

Pricing Power in Advertising Markets: Theory and Evidence

Matthew Gentzkow, *Stanford University and NBER**
Jesse M. Shapiro, *Harvard University and NBER*
Frank Yang, *Stanford University*
Ali Yurukoglu, *Stanford University and NBER*

September 2023

Abstract

Existing theories of media competition imply that advertisers will pay a lower price in equilibrium to reach consumers who multi-home across competing outlets. We generalize, extend, and test this prediction. We find that television outlets whose viewers watch more television charge a lower price per impression to advertisers. This finding helps rationalize well-known stylized facts such as a premium for younger and more male audiences on television. A quantitative version of our model whose only free parameter is a scale normalization can explain 35 percent of the variation in price per impression across owners of television networks, and aligns with recent trends in television advertising revenue. We use the model to quantify the impact of mergers and the effect of Netflix ad carriage on prices for linear television advertising. We then extend our analysis to social media markets where we find evidence of a premium for older audiences (who multi-home less), and we discuss implications for competition across ad formats.

JEL codes: L10, L82, M37

Keywords: platform competition, media markets

*E-mail: gentzkow@stanford.edu, jesse.shapiro@fas.harvard.edu, shuny@stanford.edu, ayurukog@stanford.edu. We thank our many dedicated research assistants for their contributions to this project. We also thank Sylvain Chassang, Matt Notowidigdo, Andrei Shleifer, and seminar participants at Stanford University, UC Berkeley, the University of Utah, Yale University, Hebrew University, the Department of Justice, the QME Rossi Seminar, Columbia University, UC Santa Cruz, the IIOC, Helsinki GSE, the Leuven Summer Event, the NBER, the Northwestern Conference on Antitrust Economics and Competition Policy, MIT, USC, the University of Chicago, Cornell University, the University of Rochester, the NYC Media Seminar, the WEAM, the MaCCI SICP, and especially discussants Jay Lee, Fiona Scott Morton, Brad Shapiro, Matthew McGranaghan, and Andrew Rhodes for helpful comments. We acknowledge support from the Eastman Professorship and the Population Studies and Training Center at Brown University, the Stanford Institute for Economic Policy Research, the National Science Foundation (SES 1260411), and the Toulouse Network for Information Technology. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the funding sources. We thank Mike Bailey for help executing advertising buys on Facebook and Hunt Allcott, Luca Braghieri, and Sarah Eichmeyer for sharing additional Facebook advertisement data. Portions of our analysis use data derived from a confidential, proprietary syndicated product owned by GfK US MRI, LLC, which is copyright MRI-Simmons 2021. Portions of our analysis use data from S&P Global Market Intelligence; SNL Financial LC. contains copyrighted and trade secret material distributed under license from SNL. The paper includes the researchers' own analyses calculated (or derived) based in part on data from The Nielsen Company (US), LLC and marketing databases provided through the Nielsen Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the Nielsen data are those of the researchers and do not reflect the views of Nielsen. Nielsen is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein. Similar caveats apply to other data sources.

1 Introduction

Both traditional and digital media receive substantial revenues from selling ads (e.g., Statista 2021). Prices for ads in these markets vary widely. On television, prices per impression can easily vary across programs or networks by a factor of three or more (e.g., Crupi 2009). Prices for online advertising exhibit similarly large variation (e.g., AdStage 2020). Which consumers' eyeballs command the highest prices is a key determinant of the incentive to produce content (Spence and Owen 1977; Wilbur 2008; Veiga and Weyl 2016). Pricing in advertising markets has become an important issue in antitrust policy (e.g., Competition and Markets Authority 2019).

Industry observers have long been puzzled by the large variation in the price of impressions across different groups of consumers. Perhaps the most famous example is the premium paid to advertise on television programs with younger audiences. The premium attached to younger audiences—who are sometimes known as the “coveted” or “target” demographic—is widely regarded as a major influence on content and scheduling, and persists despite the fact that older audiences tend to have greater purchasing power than younger audiences (Dee 2002; Surowiecki 2002; Einstein 2004; Pomerantz 2006; Goettler 2012; Gabler 2014).¹ Television advertisers also pay a premium for advertising to men relative to women (Papazian 2009) and (on a per-impression basis) for advertising on programs with larger relative to smaller audiences (Chwe 1998; Phillips and Young 2012; Goettler 2012).

In this paper, we develop an equilibrium model of an advertising market with competing outlets. The model implies that the price per viewer that an outlet charges for its advertisements in equilibrium is decreasing in the activity level of the outlet's audience, i.e., in the extent to which members of its audience visit competing outlets. We show that the model's predictions are borne out in data from the US television market, and can help explain well-known and potentially puzzling patterns such as premia for younger, more male, and (on a per-impression basis) larger audiences. A quantitative version of the model whose only free parameter is a scale normalization can explain 35.1 percent of the variation in price per impression across owners of television networks, and aligns with recent trends in television advertising revenue. An extension to social media adver-

¹Gabler (2014 pp. 3-4) writes that those over 49 “have been steadily disenfranchised by a ruthless, self-serving, myopic and ignorant dictator. That dictator is the eighteen to forty-nine demographic cohort, and it is the single most important factor in determining what we see, hear and read.” An advertisement run by the American Association for Retired Persons highlights the value of advertising to older audiences. Its text reads, “I may be gray, but my money is as green as it gets. Why is it all about 18-34, when they barely have a dime of their own?” (quoted in Newman 2012).

tising shows that older audiences are both the least active and command the highest prices online, and suggests that impressions on social media and television are imperfect substitutes from the perspective of advertisers.

Our model builds on a large theoretical literature on two-sided markets beginning with Rochet and Tirole (2003) and Anderson and Coate (2005), and extends Anderson, Foros, and Kind's (2018) model of advertising pricing in markets with multi-homing. In the model, each of a set of owners may own multiple outlets, and each outlet may have multiple advertising slots. Owners simultaneously announce prices to advertise on the slots they own, after which each of a set of advertisers decides which slots to purchase. Advertisers have homogeneous value functions that are submodular in the set of outlets on which they advertise. The number of slots on each outlet exceeds the number of advertisers, so slots are not rationed in equilibrium, and because advertisers are homogeneous, equilibrium is efficient. In particular, equilibrium follows the incremental pricing principle of Anderson, Foros, and Kind (2018): the price an owner commands for its slots is determined by the difference in an advertiser's value from advertising on all outlets versus all outlets except those controlled by the given owner.

An important special case of a submodular value function arises when advertisers face diminishing returns from multiple impressions to a given viewer, and viewers multi-home in a pattern that is invariant to advertisers' choices. In this case, the incremental value of an owner's advertising slots is determined by the overlap of its audience with those of other owners. In the special case of perfect diminishing returns, where advertisers value only the first impression to a given viewer, each owner's price per impression is determined solely by the fraction of its audience that is exclusive to that owner. The price per viewer that an owner can charge in equilibrium is then decreasing in the overall activity level of its audience, and increasing in the overall size of its audience.

We study these predictions empirically using data on television audiences and advertising prices from Nielsen's Ad Intel database, and audience survey data from GfK MRI. Consistent with the predictions of the model, we show that outlets whose audiences watch more television charge a lower price per impression for their ads. Turning to the demographic patterns that have received significant attention in the industry, we find that the younger, more male audiences that command a price premium are also those that watch the least television. We also find, consistent with prior evidence and with the predictions of the model, that outlets with larger audiences command higher prices per impression, even after accounting for the viewing intensity of their audiences.

We evaluate the fit of a quantitative version of the model. We consider a specification with

perfect diminishing returns in which a given viewer's probability of seeing an ad on a given outlet is proportional to the time that the viewer spends on the outlet. We consider a baseline specification in which viewers differ only in their viewing behavior, and one in which higher-income viewers are intrinsically more valuable to advertisers. Based on these specifications we use the audience survey data to calculate the incremental value of advertising on each owner's outlets, which in turn yields a prediction for the equilibrium price charged by each owner for its advertising slots. We find that the model's predictions are a good fit to observed prices. Predicted prices explain 34.0 to 35.1 percent of the variation in price per impression across owners, with a slope close to unity in the specification that incorporates viewer income. The model also rationalizes the fact that television advertising revenues have risen slightly in the last several years despite a decline in audience and impressions. This is true despite the fact that the model's quantitative predictions for relative prices across owners, and for trends in revenues over time, are based only on the audience survey data and therefore do not use any information on observed advertising prices. We use the quantitative model to show how competition shapes the incentive to attract viewers from different demographic groups.

We apply the quantitative model to two further questions. First, we study the effects of several recent mergers of television network owners on the combined advertising revenues of the merging entities. The model-predicted effects vary widely in ways that would be difficult to predict using standard concentration measures such as the Herfindahl-Hirschman Index (HHI), but are well-approximated by measures based on the overlap in the merging entities' audiences. Second, we study the effect of Netflix carrying advertising on the price of advertising on linear television. In a scenario where Netflix carries ads across its platform, and there is no change in audience behavior, we estimate a decline in price per impression of between 0.38 and 2.7 log points across television owners, with owners whose audience overlaps more with Netflix tending to experience larger declines in price per impression.

In the final part of the paper, we extend our analysis to social media. We specify a model in which ads may be sold and targeted at the viewer level on social media. The model allows diminishing returns to operate more strongly within television or within social media than between the two, such that ads on the two formats may be imperfect substitutes from the perspective of advertisers. When cross-format substitutability is limited, the model predicts that the age-price relationship on social media should be reversed relative to television, with the young—the heaviest users of social media—commanding the lowest prices, and the old—the lightest users—being the

“coveted” group. Alternative explanations for the youth premium, such as the young having more malleable preferences (Surowiecki 2002), would not all share this prediction. Using data on prices of Facebook advertisements collected as part of a series of experiments including our own, we show that ads targeting the oldest users indeed command the highest prices. We compare model fit as we vary a parameter governing cross-format diminishing returns and find that the best-fitting parameter implies that television and social media ads are imperfect substitutes.

