, but will only exceed their lower targets because of the discrete nature of the allocations. This yields the following expression for the UD value of a given deal d:

$$v_d = \min\{V_d y_d, B^M_d, (I^L_d - I^C_d) V_d J_d\}.$$

Constraints (B.1)–(B.3) linearize this expression. Although this linearization expands the feasible region, at optimality v_d will indeed equal the minimum of those three quantities.

Constraints (B.4)–(B.6) enforce the airtime capacity constraints. The units allocated across all deals to a selling-title week should not exceed the available units (constraint (B.4)); the number of units of a specific conflict code assigned to a selling-title week cannot surpass the maximum units that can still be scheduled for that conflict code (constraint (B.5)); and the number of units allocated to a given deal on a specific selling-title week cannot exceed the maximum number of units that could be later scheduled given time-separation requirements (constraint (B.6)). Note that constraint (B.6) can simply be enforced as an upper bound on the decision variables.

Downloaded from informs.org by [64.187.132.166] on 04 July 2024, at 08:32 . For personal use only, all rights reserved.

Constraint (B.7) is an expression that defines the total number of impressions allocated to deal d, which equals the total units assigned to that deal across all selling title-weeks, weighted by the estimated audience on each selling-title week in the corresponding target demographic and ratecard.

Constraints (B.8) and (B.9) relate to the new season-to-date index for deal d. The new season-to-date index equals the current index, plus the proportion of newly allocated impressions to guaranteed impressions (constraint (B.8)), and this new season-to-date index cannot exceed the desired upper target (constraint (B.9)).

Constraints (B.10)–(B.12) impose the correct usage of surrogate selling titles; that is, a deal can be assigned units on a selling title that belongs to the same group of the original selling title in the deal if and only if all available units of the original selling title have been consumed. Constraint (B.10) states that the remaining units that are available for selling title s is the difference between all units originally available minus all units assigned across all deals. Constraints (B.11) and (B.12) are disjunctive constraints stating that either all units of selling title s have been exhausted ($\psi_{ds} = 1$) and units can be assigned to a surrogate selling title, or there are units still remaining for selling title s ($\psi_{ds} = 0$), and as such no units can be assigned to surrogate selling titles. The big-M coefficients used in these constraints are computed to be as small as possible to avoid numerical issues such as ill-conditioned matrices.

Finally, constraint (B.13) enforces selling-title packaging rules; that is, the same number of units assigned on a selling-title week must be equal for all selling titles that were sold as a package on a deal.

Phase 2A Formulation

In Phase 2A, audience-deficiency units are allocated to maximize the value of the allocated UD subject to the attainable lower targets from Phase 1.

$$\text{Maximize} \quad \lambda = \sum_{d\in\mathcal{D}} v_d \tag{O2}$$

Subject to:

$$\text{(B.1)–(B.2)}$$
$$\text{(B.4)–(B.13)}$$
$$i_d^N \geq \min(I_d^L, i_d^{N*}(1)), \quad \forall\, d \in \mathcal{D} \tag{B.14}$$
$$x_{dsw} \in \mathbb{Z}^+$$
$$y_d, z_{dpw}, i_d^N, v_d, \alpha_s \in \mathbb{R}^+$$
$$\psi_{ds} \in \{0,1\}$$

The objective function (O2) maximizes the total value of the UD allocation across all deals. Therefore, it is the same as (O1). However, the definition of v_d changes for this phase because the value is no longer capped by the lower target. For this phase, the value of the UD for a given deal d becomes

$$v_d = \min\{V_d\ y_d, B_d^M\}.$$

Constraints (B.1) and (B.2) linearize this expression; constraints (B.4)–(B.13) are still valid in this phase, and constraint (B.14) enforces the attainable lower bound on the new season-to-date index for deal d, which is the minimum between the original lower target for deal d and the season-to-date index resulting from the Phase 1 allocation $i_d^{N*}(1)$.

