Appendix C. Public Policy, Generalizations, and Implementation

This appendix discusses several implications and extensions of the proposed framework. These include influence on public policy, order effects among commercials, multiple commercial breaks, advance selling, audience bundling, additional constraints, and predicting continuation probabilities for new advertisements.

C.1. Public Policy Implications

It is important to remember that even aggregate efficiency gains still may produce winners and losers. Though it did not occur in the simulations, it is possible that some viewing segments may be left worse off under AVMA if their commercial preferences are very different from those of other segments which are much larger or have greater value to advertisers. Furthermore, advertisements with exogenously low continuation rates may find it more difficult to buy advertising time; they will be forced to pay higher prices or be excluded from the market. This, of course, is the intention of the approach, since exclusion of low-quality ads raises market efficiency.

AVMA may be particularly relevant to public television networks. These networks rank among the leaders in viewing share and advertising revenues in most countries' television markets and are typically obliged to serve the public interest. Because viewer switching indicates poor audience acceptance of commercials, public networks may have additional reasons to adopt the models proposed here beyond revenue considerations.

In some countries, public policy restricts the use of consumers' individual viewing data without express consent. It would be difficult to reliably estimate continuation rates without some form of passive viewing data. Goldfarb and Tucker (2011b) show that European Union privacy regulations reduced online advertising effectiveness by 65%. Such privacy regulations may have a similar detrimental impact on a television network's ability to design commercial schedules to reduce viewer disutility. However, the potential harm to individual consumers may be smaller in television than online, given the broadcast nature of TV advertising.

C.2. Order Effects

The model in §2 is based on the assumption that the audience watching any given ad slot depends only on the product of the prior ads' continuation probabilities. It is quite possible that externalities take on a more general form. For example, showing a funny ad before a dramatic ad may have a different impact on switching than showing the dramatic ad before the funny ad. In general, combinations of ads may have effects that are different from the ones predicted only by the individual ads. A very general model of externalities was studied by Ghosh and Mahdian (2008). The main result in that paper was that it is computationally intractable even to come near the optimum welfare if externalities can take very general forms.

A further complication can arise from interactions between an ad and the program or context within which it is shown. For instance, showing a funny ad during a dramatic program may produce a different result than showing the same funny ad during a funny program. Goldfarb and Tucker (2011a) find that consumers may rebel against ads that are both obtrusive and targeted, suggesting a possible backlash to well-placed ads. To deal with these kinds of difficulties, one approach would be to learn continuation rates for each ad within each possible context it could be shown in. The natural downside of this approach is that in order to achieve the same confidence level for the continuation rates, significantly more data on viewer behavior need to be available.

C.3. Multiple Commercial Breaks and Break Length

Television networks do not sell just one commercial break. Rather, they often offer dozens or hundreds of breaks to advertisers simultaneously. The sets of potential advertisements for different breaks will naturally overlap. Indeed, advertisers may be indifferent between specific breaks, so long as their creatives get shown a certain number of times within a particular program or time period. If AVMA were run individually for each break, it might produce similar advertisement selections and orderings for each of the breaks. This would violate the goal of the advertisers to be shown only a given number of times, and may also produce undesirably high levels of viewer anticipation of ads.

The key to allocating ads across multiple breaks would be an accurate model of how continuation probabilities change with such variables as advertiser identity, viewer demographics, past exposures, time, and context effects. If such predictions could be made, then a global algorithm could allocate potential ads to breaks such that AVMA can then be run for each individual break. An alternative approach would treat all slots of a program (or evening) as one large commercial break, running AVMA on this large "break."

C.4. Advance Selling and Yield Management

The models in §2 implicitly assume that the network knows all advertisers interested in the commercial break at the time it commits to selling particular slots. In general, however, advertisers may differ significantly in their "time of arrival" to purchase advertising. For example, a movie studio may know six months in advance that it will require an advertising blitz to promote the release of a new movie. In contrast, a politician responding to recent news reports may require last-minute advertising time to swing an election. This is an important practical issue in television markets. The May/June "upfront" market allocates the majority of advertising time (60–80%) airing during the following September–August time period. The "scatter" market then allocates the remaining ad time to advertisers up until the ads' actual air dates.

Complicating matters further, advertisers' valuations may well be correlated with their time of arrival, and networks may not even know precisely how many advertisers will be in the market for a particular commercial break. Since advertising stock is perishable, television advertising is a classic example of an industry in which yield management pricing practices may be profitably applied. Indeed, several television networks such as NBC and MTV do use yield management systems to price ad inventory.

If the number of bidders is known with certainty and bids can be held until the complete set of bids is collected, AVMA can be applied without modification. However, it would be unreasonable to require an advertiser to wait months to find out whether a particular bid has been accepted. A more reasonable protocol is to require that each advertiser, upon submitting a bid, be told immediately whether the bid has been accepted, while the decision on the precise position within the break and the price to be paid is deferred until the break airs.