The primary contribution of this paper is to show that the predictions of a model of a competitive advertising market with multi-homing are a good match, both qualitatively and quantitatively, to existing and novel facts about important real-world markets. In contrast to many prior studies of advertising markets (e.g., Kaiser and Wright 2006; Bel and Domènech 2009; Wilbur 2008; Fan 2013; Chandra and Kaiser 2014; Jeziorski 2014; Berry, Eizenberg, and Waldfogel 2016; Zubanov 2020), our quantitative model explicitly derives the price of advertising on a given outlet from a microfounded equilibrium model with a multi-homing audience.² Multi-homing is essential to the model’s implications. In contrast to prior work that incorporates audience demographics into a model of advertiser demand (e.g., Wilbur 2008; Liao, Sorensen, and Zubanov 2020), our model can explain demographic premia in advertising prices without assuming that advertisers intrinsically value certain demographic characteristics.

Our analysis provides a unified explanation of several facts, some of which are new to the literature. There is a folk wisdom in television advertising that it is more expensive to advertise to groups that are harder to reach (Surowiecki 2002; Papazian 2009; Gabler 2014).³ Some have questioned the logic of this proposition.⁴ We provide what is to our knowledge the first systematic evidence on the relationship between an outlet’s advertising prices and the activity levels of its audience, and the first depiction of this relationship grounded in a quantitative economic model. We also systematically document advertising premia related to audience age, gender, and size. In

²Gentzkow, Shapiro, and Sinkinson (2014) incorporate a microfounded model of advertising with multi-homing consumers into a structural model of newspapers’ choice of political affiliation, but allow for only a small number of outlets, and do not study the cross-sectional variation in advertising prices implied by their model. Shi (2016, Section 3.2) proposes an inverse demand function for advertising that is motivated by a model with multi-homing consumers and diminishing returns to impressions. Greenwood, Ma, and Yorukoglu (2021) calibrate a macroeconomic model in which consumers can consume multiple media goods but do not receive multiple advertisements from the same advertiser. Prat and Valletti (2022, Section 4) simulate effects of platform mergers under various assumptions about overlap in their audience, though using a microfoundation different from ours.

³Papazian (2009, p. 134) writes that, “as a rule, shows that pull higher proportions of easy-to-get heavy tube watchers come in at lower [cost per thousand impressions] than those that rely less on this preponderantly lowbrow segment and more on upscale audiences.”

⁴Surowiecki (2002) writes that, “by this logic, advertisers ought to pay top dollar to reach shepherders in Uzbekistan.”

the case of social media, while some industry sources report a premium for older audiences on social media (e.g., Ampush 2014, Strikesocial 2017), we are not aware of prior evidence in the academic literature showing that transaction prices in the US are greater for Facebook ads targeted to older users.⁵ Turning to time trends, the fact that television advertising revenues have grown despite a declining audience has been noted as a puzzle, but is predicted by our quantitative model.⁶

The paper also makes a contribution to the theoretical literature on advertising in two-sided markets with multi-homing. In particular, we generalize the incremental pricing result in Anderson, Foros, and Kind (2018) to allow for arbitrary submodular value and ownership structure. Unlike Ambrus, Calvano, and Reisinger (2016) and Anderson and Peitz (2020), we do not model the determination of the number of advertising slots. Unlike Athey, Calvano, and Gans (2018), we do not allow heterogeneity among advertisers in our baseline analysis, though in an extension we show that incremental pricing holds when the extent of heterogeneity is small or when owners can charge advertiser-specific prices. Unlike Prat and Valletti (2022), we do not focus on the effects of the ad market on competition among advertisers, though we do allow for some interactions among advertisers in an extension. As in Anderson, Foros, and Kind (2018), our model allows for a very rich description of viewers' choices of which outlets to watch, a feature we take advantage of when developing the model's quantitative implications.⁷

None of the evidence we present constitutes a pure test of the forces in the model, which would require changing the competitive environment while holding all other conditions constant. Accordingly, each individual piece of evidence is subject to alternative interpretations, some of which we highlight in the paper and test in sensitivity analysis. However, to us, the fact that a model that builds on a large body of economic theory can explain such a wide range of facts—both qualitatively and quantitatively, and across markets, outlets, and over time—suggests that the economic forces we highlight are important for understanding pricing power in competitive advertising markets.

⁵Lambrecht and Tucker's (2019, Table 7) analysis of average suggested bids for a STEM career information campaign on the Facebook platform across 191 countries indicates that average suggested bids are higher for ads targeted to females. The analysis does not show clear differences in average suggested bids by the age of the target users (columns 1 and 2) but does show evidence of interactions between age and gender (column 3). Our analysis differs in using transaction price data from campaigns in the US rather than suggested bid data from a campaign across 191 countries.

⁶The Economist (2021) writes that, "the Tokyo games illustrate a puzzle: as audiences decline, the TV-ad market is holding up."

⁷As in Prat's (2018) analysis of media outlets' political power and Armstrong and Vickers' (2022) analysis of oligopoly pricing with limited consideration, our analysis of outlets' pricing power emphasizes the importance of individual consumers' allocation of attention.

The remainder of the paper proceeds as follows. Section 2 presents our model and its implications. Section 3 describes our data and variable definitions. Section 4 presents our key findings about the determinants of advertising prices on television. Section 5 presents our quantitative implementation of the model, its fit to the data, and its applications. Section 6 presents our extension to social media. Section 7 concludes.

2 Model

There is a set of outlets \mathcal{J} . A given owner can own multiple outlets, and we define a partition \mathcal{Z} on the set of outlets that describes the ownership structure, using the notation $Z \in \mathcal{Z}$ to refer both to a cell of the partition and to the owner of the outlets in that cell. Each outlet has available K advertising slots, each of which can be sold to one of the N advertisers in the set \mathcal{N} . We assume that $N \leq K$, i.e., that advertising slots are not scarce. We let $\mathcal{P}(\cdot)$ denote the power set operator, so that $\mathcal{P}(Z)$ denotes the set of all bundles of owner Z 's outlets. A bundle $B \in \mathcal{P}(Z)$ might consist, for example, of a single outlet owned by owner Z , or it might contain all outlets owned by owner Z .

The game proceeds as follows. Each owner Z simultaneously announces, for each bundle $B \in \mathcal{P}(Z)$ of its outlets, a price p_B at which it will sell one slot on each outlet $j \in B$ to any advertiser, with $p_B = \infty$ denoting that a given bundle B is unavailable. Advertisers then simultaneously decide which, if any, bundles to buy. When all advertisers have moved, ads are shown and the game ends. Since we allow each owner to own an arbitrary number of outlets, it is without loss of generality to assume that each advertiser can purchase at most one slot on each outlet.

The payoff of an owner is given by the sum of the prices p_B of all bundles B that the owner sells. The payoff of an advertiser n that buys slots in a set of bundles $\mathcal{B}_n \subseteq \mathcal{P}(\mathcal{J})$ is given by

$$V(\{j : j \in B \in \mathcal{B}_n\}) - \sum_{B \in \mathcal{B}_n} p_B,$$

where $V(\cdot)$ is a non-negative *value function* that is monotone in the set-inclusion order.⁸ We capture the idea that there are diminishing returns to advertising by assuming that $V(\cdot)$ is submodular:

⁸Formally, for any sets of outlets $\mathcal{J}', \mathcal{J}'' \subseteq \mathcal{J}$,

$$\mathcal{J}' \subseteq \mathcal{J}'' \implies V(\mathcal{J}') \leq V(\mathcal{J}'').$$

an advertiser derives less incremental value from an outlet when adding it to a larger bundle.⁹ We assume that the value function $V(\cdot)$ is the same across advertisers, i.e., the advertisers are homogeneous.

Our main result is that each owner is able to extract the incremental value of the outlets it controls. To state this result, for each bundle $B \subseteq \mathcal{J}$, let the *incremental value* v_B be given by

$$v_B = V(\mathcal{J}) - V(\mathcal{J} \setminus B),$$

i.e., the value to an advertiser of advertising on all outlets rather than all outlets except those in B . We assume that every outlet in \mathcal{J} has positive incremental value, $v_j > 0$ for all $j \in \mathcal{J}$. We use subgame perfect Nash equilibrium in pure strategies as our solution concept and hereafter refer to it as *equilibrium*.

Theorem 1. (Incremental pricing.) *There exists an equilibrium. In any equilibrium, all advertisers buy slots on all outlets, and the payment by each advertiser to each owner Z is given by $p_Z^* = v_Z$.*

All proofs are given in Appendix A. The proof of Theorem 1 shows that, in any equilibrium, every owner Z finds it optimal to offer the maximal bundle $B = Z$ for a price given by the bundle's incremental value v_Z . The intuition can be understood as follows. Consider any equilibrium. Every owner Z has the option of offering only the maximal bundle Z at a price equal to v_Z . With this offering, regardless of what other owners' prices are, any owner Z can always sell the bundle to every advertiser. To see why, note that for any set of outlets $S \subseteq \mathcal{J} \setminus Z$ that an advertiser could have purchased from other owners, the additional value of purchasing bundle Z is

$$V(S \cup Z) - V(S) \geq V(\mathcal{J}) - V(\mathcal{J} \setminus Z) = v_Z,$$

where the inequality is due to the submodularity of $V(\cdot)$. Thus, in equilibrium, owner Z must secure a profit of at least $N \cdot v_Z$. Now, suppose toward contradiction that some owner Z earns a profit strictly higher than $N \cdot v_Z$. Then, some advertiser n must pay owner Z more than v_Z . Thus, there exists an outlet $j \in \mathcal{J} \setminus Z$ on which advertiser n does not purchase an ad slot, because otherwise, by the monotonicity of $V(\cdot)$, the advertiser could profitably deviate by not trading with owner Z

⁹Formally, for any sets of outlets $\mathcal{J}', \mathcal{J}'' \subseteq \mathcal{J}$,

$$\mathcal{J}' \subseteq \mathcal{J}'', j \in \mathcal{J} \setminus \mathcal{J}'' \implies V(\mathcal{J}' \cup \{j\}) - V(\mathcal{J}') \geq V(\mathcal{J}'' \cup \{j\}) - V(\mathcal{J}'').$$

at all. But, also because of the monotonicity of $V(\cdot)$, this implies that the owner Z' of outlet j can profitably deviate by adding the ad slot on outlet j to the existing bundle that advertiser n purchases from owner Z' and slightly increasing the total price. The existence of a profitable deviation for owner Z' contradicts that the conjectured strategy is part of an equilibrium, thus implying that in equilibrium each owner must secure a profit of no more than $N \cdot v_Z$, and hence that incremental pricing holds.