Phase 2B Formulation

In Phase 2B, audience-deficiency units are allocated to minimize the total penalty from deviations from selling-title-mix distributions and weekly distributions subject to an additional lower bound on the overall UD value that must be attained.

$$\text{Minimize} \quad \pi^G \sum_{d\in\mathcal{D}} \Delta_d^G + \pi^W \sum_{d\in\mathcal{D}} \sum_{s\in\mathcal{S}} \Delta_{ds}^W \tag{O3}$$

Subject to:

$$\text{(B.1)–(B.2)}$$
$$\text{(B.4)–(B.14)}$$
$$\Theta_d = \sum_{s\in\mathcal{S}} \sum_{w\in\mathcal{W}} x_{dsw}, \quad \forall\, d \in \mathcal{D} \tag{B.15}$$
$$\sum_{w\in\mathcal{W}} \sum_{s\in g} x_{dsw} + \delta_{gd}^{G+} - \delta_{gd}^{G-} = F_{gd}^G \Theta_d, \quad \forall\, d \in \mathcal{D}, \forall\, g \in \mathcal{G} \tag{B.16}$$
$$\Delta_d^G \geq \delta_{gd}^{G+}, \quad \forall\, g \in \mathcal{G}, \forall\, d \in \mathcal{D} \tag{B.17}$$
$$\Delta_d^G \geq \delta_{gd}^{G-}, \quad \forall\, g \in \mathcal{G}, \forall\, d \in \mathcal{D} \tag{B.18}$$
$$x_{dsw} + \delta_{dsw}^{W+} - \delta_{dsw}^{W-} = F_{dw}^W \Theta_d, \quad \forall\, d \in \mathcal{D}, \forall\, s \in \mathcal{S}, \forall\, w \in \mathcal{W} \tag{B.19}$$
$$\Delta_{ds}^W \geq \delta_{dsw}^{W+}, \quad \forall\, d \in \mathcal{D}, \forall\, s \in \mathcal{S}, \forall\, w \in \mathcal{W} \tag{B.20}$$
$$\Delta_{ds}^W \geq \delta_{dsw}^{W-}, \quad \forall\, d \in \mathcal{D}, \forall\, s \in \mathcal{S}, \forall\, w \in \mathcal{W} \tag{B.21}$$
$$\sum_{d\in\mathcal{D}} v_d \geq (1-\beta)\lambda^*(2A) \tag{B.22}$$
$$x_{dsw} \in \mathbb{Z}^+$$
$$\Theta_d, y_d, z_{dpw}, i_d^N, v_d, \delta_{gd}^{G+}, \delta_{gd}^{G-}, \Delta_d^G, \delta_{dsw}^{W+}, \delta_{dsw}^{W-}, \Delta_{ds}^W, \alpha_s \in \mathbb{R}^+$$
$$\psi_{ds} \in \{0,1\}$$

The objective function (O3) minimizes the sum of total penalties resulting from the maximum deviations across the selling-title group unit goals by deal and the penalties resulting from the maximum deviations across weekly unit goals by deal selling title.

Constraints (B.1), (B.2), and (B.4)–(B.14) are still valid in this phase and have the same interpretations as above; however, constraint (B.15) is an expression that quantifies the total number of units allocated to deal d across all selling-title weeks.

Constraints (B.16)–(B.18) quantify the maximum deviations across the selling-title group unit goals by deal. These maximum deviations are defined as

$$\Delta_d^G = \max_{d\in\mathcal{D}, g\in\mathcal{G}} \left| F_{gd}^G \Theta_d - \sum_{w\in\mathcal{W}} \sum_{s\in g} x_{dsw} \right|.$$

Downloaded from informs.org by [64.187.132.166] on 04 July 2024, at 08:32 . For personal use only, all rights reserved.