The network could use historical data to predict the likely number of ads, as well as the distribution of their associated bids and continuation rates. It could then decide whether to accept a particular bid based on its attractiveness relative to the expected set of bids likely to be entered in the future. A more detailed incorporation of yield management with continuation rate-based allocations is a promising direction for future research.

C.5. Audience Bundling

Large television advertisers often receive quantity discounts in exchange for purchasing bundles of ad slots in different programs. There are multiple motivations for such bundling, including reduction of transaction costs and market power of large buyers.

In general, an advertiser could submit bids for arbitrary bundles of slots. However, due to the large number of potential bundles, even just specifying bids for all bundles of interest is intractable. Furthermore, the computation would require solving a complex combinatorial auction, a notoriously difficult problem (Rassenti et al. 1982, Crampton et al. 2006). A simpler approach would be for a large advertiser to enter a single set of segment-specific bids to the network along with a budget constraint, and let the network select programs and ad breaks that match those preferences well. This optimal program recommendation problem remains unexplored.

C.6. Additional Constraints

Section 3.4 showed how to ensure that the algorithms find a solution that satisfies all anti-competition constraints. There are many other constraints that could be placed on the network's problem. For example, there might be advertisers who demand particular slots within a break.

If the constraints become very complex, it may be advisable to use an off-the-shelf optimization approach such as Simulated Annealing on the entire problem. Such a standard optimization algorithm could be used not only to find a feasible selection of ads, but also to determine the selected set of J ads and their ordering, with suitable local search moves defined on them.

The advantage of this approach would be its simplicity. Any number and type of constraints could be incorporated by reducing the solution space, and different ways of exploring local search moves can be implemented. However, it introduces several costs. Many off-the-shelf techniques apply fairly generic searches, with the goal of finding any acceptable solution rather than the best one. AVMA is carefully designed to exploit the structure of the ordering and selection problem, so one would expect its performance to exceed a generic algorithm. More fundamentally, the general specification of the network's problem with constraints subsumes the well-known Independent Set problem. It is thus NP-complete to decide if any valuable conflict-free sets of ads exist (Karp 1972) or to approximate the maximum possible profit (Håstad 1999), so no fast and (nearly) optimal solution can be guaranteed when constraints become sufficiently complex.

C.7. Predicting Continuation Probabilities for New Ads

In this article, we have assumed that the network knows advertisements' segment-specific continuation rates. For ads that have previously aired on the network, historical data can be used to form accurate estimates. But for new ads or new advertisers, no historical data are available. Several approaches to predicting continuation probabilities are available.

(1) An approach taken by search engines is to randomly place new advertisements within different contexts and slots. Using a multiarmed bandit model, this approach allows the search engine to learn the click-through rates of ads over time (Wortman et al. 2007). It leads to a continuous tradeoff between exploration (learning rates more accurately) and exploitation (placing the best ads according to current estimates). While exploration has relatively low opportunity cost in search engines, the opportunity

cost is much higher in the context of TV ads. The model of Scott (2010) could be applied to efficiently estimate continuation rates using such experiments.

(2) The network could measure the continuation rates with a test panel, or an online video site such as Hulu.

(3) The burden could be placed on advertisers to supply accurate estimates of continuation rates for new ads. Such estimates could again be obtained online. To encourage accurate estimates, the network could charge the advertiser a fee if the ad's continuation rate falls below a minimum threshold.

(4) For advertisers that have previously purchased slots for other ad creatives, the network can use historical data on these advertisers as an initial estimate of the new ad's continuation probability.

It is likely that some combination of these approaches would be best in practice.

References

Crampton P, Shoham Y, Steinberg R, eds. (2006) *Combinatorial Auctions* (MIT Press, Cambridge, MA).

Ghosh A, Mahdian M (2008) Externalities in online advertising. *17th Internat. World Wide Web Conf.*, 161–168.

Goldfarb A, Tucker C (2011a) Online display advertising: Targeting and obtrusiveness. *Marketing Sci.* 30(3):378–404.

Goldfarb A, Tucker C (2011b) Privacy regulation and online advertising. *Management Sci.* 57(1):57–71.

Håstad J (1999) Clique is hard to approximate within $n^{1-\varepsilon}$. *Acta Mathematica* 182:105–142.

Karp R (1972) Reducibility among combinatorial problems. Miller R, Thatcher J, eds. *Complexity of Computer Computations* (Plenum Press, New York), 85–103.

Rassenti SJ, Smith VL, Bulfin RL (1982) A combinatorial auction mechanism for airport time slot allocation. *Bell J. Econom.* 13(2): 402–417.

Scott S (2010) A modern Bayesian look at the multi-armed bandit. *Appl. Stochastic Models Bus. Indust.* 26(6):639–658.

Wortman J, Vorobeychik Y, Li L, Langford J (2007) Maintaining equilibria during exploration in sponsored search auctions. *Proc. 3rd Workshop Internet Network Econom.* (*WINE*), 119–130.