Remark 1. Appendix A.3 shows that versions of incremental pricing hold with rationing ($N > K$), with partially increasing returns, with heterogeneity in advertisers' value functions, and with alternative market institutions including unbundled pricing, bargaining, or auctioning of ad slots.

2.1 Special Cases

The general model abstracts from, but is consistent with, many types of viewer-level microfoundations. Here we define an important special case that we use to illustrate comparative statics and, later, to quantify the implications of the model.

Definition. (Viewer-level model.) In the **viewer-level model**, each viewer $i \in \mathcal{I}$ is endowed with an intrinsic value a_i to advertisers. Each viewer i sees ads on outlet j with probability $\eta_{ij} \in [0, 1]$, independently across outlets. Owners are partitioned into formats F with \mathcal{F} the set of all formats. An advertiser who reaches a given viewer i exactly $M \in \mathbb{N}^{|\mathcal{F}|}$ times in each format obtains value $a_i \cdot u(M)$ where $u(\cdot)$ is monotone and submodular in M , and has decreasing differences in each argument with $u(0) = 0$.¹⁰ The **reach-only model** is a special case of the viewer-level model in which there is a single format and $u(M) = \mathbf{1}_{M>0}$.

Appendix A proves that the viewer-level model induces a value function that is monotone and submodular.

In the viewer-level model, each viewer i is endowed with an intrinsic value a_i that might reflect, for example, differences in income across households. Each viewer is also endowed with viewing behavior represented by the probability η_{ij} of seeing ads on outlet j . Outlets are partitioned into formats, for example television and social media. Advertisers' value for ads exhibits diminishing returns, with initial impressions being more valuable than subsequent ones. Diminishing returns may operate more strongly within than across formats. In the **reach-only model**, there is a single

¹⁰That is, $M \geq M' \implies u(M) \geq u(M')$; $u(M) + u(M') \geq u(\max\{M, M'\}) + u(\min\{M, M'\})$ for all $M, M' \in \mathbb{N}^{|\mathcal{F}|}$ where \max and \min denote the componentwise maximum and minimum; and $u(M + e(m+1)) - u(M + em)$ is nonincreasing in $m \in \mathbb{N}$ for all $M \in \mathbb{N}^{|\mathcal{F}|}$ and all standard basis vectors e .

format and only first impressions matter. Theorem 1 implies the following characterization of revenues in the reach-only model.

Corollary 1. *In the reach-only model the price per viewer of owner Z is given by p_Z^*/λ_Z where*

$$p_Z^* = \sum_{i \in \mathcal{I}} a_i \left(\eta_{iZ} \prod_{Z' \neq Z} (1 - \eta_{iZ'}) \right)$$

is the revenue of owner Z ,

$$\lambda_Z = \sum_{i \in \mathcal{I}} \eta_{iZ}$$

is the expected number of viewers seeing an ad on an outlet owned by owner Z , and

$$\eta_{iZ} = 1 - \prod_{j \in \mathcal{Z}} (1 - \eta_{ij})$$

is the probability that viewer i sees an ad on an outlet owned by owner Z .

Corollary 1 states that, in the reach-only model, the price per viewer p_Z^*/λ_Z is given by the weighted share of owner Z 's audience that is exclusive to that owner, in the sense of not seeing ads on any other owner's outlets.

2.2 Comparative Statics

We provide two comparative statics results in the reach-only model. For these comparative statics we suppose that all viewers are intrinsically equally valuable; that is, that $a_i = a > 0$ for all i . We also suppose that viewers can be partitioned into a set G of mutually exclusive demographic groups g , with viewing behavior homogeneous within groups, i.e., with $\eta_{iZ} = \eta_{i'Z} = \eta_{gZ}$ for all Z for any two viewers i, i' in the same group g . We define a group $g \in G$ to be *less active* than group $h \in G$ if $\eta_{gZ} \leq \eta_{hZ}$ for all $Z \in \mathcal{Z}$. We let

$$\sigma_{gZ} = \frac{\sum_{i \in g} \eta_{gZ}}{\lambda_Z}$$

denote the share of owner Z 's audience that comes from group g .

We first show that, all else equal, an owner commands a larger price premium for its viewers if its viewers come from less active groups.

Proposition 1. *Suppose that owner Y draws a larger share of its audience from the less active group g and a smaller share of its audience from the more active group h than owner Z , in the sense*

that $\sigma_{gY} \geq \sigma_{gZ}$ and $\sigma_{hY} \leq \sigma_{hZ}$, and that the two owners have equal total audience sizes, $\lambda_Y = \lambda_Z$, and equal shares of audience from groups other than g and h , $\sigma_{g'Y} = \sigma_{g'Z}$ for all $g' \neq g, h$. Then owner Y has a higher equilibrium price per viewer than owner Z , $p_Y^*/\lambda_Y \geq p_Z^*/\lambda_Z$.

The inequality in the conclusion of Proposition 1 is strict if $\eta_{gZ'} < \eta_{hZ'}$ for some $Z' \notin \{Y, Z\}$, $\sigma_{gY} > \sigma_{gZ}$, and $\eta_{gZ} \in (0, 1)$ for all Z . To see the intuition for Proposition 1, consider the pair of owners Y and Z . A more active viewer is more likely to view outlets owned by any other owner $Z' \notin \{Y, Z\}$, resulting in a lower incremental value for both owners Y and Z . Thus, because a smaller share of its audience comes from the more active group, owner Y can command a higher price per viewer than owner Z .

Remark 2. Multi-homing is essential to the result in Proposition 1. If there were only a single outlet ($J = 1$), or if each group were to watch only one outlet with positive probability, then by Corollary 1 each owner's price per viewer would be invariant to the group composition of its outlets' audiences.

Remark 3. Competition is essential for the result in Proposition 1. If a single owner were to own all outlets, then by Corollary 1 the owner's price per viewer would be invariant to the group composition of its outlets' audiences.

We next show that, all else equal, an owner commands a larger price premium for its viewers if the owner attracts a larger share of the total audience.

Proposition 2. *Suppose that owner Y draws a larger audience from every group than owner Z in the sense that for some $\delta \geq 1$, we have $\eta_{gY} = \delta \eta_{gZ}$ for all $g \in G$. Then owner Y has a higher price per viewer than owner Z , $p_Y^*/\lambda_Y \geq p_Z^*/\lambda_Z$.*


The inequality in the conclusion of Proposition 2 is strict if $\delta > 1$ and $\eta_{gZ} \in (0, 1)$ for all Z . To see the intuition for Proposition 2, consider the pair of owners Y and Z and suppose that there is only one group of viewers. If a viewer sees ads on an outlet of owner $Z' \notin \{Y, Z\}$, then by Corollary 1 the viewer does not contribute to the revenue of owners Y or Z . Among the remaining viewers, the share of those watching Y 's outlets who also watch Z 's outlets is smaller than the reverse (since for any viewer the probability of watching Y is greater than the probability of watching Z). This implies that Y has a larger share of exclusive viewers, and so commands a higher price per viewer.

Remark 4. Appendix A.3 shows that, when each owner owns a single outlet, statements analogous to Proposition 1 and Proposition 2 hold in the viewer-level model with a single format and general diminishing returns.

2.3 Illustrative Example

To illustrate the intuition for the comparative statics, consider an example of the setting in Section 2.2 with one advertiser, three outlets $\{1, 2, 3\}$ owned separately by three owners, three groups of viewers $\{f, g, h\}$ with equal size, and intrinsic value $a = 1$. Figure 1 shows four cases.

Figure 1: **Illustrative Example**



Notes: Each plot represents one of the examples discussed in Section 2.3. A solid arrow pointing from a group in $\{f, g, h\}$ toward an outlet in $\{1, 2, 3\}$ means that any viewer from the group watches the outlet with probability 1. A dashed arrow pointing from a group in $\{f, g, h\}$ toward an outlet in $\{1, 2, 3\}$ means that any viewer from the group watches the outlet with probability $\frac{1}{2}$. Nodes in darker color are the focal outlets for the discussion in the text.

First, suppose that all viewers watch exactly one outlet, as in Panel (a) of Figure 1. In this case, each owner acts as if it is a monopoly seller by charging a price per viewer equal to 1, because for each viewer of an owner’s ads, there is no other owner that can show ads to the same viewer.

Second, suppose that all viewers watch two different outlets, as in Panel (b) of Figure 1. In this case, no owner can command a positive price for its ads because, for each viewer of an owner’s ads, there is another owner that can show ads to the same viewer. In both Panel (a) and Panel (b), each owner commands one third of the total audience. The difference between the two settings is thus not in the concentration of the audience sizes, but rather in the extent of competition to reach any given viewer.

