Scheduling Spots on Television

Sridhar Seshadri
Indian School of Business, Hyderabad, India 500 032, sridhar_seshadri@isb.edu

Sriram Subramanian
RSG Media Systems, Inc., New York, NY 10016, sriram.subramanian@rsgmedia.com

Sebastian Souyris
McCombs School of Business, The University of Texas at Austin, Austin, TX 78712, sebastian.souyris@utexas.edu

Draft: June 17, 2015

The scheduling of advertisements, or *spots*, is an essential operational process of the television business that must be conducted daily. An efficient distribution of viewers among advertisers allows the television network to satisfy contracts as planned and also increase ad-sales revenue. Spot scheduling is a very hard multi-period scheduling problem. Schedules have to be created such that advertiser's campaign goals are met and ad-sales revenue is maximized. Each campaign has a specific target group of viewers and a unique set of constraints that have to be met. In addition, the number of viewers is uncertain. In this paper, we describe a practical approach that combines mathematical programming and time series methods to create daily schedules that are ready for broadcasting. This approach generates high quality schedules, according to standard business metrics and in comparison with the mathematical optimal bound. Our methods are used by leading networks and they produce substantial increases in revenue.

Key words: scheduling, optimization, advertising, television–business.

---

1. Introduction

In this paper, we address a revenue optimization problem faced by a cable television network with nationwide viewership in the United States. The term *cable television network* refers to an entity that offers programming on a single channel. This term is usually shortened to *cable network* or simply *network*. A major source of revenue for a network is the placement of advertisements, or *spots*, within its programming. The contracts between the advertisers and the network require that the network delivers a target viewership—specifically, a number of impressions of a particular demographic during the course of the advertising campaign. The demographics that advertisers seek are defined according to two dimensions: gender and age group. For example, *M18-34* refers to males between the ages of 18 and 34, *F18-34* refers to females between the ages of 18 and 34, and *P18-34* refers to persons (males and female) between the ages of 18 and 34. The contracts also specify the number of spot-seconds used to deliver the required number of impressions. The contracted number of spot-seconds must be delivered even if the target viewership has already been achieved. Advertisers also require the network to deliver the target demographic viewership according to an agreed upon pace during the course of the campaign. The contract between the advertiser and the network is referred to as a *deal*. The particular type of deal described here is referred to as *guaranteed* deal because the number of spots must be telecast and the target viewership must be achieved as promised. In addition, the advertiser places *constraints* on scheduling spots, which could include requirements that spots must air at a specified day of the week and time of day, requirements that spots are shown during specific programs, or requirements that spots air in the first or last position of the commercial break. These guaranteed deals comprise most of the revenue for a typical North American cable network. There are also deals that are not guaranteed. For these types of deals, the network is allowed to decide whether to schedule the spots, and no viewership targets are required. In this case, the advertiser pays the network on a per–telecast spot basis.

The network typically sets aside a certain amount of time every hour for airing spots. The network makes this determination partly based on regulations imposed by the government. The network also aims to maintain the right balance between content and spots so as to retain audience engagement. The finite supply of air time available for the network to air spots is referred to as the available *inventory*. From the network's perspective, in order to make the most effective use of their limited amount of inventory, spots must be scheduled within a time slot so that the following goals are met:

(i) The viewership is delivered in the target demographic requested by the corresponding advertiser at the required pace, which is the rate at which spots must be scheduled (e.g., the number of spots per day). In addition, the viewership must not exceed the number of impressions requested by the advertiser, because impressions beyond the contracted quantity do not generate additional revenue.

(ii) The viewership that is delivered in any demographic outside the target must be minimal, as these impressions are also not monetized.

The inventory is further segmented into short time intervals, or *commercial breaks*, that are distributed across many programs scheduled to air. The viewership for the break usually depends on the program containing the break, the time of day, and the day of the week. The projections for national viewership accepted by networks and advertisers are those issued by Nielsen Holdings N.V., referred to as Nielsen ratings. These ratings are published a few days after programs are broadcast. Therefore, one sub-task is to forecast the viewership based on ratings data from the past and schedule spots using the forecast so that the goals defined above are met. The quality of the schedule is obviously dependent on the quality of the forecast.

For each day, the qualifying spots are scheduled to air within the breaks to produce a *log*, or a 24–hour (one-day) schedule with programs interspersed among the breaks that contain spots. To schedule these effectively, the decision-maker must consider several other factors in addition to the scheduling constraints requested by the advertiser. At any given point in time, each deal has already delivered some of the target impressions and has a specific number of spots to run, a number of impressions to deliver, and a number of days remaining until completion. Some deals may exceed the required pace (i.e., they are over-performing) and some may fall bellow the required pace (i.e., they are under-performing). The network's ability to correct any mismatch will depend on the time remaining on the deal. The advertiser pays only for the contracted number of impressions: Any impressions beyond this number would be a service that the network has given away, and any shortfall of impressions requires the network to pay a penalty, either in the form of additional spots or in cash. Therefore, it is important for the network to ensure that the deals track as closely as possible to the required number of impressions.

Achieving the promised viewership is central to the business model of the networks. Nevertheless, existing software systems can only schedule spots so that advertiser's scheduling constraints are met. The network's software schedules logs up to a year in advance, but this software does not take into account the current state of the deal or the expected viewership of the breaks in which the software schedules the spots. Therefore, at the outset, the decision-makers realized that there was significant scope for improving on the current solution. One complication in the current process is the fact that the log for the following day is being developed until almost the day immediately preceding the airing, due to changes in the programming schedule, changes in the layout and duration of breaks, and changes in the advertiser's requirements. The availability of the corresponding videos for the spots to be scheduled is often confirmed only the preceding day. Therefore, the decision-makers considered it practical to optimize the schedule of spots only for the next day's log and not any further. Of course, they desired that the performance of deals remained steady over the contracted duration of the campaign, which means that the solution should take into account past performance and the time remaining for each deal when scheduling each day.

In this paper, we present a practical approach for solving the spots scheduling problem described above. The scheduling is achieved in three stages. We start in Stage 1 with a multi-day optimization model that takes the aggregate demand for spots as input to generate weights for the deals in play. These weights indicate the relative importance of deals after taking into account the number of impressions still due to be delivered for the deal during the remaining lifetime of the campaign. The more the deal is lacking in the required pacing (i.e., the more significant amount of impressions pending relative to the time left on the deal) the higher the weight of the deal. These weights are used in a single-day optimization model that assigns spots to breaks so as to make the best use of the viewership, given the existing inventory and the available spots. This is done so as to deliver a higher share of the impressions to the deals with larger weights. For computational efficiency, we split the single-day optimization model that schedules spots in breaks in two stages. In Stage 2, the spots are allocated to the breaks. Then, the precise positioning of the allocated spots within each break is completed in Stage 3. This approach generates a schedule for one day at a time, preferably at the latest possible point in time so that the layout of programs and commercial breaks (i.e., time and durations) as well as the set of spots available to be scheduled are as close as possible to the final version. Delaying the scheduling until the latest possible point in time has the additional benefit of generating the best possible forecast for viewership ratings. These models are supported by a ratings forecasting model developed in conjunction with the optimization model.

We mainly focus on describing the solution approach and the mathematical modeling. This includes several modeling ideas and features that can guide similar efforts. However, we also touch upon several practical and business considerations and describe how these are managed in practice. As with any first time implementation, one measure of success is the potential for revenue generation. In this regard, this application, which is being used daily at leading networks, is providing an increase in revenue of the magnitude of tens of millions dollars annually. The other measure of success is whether the application has led to more problems for analysis and solution. In that respect, several new problems were identified and are currently being solved to improve both the quality of the solutions as well as the analysis of the results for use in the ad-sales process.

The organization of the paper is as follows: Section 2 describes the related academic literature. Section 3 gives an overview of the problem and the solution approach. Section 4 describes in detail the optimization models. Section 5 specifies the solution strategy. Section 6 briefly describes the statistical models used to forecast the ratings. Section 7 presents numerical results, and Section 8 concludes.

2. Literature Review

Advertising on television is a multi-billion dollar industry in which revenue management (RM) techniques can have a considerable effect on the bottom line. Despite this, as shown Talluri and Van Ryzin (2005), only a few cases of RM applications in the television industry had been documented in the operations literature.

However, in recent years there has been a growing need in the operations management community for applications in the television industry. For a full description of the television industry, see Blumenthal and Goodenough (2006). To classify the literature, we have divided out the cable network business processes broadly into four process. Each of these process has significant opportunities for yield management and has attracted the attention of researchers over time. We describe the four processes and the most important contributions found in the operations management and related literature.

Planning of creative goods: Creative goods help to attract the audience that can be sold to advertisers. Therefore, a thoughtful generation of this material is essential and is at the core of the media and broadcasting business. This first process is a complex task that involves creative people, project directors, and network managers who decide which creative goods to air and how to produce them. This process is complex because it involves a large number of agents with sometimes conflicting interests. For example, they may be forced to decide whether to air a reality show, which brings a large audience but only provides nothing but entertainment, or air a cultural show, which could have a smaller audience but provide valuable information. Also, there is high uncertainty in various types of programs. For example, there is significant uncertainty associated with predicting the viewership of a new season of a popular show or the viewership of a new show which is believed to have significant audience potential. Despite its importance, this process of planning creative goods has not received attention from the management science literature. Caves and Guo (2009) review this first process in detail.

Scheduling shows: The objective of this process is to organize the network's schedule of shows to maximize its audience. For marketers, two important measures are *reach* and *frequency*, which quantify how many different viewers view a specific commercial and how many times a particular viewer see the commercial, respectively (Headen et al. (1977)). Hence, to have projections on how many viewers would watch a show is as important as having projections on how many viewers would watch two different shows. Goodhardt and Ehrenberg (1969) find empirically a linear relation between the audience of two different shows and the audience that watches both shows. The authors refer to this common pattern as the *Duplication of Viewing Law*. Several papers contain in-depth studies regarding the shared audience between shows (e.g, Headen et al. (1979), Henry and Rinne (1984), Webster (1985), Rust et al. (1986), Danaher (1991), and Lees and Wright (2013)). Horen (1980) and Reddy et al. (1998) use linear regression to predict the audience based on past ratings and, in a second step, propose integer programming models to schedule shows to maximize the total audience. Rust and Alpert (1984), Rust and Eechambadi (1989) and Rust et al. (1992) develop heuristics that take segmentation into account using a viewer choice model. In a different stream of research, Cancian et al. (1996) study the Nash equilibrium allocation of competitor channel shows when viewers are distributed in a Hotelling fashion with respect to air time. Kelton and Schneider Stone (1998) also study a competitive environment, but in a more realistic setting using industry data to forecast audience with a linear regression model in which each competitor uses an integer programming model to schedule shows. Danaher and Mawhinney (2001) use a choice model to estimate the audience for a given show schedule. The authors design several show schedules with potentially large audiences by using heuristics rules, and then choose the show schedule with the larger audience according to the choice model. Goettler and Shachar (2001) use an empirical structural estimation approach to study the scheduling strategies used by the major television networks in the US, and they conclude that the actual show schedules are very similar to the optimal schedules predicted by their model of competition. To answer a different question, Wilbur et al. (2008) empirically investigate the effect of advertisement time upon audience volume and the preference of viewers and advertisers according to the type of show. Their results can be used as input to schedule shows.

Inventory sales: The sales of audience inventory occurs in two periods (Phillips and Young 2012), the *upfront market* and the *scatter market* (*forward* and *spot* markets in finance parlance). The upfront market takes place approximately four months before the broadcast season, when the networks publicize the show schedule. At this time, the advertisers buy rating points that are guaranteed. All the remaining inventory that is unsold during the upfront market (about 20%–40%), is sold on the scatter market during the broadcasting season. To sell inventory during the upfront market, it is necessary to have projections of the ratings inventory and the advertisers demand. Bollapragada et al. (2002) develop a decision support system based on integer programming and heuristics to support the sales process at the NBC's television network. In subsequent work devoted to advertiser demand forecasting, Bollapragada et al. (2008) implement a combination of the Delphi and Grass Roots techniques to forecast ratings. Bollapragada and Mallik (2008) propose a chance constrained model to decide how much inventory to sell in each market period. Their model maximizes the expected revenue, subject to a lower bound of the probability of generating revenue greater or equal to a given revenue threshold. Zhang (2006) combines the two markets into one, models it as a winner determination problem, and solves it using column generation. A similar formulation is presented and heuristics are developed in Kimms and Muller-Bungart (2007). Araman and Popescu (2010) define the allocation problem of stochastic ratings between upfront and scatter markets as the *media revenue management* or capacity planning problem, and formulate it as a random yield multiple lot-sizing in production to order system (Grosfeld-Nir and Gerchak 2004). The authors obtain structural results that define the optimal capacity allocation depending on the contract parameters, audience, and time. Popescu and Seshadri (2013) study a similar but complementary problem. The authors assume a deterministic audience per show, and a stochastic arrival of advertisers per demographic who are willing to pay a price dependent on the demographic and the final allocation. Their formulation is inspired by classic network revenue management problems (i.e., the model determines the optimal allocation of audience to advertiser as time goes by). Using a different approach, Banciu et al. (2010) examine network bundling strategies to sell to advertisers. These bundles can be composed of different demographic audiences or air times.

