

INFORMS Journal on Applied Analytics

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

Turner Blazes a Trail for Audience Targeting on Television with Operations Research and Advanced Analytics

José Antonio Carbajal, Peter Williams, Andreea Popescu, Wes Chaar

To cite this article:

José Antonio Carbajal, Peter Williams, Andreea Popescu, Wes Chaar (2019) Turner Blazes a Trail for Audience Targeting on Television with Operations Research and Advanced Analytics. INFORMS Journal on Applied Analytics 49(1):64-89. <https://doi.org/10.1287/inte.2018.0971>

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 © 2019, 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>



THE FRANZ EDELMAN AWARD
Achievement in Operations Research

Turner Blazes a Trail for Audience Targeting on Television with Operations Research and Advanced Analytics

José Antonio Carbajal,^a Peter Williams,^a Andreea Popescu,^a Wes Chaar^b

^a Turner Broadcasting System, Inc., Atlanta, Georgia 30318; ^b Catalina Marketing Corporation, Inc., St. Petersburg, Florida 33716

Contact: antonio.carbajal@turner.com,  <http://orcid.org/0000-0003-4025-1755> (JAC); peter.williams@turner.com (PW);

andreea.popescu@turner.com,  <http://orcid.org/0000-0003-4027-0462> (AP); wes.chaar@catalina.com,

 <http://orcid.org/0000-0002-2743-8956> (WC)

Accepted: August 24, 2018

<https://doi.org/10.1287/inte.2018.0971>

Copyright: © 2019 INFORMS

Abstract. The novel concept of audience targeting on television poses business and technical challenges that involve disrupting decades-old paradigms about transacting and executing television advertisement deals. Turner Broadcasting System, Inc., has leveraged operations research and advanced analytics to take the lead in designing and implementing innovative and integrated forecasting and optimization models that forecast (granular) targeted and (traditional) demographic audiences in the 24/7 programming schedule, generate media deals across all of Turner’s networks, and holistically schedule commercials, balancing the objectives of all of the different types of advertising spots. These scalable and data-source-agnostic methods power Turner’s audience-targeting solutions: TargetingNOW and AudienceNOW. To date, Turner has completed more than 175 targeted media deals and is on track to sell 50% of its inventory through audience targeting by 2020, representing billions in ad revenue for the company. Every TargetingNOW deal has delivered a lift in target audience for advertisers, with a 27% average lift. AudienceNOW has delivered a decrease of at least 20% in target cost per impression for advertisers.

Keywords: audience targeting • TV advertising • forecasting • optimization • multilevel regression • mixed-integer programming

Introduction

Turner’s corporate history is shaped by no one more than Ted Turner, its brash, enigmatic, and innovative founder. Ted Turner started out in the media business after his father’s death, when he took over his father’s billboard business, which was worth about \$1 million in 1961. Turner quickly grew this advertising business and expanded his media empire by purchasing a local ultra-high-frequency (UHF) broadcasting station and the Atlanta Braves in the 1970s. Thus, began the Turner Broadcasting System, with its flagship network branded SuperStation WTBS.

Turner quickly grew the distribution of WTBS by offering Atlanta Braves baseball games, and by helping to revive interest in professional wrestling by buying and televising World Championship Wrestling (WCW). Turner managed to expand WTBS, the local station that later became the TBS Superstation, by offering the network as a satellite transmission to cable operators, enabling a national cable presence. Turner’s achievements also include creating the first 24/7 cable TV news network, CNN, which revolutionized news media through coverage of the Space Shuttle Challenger disaster in 1986 and the Persian Gulf War in 1991.

Turner has become a global entertainment, sports, and news company that creates premium content and delivers exceptional experiences to fans whenever and wherever they consume content. These efforts are fueled by data-driven insights and industry-leading technology. Turner owns and operates some of the most valuable brands in the world, including Adult Swim, Bleacher Report, Boomerang, Cartoon Network, CNN, ELEAGUE, FilmStruck, Great Big Story, HLN, iStreamPlanet, Super Deluxe, TBS, Turner Classic Movies (TCM), TNT, truTV, and Turner Sports. Turner’s domestic business generated \$11.3 billion in revenue in 2016, of which approximately half is from advertising revenue. Although Turner’s portfolio spans a multitude of platforms and includes several digital and mobile properties, this paper addresses only the television side of the business.

Most broadcasting and cable television networks in the United States sell advertisements and deliver *impressions* (audience exposures) to advertisers. For decades, television advertisement deals were guaranteed on a number of impressions with demographic characteristics specified only by gender and an age bracket, such as 25- to 54-year-old people of either gender

(P25-54). In the past few years, data fusion has allowed viewership data to be fused with frequent-shopper card data, credit card data, or even custom survey information to construct targeted audience segments such as “cereal buyers” or “automobile intenders.” These new, more granular audience segments have challenged the traditional ways of forecasting audience impressions, creating media deals, and scheduling commercials within the media industry. Turner has leveraged operations research and advanced analytics to take the lead in establishing the rules of engagement in this new business paradigm. To understand the relevance of this breakthrough, we will briefly describe the evolution of advertising in television.

Brief Chronology of Advertising in Television

It is hard to imagine a time when television did not air advertisements. Once television was perceived as an instrument of massive reach, it became the preferred medium for advertisers to use to send their branding messages. The first TV ad was aired on July 1, 1941, at 2:29 p.m., during a baseball game on a local New York channel, WNBT-TV (now WNBC). It advertised Bulova watches during a 10-second commercial, cost \$9, and marked the beginning of commercialized television (Colman 2013) (Figure 1).

After the second World War, advances in technology made television a desirable entertainment method. By 1952, nearly one-third of all households in the United States owned a television set, and by 1955, half of all U.S. homes had one (Stephens 2018, Quality Logo Products 2018). With the explosion in TV popularity, advertisements seemed a natural way for the broadcasters to monetize on the new medium.

In the 1950s, sponsorships (e.g., “this program is brought to you by...”) were the main form of advertisement; an advertiser would sponsor an entire program that directly or indirectly showcased its brands. However, both broadcasters and advertisers soon realized the drawbacks of this advertising option: on

the one hand, advertisers wanted more flexibility than having to produce an entire show to get their message across; on the other hand, broadcasters wanted to exploit opportunities for additional revenue streams by having multiple advertisers in one program.

In the 1960s, NBC executive Sylvester “Pat” Weaver devised a novel solution: NBC would produce its own programs and sell brief time slots airing commercials, also known as advertisement spots or simply spots, within program breaks to multiple advertisers. It was a win-win solution, much more cost-effective for advertisers and more efficient for monetizing airtime at scale for broadcasters. Scale brought up a new need in the industry, the need to develop a trade currency. Advertisers were seeking a count of viewers exposed to their ads, and broadcasters were trying to produce content that attracted as large an audience as possible; however, no measurement was in place to gauge the popularity of a show.

Nielsen Media Research was founded by Arthur Nielsen in the 1920s to analyze brand advertising. It expanded into radio market analysis during the 1930s, culminating in Nielsen ratings of radio programming, which became the authority in radio-audience measurement by the 1950s. When the need for television-audience measurement arose, Nielsen developed a ratings system using the methods developed for radio. The rating system was based on a sampling of more than 1,000 television homes scattered around the country. Each household in the sample had a small box, called an audimeter, which was attached to the television set and recorded when the set was on and the channel to which it was tuned. The data from the audimeters and diaries in sample households were centralized in a computer center. Using all of this information, Nielsen raters projected a total audience for each program, as well as the age and gender of the viewers (Britannica 2018). A Nielsen rating point is defined as the percentage of all households owning a TV set that is tuned into a program; for example, a Nielsen rating of 10 for a specific program denotes that 10% of the total U.S. households owning a TV set are tuned into that program.

Also in the 1960s, ABC executive Oliver Treyz was looking for ways to compete with far-more-established rivals CBS and NBC. He came up with the idea of selling airtime for advertisements, which would be measured in exposures to age segments of the population. He realized that it would be attractive for advertisers to home in on more suitable buyers according to the brands advertised—for example, men between the ages of 18 and 24 would be more likely to respond to a Corvette ad than their older counterparts, who might be more inclined to respond to Buick and Cadillac ads. So, from a sales strategy, the Nielsen demographic ratings, and later impressions (i.e., individual audience exposures), became the standard measurement used in television advertising transactions.

Figure 1. The First Televised Commercial Was a Bulova Watch Advertisement



More than half a century later, broad demographic-based metrics, defined solely on the basis of gender and age brackets, remained the currency of trade between advertisers and media broadcasting companies. In the age of great technological advances, such as the internet, smartphones, streaming services, smart TVs, and self-driving cars, little progress had been achieved in audience measurement for television advertisement deals, and the numbers of people in the P18-49 or P25-54 demographics exposed to the advertisements were typically the metrics chosen to evaluate deal performance. To quote Dan Aversano, the senior vice president of ad innovation and programmatic solutions at Turner: “I’ve tried explaining this to people outside of our industry—that the fate of billions of dollars rests on the notion of age- and gender-based metrics. For example: What do Ted Cruz and Justin Timberlake have in common? They are both men, aged 25–54. Time and time again, these ad land outsiders are perplexed by why they are treated as one and the same” (Aversano 2017).

Turner Breaks the Mold

Turner has disrupted the status quo, decades-old paradigms of transacting and executing on demographic metrics only, by introducing to the market advertising products that capitalize on more granular audience segments constructed by fusing viewership data with frequent-shopper card data, credit card data, or even custom survey information. Through this data fusion, the resulting audience segments, such as “yogurt buyers” or “exotic-vacation seekers,” provide advertisers with the capability of concentrating their marketing efforts, without the restraints of gender or age brackets, on the audiences identified as most likely to consume their products or services. This capability is referred to as *audience targeting*, and the granular audience segments are called *targeted audiences* or simply *targets*. Turner is blazing a trail with these breakthroughs in television campaign measurement and execution, moving the industry forward, and changing its overall perception as a data-driven organization.

Turner has developed two main television audience-targeting solutions: TargetingNOW and AudienceNOW. TargetingNOW enables an advertiser to take an existing media deal, guaranteed on a primary demographic, and optimize its spot placements to increase the delivery of a secondary targeted segment. The deal is still guaranteed on demographic impressions and maintains its original media mix (i.e., the proportion of impressions on the different networks and shows); in addition, an agreed-on benchmark reflecting the expected targeted impressions that the deal would receive without optimization is used to measure the lift (i.e., additional delivery) in targeted-segment viewership. AudienceNOW revolutionizes audience targeting by relaxing many of the traditional mix constraints to

produce fully optimized deals and spot placements that maximize targeted audience delivery across the entire portfolio of Turner’s networks (TBS, TNT, CNN, HLN, Cartoon Network, Boomerang, NBA TV, Adult Swim, and truTV). AudienceNOW deals are guaranteed and priced based on targeted segments, a major change from the traditional approach of pricing and guaranteeing deals based solely on demographics. For both solutions, Turner enables clients to choose virtually any data set to define their targeted segment. The data can originate from syndicated marketing research suppliers or can be custom fused or matched to television data (from panel or set-top-box sources).

The main methodological contributions of this paper are a series of integrated forecasting and optimization methods that Turner designed and implemented to allow it to operationalize its advanced targeting products. On the forecasting side, a scalable, accurate, and data-source agnostic forecasting method called competitive audience estimation (CAE) was created to forecast targeted and demographic audiences across Turner’s properties. CAE is inspired by consumer-choice modeling and uses a generalized linear mixed-effects model that builds audience estimates for television based on several factors, including time-dependent attributes, program attributes, and competitors’ program attributes for shows being aired on other networks. CAE can build estimates for virtually any audience segment based on any data set. These estimates are computed at the most granular level in the industry today, 30-minute blocks by day. For the short-term forecasting needed in operational decision making, an ensemble framework was developed to combine CAE estimates with smoothing methods that use additional data that become available close to airing. On the optimization side, large-scale mathematical programming models were created to optimize deal proposals and spot placements. A mixed-integer programming (MIP) sales-proposal builder using CAE estimates was designed to generate proposals and price them based on remaining available commercial inventory (i.e., commercial airtime measured in 30-second slots) across all of Turner’s networks, honoring agreed-on restrictions and parameters. A multistage, MIP spot scheduler using short-term audience forecasts was also designed to optimize the placement of tens of thousands of commercial spots across all domestic networks, including TargetingNOW spots, AudienceNOW spots, traditional demographic-guaranteed spots, and any other type of spot that needs to be scheduled, subject to several constraints.

Television Ad Sales: Markets, Deals, and Revenue Management

Commercial airtime is typically sold in 30-second slots through media deals (also referred to as plans or schedules), which start as proposals that are discussed,

negotiated, and revised between networks and advertisers, who are usually represented by advertiser agencies. A *selling title* is an interval of programming that networks use to sell their commercial airtime, and it can refer to a specific program—for example, *The Detour* on TBS or *Animal Kingdom* on TNT—or it can refer to a block of time, such as *Comedy Block 1*—Monday through Saturday, 3:00 p.m. to 5:30 p.m. on TBS or *Latenight Dramas*—Monday through Thursday, midnight to 3:00 a.m. on TNT. The term *daypart* is typically used to describe an aggregation of selling titles. Each network defines its own dayparts for its internal inventory management, but media rating agencies such as Nielsen provide guidelines on generic, time-based dayparts that can be used to compare media plans across networks.

Television networks operate under the so-called broadcast year, which typically spans the fourth quarter of a current calendar year through the third quarter of the following year. The TV ad sales market starts around May or June when networks announce their programming for the upcoming broadcast year. These announcements are followed by an intensive sales period, called the upfront market, in which networks sell 60%–80% of their commercial airtime. The remaining capacity not sold during the upfront market is sold in the scatter market through the remainder of the broadcast year.

In addition to upfront and scatter markets whose deals provide audience guarantees to advertisers, a third type of market, known as filler sales, also occurs throughout the broadcast year. Filler sales refer to the selling of distressed inventory close to airing, and filler deals usually specify a maximum number of commercials to be aired within a particular selling title in specific weeks, but provide no audience guarantees.

The Components of an Audience-Guaranteed Media Deal

Figure 2 shows a simplified upfront deal that illustrates some elements common to all upfront and scatter deals:

- The deal’s budget is the total advertisement expenditure in the deal.
- The deal’s flight is an industry term used to describe the range of dates in the deal.
- The ratecard type determines the components that constitute a valid audience for the deal. Examples of ratecard types include live (average number of people watching a particular show when it airs, including both programming and commercial airtime), ACM + 3 (average audience watching only commercial airtime either live or within three days of airing via digital video recording devices), and ACM + 7 (the same as ACM + 3 but with seven days of delayed viewing).
- The primary demographic, specified by gender and an age bracket, reflects the audience that an advertiser is trying to reach and is typically correlated with the consumers of interest—that is, the population segment most likely to purchase its products or services.
- The gross impressions are the guaranteed audience exposures, in thousands, within the primary demographic that the deal is committed to deliver.
- The cost per thousand impressions (CPM) is equal to the budget divided by the total number of guaranteed impressions, in thousands. Media plans are usually evaluated in terms of their CPMs.
- The brand of the deal specifies the brand or group of brands to be advertised in the deal. A various brand indicates that a multibrand advertiser, such as Unilever or General Motors, will later specify the distribution of its spots across its different brands.
- The conflict indicates the industry category of the advertiser (e.g., automotive, financial and business

Figure 2. (Color online) Audience-Guaranteed Media Deal Between a Network and an Advertiser

DEAL #: D201332P3	BUDGET: \$2,000,000
ADVERTISER: XYZ	CPM: \$16.61
BRAND: VARIOUS	GROSS IMPRESSIONS (000): 120,409
CONFLICT: AUTOMOTIVE	PRIMARY DEMOGRAPHIC: P25-54
FLIGHT DATES: 09/25/2017 - 09/30/2018	RATECARD TYPE: ACM+3

NETWORK	SELLING TITLE	2017 Q4	2018 Q1	2018 Q2	2018 Q3	TOTAL
	Conan				6	6
	Full Frontal with SB		5	6	5	16
	Search Party	4	4			8
	...					
	People of Earth	1	1	1	1	4

services, toiletries, and cosmetics) and is used to enforce competition-avoidance requirements such that two spots sharing the same conflict do not air in the same commercial break.