These expressions can be linearized in two steps. First, the absolute value is linearized by modeling it as the difference of two nonnegative variables, δ_{gd}^{G+} and δ_{gd}^{G-}. Constraint (B.16) represents this transformation. Then, the maximum is linearized by noting that the maximum deviation must be at least as large as each of its elements. Constraints (B.17) and (B.18) enforce this logic. Since the objective is minimization, at optimality, one of these constraints will be binding, which completes the correctness of the reformulation.

Constraints (B.19)–(B.21) quantify the maximum deviations across weekly unit goals by deal selling title. These maximum deviations are defined as

$$\Delta_{ds}^{W} = \max_{d\in\mathcal{D}, s\in\mathcal{S}, w\in\mathcal{W}} |F_{dw}^{W}\Theta_{d} - x_{dsw}|.$$

The logic to linearize these expressions parallels the one we present earlier. Additional details on these reformulations can be found on Bertsimas and Tsitsiklis (1997).

Finally, constraint (B.22) enforces a lower bound on the overall UD value that results from the allocation. This value is a proportion of the optimal value $\lambda^{*}(2A)$ from Phase 2A, in which no selling-title-mix distribution goals or weekly distribution goals were considered.

Acknowledgments

The authors thank Nicole Slavin and David Benoit at Turner for their help and support to implement this work. They also thank the anonymous reviewers for their valuable comments and suggestions in revising the paper.

References

Araman VF, Popescu I (2010) Media revenue management with audience uncertainty: Balancing upfront and spot market sales. *Manufacturing Service Oper. Management* 12(2):190–212.

Bai R, Xie J (2006) Heuristic algorithms for simultaneously accepting and scheduling advertisements on broadcast television. *J. Inform. Comput. Sci.* 1(4):245–251.

Bertsimas D, Tsitsiklis J (1997) *Introduction to Linear Optimization* (Athena Scientific, Belmont, MA).

Bollapragada S, Garbiras M (2004) Scheduling commercials on broadcast television. *Oper. Res.* 52(3):337–345.

Bollapragada S, Mallik S (2007) Managing on-air ad inventory in broadcast television. *IIE Trans.* 40(12):1107–1123.

Bollapragada S, Bussieck MR, Mallik S (2004) Scheduling commercial videotapes in broadcast television. *Oper. Res.* 52(5):679–689.

Bollapragada S, Cheng H, Phillips M, Garbiras M, Scholes M, Gibbs T, Humphreville M (2002) NBC's optimization systems increase revenues and productivity. *Interfaces* 32(1):47–60.

Brusco MJ (2008) Scheduling advertising slots for television. *J. Oper. Res. Soc.* 59(10):1363–1372.

Danaher P, Mawhinney D (2001) Optimizing television program schedules using choice modeling. *J. Marketing Res.* 38(3): 298–312.

Danaher PJ, Dagger TS (2012) Using a nested logit model to forecast television ratings. *Internat. J. Forecasting* 28(3):607–622.

Danaher PJ, Dagger TS, Smith MS (2011) Forecasting television ratings. *Internat. J. Forecasting* 27(4):1215–1240.

Gaur DR, Krishnamurti R, Kohli R (2009) Conflict resolution in the scheduling of television commercials. *Oper. Res.* 57(5):1098–1105.

Gensch D, Shaman P (1980) Models of competitive television ratings. *J. Marketing Res.* 17(3):307–315.

Giallombardo G, Jiang H, Miglionico G (2016) New formulations for the conflict resolution problem in the scheduling of television commercials. *Oper. Res.* 64(4):838–848.

Horen JH (1980) Scheduling of network television programs. *Management Sci.* 26(4):354–370.

Kimms A, Muller-Bungart M (2007) Revenue management for broadcasting commercials: The channel's problem of selecting and scheduling the advertisements to be aired. *Internat. J. Revenue Management* 1(1):28–44.

Meyer D, Hyndman RJ (2006) The accuracy of television network rating forecasts: The effects of data aggregation and alternative models. *Model Assisted Statist. Appl.* 1(3):147–155.