Third, suppose that viewers from group f only watch outlet 1, viewers from group g always watch outlets $\{1, 2\}$ and watch outlet 3 with probability $\frac{1}{2}$, and viewers from group h watch outlet 3 with probability $\frac{1}{2}$ and no other outlet, as in Panel (c) of Figure 1. Compared to outlet 2, outlet 3 has the same audience size but has a larger share of its audience from the less active group h . Consistent with Proposition 1, owner 3 commands a higher price per viewer than owner 2, because all viewers of outlet 2 also watch at least one other outlet, whereas half of the viewers of outlet 3 exclusively watch outlet 3.¹¹

Fourth, suppose that viewers from group g only watch outlet 2, and viewers from both of the other two groups always watch outlet 3 and watch outlet 1 with probability $\frac{1}{2}$, as in Panel (d) of Figure 1. Compared to outlet 1, outlet 3 has the same group composition but has a larger audience size. Consistent with Proposition 2, owner 3 commands a higher price per viewer than owner 1, because all viewers of outlet 1 also watch another outlet, whereas half of the viewers of outlet 3 exclusively watch outlet 3.

2.4 Discussion

Bundling. We assume that each owner can offer an arbitrary menu of bundles of its outlets, which includes selling the slots on all outlets as a single bundle (pure bundling) and selling only individual slots (no bundling). Bundled ad pricing is common in practice.¹² Appendix A.3 shows that, when some owners may not be able to bundle their outlets, a version of incremental pricing

¹¹Proposition 1 applies with strict inequality in this example because viewers from group h watch outlet $1 \neq 2, 3$ with strictly smaller probability than viewers from group g .

¹²In the context of television advertising, for example, Weprin (2015) reports that network owners arguing for aggregating audiences across their outlets is “a theme that became particularly prevalent last year, with programmers pitching all of their channels together in big ad deals, rather than focusing on individual channels.” See also Geskey (2016, p. 525), WARC (2001), Patel (2015), and Garland (2002).

holds in any equilibrium, though our results do not guarantee existence of an equilibrium in this case.

Alternative market institutions. In the model, we assume that owners simultaneously post prices. In practice, ad sales can take different forms including bargaining and auctioning. Appendix A.3 shows that incremental pricing holds if owners bargain with advertisers a la Nash-in-Nash (Lee, Whinston, and Yurukoglu 2021), or compete via reserve prices in first-price auctions.

Ad effectiveness and partially increasing returns. We assume that the value function $V(\cdot)$ is commonly known. If instead advertisers are uncertain about the returns to their advertisements (Aral 2021), we can instead interpret $V(\cdot)$ as an expected return. We also assume that $V(\cdot)$ is submodular, which implies a form of diminishing returns. Appendix A.3 extends the comparative statics in Section 2.2 to a setting in which there are increasing returns to advertising at small numbers of impressions (as in, e.g., Dubé, Hitsch, and Manchanda 2005).

Heterogeneous advertisers. In the model, we assume that all advertisers have the same value function. In practice, different advertisers can have different values for the same set of outlets. If owners can post advertiser-specific prices, then the result is parallel to that in Theorem 1 in the sense that each advertiser pays each owner its own incremental value for the owner’s outlets. If owners cannot post advertiser-specific prices, Appendix A.3 shows that incremental pricing holds if heterogeneity among the advertisers is sufficiently small, in a precise sense, compared to the incremental value of a single outlet. When heterogeneity among the advertisers is large and advertiser-specific prices are not allowed, an owner might profitably engage in screening and price discrimination by posting a menu of bundles, and incremental pricing might not hold.

Content pricing. Content owners such as television networks sometimes charge fees to viewers, either directly via “over the top” subscriptions or indirectly via bundlers like cable networks. The implications of our model are invariant to including such fees in the owners’ payoffs, provided that the fees are invariant to the outcome of the advertising game. This would be true if, for example, fees to viewers are set prior to the advertising game, or prior to viewers’ knowledge of its outcome.

Ad inventory and endogenous response of viewers. In the viewer-level model, viewers’ choices of which outlets to view are not affected by the outcome of the game. We may interpret this ei-

ther as a scenario in which viewers do not care about advertising or, following Anderson, Foros, and Kind (2018), as a scenario in which viewers make viewing decisions without knowing the outcome of the game. Appendix A.3 shows that our model can accommodate settings in which viewers care about advertising and make viewing decisions knowing the outcome of the advertising game provided that the impact of advertising on viewing is additively separable across ad slots. When this assumption is violated, an owner might profitably withhold some ad slots to increase the viewership of the remaining ad slots.

3 Data

We conduct our main analysis of the television market using data from 2015. Here we describe the concepts we measure for 2015. For sensitivity analysis and extensions we use data from 2014 through 2019, with concepts defined and calculated in an analogous manner to those we describe below for 2015. Appendix Figure 1 reports sensitivity analysis replacing data from 2015 with data from 2014 or 2016.

3.1 Television Advertising Prices, Audience, and Ownership

We obtain data on broadcast and cable television viewership and advertisement pricing in 2015 from Nielsen’s Ad Intel product (The Nielsen Company 2019). For each advertisement the data includes the telecast (e.g., NBC Nightly News, June 1), program (e.g., NBC Nightly News), daypart (e.g., early fringe), and network (e.g., NBC). It also includes the duration (e.g., 30 seconds) of the advertising spot, an estimate of its cost, and an estimate of the number of impressions (live viewers) for the associated telecast. We omit from all calculations any advertisements with zero cost or duration. We standardize the cost to a 30-second-spot basis by dividing the cost by the duration of the advertisement (in seconds) and multiplying by 30.

Advertising cost estimates in the AdIntel data are based primarily on information obtained from SQAD at the month-network-daypart level for cable television and from networks at the month-program level for broadcast television (The Nielsen Company 2017).¹³ For consistency we define our notion of an outlet j to be a network-daypart.¹⁴ Appendix Figure 2 reports results when using

¹³Hristakeva and Mortimer (forthcoming) use data from SQAD to analyze price dispersion and discounting in the television advertising market.

¹⁴In cases where a telecast spans multiple dayparts, we assign it to the daypart that contains the largest share of broadcast time.

network as our notion of an outlet, and also (for broadcast television) when using program.

For each outlet, we calculate total impressions across all advertisements and divide by the number of hours in the corresponding daypart in a 52-week year to get a measure of total impressions per hour, which we may think of as an outlet-level analogue λ_j of the concept λ_Z defined in Section 2.2. For each outlet, we also calculate the total (standardized) cost of all advertisements and divide by the number of hours in the corresponding daypart in a 52-week year to get a measure of total cost per hour, which we may think of as an outlet-level analogue p_j^* of the concept p_Z^* defined in Section 2.2. Finally, for each outlet we divide total cost per hour by total impressions per hour to obtain the average price per impression of a 30-second spot on the outlet, which we may think of as an outlet-level analogue p_j^*/λ_j of the concept p_Z^*/λ_Z defined in Section 2.2.

For each advertisement we also have information on the number of impressions by age (in bins) and gender for the associated telecast.¹⁵ From this information we compute the share of each outlet's impressions that are to adults (aged 18 and over) and the share among impressions to adults that are to females. We also compute the average age of each outlet's adult impressions by imputing each bin to its midpoint value and imputing the oldest bin (65+) to age 75.¹⁶

For a subset of advertisements representing 99.9 percent of all impressions, we also have information on the distribution of impressions across household income bins for the associated program.¹⁷ From this information we compute the average household income of each outlet's adult impressions (among those for which we measure income) by imputing each bin to its midpoint value, and imputing the highest-income bin (\$125,000+) to \$175,000.

We obtain from SNL Kagan, a product of S&P Global Market Intelligence, information on the ownership of cable networks in 2015 (S&P Global Market Intelligence 2019). We supplement this with other publicly available information, including on the owners of broadcast networks. We form the ownership partition \mathcal{Z} by assigning each outlet to its majority owner, treating joint ventures as independent ownership groups. We perform analogous calculations to those at the outlet level to compute the price per impression p_Z^*/λ_Z and audience demographics of each owner $Z \in \mathcal{Z}$.

¹⁵The age bins are 2-5, 6-8, 9-11, 12-14, 15-17, 18-20, 21-24, 25-29, 30-34, 35-39, 40-44, 45-49, 50-54, 55-64, and 65+ years.

¹⁶Using information on each advertisement's advertiser, we also compute the share of each outlet's adult impressions that are to advertisements in each of a set of industry categories, which we use in sensitivity analysis.

¹⁷The bins are 0-20, 20-30, 30-40, 40-50, 50-60, 60-75, 75-100, 100-125, and 125+, all in thousands of dollars. For programs representing 96.2 percent of all impressions, we have information on the distribution of impressions across household income bins for each month, from which we compute an annual average for the program. For programs representing 3.7 percent of all impressions, we have information on the distribution of impressions across household income bins for each of a subset of the program's telecasts, from which we compute an average for the program. We associate each advertisement with the average distribution of impressions for its respective program.

3.2 Audience Survey

From GfK MRI’s 2015 Survey of the American Consumer we obtain, for each of 23978 adult respondents, information on times of day spent watching television in the form of a week-long diary, as well as the implied total weekly television viewing time (GfK Mediamark Research and Intelligence 2017). We compute a measure of each respondent’s total viewing time in each daypart by allocating viewing time in each time slot to AdIntel dayparts in proportion to the share of the time slot that is contained within each daypart. We also obtain measures of each respondent’s viewership of each of 227 broadcast television programs,¹⁸ and time spent watching each of 115 cable television networks in the preceding week. We successfully match 173 broadcast programs and 97 cable television networks to their counterparts in AdIntel.¹⁹

We use the data on viewership by daypart, broadcast program, and cable network to construct a measure of the time that each respondent i viewed each outlet (network-daypart) j . To do this, we first allocate the viewing time of broadcast programs to their respective network-dayparts.²⁰ If in a given daypart there is viewing time that cannot be attributed to broadcast programs, we allocate that time to the cable networks in proportion to the respondent’s reported viewing time of each network.²¹

We thus arrive at a measure T_{ij} of the time each respondent i viewed each outlet j . We compute each respondent’s total weekly viewing time by summing over outlets. For each outlet j , we compute the weighted average log of total weekly viewing time of its viewers, weighting each viewer by her viewing time on outlet j .²² We treat average log total weekly viewing time as a

¹⁸The data record the number of times a respondent watches a broadcast program in a typical week (for some broadcast programs) or month (for others). We convert the latter into weekly viewing by allocating monthly viewing time evenly across weeks.