- The number of spots that will air on different selling titles across the deal's flight is usually expressed in equalized 30-second units (EQ30s). When spots are equalized, spots of a length other than 30 seconds have their impressions and rates adjusted upward or downward in proportion to the ratio of their length to 30 seconds. EQ30s are often referred to as units.

A broadcast quarter typically spans 13 weeks. Although Figure 2 presents an aggregation of units at the selling title-quarter level, media deals are usually specified more granularly at the selling title-week level because advertisers typically have requests about the distribution of their spots, impressions, or budget within a quarter; for example, an advertiser might specify that 30% of its budget should be consumed during the first 3 weeks of the flight, and the remaining 70% should be distributed equitably across the other 10 weeks in the quarter. The units are also expected to follow a certain mix across selling titles; for example, an advertiser might specify that its deals should contain at least 30% of their units or impressions from prime selling titles, or at most 20% from overnight selling titles. Finally, although not shown in Figure 2, units have different rates (i.e., prices per EQ30) for deals with different primary demographics. A unit in the same selling title commands higher rates for deals with primary demographics associated with smaller audiences.

Figure 3 illustrates the life cycle of an upfront or scatter TV media deal, which comprises three main stages: planning, execution, and posting. During the planning stage, media deals are created and specify a commitment to air a number of commercial spots at

an aggregate level, typically selling title week. This is performed months in advance of the date when the spots will air. During the execution stage, the specific placements of spots are determined—that is, the exact moments in time when the spots will air. This is determined shortly before airing, at most a few days in advance. Furthermore, another part of the execution stage performed throughout the deal's flight is deal stewardship, which involves monitoring the deal's performance to evaluate whether it is on track to deliver its guaranteed audience, and if not, allocating additional units called audience-deficiency units to decrease impression shortfall. Finally, during the posting stage, the actual audiences delivered with the plan are reported to the advertiser. This occurs after all spots have aired (although partial quarterly reports and invoicing are customary for multi-quarter deals), based on actual audiences that are reported by media rating agencies. Subsequently, the cycle begins anew and the posted results from past deals are used to inform the creation and negotiation of new deal proposals.

Related Literature and Positioning of This Work

Television networks face a challenging revenue management process that involves several interconnected planning and control problems, but is usually addressed in practice in a hierarchical fashion that involves strategic, tactical, and operational decisions. Carbajal and Char (2017) provide an overview of this process. Table 1 summarizes the process components. Our work relates mostly to audience forecasting, proposal creation on both upfront and scatter sales, and spot scheduling. All previous works have addressed forecasting and planning as disparate problems as opposed to using our integrated approach. Our work also differs from previous literature in the following ways. (1) All published

Figure 3. (Color online) The Life Cycle of an Audience-Guaranteed Deal Comprises Planning, Execution, and Posting

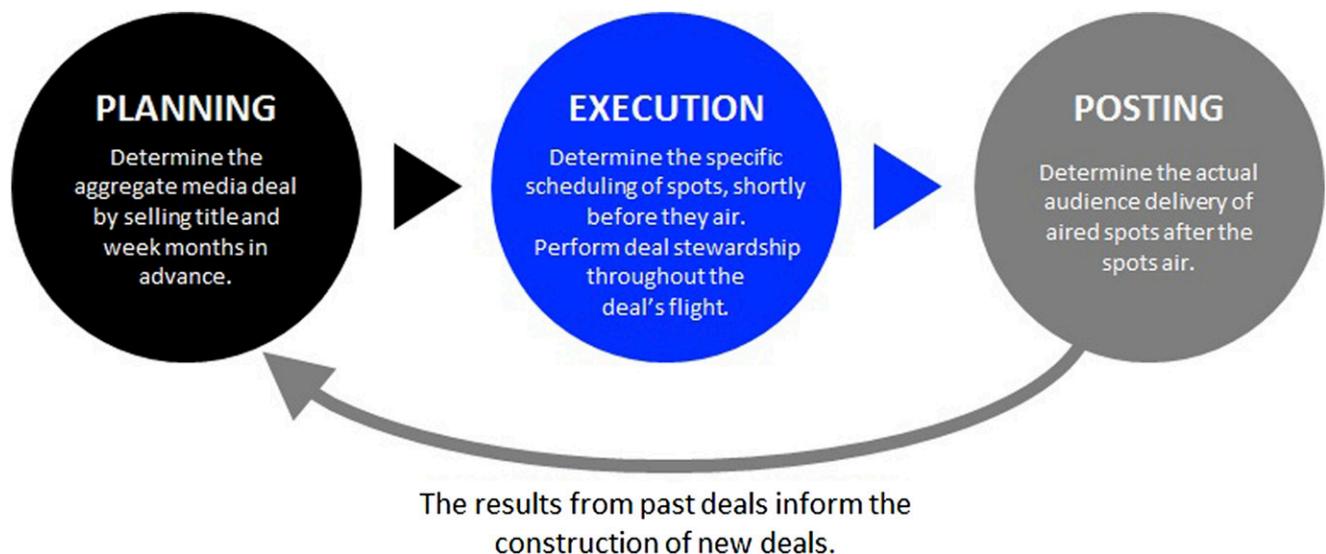


Table 1. Components of the Media Revenue Management Process and Relevant Literature

Hierarchy	Planning and control problems	Related literature
Strategic	Programming decisions: multiyear impact decisions, such as content acquisitions and original series development, program scheduling for the upcoming broadcast year	Horen (1980), Reddy et al. (1998), and Danaher and Mawhinney (2001)
	Yearly capacity management: allocation of commercial airtime between upfront and scatter markets	Bollapragada and Mallik (2007) and Araman and Popescu (2010)
	Upfront sales: creation of proposals (unit allocation and pricing) during the upfront market	Bollapragada et al. (2002) and Popescu and Seshadri (2016)
Tactical	Inventory-mix distribution: allocation of nonprogramming capacity into scatter sales, promotional (promo) campaigns, and audience-deficiency units	—
	Scatter sales: creation of proposals (unit allocation and pricing) during the scatter market	Bollapragada et al. (2002) and Popescu and Seshadri (2016)
	Promotional planning: allocation of units across promotional campaigns in the same or other sister networks (i.e., networks owned and operated by the same company), and purchasing of off-channel units on outside networks	Pereira et al. (2008)
	Allocation of audience-deficiency units: allocation of zero-rate units across deals to decrease their expected impression shortfall	Carbajal and Chaar (2017)
Operational	Brand allocation: distribution of units across brands for multibrand advertisers	—
	Commercial spot scheduling: determination of the placements of spots that advertise products, services, or events	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 (2015), and Giallombardo et al. (2016)
	Promo scheduling: determination of the placements of spots that promote network content	—
All levels	Filler sales: management of demand of filler spots and their pricing	—
	Audience forecasting: estimation of viewership of different audience segments across all programming content at different granularities	Gensch and Shaman (1980), Meyer and Hyndman (2006), Danaher et al. (2011), and Danaher and Dagger (2012).

works on forecasting have been theoretical in nature and lack practical validation (i.e., no implementation has been reported), address only demographic audiences, have a limited scope (i.e., they typically focus on programs airing between 6 p.m. and 11 p.m.), and do not address short-term, granular estimates, which are needed in operational decision making. We report on implemented forecasting models that address both demographic and (granular) targeted audiences, cover the entire 24/7 programming schedule, and generate granular and (or) aggregated estimates depending on the level of decision making. (2) Previous works on proposal creation report on models that handle single-network, demographic-based proposals with no pricing component. Our proposal-creation models allocate units and set price rates across an entire portfolio of networks, considering demographic and targeted audiences, and can also accommodate metrics other than gross impressions (total impressions regardless of duplication), such as reach (unduplicated audiences), composition (ratio of target impressions to demographic impressions), and response (consumer actions attributed to aired spots; for example, the numbers of visits to a web page that

result from a TV spot). (3) Previous works on spot scheduling have addressed mostly feasibility problems focusing particularly on product conflict and (or) spot-separation restrictions and ignoring the impact of audiences on the schedule. Our spot-scheduling models consider estimated impressions for both demographics and targets, honor all constraints that appear in practice, and balance the goals of all of the different types of spots that include demographic-guaranteed, targeted (both TargetingNOW and AudienceNOW), audience-deficiency, and filler spots.

Audience-Targeting Solutions: TargetingNOW and AudienceNOW

Nielsen Media Research is the primary data provider of television viewership behavior in the United States. Nielsen maintains a balanced, cross-sectional, and representative sample of over 40,000 households in the United States, and uses electronic meters installed in these households to make projections about what is viewed, and when and by whom it is viewed, based on actual, individual tuning behavior (Nielsen 2017). The audience impressions reported by Nielsen serve as

the primary currency of most television media transactions.

Collecting such granular, representative data are very expensive, and the underlying household panel requires ongoing balancing and maintenance. Directly enriching this individual data with information such as purchasing behaviors is prohibitive because of costs, complexity, and privacy concerns. Nevertheless, these challenges have been successfully addressed in recent years through data fusion, which enables the integration of viewership data with consumer purchasing data from other sources, such as frequent-shopper card data, credit card purchasing records, and surveys conducted by marketers (Nielsen 2009).

Data fusion matches disparate individuals from two distinct data sources that share the same target population, such as U.S. consumers. It uses characteristics that are common to both data sources as a basis to impute the values of variables measured in the second source onto the first source (D’Orazio et al. 2018). For example, the first data source may be a television viewership data set, such as Nielsen, which contains information about the TV viewing habits of a panel, whereas the second data source may be frequent-shopper card data that contain purchasing information of individuals who have signed up for a frequent-shopper discount card. The specific individuals in each data set are different; therefore, a deterministic matching (i.e., identifying the exact individual with a frequent-shopper card who corresponds to an individual from the viewing panel) is not possible. Nevertheless, the individuals in both data sources have characteristics that both data sources track, such as age, gender, and socioeconomic status. These common characteristics are

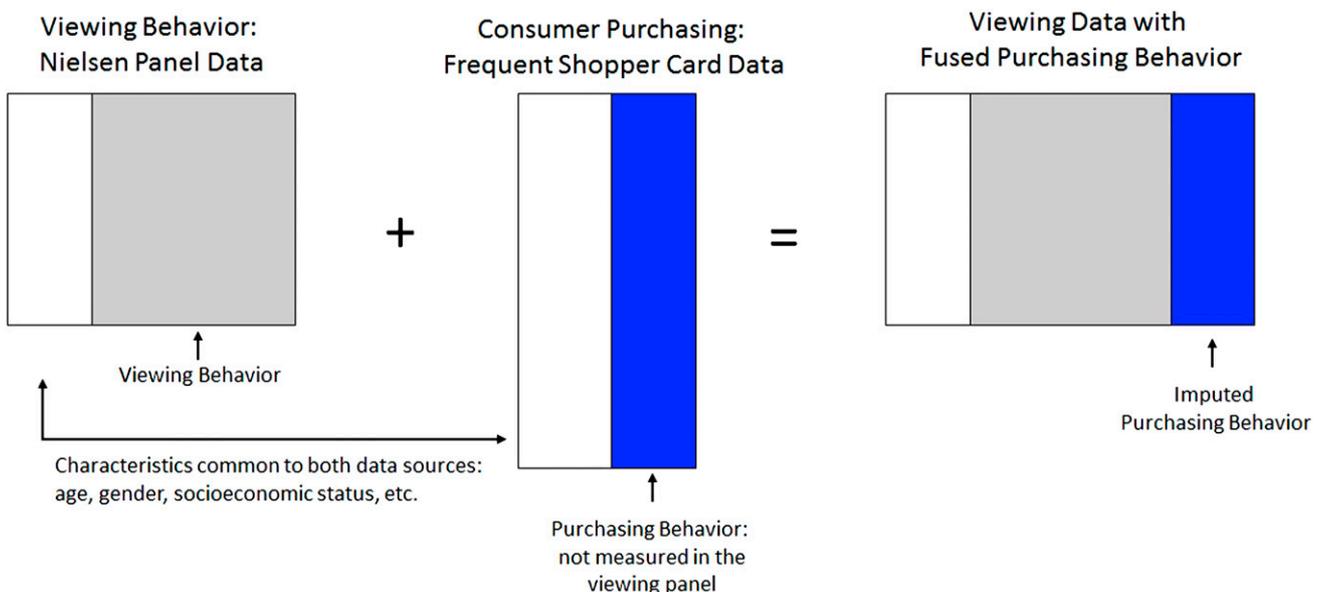
used as a basis to perform a probabilistic matching (i.e., identifying the individual with a frequent-shopper card who is most likely to match an individual from the viewership panel). After this process is performed, the purchasing behavior of each matched frequent-card shopper is imputed onto its matched individual from the viewership panel. Figure 4 illustrates this process.

Fusing viewership data with consumer purchasing data has been a catalyst for the creation of audience segments, which enable advertisers to transition from coarser, demographic-based advertising transactions toward more relevant (granular) targeted audiences. Turner’s TV audience-targeting solutions rely extensively on this capability to leverage granular targeted audience data and improve the planning and execution of media plans. The makeup of targeted audiences varies across networks, across selling titles of a network, and across airings of a selling title. Currently, targeted audience data can be produced at a 30-minute granularity, and audience differences can be exploited to create more efficient and effective media plans. Figure 5 illustrates how two target audiences—“heavy purchasing of diapers” and “typically have wine with dinner”—vary within the selling title *Daytime* on TBS. Using this type of information, proposals that yield lower target CPMs (i.e., cost per thousand impressions in the target market) can be built, and spots can be scheduled in time slots with higher concentrations of target impressions.

TargetingNOW

TargetingNOW affects the execution stage, but not the planning stage, because it involves previously agreed-on demographic deals. TargetingNOW is designed as

Figure 4. (Color online) Process of Fusing Consumer Purchasing Behavior onto Viewing Behavior



Note. This figure was inspired by an image included in van der Putten et al. (2002).