Spots scheduling: Once the content (i.e., the programming) is acquired and scheduled, and after the deals have been signed for spots, the spots must be scheduled at the right times to make the most effective use of air time. Our paper deals with this process. The first paper on this topic is by Bollapragada et al. (2004), who solve the scheduling problem of programming a given set of commercials of the same duration that should be aired a pre-specified number of times as uniformly as possible with respect to break position. The authors formulate the model as a network flow problem with a nonlinear loss objective function, and they construct ad-hoc heuristics that can produce good solutions for test instances. For the same problem and test instances, Brusco (2008) improves the solution time and optimality gap using a specialized branch and bound procedure and a simulated annealing heuristic, and also tests alternative loss functions. Brusco and Singh (2010) incorporate time separability conditions and the possibility of having spots of different durations. Bollapragada and Garbiras (2004) develop heuristic methods for scheduling spots in breaks using an integer programming model with real world conditions such as advertisers preferences for specific positions within the break. This system is used by NBC. Gaur et al. (2009) extend the previous model by adding flexibility to model commercial conflicts and by developing a specialized algorithm. Zhang (2006), as a second step, proposes a quadratic integer program to minimize the deviations from the original schedule generated in the previous process of inventory sales. Their formulation is a simplification of the real world problem and it is separated by shows, which results in one small problem per show that is easy to solve. Our paper differs from previous ones in that we develop a detailed spot schedule to minimize the penalty of under delivery while honoring a myriad of constraints. Our work does not separate per show; we can move spots from one show to another. In fact, several constraints apply for the entire day (and some across days); therefore, the scheduling problem is complex. Several papers are devoted to ratings forecasting; for example, Danaher et al. (2011), Danaher and Dagger (2012), Webster et al. (2013) and citations within. A related problem is scheduling spots for live television, in which the schedule the length of the breaks are uncertain, so the decision to air spots must be made in real time (Crama et al. 2012). Another challenge is posed by digital technology that allows personal TV advertisements (Adany et al. 2012).

3. An Overview of the Problem and Solution Approach

The problem of assigning spots to exact positions in the schedule is very challenging and fraught with practical issues, incomplete information and uncertainty. In any given day multiple deals are competing for spots in the schedule. Some deals have been running for a long time and have either been under-performing or over-performing. Some deals are due to end soon (i.e., next week) whereas some will not expire for some time (i.e., this quarter). Furthermore, the creative work for the spots are not ready until the day prior to airing and their relative lengths, specific requirements as well as their restrictions are unknown for future schedules. Finally, the ratings of programs are not known. Forecasts are imprecise and become less precise with increasing time lengths. Mathematically, it is possible to cast the entire problem, with all its associated uncertainty, into a very large dynamic program and solve it daily. However, solving the large program is impossible given today's state of the art machines and software within the available budget. Instead we chose to focus on the day-to-day scheduling problem. Even that problem was broken into three stages due to running time considerations. Figure 1 shows the schematic processes and information flow. The sequence adopted is as follows:

Stage 1: Estimate Deal Weights. Every week, we solve a model to determine the relative weights for deals. The inputs to this model are (i) all the contractual details of the deals for which spots are under consideration for inclusion in the coming week, (ii) how the deals have performed relative to their target audience, (iii) the number of weeks remaining in the deals, and (iv) information regarding spots yet to be scheduled for each deal. The output from this stage comprises the deal weights that are used in the next two stages to produce the daily schedule. The deal weight reflects the relative price per impression given to the deal in the desired demographic.

Stage 2: Schedule Spots in Breaks. The network constructs the details of the schedule showing the minute-by-minute airing of shows, advertisements and promotions, or logs, until nearly the evening prior to airing the schedule. These logs are created manually and provide a feasible schedule that finally is optimized. The decision to continue the practice of manually scheduling the log was made based on several considerations. The management did not want the entire schedule to be managed by a machine. The schedulers preferred to decide which spots to air and whether to give favored slots to certain deals. Moreover, the log is developed several weeks prior to airing. In the initial versions, the log is quite incomplete and filled with generic information or place-holders that become more specific as the date of airing approaches. Therefore, it is practical to allow the schedulers to develop the log gradually and deliver it for optimization at the final step. The log is downloaded from the production system on to the optimization server. Several pre-checks are performed to determine data accuracy as well as automatic detection of rule violations by schedulers who may occasionally relax rules to satisfy a client requirement or make a human decision to override the constraint. These violations are frozen in the final schedule. The initial schedule as well as the final schedule produced by the optimization are vetted by a traffic system. This system has a scheduling engine which automatically checks most of the rules imposed by the network and, more importantly, checks the rules imposed by government regulators. Therefore, the initial feasibility is guaranteed unless the scheduler has revised the schedule and has overridden the scheduling engine. This "close-to-final" schedule is optimized in Stage 2 to assign spots to breaks.

Stage 3: Arrange Spots in Break Positions. This model arranges the spots inside the breaks to which they were assigned in the previous stage. To the user, Stages 2 and 3 happen simultaneously and are not visible as separate steps. Afterward, the optimized schedule is pushed back into the log of the network using integration software. The traffic system verifies the schedule suggested by the optimizer. The human scheduler may make further changes until he or she is satisfied and then pushes the schedule to the production system. The production system sends the schedule for broadcast. The ads are only one component of what is finally aired. The actual programs, the creative for the ads, the promotional advertisements for the network's own programs, the local TV station content, and the local advertisements are all assembled before reaching the subscriber.

The resulting schedule and revenue gains are scrutinized on a regular basis. The metrics of performance are both quantitative and qualitative. The evaluation frequencies are daily, weekly, and monthly. Some sample metrics are:

- The revenue difference between the original schedule and the optimized schedule helps measure the lift in advertisement revenue once the actual ratings are obtained. The accounting is straightforward and the details are omitted.

- Other metrics are used, such as the distribution of spots of an advertiser or an agency by time, the day-to-day performance of deals, the use of prime time spots, the analysis of the source of revenue lift, and the forecast performance.

Figure 1 **Processes and information flow.**

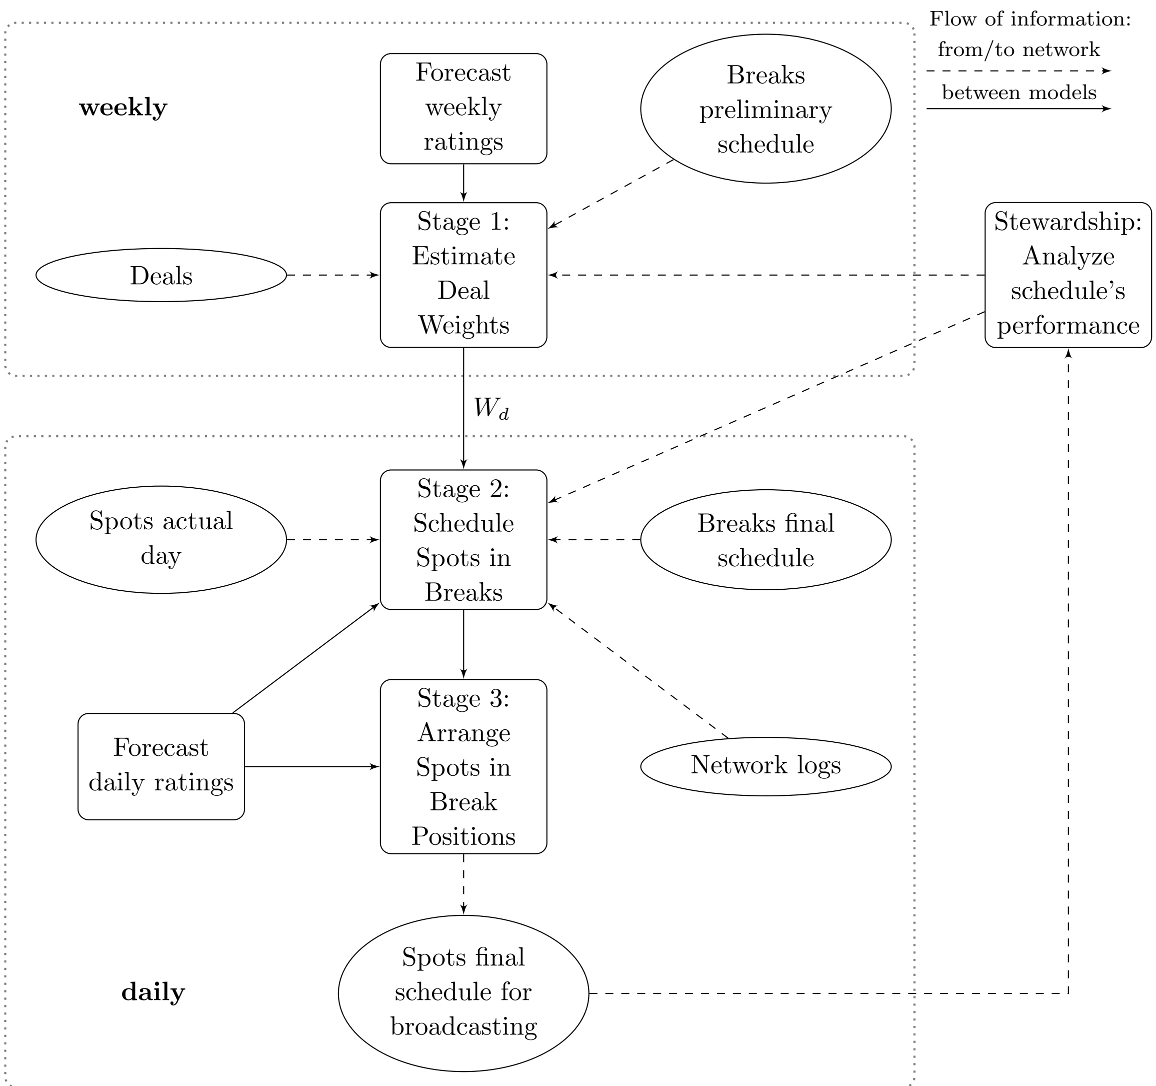

The management of the performance of deals to monitor the quality of the schedule is termed *stewardship*. The stewardship team uses the above information to analyze the effects of deal positions and the effects of the number of spots released for production on deal performance. The knowledge gained is used to fine-tune Stages 1 and 2. For example, the network may add additional spots to remediate the condition of the deal, optimization can increase weights on these deals to improve delivery, or more prime time spots can be reserved for these deals.

4. Three Stages for Scheduling Spots

As discussed in Section 3, we carry out the daily scheduling of spots into breaks, and positions within breaks, in three stages. The first stage estimates the deals' *weights*, that is, the relative importance of the deals for the next week according to the remaining number of target impressions; and the cost per thousand impressions (CPM) of the deals' target demographics. The demographic CPM is the cost of reaching 1,000 viewers the specified audience demographic, a standard measure used by the television industry. Stage 2 allocates the spots to breaks per day. This model schedules spots into the breaks on the actual day, so that the deals assigned higher weights in Stage 1 are given a larger share of the impressions, both in terms of number of spots and preferred time zones per day. Finally, Stage 3 schedules the spots to their actual positions within the breaks on the actual day. Any unscheduled spots are stored in a repository referred to as the *bin*. Stages 2 and 3 ensure that all the required constraints on the spots are satisfied. In the following subsections, we describe the mathematical formulation of each stage. For easy reference, we present summary tables of the sets and parameters used in each model. Some of the notation is used in more than one stage.

4.1. Stage 1: Estimate Deal Weights

This model is solved at the beginning of every week to define each open deal's relative importance, or *weight*, with respect to other open deals. Basically, the model creates a tentative arrangement of spots into breaks for the next week. The arrangement of spots is tentative because the length of all the video spots are not known, some of the advertiser constraints for the complete week may be not defined, and the timetable of breaks may change as the week proceeds. At the moment of solving Stage 1, the available information consists of the following:

- The set of open deals $\mathcal{D}$.
- The set of breaks Θ for the next week. Each break $b \in \Theta$ has a L_b time length.
- The set $\mathcal{Q}$ of demographics. For each demographic $q \in \mathcal{Q}$, the network provides the value for CPM_q, or the cost per thousand impressions based on historical viewership and marketing considerations.
- The deals' specifications given by the advertisers. These specifications define the $q(d) \in \mathcal{Q}$ target demographic of each deal $d \in \mathcal{D}$. Also, the set $\Theta_{\mathcal{D}}(d) \subset \Theta$ of breaks in which deal d can be

shown, for all $d \in \mathcal{D}$; and, for notational convenience, we define the set $\mathcal{D}_\Theta(b) \subset \mathcal{D}$ of deals that can be scheduled in break b, for all $b \in \Theta$. Also, each open deal has a *target impressions status* I_d for the next week, which is the difference between the promised number of impressions to be delivered on the requested demographic and the impressions already delivered for the deal from the start day, divided by the number of weeks until the deal expires. In the industry, the Deal Stewardship system closely monitors delivery and provides targets for the remaining life of the deal. In addition, each deal has a restriction on the maximum number of spots that can be broadcasted, N_d. For the next week, some of the video spots are already recorded, so their lengths are determined; however, others are not yet available. Given this lack of information, in order to create the tentative schedule, we use the historic average length of spots $\bar{H}$ in the model.

- To meet the target impressions status, our algorithm estimates a rating forecast R_{bq}, for each break $b \in \Theta$ and demographic $d \in \mathcal{Q}$, using the models presented in Section 6.

Table 1 presents a summary of the notation described above.

Table 1 Notation Stage 1.