¹⁹The broadcast programs we match span 6 networks. Some programs (e.g., those on PBS) and some cable television networks (e.g., the Disney Channel, QVC) are excluded from AdIntel because they do not carry standard advertising spots.

²⁰Specifically, we associate each program with a network-daypart following the 2014–15 United States Network Television Schedule (Wikipedia 2022). We supplement this source with information from the *Sunday News Journal* (The News Journal 2015) and other publicly available information on the program’s network and airtime, and use information on the respondent’s geographic location to adjust for time zones. If the total duration of broadcast programs allocated to a given daypart exceeds the respondent’s total viewing time of that daypart, we assume that all viewing during that daypart was to broadcast programs, and we allocate the respondent’s viewing time during that daypart to broadcast networks in proportion to the respondent’s viewing time of each broadcast network’s programs. We assume that each broadcast program viewing has the same duration, and choose that duration so that the ratio of average total broadcast viewing hours and average total cable viewing hours is equal to the one in Nielsen Local TV View (The Nielsen Company 2021).

²¹If the respondent reports zero viewing time for all cable networks, we instead allocate all viewing time during that daypart to broadcast networks in proportion to the respondent’s viewing time of each network.

²²We exclude from this calculation any respondent with zero total weekly viewing time.

measure of the overall activity level of outlet j 's audience.

We also obtain information on each respondent's gender, age (in bins),²³ reported attentiveness to different broadcast and cable programs, and attitudes towards television advertising. We obtain information on household income (in bins), which we use in some specifications as a proxy for the intrinsic value a_i in the viewer-level model.²⁴ For 2019, we additionally obtain data on time spent watching Netflix, which we use in counterfactual analysis.

4 Evidence on the Determinants of Advertising Prices

Proposition 1 predicts that more active audiences will command a lower advertising price per impression. Figure 2 shows that this prediction is borne out in a comparison of television outlets. Each panel shows a binned scatterplot of an outlet's log(price per impression) against the average log(weekly viewing time) of the outlet's audience. Panel A includes baseline controls including for daypart; Panel B additionally includes controls for log(impressions per hour).

Both panels of Figure 2 show a clear negative relationship between log(price per impression) and average log(weekly viewing time). The magnitude of the relationship is large: in Panel B, for example, moving from the bottom to the top decile of average log(weekly viewing time) corresponds to a decline in log(price per impression) of roughly 163 log points.

Proposition 1 also makes predictions about which demographic groups should command a price premium in the advertising market. Appendix Figure 4 shows that older viewers watch more television than younger viewers and that female viewers watch more television than male viewers. The logic of Proposition 1 would lead us to expect that outlets with an older audience would command a lower advertising price than outlets with a younger audience, and likewise for outlets with a more female audience. Figure 3 shows that these predictions are borne out in the data: outlets with older, more female audiences tend to exhibit both lower log(price per impression) and higher average log(weekly viewing time). The price differences between outlets with different demographics are large.

Proposition 2 predicts that outlets with a larger audience will command a higher advertising

²³The age bins are 18, 19, 20, 21, 22-24, 25-29, 30-34, 35-39, 40-44, 45-49, 50-54, 55-59, 60-64, 65-69, 70-74, and 75+ years. We impute each individual's age to the midpoint of the corresponding bin, imputing the highest bin to 77.


²⁴The household income bins are 0-4999, 5000-9999, 10000-14999, 15000-19999, 20000-24999, 25000-29999, 30000-34999, 35000-39999, 40000-44999, 45000-49999, 50000-59999, 60000-74999, 75000-99999, 100000-149999, 150000-199999, 200000-249999, and 250000+ US dollars. We impute each household's income to the midpoint of the corresponding bin, imputing the highest bin to \$300,000.

price per impression. Figure 4 shows that this prediction is borne out in the data. Panel A shows a binned scatterplot of $\log(\text{price per impression})$ against $\log(\text{impressions per hour})$ with baseline controls. Panel B additionally controls for average $\log(\text{weekly viewing time})$. Both plots show that a larger audience is associated with a higher price per impression, consistent with the logic of Proposition 2. The association is economically meaningful: in Panel B, for example, moving from the bottom to the top decile of $\log(\text{impressions per hour})$ corresponds to an increase in $\log(\text{price per impression})$ of roughly 37 log points.

Columns (1) and (2) of Table 1 summarize the patterns in Figures 2, 3, and 4. Table 1 and Appendix Figure 3 report results controlling for the average household income of an outlet's audience. Appendix Figure 3 additionally shows sensitivity to controlling for measures of the attentiveness to television and attitudes toward advertising of the outlet's audience,²⁵ and the industry mix of the outlet's advertisers.

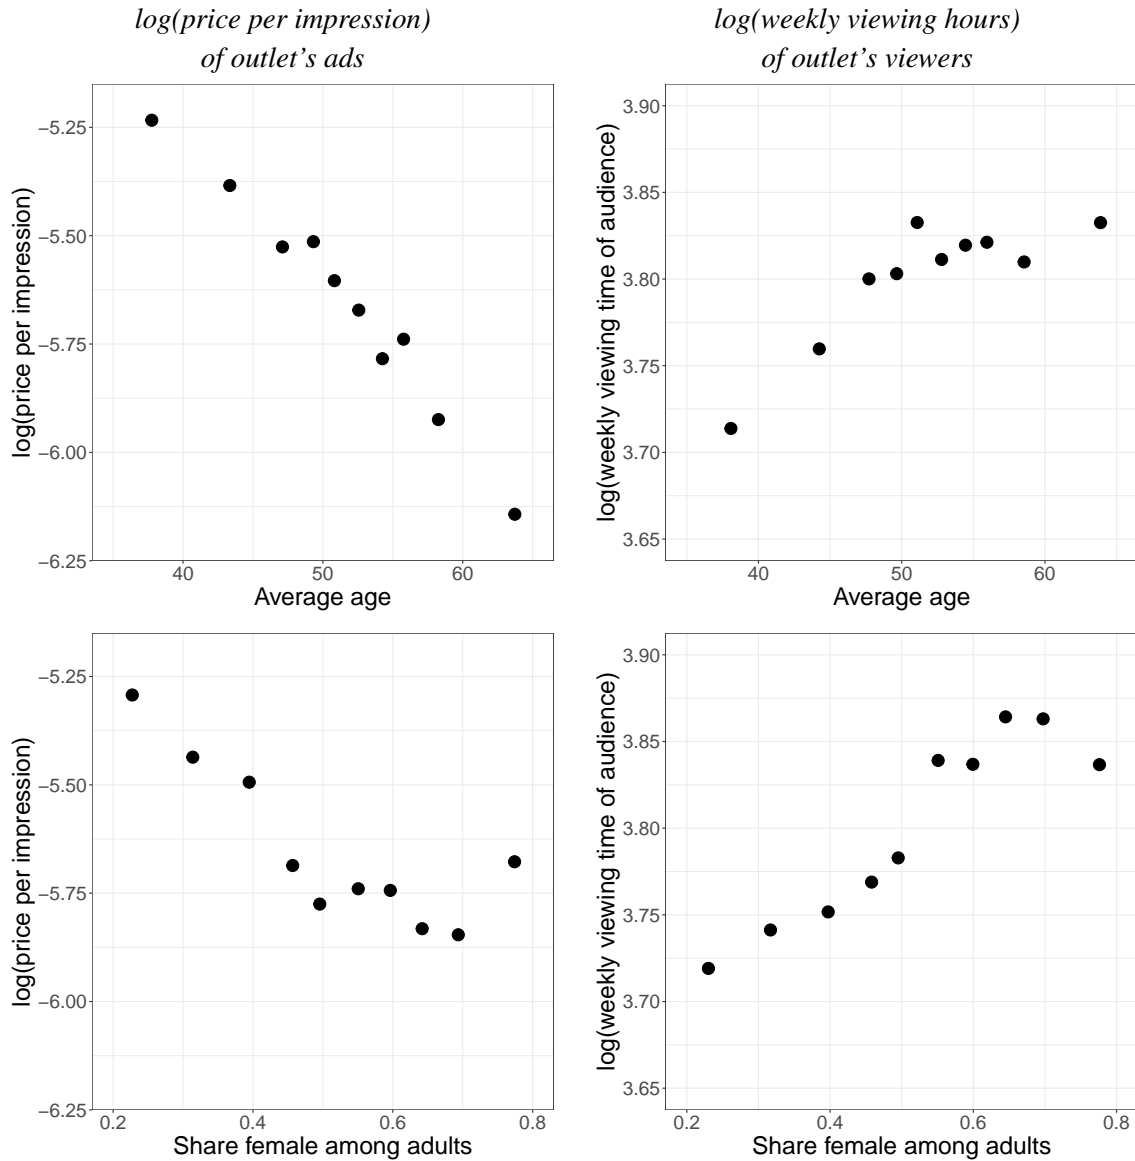
²⁵McGranaghan, Liaukonyte, and Wilbur (2022) find that younger audiences pay less attention to television advertising than older audiences. Alwitt and Prabhaker (1994) find that demographic characteristics are not strong predictors of attitudes toward television advertising.