Figure 5. (Color online) The Average Impressions (in Thousands) Vary Across Targeted Segments and Times of the Day

Broadcast Quarter: 2016 Q4						Network: TBS					Selling Title: Daytime				
Target: Heavy Purchasing of Diapers						Target: Typically Have Wine with Dinner									
Time of Day	Monday	Tuesday	Wednesday	Thursday	Friday	Monday	Tuesday	Wednesday	Thursday	Friday					
6a	5	4	4	4	4	80	61	63	65	55					
630a	7	5	5	5	5	91	80	77	88	71					
7a	8	6	7	6	7	93	80	80	88	73					
730a	8	7	8	6	8	89	77	78	89	84					
8a	9	8	9	9	10	85	90	90	88	78					
830a	11	9	11	10	11	80	78	83	94	79					
9a	12	10	10	11	12	71	66	63	92	75					
930a	13	11	11	12	11	71	66	66	95	79					
10a	12	11	11	11	12	67	79	75	95	68					
1030a	13	10	11	11	12	68	79	84	90	72					
11a	10	7	7	6	10	65	46	55	55	69					
1130a	10	8	7	6	10	65	49	52	58	68					
12p	8	9	8	8	10	72	60	58	85	82					
1230p	8	10	10	9	10	82	76	71	96	76					
1p	9	13	10	11	13	91	81	91	95	113					
130p	8	13	11	11	13	93	90	91	92	135					
2p	9	13	12	11	13	100	103	96	103	136					
230p	9	12	11	12	13	93	103	97	103	142					

Notes. Impressions represent an estimate of the in-target average viewers during the specified times. Dark gray cells represent the top 10 percentile of in-target average audience, while light gray cells represent the bottom 10 percentile.

a first step that advertisers can take to experiment with targeted audiences because it allows them to maintain their traditional media mix and guaranteed demographic impressions. The value proposition of TargetingNOW consists in taking an existing deal and, in addition to meeting its guaranteed demographic audience, optimizing its spot placements (during the execution phase) to increase impressions in a targeted audience segment (i.e., achieve a delivery lift). The delivery lift is measured against a baseline that reflects the expected impressions the deal would receive without optimization. Different options, such as an average schedule or a median schedule, can be used to estimate a baseline delivery. An average schedule assumes all units in a selling title week receive the average delivery, whereas a median schedule assumes all units in a selling title week receive the median delivery. The market has leaned toward the median schedule because it is more robust to outliers in audience delivery.

Figure 6 illustrates the TargetingNOW process. The inputs to TargetingNOW are (1) a series of demographic-guaranteed deals, each with its own primary demographic, targeted segment, and target-delivery lift goal; (2) short-term audience estimates for both demographics and targeted segments; and (3) specific commercial air-time inventory buckets in which spots can be scheduled. During the spot-scheduling process, a baseline delivery is calculated and an attainable lift goal is determined for each deal. The attainable lift goal is less than or equal to the original deal lift goal depending on the audience

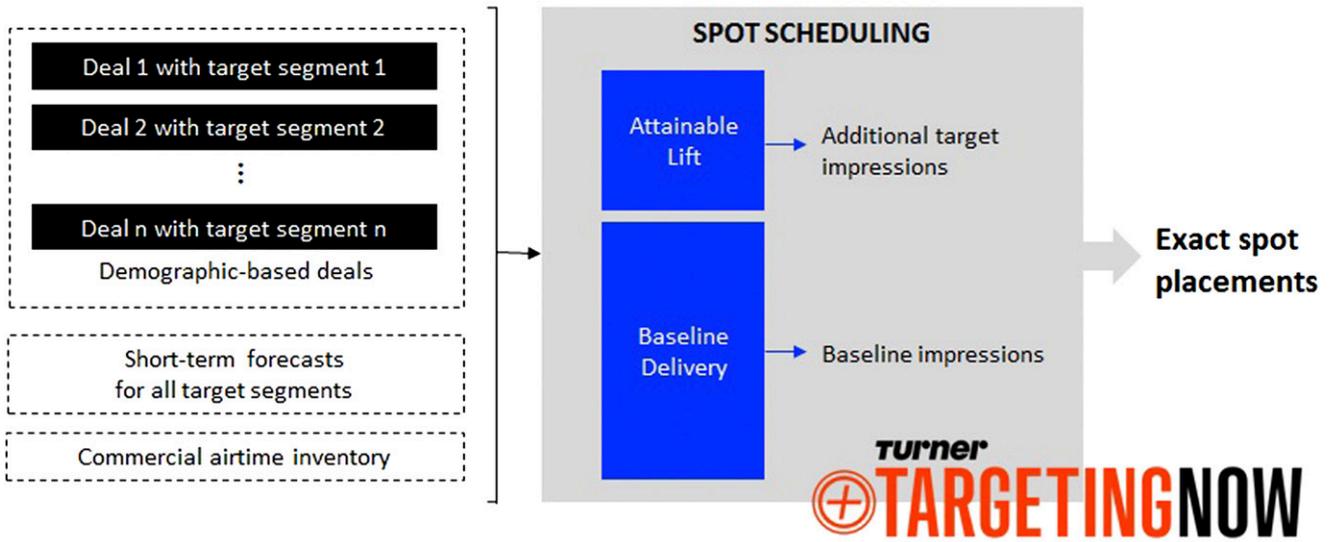
variability across eligible time-slot assignments. For example, although the lift goal of a deal may be 15%, the audience variability of time slots may allow only a 5% lift.

AudienceNOW

AudienceNOW affects both the planning and the execution stages of a media deal. AudienceNOW takes audience targeting further by removing several of the traditional media-mix limitations. The value proposition of AudienceNOW is that the advertiser can choose any data and audience, and Turner guarantees the targeted delivery across all of its networks. This allows advertisers to use virtually any audience-defining data set and target and negotiate deals directly in terms of targeted segments and target CPMs.

Figure 7 depicts the overall AudienceNOW process. During the planning stage, proposals are created to determine the allocation and pricing of units across the entire portfolio of networks. The goal of AudienceNOW at this stage is to create a proposal that is more efficient for an advertiser in terms of the targeted audience segment that it wants to reach. As such, an AudienceNOW proposal is built so that it produces a higher number of targeted impressions, which reduces the target CPM. Although demographic impressions are not guaranteed under the AudienceNOW paradigm, an advertiser uses historical proposals as benchmarks to evaluate new deals; as such, the advertiser is willing to receive a proposal that increases its demographic CPM, but only up to an agreed-on

Figure 6. (Color online) The TargetingNOW Process Schedules Commercial Spots to Produce a Lift in Targeted Audience Delivery



percentage. The inputs during this stage include (1) advertiser inputs (delivery metric to optimize, budget, overall daypart mix, network and selling title exclusions), sales inputs (target CPM reduction goals, demographic CPM increase caps, rate-increase caps), and strategic planning inputs (packaging guidelines, minimum rates to be charged for 30-second spots on the different selling titles, inventory of available commercial airtime); (2) long-term audience estimates; and (3) aggregate capacities for commercial airtime. The outputs of this stage are the number of units at the network selling title-week level and their corresponding rates.

Later, during the execution stage, the specific placements of AudienceNOW spots must be determined. The inputs for this stage include (1) the aggregate units from the AudienceNOW proposals, (2) short-term audience estimates, and (3) specific commercial airtime inventory buckets. As we can see, the inputs to the spot-scheduling stage are the same as the ones we show above for the TargetingNOW spot-scheduling stage. The spot-scheduling process for all advertising products, such as traditional demographic-guaranteed spots, TargetingNOW spots, and AudienceNOW spots, is performed concurrently because all spots are competing for the same commercial airtime.

Figure 7. (Color online) AudienceNOW Allocates Aggregate Commercial Airtime to Create Media Proposals in the Planning Stage and Follows with the Scheduling of Individual Commercial Spots in the Execution Stage

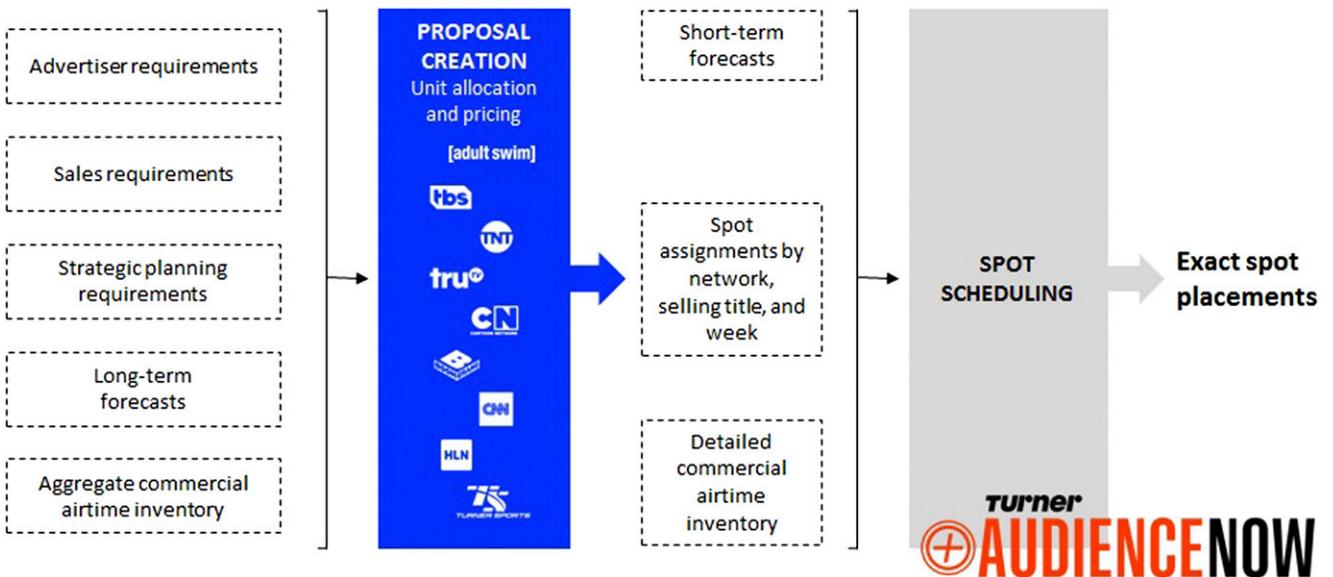


Figure 8. (Color online) Using Inputs, Including Estimates of Demographic and Targeted Audience Impressions and Floor Rates, AudienceNOW Generates Deal Proposals for Aggregate Units of Airtime by Selling Title and Week

NETWORK	SELLING TITLE	W1	W2	...	W13	TOTAL
	Conan					
	Full Frontal with SB					
	Search Party					
	...					
	People of Earth					
	The Last Ship					
	Good Behavior					
	Animal Kingdom					
	...					
	Claws					
	Rick and Morty					
	Bob's Burgers					
	Family Guy					
	...					
	Robot Chicken					



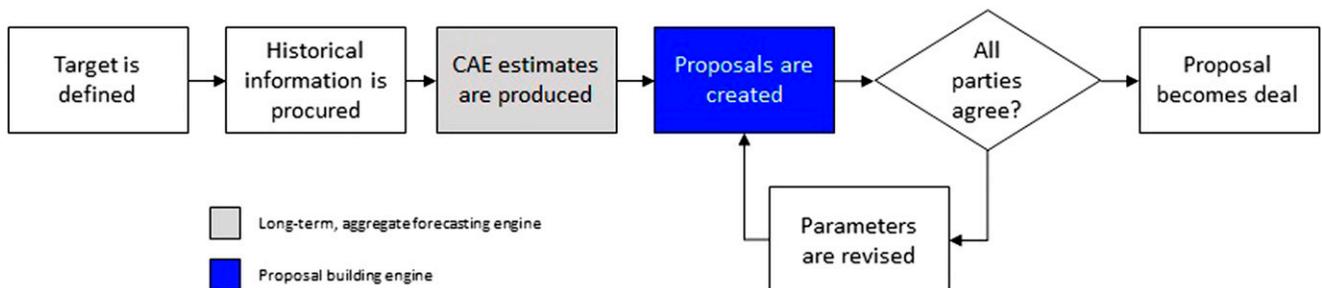
The Planning Stage: Long-Term Forecasting and Proposal Creation

During the planning stage, proposals are created, negotiated, and revised between Turner and the advertisers. The main problem addressed during this stage is to determine for a deal the number of units to allocate to each selling title-week combination during the deal's flight across all Turner networks and the prices that should be charged for each unit. Figure 8 shows an excerpt of a grid associated with the proposal-creation problem for a single quarter (although the multiple-quarter extension is straightforward). Each cell in this grid is associated with a selling title week for which estimates on demographic audiences, targeted audiences, and *floor rates* are available. The floor rates are the prices that would be charged to the advertiser under the traditional demographic paradigm and constitute the lowest rates to charge in an AudienceNOW deal for a unit in that selling title week. Therefore, the proposal-creation process can be regarded as

filling out this grid with units such that the advertiser objective is maximized subject to several requirements.

Figure 9 depicts the overall flow of the proposal-creation process, which is a time-sensitive, iterative process that is performed up to several months in advance of the airing of spots. Once the target audience has been defined, historical information is gathered and CAE forecasts are created and aggregated at the selling title-week level. A proposal-building engine is used to create alternative proposals using different parameters, based on negotiations occurring across the advertiser and Turner's organizations responsible for sales and strategic planning. Additional runs of the engine are used to refine the proposal until an agreement is reached, at which point the proposal becomes a deal. Throughout its flight, the deal is broken into spot orders, which include additional information about the spots, such as spot-length distribution, eligible airing days and times, and specific commercials to be aired. The information in these orders is then transferred to a

Figure 9. (Color online) AudienceNOW Proposals Are Built Using an Iterative Process Until All Parties Agree



spot-scheduling system to determine the specific times when spots will air.

Competitive Audience Estimation

Program announcements for an upcoming broadcast year are made before the upfront market, and program scheduling decisions are often made long in advance of airing. The lengths and types of programs (e.g., sports, series, and movies) dictate how much commercial airtime inventory will be available. Therefore, to support the proposal-creation process, audience forecasts are needed over a long-term horizon, often three months to over a year in the future. Various business groups support the audience forecasting task by providing information about which television programs the television networks will schedule at which times. An effective long-term forecasting approach should accommodate a number of business needs that include the following.

- The model should include explicit characteristics of a television programming schedule and be adaptable to programming schedule changes. The modeling approach should also provide business users with insights on the potential impact to audience impression capacity given programming schedule changes.
- The modeling approach should be versatile enough to produce forecasts across different audience segments, and at different levels of granularity, depending on planning needs.
- The model should consider the impact of the programs available on other cable, premium, and broadcast networks, and estimate the effect of concurrent competing programming on Turner's audiences.

To address these various business needs, we developed a scalable, long-term forecasting approach that we called competitive audience estimation (CAE); the term “competitive” refers to the use of variables that measure historical competitive programming at the genre level across our networks and competing networks. CAE forecasts targeted and demographic audiences at the network, date, and half-hour level across all Turner networks. CAE models the expected audience in a half hour using a mixed-effects regression model; see Gelman and Hill (2006) for a discussion of mixed-effects models, which are also referred to as

multilevel models or hierarchical models. Some effects are estimated at the observation level (i.e., fixed effects) and others at the program level (i.e., random or grouping effects). Table 2 shows a high-level description of the factors considered in the model; these include program factors, time-dependent factors, and competing program factors.

Including program-level effects is necessary because, in practice, changes to future program schedules are a constant part of internal business processes; therefore, a good estimation method needs to quantify the impact of programming changes. Furthermore, the business users of these models typically have a good understanding of the relative audience sizes of different programs; therefore, estimating program-level effects helps with user acceptance because these effects can be used to explain audience variations across programs. Because cable networks extensively repeat airings of programs, the audience rating data can often support good audience estimates of programs when they also change time periods. Additional effects associated with programs, such as whether a program is a repeat or live airing, are also critical, because repeat episodes are usually viewed by fewer people, and live programming, such as sports, typically draws larger audiences.