Sets and indices:	
$\mathcal{D}, d$	All open deals.
Θ, b	All breaks of the following week. The schedule of breaks is tentative and can change as the week proceeds.
$\mathcal{Q}, q$	Demographics.
$\mathcal{D}_\Theta(b) \subset \mathcal{D}$	Deals that can be scheduled in break b, for all $b \in \Theta$.
$\Theta_\mathcal{D}(d) \subset \Theta$	Breaks in which deal d can be shown, for all $d \in \mathcal{D}$.
$q(d) \in \mathcal{Q}$	Target demographic for deal d, for all $d \in \mathcal{D}$.
Parameters:	
CPM_q	CPM of demographic q, for all $q \in \mathcal{Q}$.
L_b	Length in seconds of break b, for all $b \in \Theta$.
$\bar{H}$	Average spot length in seconds.
R_{bq}	Ratings forecast of break b and demographic q, for all $b \in \Theta, q \in \mathcal{Q}$.
I_d	Target number of impressions for deal d for the following week, for all $d \in \mathcal{D}$.
N_d	Maximum number of deal d spots planned for delivery in the following week, for all $d \in \mathcal{D}$.
Decision Variables:	
c_{db}	Number of spots of deal d in break b, for all $d \in \mathcal{D}$, and $b \in \Theta_\mathcal{D}(d)$.
e_d	Shortfall decision variable for deal d , for all $d \in \mathcal{D}$.

As explained in above, the schedule cannot be determined exactly because the description of the spots to be scheduled are often incomplete, as several detailed scheduling constraints are missing for spots that are to be aired in the future. Also, the breaks, programs, and lengths are defined but can change significantly over the next few days. Given the available information at this stage, we formulate a linear programming model to generate a tentative multi-day schedule. A set of weights for the day-to-day scheduling problem of Stage 2 is generated as a by-product. The alternative approach would have been to rely upon decision-makers to prioritize deals in a quantitative manner to enable making trade-offs among spots from different deals in Stage 2.

Weight Problem: The decision variables of the Stage 1 model are the c_{db} number of spots of deal d in break b, for all $d \in \mathcal{D}$, and $b \in \Theta_{\mathcal{D}}(d)$; and the e_d shortfall or number of impressions that cannot be delivered of deal d in the case of a lack of forecasted viewership in the target demographic, for all $d \in \mathcal{D}$. We relax the integrality restriction of the number of spots and plan on an aggregate basis at the break level. The Stage 1 linear programming model is as follows:

$$\min \quad \sum_{d \in \mathcal{D}} e_d CPM_{q(d)} \tag{1}$$

$$\text{s.t.} \quad \sum_{d \in \mathcal{D}_\Theta(b)} c_{db} \leq \frac{L_b}{\bar{H}} \qquad \forall b \in \Theta, \tag{2}$$

$$\sum_{b \in \Theta_{\mathcal{D}}(d)} c_{db} \leq N_d \qquad \forall d \in \mathcal{D}, \tag{3}$$

$$\sum_{b \in \Theta_{\mathcal{D}}(d)} c_{db} R_{bq(d)} + e_d \geq I_d \qquad \forall d \in \mathcal{D}, \tag{4}$$

$$e_d, c_{db} \geq 0 \qquad \forall d \in \mathcal{D}, b \in \Theta_{\mathcal{D}}(d). \tag{5}$$

The objective function (1) minimizes the value of under-delivered impressions. Constraint (2) ensures that the number of spots that are delivered in each break is less than or equal to the number of spots available. Constraint (3) ensures that the planned number of spots are allocated as required by the corresponding advertiser of each deal. Constraint (4) links the number of spots delivered, the forecasted ratings in the target demographic, the shortfall of impressions, and the target number of impression for each deal. And constraint (5) is the nature of variables.

Let W_d be the shadow price of constraint (4), for all $d \in \mathcal{D}$. It can be interpreted as the price the network is willing to pay for decreasing the target I_d of deal d by one impression of the target demographic. This equivalent to consider W_d as the price of each deal-d-impression. We use this price to determine the penalty in the single-day problem that is solved in Stage 1 that we present bellow.

4.2. Stage 2: Schedule Spots in Breaks

As described in Section 3, we optimize the schedule created by the schedulers. For the actual day to be scheduled, the network allows the model to move only a given number of spots to avoid disturbing the elegance of the schedule. Moreover, we ignore the scheduling rules within each break to maintain a manageable computational complexity. We focus on optimizing the scheduling at the break level in Stage 3.

Given the Stage 1 deal weights, Stage 2 assigns the spots to breaks in the actual day so that a series of constraints are met, and the ratings per demographic allocated to spots are maximized according to the deal weights. The formulation is a mixed integer linear programing model.

Table 2 presents the sets and the indices, while Table 3 summarizes the parameters used in Stage 2. Table 4 presents the decision variables. Where possible, brief notes are given to make the notation understandable without reference to the model.

Table 2 Sets and Indices used in Stage 2.

Supply sets and indices:	
$\mathcal{B}, b$	Breaks within the actual day to be scheduled.
$\mathcal{H}, h$	Hours. $\mathcal{H} = 1, \ldots, 24$
$\mathcal{B}_{\mathcal{H}}(h) \subset \mathcal{B}$	Breaks in hour h, for all $h \in \mathcal{H}$.
$\mathcal{Q}, q$	Demographics.
Demand sets and indexes:	
$\mathcal{A}, a$	Advertisers.
$\mathcal{K}, k$	Brands.
$\mathcal{D}, d$	Open deals that can be scheduled in at least one of the breaks $b \in \mathcal{B}$.
$\mathcal{S}, i, j$	Spots.
$\mathcal{P}, p$	Product categories.
$\mathcal{S}_{\mathcal{A}}(a) \subset \mathcal{S}$	Spots that belong to advertiser a, for all $a \in \mathcal{A}$.
$\mathcal{S}_{\mathcal{K}}(k) \subset \mathcal{S}$	Spots of brand k, for all $k \in \mathcal{K}$.
$\mathcal{S}_{\mathcal{D}}(d) \subset \mathcal{S}$	Spots that belong to deal d, for all $d \in \mathcal{D}$.
$d(i) \in \mathcal{D}$	Deal to which spot i belongs, for all $i \in \mathcal{S}$.
$\hat{\mathcal{S}}_{\mathcal{D}}(d) \subset \mathcal{S}_{\mathcal{D}}(d)$	Spots of deal d that are *guaranteed* (i.e., must be aired and together must achieve the viewership target), for all $d \in \mathcal{D}$.
$\mathcal{S}_{\mathcal{P}}(p) \subset \mathcal{S}$	Spots that belong to product category p, for all $p \in \mathcal{P}$.
Sets that define constraints:	
$\mathcal{S}_{\mathtt{A}} \subset \mathcal{S}$	Spots that must be scheduled at the first position within a break*.
$\mathcal{S}_{\mathtt{Z}} \subset \mathcal{S}$	Spots that must be scheduled at the last position within a break*.
$\mathcal{S}^2_{\mathtt{Sep}} \subset \mathcal{S} \times \mathcal{S}$	Pair of spots that must be separated by time*.
$\mathcal{S}^2_{\mathcal{A}\mathtt{Sdw}}(a) \subset \mathcal{S}_{\mathcal{A}}(a) \times \mathcal{S}_{\mathcal{A}}(a)$	Pair of spots that must satisfy sandwich constraint* for advertiser a, for all $a \in \mathcal{A}$.
$\mathcal{S}^2_{\mathcal{A}\mathtt{Pig}}(a) \subset \mathcal{S}_{\mathcal{A}}(a) \times \mathcal{S}_{\mathcal{A}}(a)$	Pair of spots that must satisfy piggyback constraint* for advertiser a, for all $a \in \mathcal{A}$.
$\mathcal{S}^2_{\mathcal{A}\mathtt{Con}}(a) \subset \mathcal{S}_{\mathcal{A}}(a) \times \mathcal{S}_{\mathcal{A}}(a)$	Pair of spots that must be scheduled in consecutive breaks* for advertiser a, for all $a \in \mathcal{A}$.
Demand-Supply sets and indices:	
$\mathcal{B}_{\mathcal{S}}(i) \subset \mathcal{B}$	Valid breaks for spot i (i.e., breaks within which the spot can be aired), for all $i \in \mathcal{S}$.
$\mathcal{S}_{\mathcal{B}}(b) \subset \mathcal{S}$	Valid spots for break b, for all $b \in \mathcal{B}$.
$\hat{\mathcal{S}}_{\mathcal{B}}(b) \subset \mathcal{S}_{\mathcal{B}}(b)$	Spots assigned to break b in the original schedule, for all $b \in \mathcal{B}$.
$b(i) \in \mathcal{B}$	Original break to which spot i was assigned, for all $i \in \mathcal{S} \setminus \mathcal{S}_{\mathtt{bin}}$.
$\mathcal{S}_{\mathtt{bin}} \subset \mathcal{S}$	Spots that are in the bin in the original schedule.
$\mathcal{H}_{\mathcal{D}}(d) \subset \mathcal{H}$	Set of hours within which deal d can be shown, for all $d \in \mathcal{D}$.
$q(d) \in \mathcal{Q}$	Target demographic for deal d, for all $d \in \mathcal{D}$.

Note. Latin calligraphic uppercase denotes set. * Constraints explained in the formulation. The notation $|\mathcal{C}|$ refers to the cardinality of a set $\mathcal{C}$.

Table 3 Parameters used in Stage 2.

R_{bq}	Rating forecast in 30 seconds of break b, demographic q, for all $b \in \mathcal{B}, q \in \mathcal{Q}$.
F_b	Start time of break b, for all $b \in \mathcal{B}$.
L_b	Length of break b in seconds, for all $b \in \mathcal{B}$.
$\bar{L}$	Average break length in seconds.
H_i	Length of spot i in seconds, for all $i \in \mathcal{S}$.
W_d	Weight of deal d that was obtained in Stage 1, for all $d \in \mathcal{D}$.
$P^{\texttt{B}}$	Brand separation penalty factor*.
$P^{\texttt{A}}$	Advertiser separation penalty factor*.
$P^{\texttt{U}}$	Vertical uniformity penalty factor*.
D_{ij}	Minimum separation* between spots i and j, for all $(i,j) \in \mathcal{S}^2_{\texttt{Sep}}$.
$M^{\texttt{bin}}$	Maximum number of spots allowed in the bin. $M^{\texttt{bin}} \geq \lvert\mathcal{S}_{\texttt{bin}}\rvert$.
$M^{\texttt{move}}$	Maximum number of spots that may be moved. Policy variable to ensure the original log's "beauty" is preserved.
$M^{\texttt{dev}}_d$	Maximum number of spots of deal d per hour allow to deviate from the average number of spots per hour, for all $d \in \mathcal{D}$.
$M^{\texttt{B}}_{pb}$	Maximum number of spots from the same product category p that can be placed within break b, for all $p \in \mathcal{P}$ and $b \in \mathcal{B}$.

Note. Latin uppercase denotes parameter. *The extended definition of these parameters are presented in the formulation.

Table 4 Decision Variables Stage 2.

Binary decision variables:	
x_{ib}	Equal to 1 if spot i is scheduled in break b, 0 otherwise, for all spots $i \in \mathcal{S}$, and breaks $b \in \mathcal{B}_{\mathcal{S}}(i)$.
y_i	Equal to 1 if spot i is added to the bin, 0 otherwise, for all spots $i \in \mathcal{S}$.
z_{kb}	Equal to 1 if at least one spot of brand k is scheduled in break b, 0 otherwise, for all brands $k \in \mathcal{K}$, and breaks $b \in \mathcal{B}$.
v_{ab}	Equal to 1 if at least one spot of advertiser a is scheduled in break b, 0 otherwise, for all advertisers $a \in \mathcal{A}$, and breaks $b \in \mathcal{B}$.
w_{ij}	Equal to 1 if spot i is scheduled before spot j, 0 otherwise, for all pairs of spots $(i,j) \in \mathcal{S}^2_{\texttt{Sep}}$.
Non negative decision variables:	
α_{kb}	Penalty for scheduling two spots of brand k in consecutive breaks b and $b+1$, for all brands $k \in \mathcal{K}$ and breaks $b \in \mathcal{B}$.
β_{ab}	Penalty for scheduling two spots of advertiser a in consecutive breaks b and $b+1$, for all advertisers $a \in \mathcal{A}$ and breaks $b \in \mathcal{B}$.
γ_{dh}	Penalty for vertical uniformity deviation, of deal d and hour h, for all deals $d \in \mathcal{D}$ and hours $h \in \mathcal{H}_{\mathcal{D}}(d)$.
δ_{dh}	Vertical uniformity deviation of deal d and hour h, for all deals $d \in \mathcal{D}$ and hours $h \in \mathcal{H}_{\mathcal{D}}(d)$.

Note. Latin lowercase denotes binary decision variable. Greek letter denotes non-negative decision variable.