Figure 2: Advertising Prices and Audience Activity Levels of Television Outlets



Notes: Each plot is a binned scatterplot of a dependent variable against an independent variable of interest. To construct each plot, we regress the dependent variable on indicators for deciles of the independent variable of interest and a set of controls. The unit of analysis in the regression is an outlet. The y-axis values in the plot are the coefficients on the decile indicators, recentered by adding a scalar so that their mean value is equal to the sample mean value of the dependent variable. The x-axis values in the plot are the mean values of the independent variable of interest within the corresponding decile. In both plots, the dependent variable is the log(price per impression) of a 30-second spot on the outlet; the independent variable of interest is the weighted average log(weekly viewing time) of the outlet's viewers; and the controls include the share of the outlet's impressions that are to adults, and indicators for the outlet's daypart. In Panel B, the controls additionally include deciles for the outlet's log(impressions per hour).


Figure 3: Advertising Prices and Activity Levels by Audience Demographics of Television Outlets




Notes: Each plot is a binned scatterplot of a dependent variable against an independent variable of interest. To construct each plot, we regress the dependent variable on indicators for deciles of the independent variable of interest and a set of controls. The unit of analysis in the regression is an outlet. The y-axis values in the plot are the coefficients on the decile indicators, recentered by adding a scalar so that their mean value is equal to the sample mean value of the dependent variable. The x-axis values in the plot are the mean values of the independent variable of interest within the corresponding decile. In all plots, controls include the share of the outlet's impressions that are to adults, and indicators for the outlet's daypart. In the upper row of plots, the independent variable of interest is the average age of the outlet's adult impressions, and the controls additionally include indicators for deciles of the share of the outlet's adult impressions that are to females. In the lower row of plots, the independent variable of interest is the share of the outlet's adult impressions that are to females, and the controls additionally include indicators for deciles of the average age of the outlet's adult impressions. In the left column of plots, the dependent variable is the log(price per impression) of a 30-second spot on the outlet. In the right column of plots, the dependent variable is the weighted average log(weekly viewing time) of the outlet's viewers.

Figure 4: Advertising Prices and Audience Size of Television Outlets

Panel A: Not controlling for viewing time of outlet's audience



Panel B: Controlling for viewing time of outlet's audience



Notes: Each plot is a binned scatterplot of a dependent variable against an independent variable of interest. To construct each plot, we regress the dependent variable on indicators for deciles of the independent variable of interest and a set of controls. The unit of analysis in the regression is an outlet. The y-axis values in the plot are the coefficients on the decile indicators, recentered by adding a scalar so that their mean value is equal to the sample mean value of the dependent variable. The x-axis values in the plot are the mean values of the independent variable of interest within the corresponding decile. In both plots, the dependent variable is the log(price per impression) of a 30-second spot on the outlet; the independent variable of interest is the log(impressions per hour) of the outlet; and the controls include the share of the outlet's impressions that are to adults, and indicators for the outlet's daypart. In Panel B, the controls additionally include deciles for the weighted average log(weekly viewing time) of the outlet's viewers.

Table 1: Advertising Prices, Audience Demographics, and Audience Activity Levels of Television Outlets

<i>Dependent variable:</i>	<i>Observed</i>		<i>Predicted</i>			
	<i>log(price per impression)</i>		<i>log(price per viewer)</i>			
			<i>Value homog.</i>		<i>Value prop. to income</i>	
	(1)	(2)	(3)	(4)	(5)	(6)
Average log(weekly viewing hours) of audience	-1.5556 (0.2913)		-1.6799 (0.0607)		-1.8388 (0.1027)	
Average age of impressions		-0.0285 (0.0079)		-0.0028 (0.0024)		-0.0020 (0.0029)
Share female among adult impressions		-0.4690 (0.2599)		-0.3056 (0.0933)		-0.5230 (0.1228)
log(impressions per hour)	0.0973 (0.0292)	0.1221 (0.0306)	0.0082 (0.0044)	0.0418 (0.0109)	0.0198 (0.0075)	0.0628 (0.0125)
Average household income of impressions (\$1000)	0.0124 (0.0031)	0.0152 (0.0034)	0.0002 (0.0004)	0.0057 (0.0016)	0.0102 (0.0008)	0.0152 (0.0018)
Number of networks	103	103	103	103	103	103
Number of network-dayparts	809	809	809	809	809	809

Notes: Each column reports estimates of a linear regression. The unit of analysis is an outlet (network-daypart). In columns (1) and (2), the dependent variable is the log(price per impression) of a 30-second spot observed in the data, as described in Section 3.1. In columns (3) through (6) the dependent variable is the log(price per viewer) predicted by the model, as described in Section 5. Columns (3) and (4) use log(price per viewer) predicted from the baseline model in which advertisers' value of a first impression is homogeneous across viewers. Columns (5) and (6) use log(price per viewer) predicted from the model in which advertisers' value of a first impression is proportional to a viewer's income. All models include controls for the share of the outlet's impressions that are to adults, and indicators for the outlet's daypart. The sample includes only those outlets for which all variables are available. Standard errors in parentheses are clustered by network.

5 Quantification and Applications of the Model

5.1 Implementation and Model Fit


Corollary 1 shows that, in the reach-only model, it is possible to calculate the equilibrium price per viewer p_Z^*/λ_Z given data on the probability η_{ij} that each viewer i sees ads on each outlet j , the intrinsic value a_i of each viewer i , and the ownership partition \mathcal{Z} . We operationalize this calculation in the audience survey data, letting the viewers \mathcal{I} be the set of respondents. For each viewer $i \in \mathcal{I}$, we take $\eta_{ij} = T_{ij}/T_j$ where T_j is outlet j 's total broadcast time and we recall that T_{ij} is the time that viewer i spends viewing outlet j . We implement two specifications, a specification in which $a_i = a$ for all i , and a specification in which a_i is proportional to household income for all i . We treat the value $\frac{1}{|\mathcal{I}|} \sum_{i \in \mathcal{I}} a_i$ of reaching an average viewer as an unknown scale normalization, and therefore do not use any data on advertising prices in calculating the price per viewer predicted by the model. We perform inference via a nonparametric bootstrap over survey respondents with 100 replicates.

Figure 5 shows that the model does a good job predicting relative advertising prices across owners of television networks. Panel A shows a scatterplot of the observed log(price per impression) against the predicted log(price per viewer) for the specification with homogeneous values. Predicted log(price per viewer) explains 35.1 percent of the variation in observed log(price per impression). Panel B shows the analogous scatterplot for the specification with values proportional to household income. Predicted log(price per viewer) explains 34.0 percent of the variation in observed log(price per impression). In the latter case, the slope of the relationship between observed and predicted prices is close to 1. Across both specifications, the model is able to rationalize large differences in advertising prices between owners, as illustrated in the x-axes of the scatterplots.


Appendix Table 1 evaluates the restrictiveness (Fudenberg, Gao, and Liang 2023) and completeness (Fudenberg et al. 2022) of the economic model. The restrictiveness measures the difficulty of fitting arbitrary data, and in this sense can be thought of as a measure of parsimony. Panel A shows that the economic model is highly restrictive. This is because the economic model has no free parameters that are fit to the observed price per impression. The completeness measures the model's success in capturing the systematic patterns in the data. Panel B shows that, despite being highly restrictive, the economic model outperforms both a simple regression model and a richly parameterized machine-learning model.

Figure 5: **Observed and Predicted Television Advertising Prices**

Panel A: Baseline model with homogeneous value



Panel B: Model with value proportional to income




Notes: Each plot is a scatterplot of the log(price per impression) of a 30-second spot observed in the data (y-axis), as described in Section 3.1, against the log(price per viewer) predicted by the model (x-axis), as described in Section 5. Panel A uses log(price per viewer) predicted from the baseline model in which advertisers' value of a first impression is homogeneous across viewers. Panel B uses log(price per viewer) predicted from the model in which advertisers' value of a first impression is proportional to a viewer's income. The unit of analysis is an owner Z . Variables are residualized with respect to the share of the owner's impressions that are to adults and recentered by adding a scalar so that the mean value of each recentered variable is equal to the sample mean of the log(price per impression) observed in the data. The dashed line depicts a 45-degree line. The solid line depicts the line of best fit. The box reports the slope of the line of best fit and the R^2 of the associated linear model, with standard errors in parentheses obtained via a nonparametric bootstrap over survey respondents with 100 replicates.

Appendix Figure 5 shows that moderate departures from perfect diminishing returns have little effect on the share of variation in $\log(\text{price per impression})$ that is explained by variation in predicted $\log(\text{price per viewer})$, but lead predicted $\log(\text{price per viewer})$ to underpredict the extent of price differences across owners. Under our other maintained assumptions Appendix Figure 5 suggests that diminishing returns are strong. If diminishing returns are instead weak, Appendix Figure 5 suggests that there are forces outside the model that explain variation in $\log(\text{price per impression})$ and that are correlated with predicted $\log(\text{price per viewer})$.

We can also evaluate the fit of the model to trends in television advertising revenues during recent years in which overall television viewership has fallen. Panel A of Figure 6 shows that annual revenues increased slightly between 2014 and 2019 even as total impressions fell, a pattern that some have regarded as puzzling (The Economist 2021). Panel B shows that our baseline model predicts this pattern. In the model, a decline in impressions can increase the value captured by television owners if it results in less overlap in audience across owners. Panel B shows that the model indeed predicts an increase in the price per impression, and in revenues, over this period. The patterns in Panel B provide a reasonable qualitative and quantitative match to those in Panel A, even though the revenue calculations underlying Panel B are based only on audience survey data, and in particular do not use any information on advertising prices.²⁶ Competition is important for the findings in Panel B of Figure 6: with a single monopoly owner of television networks, our model predicts declining, rather than increasing, advertising revenue over the period we study. Diminishing returns also appear to be important for the findings in Panel B of Figure 6: Appendix A.4 shows that two alternative models without diminishing returns can generate increasing prices but not increasing revenues.