CAE also considers the cyclicity of television viewing through time-dependent factors, such as time of day, daypart, and day of week. In addition, seasonality is a critical factor to consider because television viewing levels can be affected by the outdoor temperature, the weather, and whether children are in school. We account for this seasonality through trigonometric (harmonic or Fourier) regression terms, which model the daily periodicity of the data; Gensch and Shaman (1980) include an example.

CAE overcomes the lack of readily available information on competitors' programming schedules by considering genre telecast counts of competitors, lagged by one year, which are a good proxy for current competitor programming because a network's programming typically presents only small variations during two consecutive years. These factors consist of measurements by date and half hour, of how many programs that fit a particular genre criteria were aired, the top 25 cable networks (by average total audience), and premium networks, such as HBO and Showtime. To reduce the program genres considered, we include only relevant genres based on the Turner network of interest. For example, the CAE predictions for TNT, which primarily airs dramatic series and movies, consider the number of dramas that aired historically on other networks.

CAE provides granular audience forecasts at a uniform resolution (e.g., half hour), which allows for easy aggregation based on business needs. This allows the generation of forecasts at coarser levels (e.g., a network

Table 2. Audience Viewing Is Estimated Using Several Factors Including Program Attributes, Time-Dependent Attributes, and Competitor's Program Attributes

Effect category	Factors
Program	PUT (persons using television), genre, repeat, live program
Seasonal and time	Time of day, daypart, day of week, quarter in year, Fourier terms (daily periodicity)
Competitive	Historical genre telecast counts (broadcast, pay, and cable)

selling title-week level) to serve long-term sales and financial planning. It also enables Turner to conduct analysis at finer levels, such as computing the expected lift over a median schedule that a number of units slated to air in a selling title week can receive. Furthermore, because CAE can accommodate any historical average audience input, it is flexible and scalable across different demographic and target audiences. We include the model specification for CAE in the appendix.

Proposal Building

Long-term forecasts of audiences at the selling title-week level are entered into a proposal-building engine that uses a mathematical programming formulation in which the main decision variables are the allocations of units to allowable combinations of selling titles and weeks (in integer EQ30s) and the rates to charge for those units (in dollars). The objective is to maximize the overall metric of interest (e.g., gross impressions, composition, reach, response) delivered with the proposal, while honoring constraints from the advertiser and Turner's marketing and strategic planning organizations. These constraints include requirements from the advertiser and requirements from Turner's organizations in charge of strategic planning and sales.

Advertiser Requirements.

- The proposal dollar value cannot exceed the budget.
- The proposal budget distribution across Turner networks should be between some lower and (or) upper limits (e.g., at least 5% and at most 30% of the budget should be assigned to TBS).
- The proposal should not include specific networks and (or) selling titles that the advertiser wants to exclude.
- The distribution of impressions in the proposal across Nielsen dayparts should be between some lower and (or) upper limits (e.g., at most 10% of the target impressions should be in the overnight Nielsen daypart).
- The units in the proposal across weeks should follow a desired distribution. If no particular distribution is defined, a uniform distribution is typically expected.

Strategic Planning Requirements.

- The number of units in the proposal assigned to a specific selling title week cannot exceed the remaining capacity in terms of overall EQ30s and product-conflict availability. Furthermore, it cannot exceed the maximum number of units that can air without violating time-separation restrictions.
- The rates charged for the units in the proposal should be greater than or equal to their corresponding floor rates (i.e., baseline rates).

- The distribution of target impressions in the proposal across the dayparts and (or) selling titles within each network should be between some lower and (or) upper limits (e.g., at most 10% of the impressions delivered on TBS should be in the selling title *Full Frontal with Samantha Bee*).

Sales Requirements.

- The proposal target CPM (proposal dollar value divided by the total target impressions, in thousands) should be lower than a baseline target CPM by at least a specified percentage.
- The proposal demographic CPM (proposal dollar value divided by the total demographic impressions, in thousands) cannot be increased by more than a specific percentage over a baseline demographic CPM.
- The rates charged for the units in the proposal should not exceed a maximum percentage increase over the floor rates.
- The proposal should yield at least a minimum profit margin (i.e., the difference between the proposal value under the rates charged and the proposal value under the floor rates).

In general, the resulting formulation is a nonlinear programming problem (NLP) because both unit selection and unit pricing are determined simultaneously. The pricing decision can be thought of as determining how much to increase the floor rate of each unit in the proposal; for the case in which floor rates are constrained to be increased in the same proportion (i.e., all proposal unit rates equal their corresponding floor rate multiplied by a factor greater than one), the NLP can be reformulated into an equivalent MIP problem. The same-proportion floor-rate increase is the case used most frequently in practice; however, the more general case in which floor rates are allowed to increase at different proportions can be solved heuristically starting with the solution to the reformulated MIP.

The appendix shows a detailed mathematical formulation of the NLP and the reformulated MIP that is used to maximize the gross target impressions delivered in the proposal under the same-proportion floor-rate increase assumption. As we indicate above, some advertisers may have alternative objectives for their proposals, such as maximizing a response (i.e., actions taken by consumers attributed to aired spots), maximizing the proposal composition (i.e., target impressions divided by demographic impressions), or maximizing reach (i.e., unduplicated audience). The specific objective depends on what the advertiser considers to be the best match for its marketing strategy for the media deal. The metric response can be maximized using the same MIP with a different interpretation of the parameters. The metric composition can be maximized using similar logic, but with a more complex formulation that handles the

ratio in the objective function. Williams (2013) provides details on how to reformulate linear-fractional programming models, such as the one arising with the composition metric. The metric reach is currently optimized heuristically by building several proposals and running extensive simulations of spot placements to estimate their expected reach. However, analytical approaches are being evaluated to incorporate reach as a function that can be optimized directly in the mathematical programming framework.

The Execution Stage: Short-Term Forecasting and Spot Scheduling

During the execution stage, spots in booked orders from all agreed-on deals are scheduled to air at specific times in commercial breaks within selling title weeks. Examples of the different types of spots include demographic-guaranteed, TargetingNOW, AudienceNOW, audience-deficiency, and filler spots. The spot-scheduling process that considers audience estimates is performed very close to the airing time using short-horizon forecasts at the most granular level to differentiate the value of spot placements across the entire selling title-week landscape.

Figure 10 depicts the overall flow of the spot-scheduling process. CAE forecasts are combined with recent audiences, such as the audiences during the past few weeks, to refine 30-minute-level forecasts for the current week and the next week. A spot-scheduling engine determines the specific times when spots will air and communicates them to the traffic system, which is a system that generates the daily log of programming elements (e.g., programming content, commercials, and promos) and specifies when they are planned to be aired. Subsequently, a few manual adjustments might be performed in the traffic system to complete the commercial schedule that airs on TV.

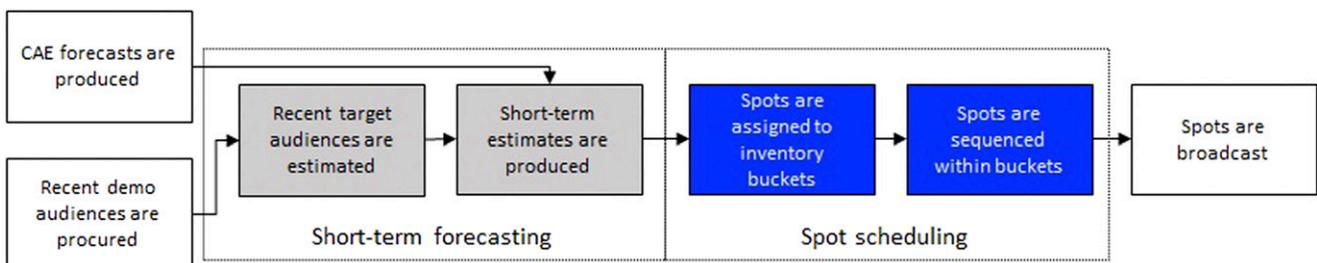
Ensemble Framework for Short-Term Forecasting

CAE estimates include time-based attributes, program attributes, and competing program attributes, but can be further refined with additional data that become available shortly before airing. We use an ensemble

approach that blends CAE forecasts with estimates produced with a single exponential smoothing model to create the short-term, granular forecasts used in the execution stage. The single exponential smoothing model (SES) is implemented at the network, day-of-week, and half-hour resolution on a time series of historical audience data. For example, if we are forecasting the audience of the program *Claws* for next Sunday between 9:00 p.m. and 9:30 p.m., we operate on a time series of historical audience data points (also called observations) on Sundays between 9:00 p.m. and 9:30 p.m. on TNT. Two important features of the SES model are that it incorporates characteristics of the specific show to forecast and it uses a capping mechanism to reduce the effects of atypically large or small audiences.

Programming can vary from week to week because of special events; therefore, when trying to forecast the audience of a specific network, day, and half hour, not all historical observations matching the same network, day, and half hour may be relevant. SES uses the historical observations that match the characteristics of the show being forecast as much as possible. A minimum number of observations is established and SES processes historical observations that match the same franchise (i.e., the specific show within a selling title); if the minimum number of observations is not met, it processes observations that match the same selling title, and if the minimum number of observations is not met again, it processes observations that match the same network, day, and half hour. For example, on May 12, 2017, when forecasting the average audience that would tune in on the following Thursday from 8:00 a.m. to 8:30 a.m. on TBS, SES identified that *Married with Children* was scheduled at that time. It analyzed historical audiences on TBS on Thursdays between 8:00 a.m. and 8:30 a.m.; however, it did not process the observation from Thursday May 4, 2017, because, to commemorate *Star Wars* day, TBS had aired *Star Wars* movies on that day instead of its regular programming. If enough observations had not been found for *Married with Children*, SES would have looked for observations matching the selling title *Daytime*, and if

Figure 10. (Color online) Scheduling Commercials Involves Producing Short-Term Audience Forecasts, Assigning Spots to Commercial Airtime Inventory Buckets, and Sequencing Spots Within the Assigned Buckets



enough observations had not been found again, it would have processed all available observations on TBS on Thursdays between 8:00 a.m. and 8:30 a.m.

Even when the same program is scheduled on two consecutive weeks, the audience for one airing might be unusually large (e.g., because of breaking news on CNN) or small (e.g., because of a sporting event on a competing network). To decrease the effect of these atypical observations on forecasts, SES uses a mechanism that keeps the value of a processed audience observation within 1.5 times an estimated absolute deviation of the observations. It uses two time series to implement this method, one for the audience observations and one for the estimated absolute deviation. When a specific observation exceeds the smoothed value of the series plus 1.5 times the corresponding absolute deviation estimate, it is replaced with a value equal to the smoothed value of the series plus 1.5 times the corresponding absolute deviation estimate. Mirror logic is used to replace the value of an observation smaller than the smoothed value of the series minus 1.5 times the corresponding absolute deviation estimate. We provide additional details on SES in the appendix.

Actual demographic audience figures are produced overnight; therefore, except for a processing lag of a few days, the most recent demographic audience information is available to be processed using SES. However, actual target-audience figures have a longer production lag of up to eight weeks because of the involved fusion process. Nevertheless, we impute recent target actuals with a linear regression model that uses the available actual demographic figures and other time- and program-based factors. SES is then applied on the combination of actual and imputed values for target audiences. Finally, our ensemble model uses a weighted average of the CAE

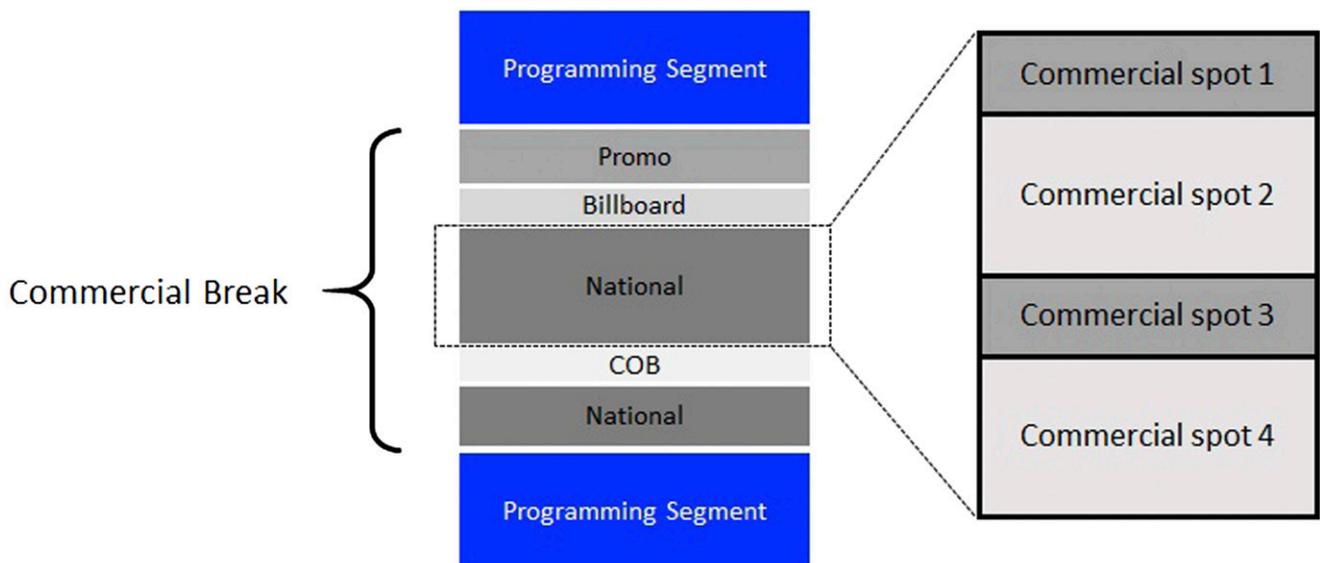
forecasts and the SES forecasts to generate the audience estimates at the network, date, and half-hour level used in spot scheduling.

Spot-Scheduling Engine

In cable networks, commercial breaks comprise different time intervals (called inventory buckets) over which different types of spots can air. Examples of inventory buckets include the following: national (airing national commercial spots), promo (airing national spots that promote TV shows), cable operator break (COB; airing locally inserted commercials, but with backup filler commercials in case the cable operator does not insert commercials), and billboard (airing billboard spots, which are short spots that provide some information about an advertiser relating it to the programming content; for example, closed captioning sponsorship). Figure 11 provides an example of a composition of a commercial break.

A spot-scheduling engine is used to determine the exact times when commercial spots air within the commercial breaks. Promos are scheduled separately under a different business process. Because of audience variations across a selling title week (estimated with the short-term ensemble framework), different placements provide distinct values across spots. Furthermore, different types of spots have diverse scheduling goals; demographic-guaranteed spots seek to achieve a demographic audience guarantee, AudienceNOW spots seek to achieve a target-audience guarantee, TargetingNOW spots try to attain a delivery lift, audience-deficiency spots aim to reduce impression shortfalls, and filler spots seek to maximize their monetary value. Audience-related deals (i.e., all deals except for filler) have an expected and an actual pacing—that is, the

Figure 11. (Color online) Commercial Breaks Are Divided into Several Types of Time Buckets that Air Different Types of Spots



percentage of guaranteed impressions that should have been delivered by now and the percentage that has been delivered, respectively. Therefore, two important inputs to the engine are the expected impression delivery goal at the end of the scheduling period for audience-related deals and the dollar values associated with the impressions of audience-related deals and with the spots of filler deals. These data enable the impression values resulting from different placements to be traded off, and the schedule of spots to be optimized holistically.