Break Problem: Let $\mathcal{B}$ be the set of breaks on the actual day to be scheduled; let $\mathcal{D}$ be the set of deals with spots that can be scheduled in the day; and let $\mathcal{Q}$ be the set of audience demographics. For each deal $d \in \mathcal{D}$, $q(d) \in \mathcal{Q}$ is the targeted demographic and $\mathcal{S}_{\mathcal{D}}(d)$ is the set of spots that belong to deal d. From this set, only $\hat{\mathcal{S}}_{\mathcal{D}}(d) \subset \mathcal{S}_{\mathcal{D}}(d)$ are guaranteed to be shown on the actual day due to contract considerations. The set of all spots is $\mathcal{S}$. For each spot $i \in \mathcal{S}$, $d(i)$ is the deal to which spot i belongs. Given the deal specifications of $d(i)$, $\mathcal{B}_{\mathcal{S}}(i) \subset \mathcal{B}$ is the set of breaks in which spot i can be scheduled. Reciprocally, $\mathcal{S}_{\mathcal{B}}(b) \subset \mathcal{S}$ is the set of spots that can be scheduled in break b, for each break $b \in \mathcal{B}$. The main decision variable of the model is x_{ib}, a binary indicator of whether spot i is in break b, for each spot $i \in \mathcal{S}$ and break $b \in \mathcal{B}_{\mathcal{S}}(i)$. If a spot i cannot be scheduled in any break because of infeasibility, then it is assigned to the bin. Let y_i be the binary decision variable that indicates bin assignation. Recall that Stage 2 receives an original schedule that can be modified by a given maximum number of moves, namely $M^{\texttt{move}}$. Let $\hat{\mathcal{S}}_{\mathcal{B}}(b) \subset \mathcal{S}_{\mathcal{B}}(b)$ be the set of spots scheduled in break b, for each $b \in \mathcal{B}$, and let $\mathcal{S}_{\texttt{bin}} \subset \mathcal{S}$ be the set of spots that are originally in the bin. Accordingly, let $b(i)$ be the scheduled break of spot i, for each $i \in \mathcal{S} \setminus \mathcal{S}_{\texttt{bin}}$. In this given schedule, all the guaranteed spots are scheduled in a break, that is $\hat{\mathcal{S}}_{\mathcal{D}}(d) \cap \mathcal{S}_{\texttt{bin}} = \emptyset$ for all $d \in \mathcal{D}$. At a demand level, the other three relevant dimensions at this stage are the advertisers, the brands, and the product categories. Let $\mathcal{A}$ be the set of advertisers who sign the deals; and, for each advertiser $a \in \mathcal{A}$, let $\mathcal{S}_{\mathcal{A}}(a) \subset \mathcal{S}$ be the set of spots that belong to advertiser a. Each advertiser can have several brands. Let $\mathcal{K}$ be the set of brands; and, for each brand $k \in \mathcal{K}$, let $\mathcal{S}_{\mathcal{K}}(k) \subset \mathcal{S}$ be the set of brand k spots. Finally, across brands, $\mathcal{P}$ is the set of product categories; and, for each $p \in \mathcal{P}$, $\mathcal{S}_{\mathcal{P}}(p) \subset \mathcal{S}$ is the set of spots that belong to product category p. At a supply level, the day is broken into hours $\mathcal{H} = 1, \ldots, 24$, such that $\mathcal{B}_{\mathcal{H}}(h)$ is the set of breaks to be air in hour h, for all $h \in \mathcal{H}$. Finally, for modeling purposes that will become clear later, let $\mathcal{H}_{\mathcal{D}}(d)$ be the set of hours in which the spots of deal d can be shown for all $d \in \mathcal{D}$. Notice that $\mathcal{H}_{\mathcal{D}}(d)$ can be inferred from the set of spots that belong to deal d and the set of hours in which these spots can be shown. Formally, $\mathcal{H}_{\mathcal{D}}(d) := \bigcup_{i \in \mathcal{S}_{\mathcal{D}}(d), b \in \mathcal{B}_{\mathcal{S}}(i)} \left\{h : \underline{h} \leq F_b \leq \bar{h}\right\}$, where $\underline{h}$ and $\bar{h}$ are the start and end day-minutes of hour h, and F_b is the start day-minute of break b.

Objective function (6) maximizes the allocated ratings minus four penalties. The first term is the sum over deals $d \in \mathcal{D}$, guaranteed spots $i \in \hat{\mathcal{S}}_{\mathcal{D}}(d)$, and breaks $b \in \mathcal{B}_{\mathcal{S}}(i)$ of the allocated ratings in the targeted demographic $q(d)$. This means that the rating per 30 seconds $R_{bq(d)}$ factored by the proportion of this rating that would be assigned to spot i $\frac{H_i}{30}$, multiplied by the allocation variable x_{ib} and factored by the weight of the deal W_d, which is obtained by the weight model presented in Subsection 4.1. The second term is the penalty for allocating guaranteed spots to the bin, which is the sum of spots assigned to the bin, y_i, multiplied by the ratings allocated in the original schedule $\frac{H_i}{30} R_{b(i)q(d)}$, and weighted by W_d. The third term is the sum of decision variable penalties α_{kb} for placing spots of the same brand k scheduled in consecutive breaks b and $b+1$. These penalties are defined in the block of constraints (8). Similarly, the fourth term is the sum of decision variable penalties β_{ab} for placing spots of the same advertiser a scheduled in consecutive breaks b and $b+1$, which are defined in the block of constraints (9). The fifth term is the sum of decision variable penalties γ_{dh} for not having a uniform scheduled distribution of deal d spots at hour h, defined in the block of constraints (10). The parameters $P^{\mathtt{B}}, P^{\mathtt{A}}$ and $P^{\mathtt{U}}$ are input multipliers that calibrate the importance of the penalties in the objective function, and are chosen based on trial and error. The decision variable γ_{duh} is an artifact to model an idea that is somewhat subjective. For example, when an advertiser requires an hour of separation between spots, the advertiser often can tolerate a separation of 59 minutes, but a separation of 45 minutes is less likely to be accepted. The Stage 2 objective function is as follows:

$$\max \sum_{\substack{d \in \mathcal{D}, i \in \hat{\mathcal{S}}_{\mathcal{D}}(d), \\ b \in \mathcal{B}_{\mathcal{S}}(i)}} W_d \frac{H_i}{30} R_{bq(d)} x_{ib} - \sum_{\substack{d \in \mathcal{D}, \\ i \in \hat{\mathcal{S}}_{\mathcal{D}}(d)}} W_d \frac{H_i}{30} R_{b(i)q(d)} y_i - P^{\mathtt{B}} \sum_{\substack{k \in \mathcal{K}, \\ b \in \mathcal{B}}} \alpha_{kb} - P^{\mathtt{A}} \sum_{\substack{a \in \mathcal{A}, \\ b \in \mathcal{B}}} \beta_{ab} - P^{\mathtt{U}} \sum_{\substack{d \in \mathcal{D}, \\ h \in \mathcal{H}_{\mathcal{D}}(d)}} \gamma_{duh} \tag{6}$$

For later reference in Section 7, we label the five terms as metrics M1 to M5. Hence, the Stage 2 objective function is equal to M1 - M2 - M3 - M4 - M5.

Block of constraints (7) is the *network feasibility constraints*: Constraint (7a) imposes that each spot $i \in \mathcal{S}$ is scheduled for exactly one break or goes into the bin. Constraint (7b) restricts for each break $b \in \mathcal{B}$ that the sum of the spots' length, H_i, scheduled in the break should be less than or equal to the break length, L_b. Constraint (7c) imposes that the number of spots that goes into the bin must be less than or equal to $M^{\texttt{bin}}$, the maximum number of spots that are allowed in the bin. Lastly, constraint (7d) establishes that the number of spots not moved from the original schedule must be greater than or equal to the total number of spots minus the maximum number of moves allowed, $M^{\texttt{move}}$. The value $M^{\texttt{move}}$ was initially chosen equal to 10% of spots and increased steadily with growing confidence in the model to 50%. The initial fear was that a machine may be unable to generate an elegant schedule—that is a schedule appreciated by experienced schedulers.

$$y_i + \sum_{b \in \mathcal{B}_{\mathcal{S}}(i)} x_{ib} = 1, \qquad \forall i \in \mathcal{S}, \tag{7a}$$

$$\sum_{i \in \mathcal{S}_{\mathcal{B}}(b)} x_{ib} H_i \le L_b, \qquad \forall b \in \mathcal{B}, \tag{7b}$$

$$\sum_{i \in \mathcal{S}} y_i \le M^{\texttt{bin}}, \tag{7c}$$

$$\sum_{b \in \mathcal{B}, i \in \hat{\mathcal{S}}_{\mathcal{B}}(b)} x_{ib} + \sum_{i \in \mathcal{S}_{\texttt{bin}}} y_i \ge |\mathcal{S}| - M^{\texttt{move}}. \tag{7d}$$

Block of constraints (8) defines the *brand uniform dispersion constraints*: It is desirable that spots of the same brand are separated by at least one break of difference. We impose the following constraint to define penalties measured in terms of ratings times deal weights when this requirement is not satisfied. For each brand $k \in \mathcal{K}$ and break $b \in \mathcal{B}$, (8a) imposes the relation between the binary variable z_{kb}—which indicates whether at least one brand k spot is scheduled in break b—and the total number of brand k spots in break b; (8b) and (8c) define the brand separation penalty α_{kb}, which must be greater than or equal to the sum over the spots that are scheduled in consecutive breaks of the delivered ratings $(\frac{H_i}{30} R_{bq(d(i))})$, weighted by the deal weight W_d.

$$\sum_{i \in \mathcal{S}_{\mathcal{B}}(b) \cap \mathcal{S}_{\mathcal{K}}(k)} x_{ib} \le |\mathcal{S}_{\mathcal{K}}(k) \cap \mathcal{S}_{\mathcal{B}}(b)| z_{kb}, \qquad \forall k \in \mathcal{K}, b \in \mathcal{B}, \tag{8a}$$

$$\alpha_{kb} \ge \sum_{i \in \mathcal{S}_{\mathcal{B}}(b) \cap \mathcal{S}_{\mathcal{K}}(k)} W_{d(i)} \frac{H_i}{30} R_{bq(d(i))} (x_{ib} - 1 + z_{kb+1}), \qquad \forall k \in \mathcal{K}, b \in \mathcal{B} \setminus \{|\mathcal{B}|\}. \tag{8b}$$

$$\alpha_{kb} \ge \sum_{i \in \mathcal{S}_{\mathcal{B}}(b) \cap \mathcal{S}_{\mathcal{K}}(k)} W_{d(i)} \frac{H_i}{30} R_{bq(d(i))} (x_{ib} - 1 + z_{kb-1}), \qquad \forall k \in \mathcal{K}, b \in \mathcal{B} \setminus \{1\}. \tag{8c}$$

Block of constraints (9) defines the *advertiser uniform dispersion constraints*: These are similar to brand uniform dispersion constraints, but these apply to advertisers instead of brands.

$$\sum_{i\in\mathcal{S}_{\mathcal{B}}(b)\cap\mathcal{S}_{\mathcal{A}}(a)} x_{ib} \le |\mathcal{S}_{\mathcal{A}}(a)\cap\mathcal{S}_{\mathcal{B}}(b)|v_{ab}, \qquad \forall a\in\mathcal{A}, b\in\mathcal{B}, \tag{9a}$$

$$\beta_{ab} \ge \sum_{i\in\mathcal{S}_{\mathcal{B}}(b)\cap\mathcal{S}_{\mathcal{A}}(a)} W_{d(i)}\frac{H_i}{30}R_{bq(d(i))}(x_{ib}-1+v_{ab+1}), \qquad \forall a\in\mathcal{K}, b\in\mathcal{B}\setminus\{|\mathcal{B}|\}. \tag{9b}$$

$$\beta_{ab} \ge \sum_{i\in\mathcal{S}_{\mathcal{B}}(b)\cap\mathcal{S}_{\mathcal{A}}(a)} W_{d(i)}\frac{H_i}{30}R_{bq(d(i))}(x_{ib}-1+v_{ab-1}), \qquad \forall a\in\mathcal{K}, b\in\mathcal{B}\setminus\{1\}. \tag{9c}$$

Block of constraints (10) imposes the *time based uniform dispersion constraints*: Advertisers want spots to be as uniformly distributed as possible in each hour. To model this, we define penalties that measure the deviation in the number of spots per hour from the average number of spots scheduled in an hour. For each deal $d\in\mathcal{D}$ and hour $h\in\mathcal{H}_{\mathcal{D}}(d)$, constraints (10a) and (10b) define the difference δ_{dh} over the maximum allowed deviation $M_d^{\mathtt{dev}}$ from the average number of spots per hour. (10c) defines the incurred penalty (γ_{dh}) because of the extra difference.

$$\delta_{dh} + M_d^{\mathtt{dev}} \ge \sum_{\substack{i\in\mathcal{S}_{\mathcal{D}}(d),\\ b\in\mathcal{B}_{\mathcal{S}}(i)\cap\mathcal{B}_{\mathcal{H}}(h)}} x_{ib} - \frac{\sum\limits_{i\in\mathcal{S}_{\mathcal{D}}(d), b\in\mathcal{B}_{\mathcal{S}}(i)} x_{ib}}{|\mathcal{H}_{\mathcal{D}}(d)|}, \qquad \forall d\in\mathcal{D}, h\in\mathcal{H}_{\mathcal{D}}(d), \tag{10a}$$

$$\delta_{dh} + M_d^{\mathtt{dev}} \ge \frac{\sum\limits_{i\in\mathcal{S}_{\mathcal{D}}(d), b\in\mathcal{B}_{\mathcal{S}}(i)} x_{ib}}{|\mathcal{H}_{\mathcal{D}}(d)|} - \sum_{\substack{i\in\mathcal{S}_{\mathcal{D}}(d)\\ b\in\mathcal{B}_{\mathcal{S}}(i)\cap\mathcal{B}_{\mathcal{H}}(h)}} x_{ib}, \qquad \forall d\in\mathcal{D}, h\in\mathcal{H}_{\mathcal{D}}(d), \tag{10b}$$

$$\gamma_{dh} \ge \delta_{dh}W_{d(i)}\frac{\sum\limits_{b\in\mathcal{B}_{\mathcal{H}}(h)} \frac{H_i}{30}R_{bq(d)}}{|\mathcal{B}_{\mathcal{H}}(h)|}, \qquad \forall d\in\mathcal{D}, h\in\mathcal{H}_{\mathcal{D}}(d). \tag{10c}$$

Block of constraints (11) is the *A-position and Z-position constraints*: By advertisement requirements, some of the spots must be scheduled at the first positions of breaks, which in industry parlance are called A-position spots. Let $\mathcal{S}_{\mathtt{A}}\subset\mathcal{S}$ be this set of spots. The positioning inside the break is determined in Stage 3, so Stage 2 must ensure that in each break there must be assigned at most one of the spots of $\mathcal{S}_{\mathtt{A}}$. Constraint (11a) ensures that this happens. Similarly, some of the spots must be scheduled at the end of a break, or, the Z-position spots. Let $\mathcal{S}_{\mathtt{Z}}\subset\mathcal{S}$ be this set of spots. Constraint (11b) imposes that for each break $b\in\mathcal{B}$ only one of these spots can be assigned.