²⁶The ratio of price per impression in 2019 to price per impression in 2014 is 1.17 in the data and 1.22 (SE = 0.01) in the model prediction. Appendix Figure 6 includes an alternative version of Panel B in which impressions are imputed from the audience survey data.

Figure 6: **Observed and Predicted Television Advertising Revenues**



Notes: Each plot depicts trends in the television advertising market over the sample period. We plot trends in total revenue, total impressions, and price per impression (total revenues divided by total impressions), all normalized relative to their 2015 value. In Panel A, all series are as observed in the data, as described in Section 3.1, and revenue is deflated to 2015 dollars using the US Consumer Price Index (Organization for Economic Co-operation and Development 2022). In Panel B, the trend in revenue is predicted by the baseline model in which advertisers' value of a first impression is homogeneous across viewers, as described in Section 5; the trend in impressions is identical to that in Panel A; and the price per impression is the ratio of the two.

5.2 Implications for Demographic Premia

The model allows us to quantify the effects of competition on the incentives of network owners to attract different kinds of audience members. Following Corollary 1, the contribution of viewer i to the revenue of owner Z is given by

$$p_{iZ}^* = a_i \eta_{iZ} \prod_{Z' \neq Z} (1 - \eta_{iZ'}).$$

Motivated by this observation we can define the television market’s total marginal willingness to pay per impression to attract viewer i ,

$$m_i = \frac{\sum_{Z \in \mathcal{Z}} p_{iZ}^*}{\sum_{Z \in \mathcal{Z}} \eta_{iZ}}.$$


Appendix A.3 relates p_{iZ}^* to the marginal willingness to pay of owner Z to attract a viewer i in a model with a content investment stage.

Figure 7 depicts the average estimated value of $\ln(m_i)$ across viewers i in different age categories under different ownership partitions, including the factual partition (“baseline”), a counterfactual partition in which a single owner owns all networks (“concentrated”), and counterfactual partitions in between “baseline” and “concentrated” in which the top two, three, or four owners by audience are merged. Panel A uses the model in which advertisers’ value a_i is homogeneous across viewers; Panel B uses the model in which a_i is proportional to viewer i ’s household income. In each panel, we normalize the reported average estimated value of $\ln(m_i)$ to be relative to the average value for the youngest age group. Therefore, y-axis values can be read as differences in the log of the market’s willingness to pay per impression for viewers in a given age group relative to those of the youngest group.


Figure 7 shows that, under the factual ownership partition, willingness to pay per impression is 82.8 (Panel A, SE = 3.21) or 88.8 (Panel B, SE = 4.38) log points lower to attract the average member of the oldest group than to attract the average member of the youngest group. These differences attenuate with reduced competition, down to 0 (Panel A, SE = 0) or 6.0 (Panel B, SE = 3.06) under concentrated ownership. Figure 7 thus suggests that television network owners would have a stronger incentive to target content to older viewers if the television market were less competitive.

Figure 7: Marginal Willingness to Pay to Attract Older vs. Younger Viewers

Panel A: Baseline model with homogeneous value



Panel B: Model with value proportional to income



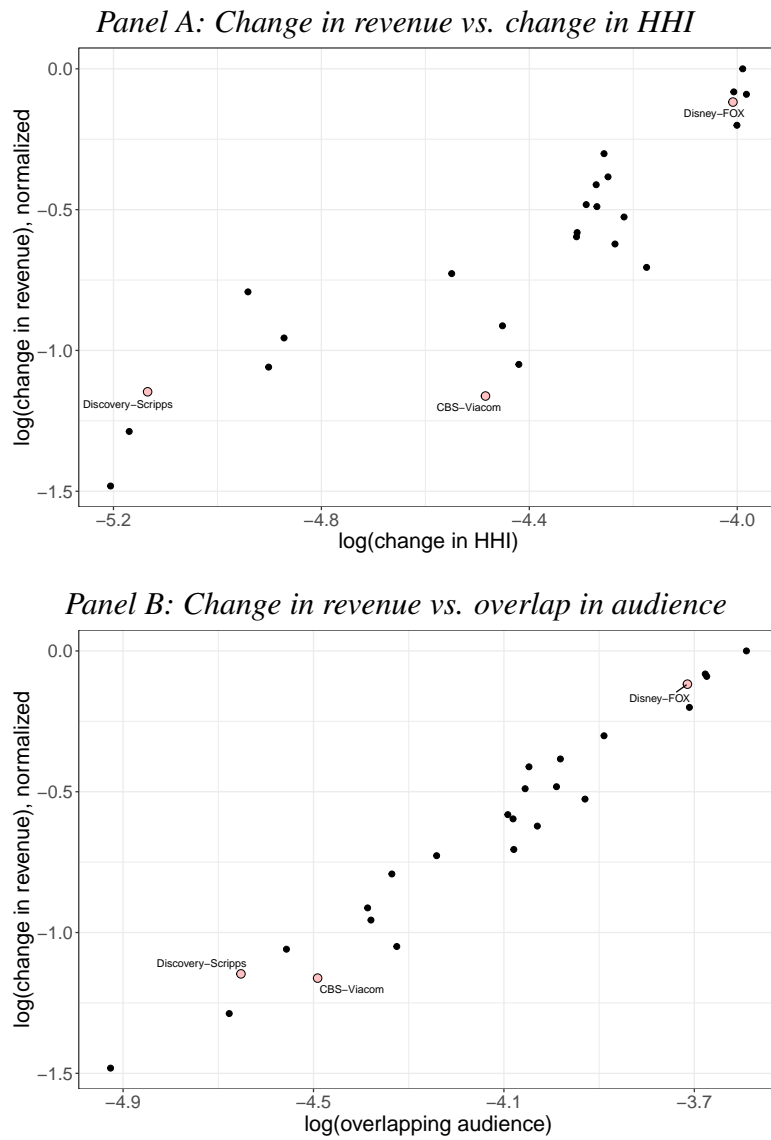
Notes: In each plot, the y-axis value corresponds to the estimated average log(total marginal willingness to pay per impression), $\ln(m_i)$, as described in Section 5.2, for viewers in the age bin listed on the x-axis, under different ownership scenarios. These scenarios include: the “baseline” ownership corresponding to the observed partition \mathcal{Z} ; the “concentrated” ownership corresponding to the counterfactual scenario in which one entity owns all television networks; and three counterfactual ownership partitions in between, in which the top two, three, or four owners by audience are merged. Panel A uses the baseline model in which advertisers’ value of a first impression is homogeneous across viewers. Panel B uses the model in which advertisers’ value of a first impression is proportional to a viewer’s income. In both plots, darker colors correspond to more concentrated ownership scenarios, and the y-axis value is normalized by adding a scalar so that its average value is zero in the youngest age group.

We can also evaluate the fit of the model to the patterns we documented in Section 4. To do this, we calculate the predicted price per viewer p_j^*/λ_j for each outlet, treating outlets as if they were independently owned. Columns (3) and (4) of Table 1 report estimates of the same regression models as in columns (1) and (2), replacing the observed $\log(\text{price per impression})$ with $\log(\text{price per viewer})$ $\ln\left(p_j^*/\lambda_j\right)$ predicted from the model in which advertisers' value a_i is homogeneous across viewers. The model matches the qualitative patterns in the data well, but predicts weaker relationships on some dimensions than those observed in the data. Columns (5) and (6) repeat the exercise in columns (3) and (4), with $\ln\left(p_j^*/\lambda_j\right)$ predicted from the model where a_i is proportional to viewer i 's household income. This specification's predictions align better on some dimensions with the patterns observed in the data. Both specifications underpredict the magnitude of the relationship between price and average age of impressions. A possible interpretation is that audience age influences advertising prices through other channels in addition to those captured in the model.

5.3 Application to Mergers of Television Network Owners

Figure 8 visualizes the implications, under the model with homogeneous values, of each possible pairwise merger among the top eight owners of television networks by audience. For each merger, we calculate the log of the predicted change in revenue, the log of the predicted change in the Hirschman-Herfindahl index (HHI) of audience shares, and the log of the size of the overlapping audience between the two merging owners. Panel A plots the log of the predicted change in revenue against the log of the predicted change in HHI. Panel B plots the log of the predicted change in revenue against the log of the overlapping audience. In both panels, we normalize the log of the predicted change in revenue relative to its maximum value, and we highlight three mergers that occurred after 2015: Discovery and Scripps (2018), CBS and Viacom (2019), Disney and Fox (2019).

Figure 8: **Predicted Effects of Mergers on Advertising Revenue**



Notes: For each of a set of simulated mergers, Panel A plots the log of the simulated change in revenue (y-axis) against the log of the simulated change in HHI (x-axis), and Panel B plots the log of the simulated change in revenue (y-axis) against the log of the overlapping audience (x-axis). Larger, more lightly shaded circles indicate mergers that occurred after 2015; smaller, more darkly shaded circles indicate hypothetical mergers that have not occurred. We construct the plots as follows. For each owner we compute the audience size as the probability of an average viewer seeing an ad on at least one of the owner's outlets, as described in Section 5. We select the top eight owners by this metric, excluding joint ventures, and form all possible pairs of these eight. For each pair, we compute the overlapping audience, defined as the share of the audience seeing an ad on both owners' outlets. For each pair we also simulate the effect of a pairwise merger on the pair's total advertising revenue, using our model of advertising-market equilibrium as described in Section 5. We also simulate the effect of the merger on the Herfindahl-Hirschman Index (HHI), where the HHI is computed with respect to the probability of an average viewer seeing an ad on at least one of the owner's outlets. We exclude from the plots any merger that changes the HHI by less than 0.001. For simulated mergers between two owners each of which owns one of the broadcast networks {ABC, CBS, FOX, NBC}, we exclude one of the two owners' broadcast networks from the simulated merger and treat it as a separate entity for all calculations. For mergers that took place we exclude the broadcast network that was excluded in practice; for other mergers we exclude the broadcast network owned by whichever owner had a smaller total pre-merger audience. We normalize the log of the simulated change in revenue by subtracting the value for the merger with the largest simulated change in revenue. For the mergers that took place, the normalized log of the simulated change in revenue is -1.15 for Discovery-Scripps (SE = 0.04), -1.16 for CBS-Viacom (SE = 0.03), and -0.12 for Disney-FOX (SE = 0), where standard errors in parentheses are obtained via a nonparametric bootstrap over survey respondents with 100 replicates.