At Turner, thousands of commercial spots need to be scheduled every week across all of the selling titles in Turner's networks. Scheduling all spots at once is prohibitive; however, because deals specify a fixed number of units per selling title week, the spot-scheduling task can be decomposed by selling title week within each network. Furthermore, because the most granular level of posting is typically 30 minutes and almost all inventory buckets within commercial breaks fall into a single half hour, the spot-scheduling task can be further decomposed, without loss of optimality, into an assignment phase that assigns spots to inventory buckets and a sequencing phase that determines the specific order of the spots within each bucket.

The assignment phase of the spot scheduler is solved using a mathematical programming formulation in which the main decision variables are binary variables that represent whether a specific spot is assigned to a particular inventory bucket. The objective is to maximize the overall weighted dollar value of the schedule (i.e., the dollar value of the different types of orders is weighted to reflect their relative business importance and priority), while honoring the advertiser requirements and operational restrictions that we list below.

- The total length of spots assigned to an inventory bucket cannot exceed the bucket's duration.
- Some spots have specific position requirements within the commercial break. Typically only the first, second, second-to-last, and last positions within a break are distinguished, and each break can be assigned at most one spot in each of those positions.
- For each product conflict, every commercial break can air a maximum number of spots that share that conflict. Although the maximum per-product conflict is typically 1, there are exceptions, such as the product conflict "Toys" in Cartoon Network.
- All spots from the same advertiser should honor an agreed-on minimum time separation. The most common time separation is 30 minutes; however, other lengths, such as 15 minutes or 60 minutes, are also possible.
- Franchise and title exclusions should be honored; that is, spots should not be scheduled within specific shows (franchises) or within specific episodes or movies (titles), which an advertiser has requested to exclude.

- Spots should air within specified dates and times within the selling title. For example, a spot advertising a movie might be eligible to air only on Thursday, Friday, and Saturday within a selling title that airs the entire week, or a spot advertising breakfast options may be limited to air before 11:00 a.m. within a selling title that runs from 6:00 a.m. to 3:00 p.m.

- Spots should air in inventory buckets that match their type. For example, national spots should not air on COB buckets.

- Some spots have associative requirements; examples include piggyback (two spots that should air back to back), bookend (two spots that should air in the same break with at least one unrelated spot between them), sandwich bookend (three spots that should air in the same break with at least one unrelated spot between pairs of them), and billboard (the billboard should air immediately before a commercial spot from that advertiser).

After preprocessing input data, the formulation for the assignment phase is a mixed-integer programming problem; we provide details in the appendix. Data preprocessing includes determining eligible inventory buckets for each spot based on days and times, inventory-bucket types, and franchises and titles, and consolidating spots with associative constraints so that they can be treated as single spots during this phase. It also involves building clusters of buckets to enforce time-separation requirements without explicitly modeling precedence relationships among spots from the same order, and determining impression delivery goals for all orders based on their type. The main idea behind time-separation clusters is to create groups of inventory buckets that can be assigned at most one spot from an order without violating time-separation requirements. The key point of the process of setting impression delivery goals is to translate the objectives from the different types of spots into impressions. For example, for a TargetingNOW order, this process consists of finding the selling title-week baseline delivery for the order and multiplying this baseline delivery by the order's lift goal to obtain the impression delivery goal for the order.

The sequencing phase of the spot scheduler determines the order of the spots within their assigned buckets, explicitly honoring constraints that are only implicitly enforced during the assignment phase. These include break-position requirements and relative positions of spots with associative constraints. This phase is solved with another mixed-integer programming formulation; we omit the explicit details of this formulation.

The spot-scheduling engine with audience considerations is run in a shrinking rolling horizon. An entire broadcast week, Monday through Sunday, is run; however, only the placements on Monday are finalized. The next day, Tuesday through Sunday is run (with

visibility of the placements scheduled on Monday); but only the placements on Tuesday are finalized. This process repeats until the last two days, Saturday and Sunday, are run together and their spot placements are finalized.

The Need for Operations Research and Advanced Analytics

A seemingly unattainable goal of any media broadcasting company is to build a fully automated system that automatically receives deals, pushes them into the inventory system containing all available commercial minutes across all networks, schedules every unit according to its requirements (e.g., desired selling titles across weeks, competition avoidance, desired time separation, and equitable distribution of spots from a deal across days and within days), performs deal stewardship keeping track of the delivery on each deal, proactively selects audience-deficiency units to reduce impression shortfall, prepares deal postings at the end of their flights, and invoices clients. Such a solution is considered an unrealistic expectation because of the practical limitations in which the industry operates; nonetheless, its ideation provides an ultimate goal toward which the industry should strive.

Imagine multiple disconnected systems under the ownership of several distinct business units that handle most of their communications through spreadsheets, with a semiautomated way of placing commercial advertisements that can honor only simple constraints, such as availability in terms of day and time, and cannot move a commercial once its placement has been decided. This was the status in 2005 when Turner started its quest for a fully automated media broadcasting system. The far-reaching initiative not only addressed consolidation and automation across the board, but also prepared Turner for adding analytics, data, and decision sciences at the core of its business.

In 2006, Turner employed its first operations research analyst to develop an intelligent spot-scheduling engine as part of the larger effort that involved the centralization and automation of its entire traffic system from order intake to commercial ad placement across all Turner networks. The result, Crossroad, has been praised as “the industry’s most advanced traffic system” (CISION 2011), which “will transform the technology landscape of the media industry” (Spangler 2010). In addition to enabling Turner as a frontrunner among its peers, Crossroad proved to be an additional revenue stream. In 2010, Turner gave exclusive worldwide distribution rights to Invision, a well-established provider of sales solutions to the media industry, whose goal was to bring to the industry an “advanced, streamlined, and modern solution that delivers a strategic advantage.”

The success of Crossroad opened the door to more analytical solutions within Turner. However, as in most

industries, adoption of these tools was slow. Decades-long industry practices are not trivial to change, and some of our tools had to be shelved to await a better chance for acceptance. When the industry’s status quo was shaken by allowing viewership data to be enriched with consumer-purchasing information, the time was ripe for Turner to expand its set of analytical tools, because no established industry standard was in place.

Although the availability of targeted audience segments was a necessary condition for a breakthrough in TV ad campaign measurement and execution, it was not sufficient. Data that could be used to create targeted audience segments had been available for years, and some networks and advertisers may have reviewed these data to gain an understanding beyond demographics of the audiences they were reaching; however, the systematic use of these data or the advertising products leveraging them had not been put in place. The development of advanced forecasting and optimization models was the catalyst that enabled the creation of audience-targeted solutions.

Broadcast and cable networks have typically focused on content as their main competitive strategy to increase the audiences that watch their programming. Although content remains important, the other side of the competitive equation, which had not received enough attention, is improving the valuation of their audiences with concomitant changes to the value of their inventory across their programming landscape.

In the traditional television advertising world, units in selling titles are priced based on their expected demographic audiences. A unit will command a higher price if the audience guarantee is expressed in terms of a narrow demographic such as M18-34 as opposed to a broad demographic such as P18+, because the former demographic is scarcer than the latter. The state of the practice in the industry is to produce demographic audience forecasts at the selling title-quarter level (called ratecards); that is, a single estimate is used to measure the expected audience for any commercial placed in any airing of a particular selling title across the entire quarter. Under this limited view of audiences, advertisers become biased toward not advertising on certain networks and (or) dayparts because they may not on average deliver larger audiences in the advertiser’s demographic of interest.

The advent of targeted audiences resulted in a renewed need to measure audiences beyond the coarse selling title-quarter level, to be able to identify and exploit audience variations across networks, selling titles of a network, and airings of a selling title, and to thus create more efficient and effective media plans. However, traditional audience-estimation methods could not be applied to yield these granular forecasts; in addition, the simple approaches used by most vendors, including (1) the use of indices (i.e., ratios) of targeted audiences to

demographic audiences, and (2) the use of time-series-based models, such as naive averages of prior-quarter audiences or prior-year audiences, fail to recognize that targeted audiences are not just a proportion of demographic audiences and that they vary depending on time attributes, programming attributes, and competitors' programming attributes. Therefore, the development of advanced forecasting models that capture all of these effects was a critical step.

Furthermore, the proposal-building task has typically been a largely manual process relying mostly on the expertise of planners and their knowledge of advertisers to produce single-network proposals that vary little from year to year. The value proposition of AudienceNOW is that the advertiser can choose any data and audience, and Turner guarantees the targeted delivery across all of its networks; therefore, Turner saw a fundamental need to create cross-network proposals with methods that can exploit audience variations and create campaigns that deliver simultaneous value for the company, through the increase of unit rates, and for its advertising partners, through the reduction of their target CPMs. Therefore, developing sophisticated mathematical programming models that could ingest all of the different business and advertiser requirements was also essential.

Thus, operations research and advanced analytics have helped Turner take targeted audience segments from concept to reality, from information that is nice to know to the actual planning and execution of successful media campaigns that benefit Turner and its clients. They have broken paradigms, such as the previous presumption that advertisers should stay away from entire networks or selling titles because they lack their audiences on average.

Change Now or Be Left Behind: Impact and Challenges of Audience-Targeting Solutions

The \$68 billion TV ad marketplace is on the verge of overturning traditional buying. By enabling audience-based transactions in the TV advertising space and by taking the lead in establishing the rules of engagement for audience-targeting deals, Turner has moved the industry forward and it has changed its overall perception as a data-driven organization.

Turner's audience-targeting solutions have been a game changer for the advertising industry. Turner was the first to market in offering end-to-end audience-targeting solutions (from proposal creation to posting) in the television advertising space. These solutions represented a significant competitive advantage because this leadership status has given Turner enormous leverage in determining the rules of engagement for advertising sales under audience targeting. Turner

shifted the focus of the advertising proposal process from being centered around an advertiser's budget and flight to being centered around defining the right audience and targeting criteria, setting benchmarks, and agreeing on key performance indicators for the campaign.

Qualitative Benefits

In addition to revolutionizing traditional advertising sales processes, audience targeting has changed the executive perception of the value of data and analytics in the business and has helped facilitate a positive emphasis on being data driven within our corporate culture. Further, targeting initiatives have facilitated technological disruption in the company. Visibility into how data are monetized is much greater in the targeting space. Because so much of the data available in the media space are delivered by third parties, the granularity of targeting data helps our finance teams to better understand the value of data initiatives, their incremental revenue, and the cost savings targeting brings. This has also led to a cultural impact on the business as more internal communications focus on our business as driven by data and technology.

Targeting has also facilitated technological disruption across the company. To support audience-based advertising transactions, operations research and advanced analytics technologies that service the entire life cycle of an advertising deal were needed. In addition to the tools we present in this paper, new analytics tools were developed to support presale and sales-prospecting initiatives, deal pacing, and postsales reporting efforts. The emphasis and use of data visualization and self-service business intelligence technologies have grown significantly with Turner's use of audience targeting. The need to support a growing number of practitioners who are not knowledgeable about analytics, but who need to interact with data in the business, has grown as the analytic sales process has matured.

This adoption of analytic technologies for audience-based selling has led to faster adoption of advanced analytics in other business functions, such as content scheduling, financial planning, and marketing, as we explain below in the Transportability section.

Quantitative Benefits

Although we cannot disclose specific revenue numbers for confidentiality reasons, we can present some numbers that quantify the large financial impact that audience-targeting solutions have brought to Turner and its advertising partners:

- Generation of advertiser demand for new products: Television has gradually become less attractive to advertisers since the inception of the internet and mobile telephony. Digital advertising is easy to customize to reach the right person at the right moment on

the right device. According to industry studies, digital advertising revenue surpassed television advertising revenue for the first time in 2017 (Slefo 2017). Nevertheless, Turner has been able to maintain a growing ad revenue business despite these trends. This has been possible only because of the customized targeting capabilities it has offered to advertisers. Two factors have been instrumental in increasing revenue: (1) revenue retention due to Turner's targeting capabilities (i.e., the revenue from advertisers that would have stopped advertising with Turner under the traditional demographic-based paradigm), and (2) the premium Turner has been able to charge for its new products. To date, more than 175 targeted deals have been executed, representing hundreds of millions in new ad revenue, and the year-over-year revenue growth rates for these deals has increased from 80% in the first two years (2015 to 2017) to 186% in the period 2017 to 2018, based on advanced bookings. Targeted deal volume has increased year over year every quarter since the inception of TargetingNOW in 2014 and AudienceNOW in 2015. This is reflective of the industry's desire to move away from demographic-based guarantees.

- Advertiser return in media investment: Some advertisers using AudienceNOW and TargetingNOW have engaged Nielsen to do a study to link their sales data to their respective campaign spend. The study considered a random panel of customers reflective of the national population; the panel included people who were exposed to the respective campaign and people who were not. Nielsen then recorded the money spent on the advertised good from the two categories of people on the panel. It used statistics to project the findings to the national population. These findings showed:

- 25% average lift in target audience across hundreds of linear TV audience-targeting deals (increased efficiency);

- \$118 million in incremental sales, 4%–15% sales lift, and 21% average lift in online sales behavior for advertisers who requested this study.

- Simultaneous creation of advertiser and sales efficiencies: optimized audience-targeted media schedules generate value for advertisers by pairing commercial inventory to audiences of interest, thus achieving substantive cost-per-impression reductions. It also allows Turner to use airtime inventory more efficiently (i.e., generate more revenue) by considering its total impression capacity and holistically prioritizing the inventory it should pair with each advertising deal. Of all of the TargetingNOW campaigns that have been completed so far, each received a lift in target delivery across categories and sizes of targets; the average lift was 27%, and the high was 51% in delivered targeted impressions when compared with a benchmark. AudienceNOW has built on the success of

TargetingNOW; so far, the target CPM of all schedules created has decreased by at least 20%; in most cases, their demographic CPMs also decreased. Because these products enable an advertiser to achieve both audience and pricing efficiencies, we also see evidence of increased advertiser loyalty; for example, all advertising partners that beta tested TargetingNOW and AudienceNOW renewed their targeted deals, and some increased their investments in the products in subsequent quarters.

Implementation Challenges

In implementing audience-targeting solutions, we had to overcome some business, data, and technology challenges. From a business perspective, although some forecasting and proposal-building models had been presented in the literature prior to our work, the state of the practice at Turner (and, to the best of our knowledge, the current state for most media companies) involved (1) producing demographic audience forecasts at the selling title-quarter level (i.e., using a single estimate to forecast the audience for all potential commercial placements on a selling title throughout an entire quarter), and (2) manually generating single-network proposals. Turner's research department dedicates groups of individuals to generating demographic audience estimates at the selling title-quarter for each network. These coarse forecasts are used in long-term planning, but are unsuitable for tactical and operational planning because they do not identify audience variations across weeks and specific times of the day. The strategic planning department, which relies mostly on the expertise of planners and their knowledge of advertisers, performed the proposal-building task but generated proposals with media mixes that changed very little from year to year. Cross-network media plans were mostly created, negotiated, and managed as independent, single-network plans. Thus, developing and deploying our models required major transformations in business processes because the traditional approaches were unsuitable for scaling under the targeting paradigm, which involves an extremely large number of potential targeted audience segments and proposals that cross the entire portfolio of networks.