The other specific positions inside the breaks are determined by Stage 3, which is explained in Subsection 4.3.

$$\sum_{i \in \mathcal{S}_{\mathtt{A}} \cap \mathcal{S}_{\mathcal{B}}(b)} x_{ib} \leq 1, \qquad \forall b \in \mathcal{B}, \tag{11a}$$

$$\sum_{i \in \mathcal{S}_{\mathtt{Z}} \cap \mathcal{S}_{\mathcal{B}}(b)} x_{ib} \leq 1, \qquad \forall b \in \mathcal{B}. \tag{11b}$$

Block of constraints (12) defines the *minimum separation constraints*: The contracts define that certain pairs of spot must be scheduled in breaks separated by at least a given amount of time. For example, an advertiser may want to have their spots separated by at least 30 minutes, or may want a separation of 45 minutes from other spots of the same product category, or may want a separation of one hour from spots of a rival brand. Let $\mathcal{S}^2_{\mathtt{Sep}} \subset \mathcal{S} \times \mathcal{S}$ be the set of pair of spot that must be separated; and, for each $(i,j) \in \mathcal{S}^2_{\mathtt{Sep}}$, let D_{ij} be the requested separation time, in minutes. Because at this stage we only schedule at a break level and not at a position inside the break, the constraints must ensure that the start time of the breaks where spots i and j are scheduled must be separated by at least D_{ij} plus the average length of a break in minutes, $\bar{L}$. For that, we define the binary decision variable w_{ij} that is equal to 1 if spot i is scheduled before spot j. Constraints (12a) and (12b) impose that the difference between the start time of the spot that is scheduled later in the day and the spot that is scheduled earlier must be at least $D_{ij} + \bar{L}$, where F_b is the start day-minute of break b.

$$\sum_{b \in \mathcal{B}_{\mathcal{S}}(j)} F_b x_{jb} - \sum_{b \in \mathcal{B}_{\mathcal{S}}(i)} F_b x_{ib} \geq (D_{ij} + \bar{L}) w_{ij} - (24 \times 60)(1 - w_{ij}), \qquad \forall (i,j) \in \mathcal{S}^2_{\mathtt{Sep}}, \tag{12a}$$

$$\sum_{b \in \mathcal{B}_{\mathcal{S}}(i)} F_b x_{ib} - \sum_{b \in \mathcal{B}_{\mathcal{S}}(j)} F_b x_{jb} \geq (D_{ij} + \bar{L})(1 - w_{ij}) - (24 \times 60) w_{ij}, \qquad \forall (i,j) \in \mathcal{S}^2_{\mathtt{Sep}}. \tag{12b}$$

The set $\mathcal{S}^2_{\mathtt{Sep}}$ is defined for three different pairs of spot types: 1) separation of spots from the same advertiser or spots from a specific set of advertisers, 2) separation of spots from the same brand or spots from a specific set of brands, and 3) separation of spots from the same product category or spots from a specific set of product categories. If a pair of spots has multiple separation requests, only the constraint for the maximum separation is created.

Block of constraints (13) is the *association constraints*: These are three constraint types required by the advertisers for specific spot pairs related by their content, so they must the broadcasted in a specific order and positions inside a break. The first type are the so called *sandwich constraints*. These spot pairs must be shown in the same break but should be separated by at least one spot within the break, hence the name sandwich. Equivalently, the second type are the *piggyback constraints*. These spot pairs must be shown in the same break in consecutive positions, hence the name "piggyback". For each advertiser $a \in \mathcal{A}$, let $\mathcal{S}^2_{\mathcal{A}\texttt{Sdw}}(a) \subset \mathcal{S}_{\mathcal{A}}(a) \times \mathcal{S}_{\mathcal{A}}(a)$ be the set of spot pairs that must satisfy sandwich constraints, and let $\mathcal{S}^2_{\mathcal{A}\texttt{Pig}}(a) \subset \mathcal{S}_{\mathcal{A}}(a) \times \mathcal{S}_{\mathcal{A}}(a)$ be the spot pairs set that must satisfy piggyback constraints. Constraint (13a) imposes that each sandwich and piggyback spot pair must be in the same break. Additionally, constraint (13b) ensures that if a sandwich pair (i,j) is scheduled in break $b \in \mathcal{B}_{\mathcal{S}}(i)$, then at least one other spot that is not type A-position or type Z-position must be assigned to the break, in order to have a spot between i and j for the model of Stage 3. Last, the third constraints type are the *consecutive breaks constraints* (13c), for which each pair of spots in $\mathcal{S}^2_{\mathcal{A}\texttt{Con}}(a) \subset \mathcal{S}_{\mathcal{A}}(a) \times \mathcal{S}_{\mathcal{A}}(a)$ must be assigned to consecutive breaks.

$$x_{ib} = x_{jb}, \qquad \forall a \in \mathcal{A}, (i,j) \in \mathcal{S}^2_{\mathcal{A}\texttt{Sdw}}(a) \cup \mathcal{S}^2_{\mathcal{A}\texttt{Pig}}(a), b \in \mathcal{B}_{\mathcal{S}}(i), \tag{13a}$$

$$\sum_{i' \in \mathcal{S}_{\mathcal{B}}(b)/\{i,j\} \cup \mathcal{S}_{\texttt{A}} \cup \mathcal{S}_{\texttt{Z}}} x_{bi'} \geq x_{ib}, \qquad \forall a \in \mathcal{A}, (i,j) \in \mathcal{S}^2_{\mathcal{A}\texttt{Sdw}}(a), b \in \mathcal{B}_{\mathcal{S}}(i), \tag{13b}$$

$$x_{ib} = x_{j(b+1)}, \qquad \forall a \in \mathcal{A}, (i,j) \in \mathcal{S}^2_{\mathcal{A}\texttt{Con}}(a), b \in \mathcal{B} \setminus \{|\mathcal{B}|\}. \tag{13c}$$

Another type of business constraints required by advertisers are called the *product category constraints*: In each break, a specific maximum number of spots of the same product category can be shown. Let $\mathcal{P}$ be the set of product categories; and, for each $p \in \mathcal{P}$, let $\mathcal{S}_{\mathcal{P}}(p) \subset \mathcal{S}$ be the set of spots that belong to product category p. Constraint (14) imposes that a maximum of $M^{\texttt{B}}_{pb}$ product category p spots can be assigned to break b.

$$\sum_{i \in \mathcal{S}_{\mathcal{P}}(p) \cap \mathcal{S}_{\mathcal{B}}(b)} x_{ib} \leq M^{\texttt{B}}_{pb}, \qquad p \in \mathcal{P}, b \in \mathcal{B}, \tag{14}$$

Last, block of constraints (15) defines the nature of variables.

$$x_{ib} \in \{0,1\}, \qquad \forall i \in \mathcal{S}, b \in \mathcal{B}_{\mathcal{S}}(i),$$

$$
\begin{aligned}
& y_i \in \{0,1\}, && \forall i \in \mathcal{S}, \\
& z_{kb} \in \{0,1\}, \quad \alpha_{kb} \geq 0, && \forall k \in \mathcal{K}, b \in \mathcal{B}, \\
& v_{ab} \in \{0,1\}, \quad \beta_{ab} \geq 0, && \forall a \in \mathcal{A}, b \in \mathcal{B}, \\
& \delta_{dh}, \gamma_{dh} \geq 0, && \forall d \in \mathcal{D}, h \in \mathcal{H}_{\mathcal{D}}(h), \\
& w_{ij} \in \{0,1\}, && \forall (i,j) \in \mathcal{S}^2_{\mathtt{Sep}}.
\end{aligned} \tag{15}
$$

Formulation (6) to (15) is flexible in terms of including other considerations that may be required by the advertisers or the network. For example, the user may wish to force the scheduling of certain spots that are not guaranteed but must be shown for business considerations, or the user may wish to impose a given number of spots on the log and add constraints that prevent under-delivered deals from becoming worse than before.

4.3. Stage 3: Arrange Spots in Break Positions

As we discussed in Section 3, after the spots are assigned to breaks, it is necessary to sort them to satisfy the internal break constraints. We can solve the arrangement problem by break. For each break $b \in \mathcal{B}$, let $\tilde{\mathcal{S}}(b)$ be the set of spots that are assigned to break b in Stage 2. Therefore, the number of positions in break b is $|\tilde{\mathcal{S}}(b)|$. Let $\hat{x}_{il}$ be a binary variable that is equal to 1 if spot i is scheduled in position l, for all $i \in \tilde{\mathcal{S}}(b), l \leq |\tilde{\mathcal{S}}(b)|$. To model the piggyback and sandwich constraints, let $\tilde{\mathcal{S}}^2_{\mathtt{Pig}}(b)$ be the set of piggyback spot pairs, and let $\tilde{\mathcal{S}}^2_{\mathtt{Sdw}}(b)$ be the set of sandwich spot pairs that were assigned to break b in Stage 2. Let $i_{\mathtt{A}}(b)$ and $i_{\mathtt{Z}}(b)$ be the A-position and Z-position spots that were assigned to break b in Stage 2, if any. Because of all these constraints, there is a chance that the allocation of all the spots to positions is infeasible. In that case, we assign spots to the bin and maximize the weighted ratings allocated, as in Stage 2. Let $\hat{y}_i$ be the binary variable that indicates whether spot $i \in \tilde{\mathcal{S}}(b)$ goes into the bin. Table 5 summarizes the notation described above.

The Stage 3 integer programming model, or the *Position Problem*, is as follows:

$$
\max \sum_{i \in \tilde{\mathcal{S}}(b), l \leq |\tilde{\mathcal{S}}(b)|} W_{d(i)} \frac{H_i}{30} R_{lq(d(i))} \hat{x}_{il} - \sum_{i \in \tilde{\mathcal{S}}(b)} W_{d(i)} \frac{H_i}{30} R_{b(i)q(d(i))} \hat{y}_i \tag{16}
$$

$$\text{s.t.} \quad \sum_{i\in\tilde{\mathcal{S}}(b)} \hat{x}_{il} \le 1 \qquad \forall l \le |\tilde{\mathcal{S}}(b)|, \tag{17}$$

$$\sum_{l\le|\tilde{\mathcal{S}}(b)|} \hat{x}_{il} + \hat{y}_i = 1 \qquad \forall i \in \tilde{\mathcal{S}}(b), \tag{18}$$

$$\hat{x}_{i_{\mathtt{A}}(b)1} = 1, \qquad \hat{x}_{i_{\mathtt{Z}}(b)|\tilde{\mathcal{S}}(b)|} = 1, \tag{19}$$

$$\hat{x}_{il} = \hat{x}_{j(l+1)} \qquad l \le |\tilde{\mathcal{S}}(b)|-1, \forall (i,j) \in \tilde{\mathcal{S}}^2_{\mathtt{Pig}}(b), \tag{20}$$

$$\hat{x}_{jl} = 0 \qquad l \in \{1,2\}, \forall (i,j) \in \tilde{\mathcal{S}}^2_{\mathtt{Sdw}}(b), \tag{21}$$

$$\hat{x}_{il} + \hat{x}_{j(l+1)} \le 1 \qquad l \le |\tilde{\mathcal{S}}(b)|-1, \forall (i,j) \in \tilde{\mathcal{S}}^2_{\mathtt{Sdw}}(b), \tag{22}$$

$$\sum_{t=1}^{l} \hat{x}_{it} \ge \hat{x}_{j(l+2)} \qquad l \le |\tilde{\mathcal{S}}(b)|-2, \forall (i,j) \in \tilde{\mathcal{S}}^2_{\mathtt{Sdw}}(b), \tag{23}$$

$$\hat{x}_{il}, \hat{y}_i \in \{0,1\} \qquad i \in \tilde{\mathcal{S}}(b), l \le |\tilde{\mathcal{S}}(b)|. \tag{24}$$

The objective function (16) maximizes the weighted ratings allocated to spots as in Stage 2. Constraint (17) ensures that no more than one spot is assigned per position, and constraint (18) imposes that each spot must be assigned to only one position or to the bin. Constraint (19) imposes A-position and Z-position constraints. This means that if Stage 2 assigns the A-position spot $i_{\mathtt{A}(b)}$ to break b, then that spot must be scheduled in the first position of the break. Similarly, if the Z-position spot $i_{\mathtt{A}(b)}$ is assigned to break b by Stage 2, then that spot must be scheduled in the last position of the break. Constraint (20) imposes the piggyback condition, meaning that the spots $(i,j) \in \tilde{\mathcal{S}}^2_{\mathtt{Pig}}(b)$ must be scheduled in consecutive positions. Constraints block (21) to (23) are the sandwich constraints: For each pair of sandwich spots $(i,j) \in \tilde{\mathcal{S}}^2_{\mathtt{Sdw}}(b)$, (21) imposes that j can neither be in position 1 nor position 2, constraint (22) ensures that spots i and j are separated by at least one position, and constraint (23) establishes that spot i must be scheduled before spot j. Finally, constraint (24) is the nature of variables.