Popescu DG, Crama P (2016) Ad revenue optimization in live broadcasting. *Management Sci.* 62(4):1145–1164.

Popescu DG, Seshadri S (2016) Managing revenue from television advertising sales. Accessed April 28, 2017, http://faculty.insead.edu/dana-popescu/documents/Ad%20Sales.pdf.

Reddy SK, Aronson JE, Stam A (1998) SPOT: Scheduling programs optimally for television. *Management Sci.* 44(1):83–102.

Zhang X (2006) Mathematical models for the television advertising allocation problem. *Internat. J. Oper. Res.* 1(3):302–322.

Verification Letter

Nicole Slavin, SVP, Strategic Planning, Turner Broadcasting Sales, One Time Warner Center, New York, NY 10019, writes:

"The purpose of this letter is to verify the use and impact of the Liability Efficiency Optimizer (LEO), a solution based system which allocates audiences deficiency units against liability incurrence on media deals. This system was developed by members of our Analytics, Data, and Decision Sciences (ADDS) team at Turner Broadcasting.

"Our relationship with the ADDS team started several years ago and has proven to be extremely beneficial. The team has developed and implemented several advanced analytics, operations and research decision support systems that help Strategic Planning, and other divisions in our organization optimize revenue and improve productivity. Examples include a spot scheduler engine that schedules thousands of commercial spots every day satisfying a plethora of constraints; a brand allocation engine that makes the best use of our inventory in assigning branded spots for large advertiser orders; and a competitive audience estimation framework that forecasts audiences at different horizons for different demos and targeted audience segments.

"The development and implementation of LEO is another example of the amazing work of the ADDS team. Since its deployment in the fourth quarter of 2013, LEO has significantly improved the way in which the Strategic Planning team allocates audience deficiency units (ADU) on all domestic networks, which has resulted in more accurate and efficient inventory and revenue management. Each broadcast quarter, audience deficiency units need to be allocated to make good on audience guarantees across hundreds of deals, multiple selling titles with an extensive number of options, rules, client restrictions, and business objectives that need to be balanced. Given the size and complexity of the ADU

Downloaded from informs.org by [64.187.132.166] on 04 July 2024, at 08:32 . For personal use only, all rights reserved.

allocation process, there are limitations on how good a manual solution can be, despite the skills and expertise of the team analysts. LEO has taken what used to be a manual, time consuming process that would take several weeks to complete, into an allocation process which only takes minutes.

"In addition, this system allows analysts to run different scenarios (something that was simply prohibitive before), and choose the allocation that best meets the current business needs. The expertise of the analysts is used to set up and run different scenarios and it is up to LEO to determine the optimal allocation that maximizes the liability reduction within each scenario. In short, LEO ensures that we are making the best use of the limited inventory that each network has available to clear deal liability. The financial impact of LEO has been conservatively estimated to be \$10M of annual liability reduction with respect to the manual processes that it replaced.

The outstanding results obtained by using advance analytics in our media planning and control problems has encouraged us to expand the use of advanced analytics/operations research other areas in our organization. Overall, our satisfaction with the system is very high and we expect to continue its use and enhance it further with additional functionality."

José Antonio Carbajal is a Director of Operations Research and Advanced Analytics at Turner, where he leads an Optimization and Revenue Management team that develops optimization-based decision support systems that assist Turner management with decision making, policy formulation, and other managerial functions. He received both his PhD in industrial engineering and an MS in operations research from Georgia Institute of Technology. He earned his BS in industrial and systems engineering from Monterrey Institute of Technology.

Wes Chaar is the Senior Vice President of Analytics, Data and Decision Sciences at Turner. Wes attained his PhD from the University of Texas at Austin. He has more than 20 years of expertise in the Analytics and Data and Science field. Before joining Turner, he held various research, consulting, management, and leadership positions, working on problems in the airlines, hospitality, and online retail sectors.