Many of the data and technology implementation challenges we encountered when deploying and supporting a robust forecasting system relate to data development and business-specific use cases. Some of these challenges included the following:

- Acquiring detailed programming information early in the process: Previously, long-term programming schedules were typically managed and shared among internal business groups through emails and complex spreadsheets. The need to collect standardized program schedule information at the needed resolution, date, and

half hour over a long horizon required substantial data development.

- Enabling internal software systems to retrieve forecasts as a service, on demand: this required development of application program interfaces (APIs) that allow internal software systems to request and present forecasts to users by leveraging standard data stores of previously produced estimates and by dynamically generating new audience forecasts, given changes to programming schedules and (or) audience target definitions.

- Customizing forecasting models for specific networks and special content: Although the application of one general audience forecasting model across various cable networks typically provided unbiased forecasts, it generated estimates that sometimes did not suit the intuition of business users. Leveraging alternative features and (or) developing separate models for specific networks or special programs served to improve forecast accuracy, provide more granular, program-specific forecasts, and facilitate user acceptance. This was especially apparent for news and sporting events, such as telecasts of NBA and MLB games and NASCAR races.

Moreover, the data and technology implementation challenges encountered when deploying and supporting our optimization models relate mostly to a flexibility to quickly accommodate audience-targeting business rules that are continuously evolving based on feedback from the market and internal business users. Some of these challenges included the following:

- Developing rapid prototypes of optimization models: The traditional process to develop and deploy optimization-based decision support systems involved creating a proof of concept (POC) to show the business value and then integrating the models with the rest of Turner's systems. The need to respond quickly to the market demands necessitated the acceleration of this process, bypassing the POC stage. Thus, we had to quickly demonstrate the value of our optimization models while already operating on real media deals as they were being negotiated with advertisers. This necessitated a high level of communication with our business users such that we could quickly iterate and made modifications to the models as new requirements were identified during the negotiation process.

- Refining suitable goals and requirements in flux: In the past, proposals had specific media-mix requirements and only addressed demographic audiences. With the advent of targeted audiences, new goals and benchmarks had to be incorporated in a flexible way that would enable iterating with business users to experiment and refine different types of objective functions and constraints.

- Implementing spot-scheduling changes gradually: The first implementation of TargetingNOW placements was performed by time locking spots to specific

half hours, thus ensuring that these spots would receive preferential placements, which would guarantee that they exceeded their desired impression lift; later, changes were made to the spot scheduler such that these types of spots attained only their desired impression lift. The first implementation of AudienceNOW placements leveraged the TargetingNOW logic as a proxy; later, a reengineering of the entire spot-scheduler logic was performed to trade off the value of all different types of spots on a consistent basis.

Transportability

After successfully deploying its audience-targeting solutions on the domestic commercial side, Turner is exploring alternatives to leverage these methods in different areas of its business. So far, we are transporting the audience-targeting methods to promotional business processes, international markets, and additional components of the revenue management process in the domestic business.

The promotional processes (i.e., the planning and scheduling of spot-advertising network content) are a natural additional application of our audience-targeting methods because their business processes parallel those of the commercial side, although they are performed by different business units. Traditionally, promotional plans, which are similar to proposals, have been developed to reach a desired audience in the P18-49 demographic. However, the use of targeted audience segments, such as “fantasy and imaginative drama lovers” or “reality competition enthusiasts,” enables marketing to efficiently concentrate on its desired audience for each promotional campaign. The POC of audience-targeting solutions on the promotional side has shown promise to reduce promotional inventory needs by up to 50%. This reduction is possible because the forecasting and optimization models allow Turner to differentiate commercial airtime slots with the same demographic audiences, which used to be considered equivalent under the demographic paradigm, such that promos can be allocated to the slots with higher concentrations of their desired targeted audiences; therefore, audience goals for promotional plans can be attained using fewer promos. Transporting our models to the promotional side involved additional challenges, the redesign of processes, and the implementation of change management. An entire process redesign was not possible immediately; therefore, the changes will be incorporated gradually into the promotional planning. We first reengineered a promo-scheduling engine so that it can accommodate CAE estimates for both demographic and targeted audiences, and we deployed this promo scheduler in production in the fourth quarter of 2017. In parallel, we have been running optimization models similar to the proposal-creation model to evaluate the aggregate promotional

inventory required for the promotional campaigns slated to run during a planning period, and we have been providing comparisons with the current manual approaches.

Moreover, our audience-targeting methods have been customized to conduct business in the international sales arm of Turner, considering the nuances of local selling and commercial guarantees. Turner Latin America was first to market with audience-targeting capabilities in the Latin American region, taking the lead in establishing the sales rules under the targeting paradigm. In January 2018, Turner introduced its audience-targeting solutions during its upfront-marketing event in Mexico. The first targeted campaign in the Mexican market was a huge success. First, the advertiser doubled its regular advertising campaign investment, a sign of the excitement in the market about the promise and value of audience targeting. Second, the results produced provided a lift of 32% in demographic audience and 43% in targeted-segment audience; that is, the advertiser received 32% more viewers in its chosen demographic and 43% more viewers in its targeted segment, which translates to a significant reduction in cost per impression for the advertiser. Furthermore, an additional benefit to Turner is that audience targeting provides an alternative to advertisers to induce them to advertise on networks and (or) at times that they have been traditionally reluctant to try; thus, Turner can capitalize on networks and times that have typically been undersold. Turner plans to roll out its audience-targeting capabilities to the other six Latin American countries in which it operates.

Finally, as targeted audiences permeate the industry, all other related revenue management processes need to be redesigned under our audience-targeting framework. Currently, Turner is developing proofs of concept to transport our methods into other areas such as programming optimization, the allocation of audience-deficiency units, and the distribution of branded units (breaking down spots purchased by large advertisers into smaller branded deals). We will provide some additional information for the programming optimization application only.

In the TV industry, the business function of acquiring and organizing television programs into a 24-hour schedule, which is referred to as TV programming, has significant challenges that affect the company's financial performance. These challenges include the following.

- The process of content scheduling is often tedious and complex because of the detailed and burdensome contractual obligations associated with licensing a syndicated program.
- Internally, the sales, programming, and finance organizations may have different views about content scheduling; examples of these differences include the following:

- How appealing or advertiser friendly is the program?

- Which program will attract the largest audience?

- What are the tax and financial implications related to the licensing of a show?

The core forecasting models utilized for audience-based selling have been modified for this specific business task, and an optimization model that considers licensing constraints, program airing commitments, and program licensing costs, can be leveraged using different objectives, such as maximizing the expected audience or rating of a program schedule.

Beyond these proven transportability applications, the core concepts behind our audience-targeting solutions will be the cornerstone of additional applications such as inventory-mix optimization (determining the proportions of nonprogramming airtime that should be devoted to sales, promotional efforts, and allocations of audience-deficiency units), cross-platform advertising solutions that offer advertisers holistic deals across our cable networks and digital and mobile properties, and the deployment of these solutions to the other international regions in which Turner operates: Europe, Middle East, Africa (EMEA) and the Asia Pacific.

Acknowledgments

The authors thank Dan Aversano, David Benoit, and Brad Imbach, together with their teams, for their help and support in implementing this work. They also thank the operations research analysts who have expanded upon this work: Brandon Wang, Sin-Yi Chou, Sina Faridimehr, Haifan Li, Yingjie Li, Jia Liu, Juan Ma, Xin Ma, Connor Owen, and Vivek Vasudeva. Furthermore, they are grateful to their Edelman coaches, Jeff Alden and Arnie Greenland, for their guidance throughout the competition, and to their editors, John Milne and Alice Mack, for their valuable revisions and comments on this paper.

José Antonio Carbajal's current affiliation is iHeartMedia, Inc., San Antonio, Texas; antoniocarbajal@iheartmedia.com. Peter Williams' current affiliation is TCGplayer, Roswell, Georgia; pwilliams26@gmail.com.

Appendix: Forecasting Model Specifications and Mathematical Programming Formulations

This appendix contains the specifications of the forecasting models and the mathematical programming formulations of the optimization models we describe above in the sections titled *The Planning Stage: Long-term Forecasting and Proposal Creation* and *The Execution Stage: Short-term Forecasting and Spot Scheduling*.

Competitive Audience Estimation

The CAE model is used to produce audience estimates across all of Turner's networks. CAE is a multilevel or mixed-effects regression in which observations are represented at a network, date, and half-hour level, indexed $i = 1, 2, \dots, n$. The response of interest is the average audience, which is estimated via a logarithmic transformation $y_i = \log(AA_i + 1)$,

where AA_i represents the average audience or impressions at the observation level. This choice of transformation is informed by the occurrence of natural zeroes—that is, observations with no impressions delivered.

The general form of the model used to estimate the transformed average audience is written below. The model has predictors that vary at the TV program level, indexed $j = 1, \dots, J$, and predictors that are common across all observations $i = 1, \dots, n$.

$$\begin{aligned} y_i &= X_i^0 \beta^0 + X_i B_{j[i]} + \epsilon_i, \text{ for } i = 1, \dots, n \text{ with,} \\ y_i &\sim N(X_i^0 \beta^0 + X_i B_{j[i]}, \sigma_y^2), \text{ and,} \\ B_j &\sim N(U_j G, \Sigma_B), \text{ for } j = 1, \dots, J, \end{aligned} \quad (\text{A.1})$$

where

- X is a $J \times 2$ matrix of program-level predictors, and B is a $2 \times J$ matrix of program-level coefficients. This enables a program-level intercept and slope on the PUT factor, where $j[i]$ represents the program associated with observation i that represents a half-hour program airing specific to a date and time.

- U is a $J \times 2$ matrix of program level predictors and G is a 2×2 matrix of coefficients at the program level. Their product is a vector of length 2 that is modeled as a bivariate normal, with covariance matrix Σ_B .

- X^0 is an $n \times L$ matrix of predictors and β^0 is an $L \times 1$ vector of regression coefficients that are common to all observations. These predictors include binary or dummy variables that identify whether an individual observation represents a repeat or new telecast of a TV program episode, a variable that measures the historical counts of programs of the same genre, and a combination of trigonometric functions to capture seasonality in television; see Danaher and Dagger (2012) and Gensch and Shaman (1980) for additional details on multilevel models and trigonometric regressions.

To estimate model parameters, the restricted maximum likelihood criterion is used; see Bates et al. (2015) for more details.

Proposal-Building Optimization Model

The goal of the proposal-building optimization model is to determine for a deal the number of units to allocate to each selling title-week combination during the deal's flight across all Turner networks and the prices that should be charged for each unit. First, we present a nonlinear programming model we designed for unit selection and pricing, and we then show an implemented MIP reformulation that is equivalent for the case in which rates are increased in the same proportion. For simplicity, we present the formulations that build a single proposal; however, we can extend these models to build multiple proposals for several advertisers at once. Regarding notation, we use uppercase letters to represent sets and parameters and lowercase letters to represent variables and set indices.

• Sets

\mathcal{N} = the set of networks to be included in the proposal, indexed by n .

\mathcal{G} = the set of time-based Nielsen dayparts, indexed by g .

\mathcal{H} = the set of all Turner dayparts, indexed by h .

\mathcal{S} = the set of all selling titles across all included networks, indexed by s .

$\mathcal{S}_n, \mathcal{S}_g, \mathcal{S}_h$ = the subset of selling titles associated with network n , Nielsen daypart g , and Turner daypart h , respectively.

\mathcal{W} = the set of in-flight weeks for the proposal, indexed by w .

• Parameters

B = Proposal budget.

O^T = Baseline for target CPM.

O^D = Baseline for demographic CPM.

Δ^T = Minimum proportion reduction in target CPM.

Δ^D = Maximum proportion increase in demographic CPM.

R_{sw}^F = Floor rate for units of selling title s in week w .

Π = Maximum allowable rate-increase proportion over each floor rate.

A_{sw}^T = Estimated target impressions (in thousands) on selling title s in week w .

A_{sw}^D = Estimated demographic impressions (in thousands) on selling title s in week w .

I_{sw}^e = Remaining number of EQ30s available on selling title s in week w .

I_{sw}^c = Remaining number of units with the same product conflict that can still be scheduled for selling title s in week w (i.e., number of commercial breaks on selling title s in week w minus the number of units with the same product conflict, which are already booked in other deals).

I_{sw}^t = Maximum number of units that can air on selling title s in week w without violating time-separation constraints.

N_h, N_s = The Turner network associated with Turner daypart h and selling title s , respectively.

G_s = The Nielsen daypart associated with selling title s .

H_s = The Turner daypart associated with selling title s .

$P_n^B(\min), P_n^B(\max)$ = Minimum and maximum proportions of the total proposal budget to be assigned to network n .

$P_g^A(\min), P_g^A(\max)$ = Minimum and maximum proportions of the total proposal target impressions to be assigned to Nielsen daypart g .

$P_h^A(\min), P_h^A(\max)$ = Minimum and maximum proportions of the target impressions on network N_h to be assigned to Turner daypart h .

$P_s^A(\min), P_s^A(\max)$ = Minimum and maximum proportions of the target impressions on network N_s to be assigned to selling title s .

F_w = Desired proportion of the total units in the proposal to be aired in week w .

L = Penalty per unit for deviations from the weekly goals.

M = Minimum relative profit margin that the proposal should yield.

• Decision Variables

x_{sw} = Number of units (EQ30s) in the proposal assigned to selling title s in week w .

r_{sw} = Rate (i.e., price) to be charged in the proposal for an EQ30 on selling title s in week w .

δ_w = Deviation from the goal on the number of units to be aired in week w .

δ_w^+, δ_w^- = Positive and negative deviations from the goal on the number of units to be aired in week w .

• Auxiliary Variables

b = Total dollar value of the proposal.

b_n = Dollar value of the proposal that is assigned to network n .

β = Total floor-rate dollar value of the proposal.
 β_n = Total floor-rate dollar value of the proposal that is assigned to network n .
 μ = Profit margin of the proposal (in dollars).
 a^T = Total target impressions in the proposal.
 a^D = Total demographic impressions in the proposal.
 $a_n^T, a_g^T, a_h^T, a_s^T$ = Target impressions in the proposal that are assigned to network n , Nielsen daypart g , Turner daypart h , and selling title s , respectively.