5. Solution Approach

The three stages are solved sequentially and with each stage we use a more granular formulation, as discussed in the prior sections. Stage 1 is solved once every week. This stage is solved offline

Table 5 Sets and indices for Stage 3.

$\mathcal{B}, b$	Breaks.
$\mathcal{S}, i, j$	Spots.
$\tilde{\mathcal{S}}(b) \subset \mathcal{S}$	Spots that were assigned to break b in Stage 2, for all $b \in \mathcal{B}$.
$i_{\mathtt{A}}(b)$	A-position spots that were assigned to break b in Stage 2, if any, for all $b \in \mathcal{B}$.
$i_{\mathtt{Z}}(b)$	Z-position spots that were assigned to break b in Stage 2, if any, for all $b \in \mathcal{B}$.
$\tilde{\mathcal{S}}^2_{\mathtt{Sdw}}(b) \subset \tilde{\mathcal{S}}(b) \times \tilde{\mathcal{S}}(b)$	Pairs of spots that must satisfy sandwich constraints and were assigned to break b in Stage 2, for all $b \in \mathcal{B}$.
$\tilde{\mathcal{S}}^2_{\mathtt{Pig}}(b) \subset \tilde{\mathcal{S}}(b) \times \tilde{\mathcal{S}}(b)$	Pairs of spots that must satisfy piggyback constraints and were assigned to break b in Stage 2, for all $b \in \mathcal{B}$.

so as to not impact any business operations. However, Stages 2 and 3 are solved one or more times each day. The maximum acceptable solution time is determined by the following two factors: First, all business operations related to the spots for the day in question must be suspended while optimization occurs, and cannot resume until the optimization model has been solved and the revised schedule has been created. Second, the network prefers to perform the optimization after all the work and changes related to the schedule of programs and the spots to schedule have been made, because any changes after Stages 2 and 3 can disturb the optimal solution. In this way, the network can implement the schedule immediately after the optimization. Therefore, there is a limited time for the optimization models to run. Typically, the network allows no more than 10 minutes for solving Stages 2 and 3 combined in order to have a final spot schedule. We describe the solution approach for each stage bellow.

Stage 1: Estimate Deal Weights. This model is an LP that is easy to solve using a state-of-the-art linear programming solver.

Stage 2: Schedule Spots in Breaks. Initially, we modeled and experimented computationally with Stages 2 and 3 in a combined formulation. That formulation did not produce a feasible solution after several hours of branch and bound using a state-of-the-art integer programming solver. So we divided that formulation into Stages 2 and 3. In terms of problem size and computational effort, Stage 2 is the most challenging. For a single day, the typical size of the IP model described in Subsection 4.2 includes tens of thousands of variables and constraints (see Table 7). The integer programming solver that we used for this project, Fico Xpress 7.6 (Daniel 2009), produces a near optimal solution after more than three hours, by branch and bound and default settings. Given the limited time and the large problem size, the availability of a starting solution given by the network is an advantage. After experimentation, we found that the solution approach that produces the best solutions quickly is an iterative process in which we feed an initial solution, or warm start solution, and allow a small number of changes. Recall that the network allows a maximum number of moves $M^{\texttt{move}}$, used in constraint (7d). As we described previously, that number is approximately 50% of the total number of spots. Initially, we set $M^{\texttt{move}}$ equal to 5%, feed the initial solution as warm start producing a new near optimal solution in a few minutes. We then increase $M^{\texttt{move}}$ to 10% using the previously obtained schedule as a warm start to produce a new near optimal solution. Our iterative procedure continues in this fashion until reaching the original $M^{\texttt{move}}$. This procedure produces a near optimal solution in more than 30 minutes, which is not quick enough. A second enhancement that accelerates the iterative procedure is to randomly fix certain spots in breaks; then unfixing them and fixing others; until the last iteration has all the spots unfixed.

The constraints that make the problem difficult to solve are the *minimum separation constraints* (12). Without these constraints the problem solves relatively quickly. These constraints enforce an ordering in every pair of breaks that the network imposes. This difficulty motivates the idea of including in the fixing procedure at each iteration a certain number of the w_{ij} variables that model the ordering between spots i and j, to be equal to the sequence of the warm start solution. If w_{ij} is fixed to one at any iteration, then spot i must be scheduled before spot j and constraint (12) associated to the tuple (i, j) is not needed. At the next iteration, this variable is unfixed.

Using these enhancements, Xpress computes a near optimal solution at the root node of the branch and bound at each iteration of our approach. It produces a final solution with the original $M^{\texttt{move}}$ in less than 10 minutes.

Stage 3: Arrange Spots in Break Positions. The Stage 3 model that positions the spots inside the breaks is an IP that is easy to solve because the problem can be separated by breaks; therefore, it reduces to arranging about 10 spots in an ordering that satisfies all the constraints.

In summary, by reformulating the problem in three stages and implementing the enhancements of Stage 2, our approach solves the scheduling problem while satisfying all the constraints with a near optimal solution in less than 10 minutes. Computational evidence is presented in Section 7.

6. Rating Forecasts

Recall that the rating forecasts are used in the weight problem (WP) and the break problem (BP), which we present in Subsections 4.1 and 4.2, respectively. The parameter to estimate is R_{bq}, the rating forecast in break b of demographic q, for all $b \in \mathcal{B}, q \in \mathcal{Q}$. We briefly describe the methods followed and the issues encountered below.

Time considerations: We receive the realized ratings one week after the day of broadcasting. Therefore, we have the history of ratings up to one week before the date of airing. Also, a full day is broken into 15–minute intervals; hence, there are 96 quarter hours per day. So, we forecast the ratings per quarter hour and then we map to the corresponding break(s) that are broadcasted in each quarter hour.

Demographics: As we describe in Section 1, the demographics that the networks sell are defined in two dimensions: gender and age group. The networks consider 2 possibilities for gender, Female (F) and Male (M) (some media are starting to recognize other genders; for example, Facebook). The age dimension is divided into eight groups (1-5, 6-10, 11-15, 16-20, 21-30, 31-40, 41-51, and 51+). Hence, there are 16 demographics in total ($|\mathcal{Q}| = 16$). this results in a total of 96 $\times$ 16 time series to forecast.

Rating Patterns: The ratings are highly variable, as the plots in Figure 2 show. Plots 2(a) are ratings per day of the week in a representative demographic; each day of the week follows its own pattern. Plots 2(b) are the ratings per quarter–hour from 07:00 to 23:45 on a representative day. We can see that the ratings per consecutive quarter–hours are highly correlated, as it would be expected in television programming. A typical quarter– hour–demographic time series has an approximate coefficient of variation (standard deviation/sample mean) of 0.5. Also, when a new season starts with different programming, the past ratings for different programs are often poor predictors for the future. Therefore, as a policy, we restart the forecast every season.

Rating Models: We tested a battery of time series methods (past ratings average, Holt-Winters and ARIMA); for reference, see any publication on time series (e.g., Shumway and Stoffer (2011)).

Figure 2 Examples of ratings time series.

(a) Ratings by day of the week from August 25, 2008, to July 10, 2011.

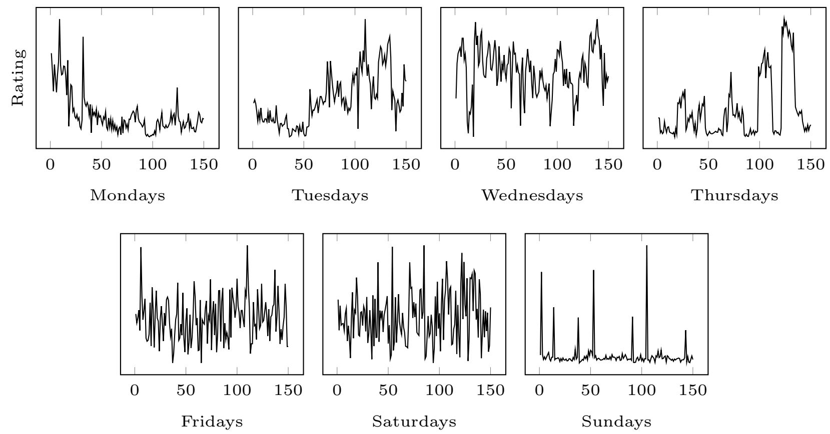

(b) Ratings per quarter hour from Mon 1st to Sun 7th of September, 2008.

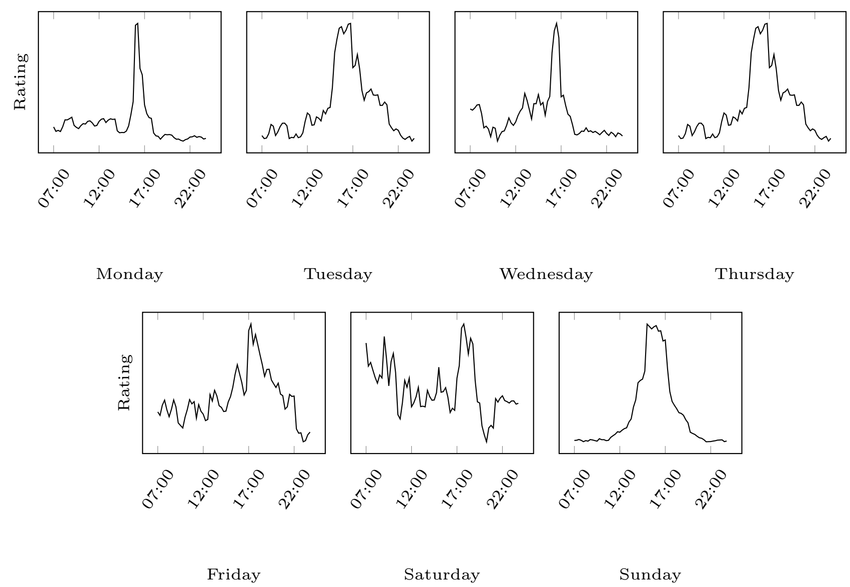

Note. Ratings are normalized to one and correspond to one particular large demographic.

Also, we evaluated different strategies of data usage: all past ratings or only a recent subset; all the days of the week combined or only the corresponding day of the week; different aggregation and disaggregation methods (aggregation per day, by show, and by demographic in one or two dimensions). Our benchmark forecast is the past ratings average, which has an approximate 44% mean absolute percentage error (MAPE). Measured by MAPE, none of the methods were found to be superior for all the series. But, for a particular combination of day of the week, time of the day, and demographic, a specific method often produced consistently superior results. We use the method that is the most competitive in terms of MAPE for test sets. The MAPEs that we obtain for the actual ratings are around 32%.

In practice, we have found that spots are moved among a small set of breaks. Therefore, it is adequate for practical purposes to focus on forecasting the difference in ratings between pairs of breaks within the smaller set. Improving the accuracy of forecast is an ongoing area of research.

7. Results

Our optimization models have been used successfully for major US cable networks since 2008, generating millions of dollars of additional revenue annually. In order to communicate with the existing networks' data bases and also to integrate the in-house software used for scheduling the logs (before passing a feasible log to the optimizer) with our scheduler, we built an ad hoc decision support system. As we mention in Section 5, we implemented the optimization models in Xpress 7.6.

To quantify the yield generated and the quality of the optimized schedules, in this section we benchmark the schedules produced by our models against the ones provided before optimization for one major network's channel. We compare the daily schedules for a full month. On average, the optimized schedules produce a conservative increase in revenue of more than $24,500 per day, which translates to almost $9 million per year.

All data, while representative of the problem, are disguised for the purpose of maintaining confidentiality.

7.1. Instances, problem size, and optimality gap

We benchmark the schedules of one major network's channel for the 31 days of August 2010. In this subsection, we present the mean, standard deviation and coefficient of variation across the 31 days of several measures. The disaggregated data is shown in the Appendix.

The average day instance has 150 breaks, 16 targeted demographics, 83 advertisers, 164 brands, 124 deals, 672 spots, and 255 product categories. See Table 6 for the descriptive statistics of the instances and Table 11, in the Appendix, for the daily data.

Table 6 **Stage 2 instances size summary statistics for the 31 days of August 2010.**

	# Breaks $\lvert\mathcal{B}\rvert$	# Demographics $\lvert\mathcal{Q}\rvert$	# Advertisers $\lvert\mathcal{A}\rvert$	# Brands $\lvert\mathcal{K}\rvert$	# Deals $\lvert\mathcal{D}\rvert$	# Spots $\lvert\mathcal{S}\rvert$	# PCat $\lvert\mathcal{P}\rvert$
Mean	150	16	83	164	124	672	255
Stdev	19.77	2.04	6.50	20.01	18.07	82.62	0.00
CV	0.13	0.12	0.08	0.12	0.15	0.12	0.00

|Set|: number of elements in Set. Stdev: standard deviation. CV: coefficient of variation (Stdev/mean). PCat: product categories.

Because the Stage 2 model is the most difficult to solve, we describe the problem size of that model. The average problem has 11,661 binary variables, 13,489 continuous variables, and 24,007 constraints. After 10 minutes of optimization, our Stage 2 iterative procedure described in Section 5 produces a nearly optimal schedule as indicated by the 0.57% MIP gap on average (100×(LP solution-IP solution)/LP solution). See Table 7 for the summary and Table 12, in the Appendix, for the daily problem size.

Table 7 **Stage 2 problems size summary statistics for the 31 days of August 2010.**

Per day	# Binary vars.	# Continuous vars.	# Constraints	MIP gap
Mean	11,661	13,489	24,007	0.57%
Stdev	7,315	5,704	9,539	0.39%
CV	0.63	0.42	0.40	0.68

Stdev: standard deviation. CV: coefficient of variation (Stdev/mean). MIP gap = ((MIP objective function - best LP bound objective function)/best LP bound objective function)× 100.

7.2. Benchmark

In this subsection, we present results showing the day-by-day yield for one month. We set benchmarks using two metrics, namely, the weighted average ratings (M1) and the increase in revenue (Value). The comparison is made between the original schedule provided by the network (ORG) and the schedule obtained after optimization (OPT). The two metrics are:

$$\text{M1} = \sum_{\substack{d\in\mathcal{D}, i\in\hat{\mathcal{S}}_{\mathcal{D}}(d),\\ b\in\mathcal{B}_{\mathcal{S}}(i)}} W_d \frac{H_i}{30} R_{bq(d)} x_{ib}; \qquad \text{Value} = \sum_{\substack{d\in\mathcal{D}, i\in\hat{\mathcal{S}}_{\mathcal{D}}(d),\\ b\in\mathcal{B}_{\mathcal{S}}(i)}} CPM_{q(d)} W_d \frac{H_i}{30} R_{bq(d)} x_{ib}.$$

We compute the two measures for the ORG and OPT schedules using forecasted as well as actual ratings.

Recall that the weight W_d is an artifact that we compute in order to assign relative importance to the deals as discussed in Section 4. However, the network assumes a weight equal to 1 for all deals at the time its original schedule is generated. Hence, in order to establish a fair comparison, in this benchmark we assume $W_d = 1$ for all deals.

Concerning penalties, we compare the sum of the penalties M2+M3+M4+M5 for the ORG and OPT schedules. These penalties are used to impose the soft constraints in the design of the schedule and they do not have an intrinsic monetary value. Therefore, we compare them only for the schedule with forecasted ratings.

Table 8 presents the summary results based on the forecasted ratings, which are used to produce the schedules. In terms of allocated ratings, or M1, the optimized schedule allocates, on average, 2.21% higher ratings to the spots. In terms of value, the average day of forecasted ratings has a total value of \$1.47 million, so each percent gain is equivalent to roughly \$14,700. For the test month using the forecasted ratings, the OPT schedules produce a 2.08% value gain. If we consider the variability of value gain across days using the forecasted ratings, we observe that the coefficient of variation is 0.62, which indicates that the gain is consistently superior. Evidently, the optimized schedule for the forecasted ratings always dominates the original schedule, which was used as starting solution. Table 13 in the Appendix shows the results disaggregated by day.

Table 8 **Allocated Ratings and Value based on forecasted ratings.**

	Weighted average ratings (M1)			Value		
Per day	ORG	OPT	gain	ORG	OPT	gain
Mean	87,193,858	89,088,245	2.21%	\$1,439,032	\$1,468,577	2.08%
Stdev	14,114,964	14,215,505	1.31%	\$279,133	\$283,092	1.30%
CV	0.16	0.16	0.59	0.19	0.19	0.62

Stdev: standard deviation. CV: coefficient of variation (Stdev/mean). ORG: original schedule. OPT: optimized schedule. gain = ((OPT-ORG)/ORG)× 100. We use a random scaling factor to maintain confidentiality, but the order of magnitudes correspond to the true values.

Table 9 displays the mean and standard deviation across the 31 days of the penalties M2+M3+M4+M5 for the ORG and OPT schedules. The penalties of ORG are much higher than OPT, as the 31.49 ratio indicates. When we examine the disaggregated data per day in Table 14, we observe that the OPT penalties are equal to 0 for many days. Therefore, the OPT schedules violate far fewer the soft constraints (on some days they violate none of the soft constraints) than the ORG schedules, which always incur penalties due to violations. Table 14 in the Appendix shows the penalties disaggregated by day.

Table 9 Penalties based on forecast ratings.

	Penalties (M2+M3+M4+M5)		
	ORG	OPT	ORG/OPT
Mean	3,734,878	118,604	31.49
Stdev	2,292,353	181,199	12.65

Stdev: standard deviation. ORG: original schedule. OPT: optimized schedule. We use a random scaling factor to maintain confidentiality, but the order of magnitudes correspond to the true values.

Table 10 shows the allocated ratings and value that were obtained by inserting the actual ratings into the schedules produced using the forecasted ratings. In terms of value gain, the OPT schedule has daily a value of \$1.85 million and is, on average, 1.34% superior to ORG, which translates into a daily average of more than \$24,500. However, the value gain coefficient of variation is 1.02, which indicates a high variation. Observing the value gain in each day in Table 15, in the Appendix, shows that the optimized schedule dominates the original schedule in 29 of the 31 days. In the two days on which ORG is better than OPT (8/11/2010 and 8/31/2010), the difference is no more than 0.34%. Based on the data, we can compute the empirical probability of a gain using the following expression:

$$\Pr(\text{having a gain}) = \frac{\displaystyle\sum_{\text{day}\in\text{days with gain}} \text{gain}_{\text{day}}\%}{\displaystyle\sum_{\text{day}\in\text{days with gain}} \text{gain}_{\text{day}}\% + \sum_{\text{day}\in\text{days with loss}} \text{loss}_{\text{day}}\%} = 0.9899$$

Therefore, although it is possible that the original schedule is superior compared to the optimized schedule for the actual ratings, the probability of that event is less than 0.0101 for the test month.

This unlikely event could occur because the forecasted ratings differ significantly from the actual ratings; and, merely due to chance, the original schedule is better than the optimized schedule. The feasibility of this event shows that the original schedules are of good quality, but because of the highly competitive nature of the US media industry, the networks are constantly searching for sophisticated methods to extract more yield from their audience.

Table 10 Allocated Ratings and Value based on actual ratings.

	Weighted average ratings (M1)			Value		
Per day	org	opt	gain	org	opt	gain
Mean	106,047,503	107,570,039	1.50%	$1,825,764	$1,849,022	1.34%
Stdev	22,259,809	22,256,253	1.56%	$500,914	$501,706	1.37%
CV	0.21	0.21	1.03	0.27	0.27	1.02

Stdev: standard deviation. CV: coefficient of variation (Stdev/mean). org: original schedule. opt: optimized schedule. gain = ((opt-org)/org)$\times$ 100. We use a random scaling factor to maintain confidentiality, but the order of magnitudes correspond to the true values. Actual ratings are obtained after the spots are aired.

8. Conclusions

The major revenue source for television networks is the selling of viewers to advertisers. Therefore, an efficient audience distribution among advertisers is essential to maximize the yield. This process involves the scheduling of advertisements, or *spots*, within commercial breaks. The goal of the schedule is to arrange the spots in the breaks so that each spot is shown to its targeted demographic. This multi-period scheduling problem is very difficult to solve, because viewers of different demographics will be watching a particular break at the same time. Thus, at the moment of producing the schedule, various spots are competing for the same break to reach different demographics. On the other hand, the advertisers impose several business restrictions on the schedule, such as minimum separation time between spots of the same product category, and uniform distribution of spots from the same brand.

We designed and implemented a combined solution based on mathematical programming and time series forecasting methods to schedule the spots within breaks in a way that maximizes the value of the audience. The scheduler arranges the spots at the level of positions inside the breaks, which is the maximum level of resolution. The optimization model is a large scale integer programming model. We solve it close to optimality by using an ad-hoc iterative procedure in less than 10 minutes, which is the time available to produce the daily schedule. The schedules are of high quality as measured by standard business metrics and when compared to the mathematical optimal bound.

Acknowledgments

Appendix. Disaggregated Results for August 2010

Table 11 Statistics of instances August 2010.

Instance-Day	# Breaks $\lvert\mathcal{B}\rvert$	# Demos $\lvert\mathcal{Q}\rvert$	# Advertisers $\lvert\mathcal{A}\rvert$	# Brands $\lvert\mathcal{K}\rvert$	# Deals $\lvert\mathcal{D}\rvert$	# Spots $\lvert\mathcal{S}\rvert$	# PCat. $\lvert\mathcal{P}\rvert$
8/1/2010	132	16	69	138	98	622	255
8/2/2010	129	18	79	160	121	624	255
8/3/2010	126	18	82	166	133	663	255
8/4/2010	129	19	93	194	155	728	255
8/5/2010	133	18	90	181	152	715	255
8/6/2010	127	17	83	173	109	746	255
8/7/2010	135	18	85	160	118	786	255
8/8/2010	135	18	76	135	100	765	255
8/9/2010	128	16	86	174	129	655	255
8/10/2010	123	14	88	172	138	683	255
8/11/2010	128	14	86	173	127	669	255
8/12/2010	129	14	95	176	152	822	255
8/13/2010	132	12	84	158	113	858	255
8/14/2010	132	13	75	136	100	697	255
8/15/2010	137	13	80	147	106	756	255
8/16/2010	171	18	88	165	128	579	255
8/17/2010	167	18	85	177	132	638	255
8/18/2010	169	16	76	167	117	636	255
8/19/2010	171	15	84	160	126	607	255
8/20/2010	172	16	81	153	112	680	255
8/21/2010	173	16	81	159	113	635	255
8/22/2010	176	15	73	121	93	464	255
8/23/2010	165	17	83	126	129	582	255
8/24/2010	161	17	86	186	143	657	255
8/25/2010	168	16	81	170	129	625	255
8/26/2010	170	18	91	183	144	621	255
8/27/2010	173	15	85	174	124	745	255
8/28/2010	173	16	85	159	115	657	255
8/29/2010	170	16	72	134	97	505	255
8/30/2010	156	21	98	210	162	717	255
8/31/2010	168	20	88	185	142	700	255

Note. Stdev: standard deviation. CV: coefficient of variation (Stdev/mean). PCat: product categories.

Table 12 Stage 2 problems size, and MIP gap.

Instance-Day	# Binary vars.	# Continuous vars.	# Constraints	MIP gap [%]
8/1/2010	7,248	11,187	20,445	0.82
8/2/2010	8,875	10,502	18,579	1.01
8/3/2010	7,495	9,955	18,022	0.03
8/4/2010	9,865	12,034	21,419	0.02
8/5/2010	11,064	11,971	21,562	0.06
8/6/2010	18,033	18,302	32,181	0.79
8/7/2010	8,560	12,889	23,007	0.62
8/8/2010	10,274	13,947	25,103	0.83
8/9/2010	6,623	9,086	16,451	0.76
8/10/2010	5,428	7,984	14,693	0.74
8/11/2010	9,224	11,723	20,847	0.04
8/12/2010	7,429	10,498	19,016	0.96
8/13/2010	6,845	10,230	19,135	0.99
8/14/2010	7,378	10,322	19,130	1.00
8/15/2010	12,968	16,241	29,659	0.99
8/16/2010	5,106	7,942	14,159	0.28
8/17/2010	12,517	14,297	24,860	0.78
8/18/2010	8,532	10,894	20,022	0.99
8/19/2010	10,669	12,101	21,686	0.08
8/20/2010	21,332	21,495	37,548	1.37
8/21/2010	17,095	17,955	31,709	0.03
8/22/2010	4,176	6,470	12,615	0.98
8/23/2010	2,188	4,728	9,335	0.21
8/24/2010	12,350	14,600	25,657	0.26
8/25/2010	24,694	23,538	39,791	0.12
8/26/2010	20,861	21,196	36,292	0.24
8/27/2010	37,418	32,206	55,563	0.43
8/28/2010	23,462	22,058	38,624	0.32
8/29/2010	6,277	9,256	16,555	0.70
8/30/2010	8,980	11,744	20,954	0.78
8/31/2010	8,512	10,822	19,610	0.48

Note. MIP gap = ((MIP objective function - best LP bound objective function)/best LP bound objective function)$\times$ 100.