• **NLP Formulation**

Maximize

$$a^T - L \sum_{w \in \mathcal{W}} |\delta_w| \quad (\text{A.O1})$$

Subject to

$$b = \sum_{s \in \mathcal{S}, w \in \mathcal{W}} r_{sw} x_{sw} \quad (\text{A.2})$$

$$\beta_n = \sum_{s \in \mathcal{S}_n, w \in \mathcal{W}} r_{sw} x_{sw} \quad (\text{A.3})$$

$$\beta = \sum_{s \in \mathcal{S}, w \in \mathcal{W}} R_{sw}^F x_{sw} \quad (\text{A.4})$$

$$b \leq B \quad (\text{A.5})$$

$$R_{sw}^F \leq r_{sw} \leq (1 + \Pi) R_{sw}^F, \forall s \in \mathcal{S}, w \in \mathcal{W} \quad (\text{A.6})$$

$$\mu = b - \beta \quad (\text{A.7})$$

$$\frac{\mu}{\beta} \geq M \quad (\text{A.8})$$

$$a^T = \sum_{s \in \mathcal{S}, w \in \mathcal{W}} A_{sw}^T x_{sw} \quad (\text{A.9})$$

$$a^D = \sum_{s \in \mathcal{S}, w \in \mathcal{W}} A_{sw}^D x_{sw} \quad (\text{A.10})$$

$$a_i^T = \sum_{s \in \mathcal{S}_i, w \in \mathcal{W}} A_{sw}^T x_{sw}, \forall i \in \{\mathcal{N}, \mathcal{G}, \mathcal{H}\} \quad (\text{A.11})$$

$$a_s^T = \sum_{w \in \mathcal{W}} A_{sw}^T x_{sw}, \forall s \in \mathcal{S} \quad (\text{A.12})$$

$$\frac{b}{a^T} \leq (1 - \Delta^T) O^T \quad (\text{A.13})$$

$$\frac{b}{a^D} \leq (1 + \Delta^D) O^D \quad (\text{A.14})$$

$$x_{sw} \leq \min\{I_{sw}^e, I_{sw}^c, I_{sw}^t\}, \forall s \in \mathcal{S}, w \in \mathcal{W} \quad (\text{A.15})$$

$$P_n^B(\min) \leq \frac{b_n}{b} \leq P_n^B(\max), \forall n \in \mathcal{N} \quad (\text{A.16})$$

$$P_g^A(\min) \leq \frac{a_g^T}{a^T} \leq P_g^A(\max), \forall g \in \mathcal{G} \quad (\text{A.17})$$

$$P_h^A(\min) \leq \frac{a_h^T}{a_{N_h}^T} \leq P_h^A(\max), \forall h \in \mathcal{H} \quad (\text{A.18})$$

$$P_s^A(\min) \leq \frac{a_s^T}{a_{N_s}^T} \leq P_s^A(\max), \forall s \in \mathcal{S} \quad (\text{A.19})$$

$$\delta_w = F_w \sum_{s \in \mathcal{S}, w_2 \in \mathcal{W}} x_{sw_2} - \sum_{s \in \mathcal{S}} x_{sw}, \forall w \in \mathcal{W} \quad (\text{A.20})$$

$$x_{sw} \in \mathbb{Z}^+, \quad r_{sw} \in \mathbb{R}^+, \quad \delta_w \in \mathbb{R}.$$

The objective function (A.O1) maximizes the total target impressions in the proposal minus a penalty for the total deviations from weekly unit goals. The absolute value reflects that both positive and negative deviations are penalized.

Constraints (A.2)–(A.4) enforce dollar-value definitions. Constraint (A.2) is a quadratic expression, which states that the total dollar value of the proposal equals the sum across all selling titles and weeks of the products of the units selected in that selling title week times the rate to be charged for units in that selling title week, both of which are decision variables. Constraint (A.3) is similar to Constraint (A.2), but is limited to the dollar value of the proposal that is assigned to a particular network. Constraint (A.4) is a linear expression similar to Constraint (A.2) and defines the value of the proposal in terms of the floor rates (i.e., the baseline rates).

Constraints (A.5)–(A.8) are related to dollar value, rates, and margins. Constraint (A.5) states that the proposal value should not exceed the budget. Constraint (A.6) stipulates that the rates charged for units on a given selling title week should be greater than or equal to the corresponding floor rates. Constraint (A.7) defines the proposal margin (in dollars) as the difference between the charged-rate proposal value and the floor-rate proposal value. Constraint (A.8) enforces the lower bound on the relative proposal margin as a proportion of the floor-rate proposal value.

Constraints (A.9)–(A.12) enforce target and demographic impression aggregations at different levels of granularity. Constraint (A.9) states that the total target impressions in the proposal equal the sum across all selling titles and weeks of the products of the units selected in that selling title week times the corresponding target audience. Constraint (A.10) states that the total demographic impressions in the proposal equal the sum across all selling titles and weeks of the products of the units selected in that selling title week times the corresponding demographic audience. Constraint (A.11) aggregates the target impressions in the proposal assigned to specific networks, Nielsen dayparts, and Turner dayparts. Constraint (A.12) aggregates the target impressions by selling title.

Constraints (A.13) and (A.14) relate the proposal dollar value and impressions. Constraint (A.13) states that the target CPM (i.e., proposal dollar value divided by the total target impressions, in thousands) should be reduced by at least a desired proportion of the baseline target CPM. Constraint (A.14) states that the demographic CPM (i.e., proposal dollar value divided by the total demographic impressions, in thousands) should not increase beyond a specific proportion of the baseline demographic CPM.

Constraint (A.15) establishes that the number of units in the proposal assigned to a particular selling title week cannot exceed the remaining EQ30s available, the remaining units with the same product conflict that can still be scheduled, or the maximum number of units that can air without violating time-separation restrictions.

Constraints (A.16)–(A.19) define mix restrictions on the units selected. Constraint (A.16) specifies lower and upper bounds on the proportion of dollar value assigned to each network. Constraints (A.17), (A.18), and (A.19) enforce lower and upper bounds on the proportions of impressions that are assigned to specific Nielsen dayparts, Turner dayparts, and selling titles, respectively.

Finally, Constraint (A.20) establishes that the deviation from a weekly unit goal equals the difference of the desired weekly goal (i.e., desired weekly proportion times the total number of units in the proposal) minus the units that are scheduled in the corresponding week.

• **MIP Reformulation**

As we indicate earlier, the pricing decision can be viewed as determining how much to increase the floor rate of each unit in the proposal. For the case in which floor rates are constrained to increase in the same proportion (i.e., $r_{sw} = (1 + \alpha)R_{sw}^F, \forall s \in \mathcal{S}, \forall w \in \mathcal{W}$, for some $\alpha \geq 0$), the NLP can be reformulated as the MIP shown below, in which Constraints (A.4), (A.8)–(A.12), and (A.15) are still valid and have the same interpretations as above.

Note that under proportional rate increase, from Constraint (A.2), $b = \sum_{s \in \mathcal{S}, w \in \mathcal{W}} r_{sw} x_{sw} = \sum_{s \in \mathcal{S}, w \in \mathcal{W}} (1 + \alpha)R_{sw}^F x_{sw}$. Using Constraint (A.4), it follows that $b = \beta + \alpha\mu$, and we can combine this with Constraint (A.7) to solve for α as $\alpha = \frac{\mu}{\beta}$. Thus, once the MIP is solved, the values of r_{sw} can be calculated as $r_{sw} = \left(1 + \frac{\mu}{\beta}\right)R_{sw}^F, \forall s \in \mathcal{S}, \forall w \in \mathcal{W}$.

Maximize

$$a^T - L \sum_{w \in \mathcal{W}} (\delta_w^+ + \delta_w^-) \tag{A.O2}$$

subject to

$$(A.4), (A.8)–(A.12), (A.15)$$

$$\beta_n = \sum_{s \in \mathcal{S}_n, w \in \mathcal{W}} R_{sw}^F x_{sw} \tag{A.21}$$

$$\beta + \mu = B \tag{A.22}$$

$$\frac{\mu}{\beta} \leq \Pi \tag{A.23}$$

$$\beta + \mu \leq (1 - \Delta^T) O^T a^T \tag{A.24}$$

$$\beta + \mu \leq (1 + \Delta^D) O^D a^D \tag{A.25}$$

$$P_n^B(\min) \beta \leq \beta_n \leq P_n^B(\max) \beta, \forall n \in \mathcal{N} \tag{A.26}$$

$$P_g^A(\min) a^T \leq a_g^T \leq P_g^A(\max) a^T, \forall g \in \mathcal{G} \tag{A.27}$$

$$P_h^A(\min) a_{N_h}^T \leq a_h^T \leq P_h^A(\max) a_{N_h}^T, \forall h \in \mathcal{H} \tag{A.28}$$

$$P_s^A(\min) a_{N_s}^T \leq a_s^T \leq P_s^A(\max) a_{N_s}^T, \forall s \in \mathcal{S} \tag{A.29}$$

$$\delta_w^+ - \delta_w^- = F_w \sum_{s \in \mathcal{S}, w_2 \in \mathcal{W}} x_{sw_2} - \sum_{s \in \mathcal{S}} x_{sw}, \forall w \in \mathcal{W} \tag{A.30}$$

$$x_{sw} \in \mathbb{Z}^+, \mu, \delta_w^+, \delta_w^- \in \mathbb{R}^+.$$

The objective function (A.O2) is equivalent to (A.O1), but the absolute value of the deviations from weekly unit goals are linearized in two steps: First, each of the original unrestricted in sign variables δ_w are modeled as the difference of two nonnegative variables, δ_w^+ and δ_w^- . This is enforced by Constraint (A.30). Second, under this definition, $|\delta_w| = \delta_w^+ + \delta_w^-$, which are the terms that appear in (A.O2).

Constraint (A.21) defines the dollar value of the proposal that is assigned to a particular network in terms of the floor rates. Constraint (A.22) exploits the fact that the entire budget will be exhausted to substitute Constraints (A.5) and (A.7) in defining the proposal margin. Constraint (A.23), together with the nonnegativity restriction on μ , substitutes Constraint (A.6) in establishing bounds on rate increases. Constraints (A.24)

and (A.25) and Constraints (A.26)–(A.29) are linearized reexpressions of Constraints (A.13) and (A.14) and Constraints (A.16)–(A.19), respectively. Note that $\beta_n/b_n, \forall n \in \mathcal{N}$ is constant under proportional floor-rate increase.

Short-Term Audience Forecasts

Short-term forecasts of audiences are generated through averaging forecasts from our long-term forecasting model and forecasts from the SES model. The SES model is implemented at the network, day-of-week, and half-hour level. As we note earlier, the historical data that are represented in each time series consist of observations matched at the franchise, or selling title level to ensure forecasts are relevant. For each network, there are thus 7×48 time series with historical audience data indexed by time t . We denote the actual historical time series of audience observations s , representing the audience for each network, half hour, and day of week.

To decrease the effect of atypical observations on the model in practice, the time series s' is employed in place of the actual observations s . s' represents a series of “capped observations” that are generated below based on an estimate of model deviation.

The time series $d_t = |s_t - \hat{s}_t|$ tracks the absolute deviation between the observations and their fitted values, and is smoothed with the following specification:

$$\begin{aligned} \hat{d}_t &= \beta |s_{t-1} - \hat{s}_{t-1}| + (1 - \beta) \hat{d}_{t-1} \\ &= \beta d_{t-1} + (1 - \beta) \hat{d}_{t-1} \end{aligned} \tag{A.31}$$

The time series of audience observations is smoothed with the following specification:

$$\begin{aligned} \hat{s}_{t+1} &= \alpha s'_t + (1 - \alpha) \hat{s}_t \\ &= \alpha s'_t + (1 - \alpha) [\alpha (s'_{t-1}) + (1 - \alpha) \hat{s}_{t-1}] \end{aligned} \tag{A.32}$$

As is done typically with exponential smoothing, these models are initialized with the first observed values. The series s' represents the capped observations of s based on the following rule:

$$s'_t = \begin{cases} s_t, & \text{if } L_t \leq s_t \leq U_t \\ U_t, & \text{for } s_t > U_t, \\ L_t, & \text{for } s_t < L_t \end{cases} \tag{A.33}$$

where U_t and L_t represent upper and lower bounds for an observation inclusion in the model and are defined as

$$\begin{aligned} U_t &= \hat{s}_{t-1} + 1.5 \hat{d}_t, \\ L_t &= \hat{s}_{t-1} - 1.5 \hat{d}_t. \end{aligned} \tag{A.34}$$

Both model parameters $0 \leq \alpha, \beta \leq 1$ are specific to each network, date, and half hour, and tuned to minimize the sum of squared errors (SSE) between the observed audience data and the audience smoothed values:

$$SSE = \frac{\sum_{t=0}^T (s_t - \hat{s}_t)^2}{T}. \tag{A.35}$$

The long-term and smoothing model forecasts may be combined as an ensemble that generates the final short-term audience estimates. Each long-term forecast \hat{y}_t is mapped to its corresponding smoothing-model forecast \hat{s}_{t+1} , because the output of both models can be mapped to the same resolution

(i.e., network, day, half hour). This ensemble is computed as a weighted average of predictions from both models:

$$\text{audi\^{}ence}_{t+1} = \lambda \hat{y}_{i=t+1} + (1 - \lambda) \hat{s}_{t+1}, \text{ where } 0 \leq \lambda \leq 1. \quad (\text{A.36})$$

Spot-Scheduling Optimization Model

The spot-scheduling engine is used to determine the exact times when commercial spots air within the commercial breaks. We present the MIP formulation solved during the assignment phase of spot scheduling, which assigns spots to inventory time buckets. The assignment phase is followed by a second MIP (not shown) that determines the specific times that spots air within the assigned inventory buckets. Regarding notation, we use uppercase letters to represent sets and parameters and lowercase letters to represent variables and set indices.

• Sets

\mathcal{B} = the set of commercial breaks, indexed by b .

\mathcal{I}_b = the set of airtime inventory buckets within commercial break b , indexed by i .

\mathcal{I} = the set of all inventory buckets, indexed by i .

\mathcal{I}_s = the subset of inventory buckets in which spot s can be scheduled, indexed by i . This subset already accounts for constraints on eligible days, eligible times, inventory-bucket type matching, and franchise and title exclusions for spot s .

\mathcal{O} = the set of orders with spots to be scheduled, indexed by o .

$\mathcal{O}^D, \mathcal{O}^T, \mathcal{O}^A, \mathcal{O}^U, \mathcal{O}^F$ = the subset of demographic-guaranteed, TargetingNOW, AudienceNOW, audience-deficiency, and filler orders with spots to be scheduled, respectively, all indexed by o .

\mathcal{S} = the set of spots to be scheduled, indexed by s .

\mathcal{S}_o = the set of spots in order o , indexed by s .

$\mathcal{S}^D, \mathcal{S}^T, \mathcal{S}^A, \mathcal{S}^U, \mathcal{S}^F$ = the subset of demographic-guaranteed, TargetingNOW, AudienceNOW, audience-deficiency, and filler spots to be scheduled, respectively, all indexed by s .

\mathcal{S}_i = the subset of spots that are eligible to be scheduled in inventory bucket i , indexed by s .

$\mathcal{P} = \{a, b, y, z\}$ = the set of break positions, indexed by p , where a = first in break, b = second in break, y = second-to-last in break, and z = last in break.

\mathcal{C} = the set of product-conflict codes, indexed by c .

\mathcal{H} = the set of days in the scheduling period, indexed by h .

\mathcal{I}_h = the set of all inventory buckets that are scheduled on day h , indexed by i .

\mathcal{N}_f = the set of f -minute separation clusters, indexed by n .

\mathcal{I}_n = the set of inventory buckets included in time-separation cluster n , indexed by i .

• Parameters

Δ_s = Primary demographic associated with demographic-guaranteed or audience-deficiency spot s .

\mathcal{T}_s = Audience segment (target) associated with TargetingNOW or AudienceNOW spot s .