Table 13 **Weighted average ratings (M1) and Value by day based on forecast ratings.**

	M1			Value		
Day	ORG	OPT	gain [%]	ORG [$]	OPT [$]	gain [%]
8/1/2010	115,915,000	117,609,000	1.46	1,565,770	1,589,100	1.49
8/2/2010	84,954,200	86,191,100	1.46	1,190,980	1,207,240	1.37
8/3/2010	107,126,000	108,123,000	0.93	1,962,410	1,979,820	0.89
8/4/2010	99,418,400	100,947,000	1.54	1,541,040	1,563,640	1.47
8/5/2010	75,768,000	77,270,100	1.98	1,109,250	1,131,060	1.97
8/6/2010	73,366,900	75,414,800	2.79	1,192,450	1,221,150	2.41
8/7/2010	94,228,400	95,529,900	1.38	1,503,240	1,519,120	1.06
8/8/2010	98,232,000	100,735,000	2.55	1,691,920	1,723,760	1.88
8/9/2010	78,514,900	79,433,600	1.17	1,289,620	1,302,440	0.99
8/10/2010	106,816,000	107,852,000	0.97	1,773,440	1,793,060	1.11
8/11/2010	81,675,700	82,654,400	1.20	1,347,240	1,359,650	0.92
8/12/2010	76,324,400	77,427,200	1.44	1,320,000	1,339,410	1.47
8/13/2010	67,917,900	68,933,900	1.50	1,147,950	1,165,090	1.49
8/14/2010	87,648,000	90,476,000	3.23	1,445,980	1,489,430	3.00
8/15/2010	83,754,100	85,588,900	2.19	1,356,760	1,381,780	1.84
8/16/2010	76,467,100	77,322,700	1.12	1,223,040	1,237,260	1.16
8/17/2010	98,908,900	101,478,000	2.60	1,844,080	1,894,520	2.74
8/18/2010	86,515,300	88,296,800	2.06	1,479,870	1,502,280	1.51
8/19/2010	78,213,600	79,785,100	2.01	1,412,880	1,437,550	1.75
8/20/2010	66,027,300	69,624,400	5.45	1,084,500	1,142,620	5.36
8/21/2010	76,425,900	78,742,200	3.03	1,223,300	1,258,260	2.86
8/22/2010	66,274,700	67,104,600	1.25	1,105,080	1,118,590	1.22
8/23/2010	78,573,000	78,941,700	0.47	1,274,720	1,279,520	0.38
8/24/2010	111,999,000	114,717,000	2.43	2,028,790	2,088,610	2.95
8/25/2010	89,990,800	95,485,200	6.11	1,520,670	1,606,590	5.65
8/26/2010	97,527,000	100,115,000	2.65	1,805,710	1,843,780	2.11
8/27/2010	89,293,200	93,426,700	4.63	1,459,230	1,525,290	4.53
8/28/2010	84,572,300	87,770,200	3.78	1,269,710	1,324,050	4.28
8/29/2010	66,996,400	68,641,000	2.45	1,028,270	1,048,840	2.00
8/30/2010	88,940,200	90,928,100	2.24	1,378,660	1,406,970	2.05
8/31/2010	114,625,000	115,171,000	0.48	2,033,440	2,045,400	0.59

Note. ORG: original schedule. OPT: optimized schedule. gain = ((OPT-ORG)/ORG)× 100. We use a random scaling factor to maintain confidentiality, but the order of magnitudes correspond to the true values.

Table 14 Sum of Penalties by day.

Penalties M2+M3+M4+M5		
ORG	OPT	ORG/OPT
1,609,370	0	-
6,093,750	556,003	10.96
5,387,620	142,321	37.86
9,428,160	226,893	41.55
2,687,760	0	-
11,860,200	798,251	14.86
2,399,360	14,316	167.61
3,546,980	0	-
4,968,830	186,559	26.63
6,248,010	332,281	18.80
5,823,790	198,591	29.33
3,373,700	0	-
2,462,360	219,410	11.22
4,063,760	234,747	17.31
4,324,200	215,280	20.09
3,749,290	0	-
2,438,150	44,884	54.32
2,538,320	266,903	9.51
2,385,700	0	-
2,496,930	0	-
3,313,990	68,886	48.11
741,317	0	-
2,863,230	0	-
3,418,260	151,616	22.55
2,680,790	0	-
1,791,720	0	-
3,753,100	0	-
3,712,180	0	-
982,697	0	-
3,153,530	19,800	159.27
1,484,170	0	-

Note. ORG: original schedule. OPT: optimized schedule. We use a random scaling factor to maintain confidentiality, but the order of magnitudes correspond to the true values.

Table 15 Allocated Ratings and Value by day based on actual ratings.

	M1			Value		
Day	ORG	OPT	gain [%]	ORG [$]	OPT [$]	gain [%]
8/1/2010	120,469,000	121,577,000	0.92	1,651,300	1,674,730	1.42
8/2/2010	98,553,200	99,305,600	0.76	1,375,570	1,388,250	0.92
8/3/2010	109,442,000	110,609,000	1.07	2,173,350	2,196,640	1.07
8/4/2010	123,849,000	125,280,000	1.16	2,028,000	2,050,400	1.10
8/5/2010	159,841,000	161,165,000	0.83	2,817,310	2,832,800	0.55
8/6/2010	112,041,000	113,524,000	1.32	1,820,170	1,837,830	0.97
8/7/2010	124,519,000	125,322,000	0.64	1,987,650	1,996,780	0.46
8/8/2010	92,648,400	93,947,800	1.40	1,558,510	1,576,390	1.15
8/9/2010	85,535,900	86,291,600	0.88	1,440,120	1,456,780	1.16
8/10/2010	108,325,000	109,629,000	1.20	1,884,040	1,929,940	2.44
8/11/2010	92,348,500	92,280,000	-0.07	1,596,280	1,590,920	-0.34
8/12/2010	135,347,000	137,014,000	1.23	2,661,940	2,693,500	1.19
8/13/2010	90,576,700	91,241,600	0.73	1,547,220	1,552,710	0.35
8/14/2010	89,458,100	91,624,100	2.42	1,483,750	1,520,410	2.47
8/15/2010	84,433,800	84,591,700	0.19	1,377,600	1,380,390	0.20
8/16/2010	94,496,100	94,551,500	0.06	1,464,360	1,473,050	0.59
8/17/2010	121,118,000	122,386,000	1.05	2,317,580	2,354,000	1.57
8/18/2010	109,291,000	111,762,000	2.26	1,940,750	1,978,660	1.95
8/19/2010	140,016,000	140,878,000	0.62	2,754,810	2,781,890	0.98
8/20/2010	89,270,100	94,830,900	6.23	1,459,720	1,542,500	5.67
8/21/2010	71,132,400	73,196,100	2.90	1,166,750	1,191,750	2.14
8/22/2010	71,899,700	72,694,700	1.11	1,211,100	1,225,200	1.16
8/23/2010	93,019,200	93,484,400	0.50	1,559,620	1,567,760	0.52
8/24/2010	94,542,800	94,859,100	0.33	1,739,730	1,748,140	0.48
8/25/2010	95,303,100	101,850,000	6.87	1,726,650	1,821,010	5.46
8/26/2010	151,335,000	154,379,000	2.01	3,118,450	3,141,130	0.73
8/27/2010	105,478,000	106,425,000	0.90	1,697,640	1,708,100	0.62
8/28/2010	98,512,300	102,000,000	3.54	1,522,060	1,578,000	3.68
8/29/2010	74,459,300	75,820,100	1.83	1,159,240	1,165,810	0.57
8/30/2010	110,574,000	112,516,000	1.76	1,827,100	1,836,320	0.50
8/31/2010	139,638,000	139,636,000	0.00	2,530,300	2,527,880	-0.10

Note. ORG: original schedule. OPT: optimized schedule. gain = ((OPT-ORG)/ORG)× 100. We use a random scaling factor to maintain confidentiality, but the order of magnitudes correspond to the true values. Actual ratings are obtained after the spots are aired.

References

Adany, R., S. Kraus, F. Ordonez. 2012. Allocation algorithms for personal TV advertisements. *Multimedia Systems* **19**(2) 79–93.

Araman, V. F., I. Popescu. 2010. Media Revenue Management with Audience Uncertainty: Balancing Upfront and Spot Market Sales. *Manufacturing & Service Operations Management* **12**(2) 190–212.

Banciu, M., E. Gal-Or, P. Mirchandani. 2010. Bundling Strategies When Products Are Vertically Differentiated and Capacities Are Limited. *Management Science* **56**(12) 2207–2223.

Blumenthal, H., O. R. Goodenough. 2006. *This business of television*. 3rd ed. Billboard Books.

Bollapragada, S., M. R. Bussieck, S. Mallik. 2004. Scheduling Commercial Videotapes in Broadcast Television.

Bollapragada, S., H. Cheng, M. Phillips, M. Garbiras, M. Scholes, T. Gibbs, M. Humphreville. 2002. NBC's Optimization Systems Increase Revenues and Productivity. *Interfaces* **32**(1) 47–60.

Bollapragada, S., M. Garbiras. 2004. Scheduling Commercials on Broadcast Television.

Bollapragada, S., S. Gupta, B. Hurwitz, P. Miles, R. Tyagi. 2008. NBC-Universal uses a novel qualitative forecasting technique to predict advertising demand. *Interfaces* **38**(2) 103–111.

Bollapragada, S., S. Mallik. 2008. Managing on-air ad inventory in broadcast television.

Brusco, M. J. 2008. Scheduling advertising slots for television.

Brusco, M. J., R. Singh. 2010. Assigning television commercial videotapes to time slots under alternative message spacing policies. *International Journal of Advertising* **29** 431–450.

Cancian, M., A. Bills, T. Bergstrom. 1996. Hotelling location problems with directional constraints: an application to television news scheduling.

Caves, R. E., K. Guo. 2009. *Switching channels: Organization and change in TV broadcasting*. Harvard University Press.

Crama, P., D. G. Popescu, A. S. Aravamudhan. 2012. Advertising Revenue Optimization in Live Television Broadcasting. *Working Paper* 1–31.

Danaher, P., T. Dagger. 2012. Using a nested logit model to forecast television ratings. *International Journal of Forecasting* **28** 607–622.

Danaher, P. J. 1991. A Canonical Expansion Model for Multivariate Media Exposure Distributions: A Generalization of the" Duplication of Viewing Law". *Journal of Marketing Research* **XXVIIl** 361–368.

Danaher, P. J., T. S. Dagger, M. S. Smith. 2011. Forecasting television ratings. *International Journal of Forecasting* **27**(4) 1215–1240.

Danaher, P. J., D. F. Mawhinney. 2001. Optimizing television program schedules using choice modeling. *Journal of Marketing Research* **38** 298–312.

Daniel, B. 2009. Modeling with Xpress. *FICO Xpress Optimization Suite whitepaper* .

Gaur, D. R., R. Krishnamurti, R. Kohli. 2009. Conflict Resolution in the Scheduling of Television Commercials. *Operations Research* **57**(5) 1098–1105.

Goettler, R. L., R. Shachar. 2001. Spatial competition in the network television industry. *Rand Journal of Economics* **32** 624–656.

Goodhardt, G. J., A. S. C. Ehrenberg. 1969. Duplication of television viewing between and within channels. *Journal of Marketing Research* **6** 169–178.

Grosfeld-Nir, A., Y. Gerchak. 2004. Multiple lotsizing in production to order with random yields: Review of recent advances. *Annals of Operations Research* **126** 43–69.

Headen, R. S., J. E. Klompmaker, R. T. Rust. 1979. The Duplication of Viewing Law and Television Media Schedule Evaluation. *Journal of Marketing Research* **16** 333–340.

Headen, R. S., J. E. Klompmaker, J. E. Teel. 1977. Predicting Audience Exposure to Spot TV Advertising Schedules. *Journal of Marketing Research* **14** 1–9.

Henry, M. D., H. J. Rinne. 1984. Predicting program shares in new time slots. *Journal of Advertising Research* **24** 9–17.

Horen, J. H. 1980. Scheduling of Network Television Programs. *Management Science* **26** 354–370.

Kelton, C. M. L., L. G. Schneider Stone. 1998. Optimal television schedules in alternative competitive environments.

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

Lees, G., M. Wright. 2013. Does the duplication of viewing law apply to radio listening? *European Journal of Marketing* **47** 674–685.

Phillips, R., G. Young. 2012. Television Advertisement Pricing in the U.S. Özalp Özer Phillips, Robert, eds., *Oxford Handbook of Pricing Management*, chap. 14. Oxford University Press, New York, 230–249.

Popescu, D. G., S. Seshadri. 2013. Managing Revenue from Television Advertising Sales. *Working Paper* 1–27.

Reddy, S. K., J. E. Aronson, A. Stam. 1998. SPOT: Scheduling Programs Optimally for Television. *Management Science* **44**(1) 83–102.

Rust, R. T., M. I. Alpert. 1984. An Audience Flow Model of Television Viewing Choice. *Marketing Science* **3**(2) 113–124.

Rust, R. T., N. V. Eechambadi. 1989. Scheduling network television programs: A heuristic audience flow approach to maximizing audience share. *Journal of Advertising* **18** 11–18.

Rust, R. T., W. A. Kamakura, M. I. Alpert. 1992. Viewer Preference Segmentation and Viewing Choice Models for Network Television. *Journal of Advertising* **21**(1) 1–18.

Rust, R. T., M. R. Zimmer, R. P. Leone. 1986. Estimating the Duplicated Audience of Media Vehicles in National Advertising Schedules. *Journal of Advertising* **15** 30–37.

Shumway, R. H., D. S. Stoffer. 2011. *Time Series Analysis and Its Applications*. 3rd ed. Springer.

Talluri, K. T., G. J. Van Ryzin. 2005. Media and Broadcating. *The theory and practice of revenue management*, chap. 10.5. Springer, 542–546.

Webster, J., P. Phalen, L. Lichty. 2013. *Ratings Analysis: Audience Measurement and Analytics*. 4th ed. Routledge.

Webster, J. G. 1985. Program audience duplication: A study of television inheritance effects. *Journal of Broadcasting & Electronic Media* **29** 121–133.

Wilbur, K., M. Goeree, G. Ridder. 2008. Effects of advertising and product placement on television audiences. *Working Paper* .

Zhang, X. 2006. Mathematical models for the television advertising allocation problem. *International Journal of Operational Research* **1**(3) 302–322.