$G_o^D, G_o^T, G_o^A, G_o^U$ = Impression delivery goal for demographic-guaranteed, TargetingNOW, AudienceNOW, and audience-deficiency order o , respectively.

L_s, R_s, C_s, P_s = Length, rate, product conflict, and break position, associated with spot s , respectively.

V_o = Deal CPM associated with order o .

f_o = the time separation required among spots from order o .

$\Lambda_i, A_{di}^D, A_{ti}^T$ = Length, estimated demographic- d impressions, and estimated target- t impressions associated with inventory bucket i , respectively.

M_c = Maximum number of spots of product conflict c that can air in a single break.

F_{oh} = Desired proportion of the total spots in order o to be aired on day h .

Π = Penalty per unit for deviations from the daily goals of spots.

$\Omega^D, \Omega^T, \Omega^A, \Omega^U, \Omega^F$ = Weights to trade off the dollar value of spots in demographic-guaranteed, TargetingNOW, AudienceNOW, audience-deficiency, and filler orders, respectively.

• Decision Variables

x_{si} = Boolean variable, which equals 1 if spot s is scheduled in inventory bucket i and equals 0 otherwise.

z_s = Boolean variable, which equals 1 if spot s is infeasible to be scheduled and equals 0 otherwise.

$\delta_{oh}^+, \delta_{oh}^-$ = Positive and negative deviations from the goal of the number of EQ30s in order o to be aired on day h .

• Auxiliary Variables

a_o^D = Total scheduled demographic impressions on spots in demographic-guaranteed or audience-deficiency order o .

a_o^T = Total scheduled target impressions on spots in TargetingNOW or AudienceNOW order o .

σ_o = Dollar value of scheduled spots in order o .

e_o = Number of scheduled EQ30s in order o .

e_{oh} = Number of EQ30s in order o scheduled on day h .

• Formulation

Maximize

$$\begin{aligned} & \Omega^D \sum_{o \in \mathcal{O}^D} \sigma_o + \Omega^T \sum_{o \in \mathcal{O}^T} \sigma_o + \Omega^A \sum_{o \in \mathcal{O}^A} \sigma_o + \Omega^U \sum_{o \in \mathcal{O}^U} \sigma_o \\ & + \Omega^F \sum_{o \in \mathcal{O}^F} \sigma_o - \Pi \sum_{o \in \mathcal{O}, h \in \mathcal{H}} (\delta_{oh}^+ + \delta_{oh}^-) \end{aligned} \quad (\text{A.O3})$$

Subject to

$$a_o^D = \sum_{s \in \mathcal{S}_o, i \in \mathcal{I}_s} \left(\frac{L_s}{30} \right) A_{\Delta_s i}^D x_{si}, \forall o \in \mathcal{O}^D \cup \mathcal{O}^U \quad (\text{A.37})$$

$$a_o^T = \sum_{s \in \mathcal{S}_o, i \in \mathcal{I}_s} \left(\frac{L_s}{30} \right) A_{\mathcal{T}_s i}^T x_{si}, \forall o \in \mathcal{O}^T \cup \mathcal{O}^A \quad (\text{A.38})$$

$$\sigma_o \leq V_o a_o^D, \forall o \in \mathcal{O}^D \cup \mathcal{O}^U \quad (\text{A.39})$$

$$\sigma_o \leq V_o G_o^D, \forall o \in \mathcal{O}^D \quad (\text{A.40})$$

$$\sigma_o \leq V_o G_o^U, \forall o \in \mathcal{O}^U \quad (\text{A.41})$$

$$\sigma_o \leq V_o a_o^T, \forall o \in \mathcal{O}^T \cup \mathcal{O}^A \quad (\text{A.42})$$

$$\sigma_o \leq V_o G_o^T, \forall o \in \mathcal{O}^T \quad (\text{A.43})$$

$$\sigma_o \leq V_o G_o^A, \forall o \in \mathcal{O}^A \quad (\text{A.44})$$

$$\sigma_o = \sum_{s \in \mathcal{S}_o, i \in \mathcal{I}_s} R_s x_{si}, \forall o \in \mathcal{O}^F \quad (\text{A.45})$$

$$e_o = \sum_{s \in \mathcal{S}_o, i \in \mathcal{I}_s} \left(\frac{L_s}{30} \right) x_{si}, \forall o \in \mathcal{O} \quad (\text{A.46})$$

$$e_{oh} = \sum_{s \in \mathcal{S}_o, i \in \mathcal{I}_s \cap \mathcal{I}_h} \left(\frac{L_s}{30} \right) x_{si}, \forall o \in \mathcal{O}, \forall h \in \mathcal{H} \quad (\text{A.47})$$

$$z_s + \sum_{i \in \mathcal{F}_s} x_{si} = 1, \forall s \in \mathcal{S} \quad (\text{A.48})$$

$$\sum_{s \in \mathcal{F}_i} L_s x_{si} \leq \Lambda_i, \forall i \in \mathcal{I} \quad (\text{A.49})$$

$$\sum_{i \in \mathcal{F}_b, s \in \mathcal{F}_i : P_s = j} x_{si} \leq 1, \forall b \in \mathcal{B}, \forall j \in \mathcal{P} \quad (\text{A.50})$$

$$\sum_{i \in \mathcal{F}_b, s \in \mathcal{F}_i : C_s = c} x_{si} \leq M_c, \forall b \in \mathcal{B}, \forall c \in \mathcal{C} \quad (\text{A.51})$$

$$\sum_{s \in \mathcal{F}_o, i \in \mathcal{F}_s \cap \mathcal{F}_n} x_{si} \leq 1, \forall o \in \mathcal{O}, \forall n \in \mathcal{N}_{f_o} \quad (\text{A.52})$$

$$\delta_{oh}^+ - \delta_{oh}^- = F_{oh} e_o - e_{oh}, \forall o \in \mathcal{O}, \forall h \in \mathcal{H} \quad (\text{A.53})$$

$$x_{si}, z_s \in \{0, 1\}, \quad \delta_{oh}^+, \delta_{oh}^- \in \mathbb{R}$$

The objective function (A.03) maximizes the total weighted value of the spots scheduled minus a penalty for the total deviations from the daily unit goals by order. The weights on the value of each type of spot (i.e., demographic-guaranteed, TargetingNOW, AudienceNOW, audience deficiency, and filler) reflect the relative business importance of these classes of spots.

Constraint (A.37) quantifies the total scheduled demographic impressions from a demographic-guaranteed or audience-deficiency order by adding the corresponding demographic equalized audiences of all its scheduled spots, whereas Constraint (A.38) quantifies the total scheduled target impressions from a TargetingNOW or AudienceNOW order by adding the corresponding targeted equalized audiences of all its scheduled spots.

Constraints (A.39)–(A.45) quantify the total dollar value of the spots scheduled from different order types. The dollar value of the spots scheduled from a demographic-guaranteed or audience-deficiency order is the minimum of the gross value of the scheduled demographic impressions (i.e., total demographic impressions, in thousands, times deal CPM) and the gross value of the impression delivery goal for the order. That is, $\sigma_o = \min\{V_o a_o^D, V_o G_o^D\}$ or $\sigma_o = \min\{V_o a_o^D, V_o G_o^U\}$, depending on the order type. Constraints (A.39)–(A.41) linearize these expressions. Although this linearization expands the feasible region, because of the maximization objective, at optimality σ_o will equal the minimum of the corresponding two quantities. Similarly, the dollar value of the spots scheduled from a TargetingNOW or AudienceNOW order is the minimum of the gross value of the scheduled target impressions (i.e., total target impressions, in thousands, times deal CPM) and the gross value of the impression delivery goal for the order. That is, $\sigma_o = \min\{V_o a_o^T, V_o G_o^T\}$ or $\sigma_o = \min\{V_o a_o^T, V_o G_o^A\}$. Constraints (A.42)–(A.44) linearize these expressions. Constraint (A.45) quantifies the total dollar value of the spots scheduled from a filler order as the sum of the rates charged for the spots scheduled.

Constraints (A.46) and (A.47) count the number of EQ30s scheduled from a given order across all days and by day, respectively. In addition, 30-second spots count as 1 unit, 15-second spots count as 0.5 units, etc. Constraint (A.48) establishes that a spot is either infeasible to be scheduled or it is scheduled in one of its eligible inventory buckets.

Constraints (A.49)–(A.52) honor capacity limitations in the schedule. Constraint (A.49) states that the total duration of all

of the spots scheduled in an inventory bucket cannot exceed the bucket length. Constraint (A.50) establishes that for every commercial break, at most one spot per position can be scheduled in the first, second, second-to-last, and last positions. Constraint (A.51) enforces that the total number of same-product-conflict spots scheduled in a commercial break cannot exceed the maximum per break for that product conflict. Constraint (A.52) enforces time-separation restrictions among spots from the same order.

Finally, Constraint (A.53) quantifies the positive and negative deviations from the daily EQ30 goals by order. As we explain in the proposal-building reformulation, this constraint plus the objective function terms $\delta_{oh}^+ + \delta_{oh}^-$ linearize the absolute value of the quantified deviations. Further details on absolute-value linearizations can be found in Bertsimas and Tsitsiklis (1997).

References

- Araman V, Popescu I (2010) Media revenue management with audience uncertainty: Balancing upfront and spot market sales. *Manufacturing Service Oper. Management* 2(12):190–212.
- Aversano D (2017) Stop the insanity: Why TV audience targeting matters right now. Adage (October 16), <http://adage.com/article/turner-ignite/stop-insanity-tv-audience-targeting-matters/310773/>.
- Bai R, Xie J (2006) Heuristic algorithms for simultaneously accepting and scheduling advertisements on broadcast television. *J. Inform. Comput. Sci.* 1(4):245–251.
- Bates D, Mächler M, Bolker B, Walker S (2015) Fitting linear mixed-effects models using lme4. *J. Statist. Software* 67(1):1–48.
- 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 M, Mallik S (2004) Scheduling commercial videotapes in broadcast television. *Oper. Res.* 52(5):679–689.
- Bollapragada S, Cheng H, Phillips M, Scholes M, Gibbs T, Humphreville M (2002) NBC's optimization systems increase its revenues and productivity. *Interfaces* 32(1):47–60.
- Brusco MJ (2008) Scheduling advertising slots for television. *J. Oper. Res. Soc.* 59(10):1363–1372.
- Carbajal JA, Chaar W (2017) Turner optimizes the allocation of audience deficiency units. *Interfaces.* 47(6):518–536.
- CISION (2011) Bloomberg selects Invision, Inc.'s dealmaker crossroad system. PR Newswire (December 13), <https://www.prnewswire.com/news-releases/bloomberg-selects-invision-incs-dealmaker-crossroad-system-135506633.html>.
- Colman D (2013) Watch the first commercial ever shown on American TV, 1941. *Open Culture* (August 2), <http://www.openculture.com/2013/08/watch-the-first-commercial-ever-shown-on-american-tv-1941.html>.
- 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. Forecast.* 28(3):607–622.
- Danaher PJ, Dagger TS, Smith MS (2011) Forecasting television ratings. *Internat. J. Forecast.* 27(4):1215–1240.
- D'Orazio M, Di Zio M, Scanu M (2019) Statistical matching of data from complex sample surveys. Working paper. Available at: https://www.researchgate.net/publication/266889507_Statistical_matching_of_data_from_complex_sample_surveys.

- Encyclopedia Britannica (2018) A.C. Nielsen. Retrieved February 1, 2018, <https://www.britannica.com/biography/A-C-Nielsen>.
- Gaur DR, Krishnamurti R, Kohli R (2009) Conflict resolution in the scheduling of television commercials. *Oper. Res.* 57(5):1098–1105.
- Gelman A, Hill J (2006) *Data Analysis Using Regression and Multilevel/Hierarchical Models* (Cambridge University Press, New York).
- 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 J (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.
- Nielsen (2009) Introduction to Nielsen data fusion. Retrieved February 1, 2018, <http://www.nielsen.com/content/dam/corporate/us/en/docs/solutions/Nielsen-Introduction-to-Data-Fusion.pdf>.
- Nielsen (2017) TV ratings. Retrieved February 1, 2018, <http://www.nielsen.com/us/en/solutions/measurement/television.html>.
- Pereira PA, Fontes FA, Fontes DB (2008) A decision support system for planning promotion time slots. Kalcsics J, Nickel S, eds. *Operations Research Proceedings 2007* (Springer, Berlin), 147–152.
- Popescu D, Crama P (2015) Ad revenue optimization in live broadcasting. *Management Sci.* 62(4):1145–1164.
- Popescu D, Seshadri S (2016) Managing revenue from television advertising sales. Working paper, INSEAD, New York.
- Quality Logo Products (2018) History of TV ads. Retrieved February 1, 2018, <https://www.qualitylogoproducts.com/promo-university/history-of-tv-ads.html>.
- Reddy S, Aronson J, Stam A (1998) SPOT: Scheduling programs optimally for television. *Management Sci.* 44(1):83–102.
- Slefo G (2017) Desktop and mobile ad revenue surpasses tv for the first time. Adage (April 26), <http://adage.com/article/digital/digital-ad-revenue-surpasses-tv-desktop-iab/308808/>.
- Spangler T (2010) Invision to resell Turner's crossroad traffic system. Multichannel (November 5), <https://www.multichannel.com/news/invision-resell-turners-crossroad-traffic-system-291051>.
- Stephens M (2018) History of television. Retrieved February 1, 2018, <https://www.nyu.edu/classes/stephens/History%20of%20Television%20page.htm>.
- van der Putten P, Kok JN, Gupta A (2002) Data fusion through statistical matching. Working paper, MIT Sloan, Cambridge, MA.
- Williams HP (2013) *Model Building in Mathematical Programming* (Wiley, West Sussex, UK).
- Zhang X (2006) Mathematical models for the television advertising allocation problem. *Internat. J. Oper. Res.* 1(3):302–322.
-
- José Antonio Carbajal** is the head of analytics, data & decision sciences at iHeartMedia, Inc. He has applied his expertise in optimization and advanced analytics on various scheduling, revenue management, resource allocation, pricing, and forecasting problems. His work has been recognized with awards from INFORMS and the Association of National Advertisers. He holds an MS in operations research and a PhD in industrial engineering from the Georgia Institute of Technology.
- Peter Williams** is currently vice president of analytics at TCGPlayer. He has been extensively recognized for his work in advertising and media with awards from organizations like INFORMS and the Association of National Advertisers, and he is the author of several patents. Peter earned an MBA from the Rotterdam School of Management, Erasmus University, and is currently a graduate student in the School of Mathematics at the Georgia Institute of Technology.
- Andreea Popescu** is a director of operations research, leading the statistics and predictive analytics team at Turner. She is the founder of ATLytiCS, a volunteer based nonprofit organization that leverages Atlanta area analytics talent and resources to benefit community betterment causes. Andreea attained her PhD in optimization from Georgia Institute of Technology in Atlanta and her MS in artificial intelligence and pattern recognition from the Polytechnic Institute in Bucharest.
- Wes Chaar** is the chief data & analytics officer at Catalina. Wes attained his PhD from the University of Texas at Austin. He has more than 20 years of expertise in the analytics, data, and decision sciences fields. He applied his skills to work on revenue management, pricing, channel distribution, consumer choice modeling, scheduling, etc. His work covers the airlines, hospitality, online retail, marketing, and media sectors, where he has held various research, consulting, and leadership positions.