Comparing Panels A and B of Figure 8 shows that, according to the model, the revenue effects of a given merger are more strongly related to the overlap in audience between the merging entities than to the change in HHI induced by the merger. Among the three mergers that occurred, for example, the CBS-Viacom merger is roughly midway between the Discovery-Scripps merger and the Disney-Fox merger in terms of its log impact on HHI, but is much closer to the Discovery-Scripps merger in terms of both the log of audience overlap and the log of the predicted revenue impact.²⁷


5.4 Application to Netflix Carrying Advertising

In 2022, Netflix launched a new service tier that carries advertising (Krouse 2022). Figure 9 visualizes the implications, under the model with homogeneous values, for television network owners of Netflix counterfactually adding advertising to its platform in 2019, assuming no change in audience behavior and that Netflix shows advertisements across all of its content and subscribers. Across the owners, we estimate that Netflix ad carriage would reduce the price per viewer by between 0.38 and 2.7 log points, with a mean reduction of 1.68 (SE = 0.05).²⁸ As the plot illustrates, owners whose outlets have greater audience overlap with Netflix tend to experience greater proportional declines in price per viewer in this counterfactual, though there is substantial variation in the effect of Netflix for a given level of audience overlap, owing to variation across owners in the overlap of their outlets' audience with that of other owners. Our estimates imply that Netflix itself would have a relatively high price per impression—about 24.4 log points larger than the average of the five largest TV owners—consistent with its relatively young audience.

²⁷Prat and Valletti (2022, Section 4) also simulate effects of platform mergers under various assumptions about overlap in their audience. In their setting, mergers impact ad prices through changes in the quantity of ad slots available, and audience overlap affects the revenue impact of a merger through product market interactions between advertisers. In contrast, our baseline model features a fixed quantity of ad slots and no interactions between advertisers outside of the advertising market.

²⁸Dividing the total daily US Netflix viewing minutes implied by values reported in MoffettNathanson (2022, Exhibit 11), assuming a 30.5-day average month, by the 2020 US Population (U.S. Census Bureau 2021) yields daily viewing of 22.7 minutes per capita, close to the average of 24.4 that we calculate from the audience survey data. If, due to covieing and other factors, Netflix viewing time is larger than what we estimate, then we expect our calculations to understate the effect of Netflix ad carriage on advertising prices.

Figure 9: Predicted Effect of Netflix Advertising on Advertising Prices



Notes: The plot shows a scatterplot, across owners of television networks, of the change in the owner’s log(price per viewer) predicted by the model if Netflix were to carry ads (y-axis), against the log of the overlapping audience between the owner’s audience and Netflix’s audience (x-axis). To construct the plot, we compile audience data from GfK MRI’s 2019 Survey of the American Consumer (GfK Mediamark Research and Intelligence 2021), analogous to the data described in Section 3.2 for the 2015 survey, and treating Netflix as an additional television outlet. To calculate the probability of a viewer seeing an ad spot on Netflix, we divide the number of hours the viewer reports spending watching Netflix over the last seven days (topcoded at 21 hours) by the number of hours in the week. We compute the difference in log(price per viewer) implied by the baseline model in which advertisers’ value of a first impression is homogeneous across viewers, as described in Section 5, between the scenarios with and without Netflix included in the advertising market (y-axis). We also compute the log of the share of the television audience seeing an ad on both the given owner’s outlets and Netflix (x-axis).

6 Extension to Social Media

Here we extend our analysis of the television advertising market to incorporate the social media advertising market.

6.1 Model

We first generalize the model in Section 2 to allow a subset of owners to post viewer-specific advertising prices. Suppose that there are viewers $i \in \mathcal{I}$ and that a subset $\tilde{\mathcal{Z}} \subseteq \mathcal{Z}$ of owners can post viewer-specific prices for each bundle of its outlets. Suppose that the advertisers' value function $V(\cdot) = \sum_{i \in \mathcal{I}} V_i(\cdot)$, where $V_i(\cdot)$ is the value function if viewer i were the only viewer, and $v_{iB} = V_i(\mathcal{J}) - V_i(\mathcal{J} \setminus B)$ is the viewer-specific incremental value of bundle B .

Proposition 3. *Suppose that $V_i(\cdot)$ is monotone and submodular and that $v_{ij} > 0$ for all i and j . Then there exists an equilibrium. In any equilibrium, for any viewer i , all advertisers buy slots on all outlets owned by the owners in $\tilde{\mathcal{Z}}$, and the payment by each advertiser to each owner $Z \in \tilde{\mathcal{Z}}$ is given by $p_{iZ}^* = v_{iZ}$. Moreover, in any equilibrium, the outcome for the outlets owned by the owners in $\mathcal{Z} \setminus \tilde{\mathcal{Z}}$ is identical to that in Theorem 1.*

In the setting of Proposition 3, the owners in $\tilde{\mathcal{Z}}$ are indifferent between selling the ad slots at the outlet level or at the viewer level because the advertisers are homogeneous. If the incremental value of each viewer could differ across advertisers, then the owners might benefit from viewer-level ad pricing and targeting.

For our empirical analysis, we consider the following special case of the setting of Proposition 3, analogous to the reach-only model in the single-format case.

Definition. (Cross-format reach-only model.) In the **cross-format reach-only model**, owners are partitioned into two formats, 1 and 2, corresponding to the sets $\tilde{\mathcal{Z}}$ and $\mathcal{Z} \setminus \tilde{\mathcal{Z}}$, respectively. Advertisers' valuations follow the viewer-level model where $u(M)$ depends on M only via $(\mathbf{1}_{M_1 > 0}, \mathbf{1}_{M_2 > 0})$ and $u(M)$ is invariant to exchanging the elements of M .

In the cross-format reach-only model, there are two formats, with owners in one format (e.g., social media) able to post viewer-specific prices and one format (e.g., television) unable to do so. An advertiser's value from a given viewer is determined by whether that viewer sees its ad at least once on each format, with no incremental return to multiple impressions on a given format. In this

case, it is without loss of generality to write that

$$u(M_1, M_2) = \mathbf{1}_{M_1 > 0} + \mathbf{1}_{M_2 > 0} - \phi \mathbf{1}_{M_1 M_2 > 0}$$

where $\phi \in [0, 1]$ is a parameter. When $\phi = 0$ diminishing returns operate only within and not between formats. When $\phi = 1$ the distinction between formats is irrelevant. We can characterize pricing in this model via the following generalization of Corollary 1.

Corollary 2. *In the cross-format reach-only model the equilibrium revenue p_Z^* of owner Z is given by $p_Z^* = \sum_{i \in \mathcal{I}} p_{iZ}^*$ for*

$$p_{iZ}^* = a_i \eta_{iZ} \left[\left(1 - \phi \left(1 - \prod_{Z' \in \tilde{\mathcal{Z}} \setminus F(Z)} (1 - \eta_{iZ'}) \right) \right) \prod_{Z' \in F(Z) \setminus \{Z\}} (1 - \eta_{iZ'}) \right]$$

the equilibrium revenue from viewer i , where $F(Z)$ is the set of owners sharing the same format as owner Z and where η_{iZ} is defined as in the reach-only model. For an owner $Z \in \tilde{\mathcal{Z}}$ the price per impression for viewer i is p_{iZ}^/η_{iZ} , whereas for an owner $Z \notin \tilde{\mathcal{Z}}$ the price per viewer is p_Z^*/λ_Z , with λ_Z defined as in the reach-only model.*

We will apply the cross-format model to the case of television and social media in the mid-2010s. Although we do not take a strong *a priori* stand on the strength of cross-format diminishing returns, i.e., on the value of the parameter ϕ , we think it is plausible that $\phi \in (0, 1)$ during the period we study. Advertisers coordinate campaigns across television and social media, and industry sources suggest that advertisers are attentive to audience overlap when they do.²⁹ These facts suggest that $\phi > 0$. At the same time, industry sources suggest that television and social media ads tend to serve different functions, with the former best suited to “top of funnel” brand building strategies and social media ads best suited to “low funnel” activities like acquiring customers and inducing immediate purchases (e.g., The Nielsen Company 2018, p. 38). Correspondingly, ads on social media often come in the form of product images or very short videos, in contrast to television ads which tend to be 30-second video spots.³⁰ The growth in television advertising revenue during a period of rapid expansion in social media advertising further suggests that advertisers do not view

²⁹Facebook tells advertisers that “Nielsen’s Total Ad Ratings (TAR) helps you see who you’re reaching across TV and Facebook. It also helps you identify overlap in the audience seeing your ads. This can then help you determine how to optimize your campaigns more effectively and efficiently” (Facebook 2018).

³⁰As of 2016, only 9.7 percent of digital ad spending went to video ads (MoffettNathanson 2017, Exhibit 3). Facebook only introduced embeddable videos in 2015 (MoffettNathanson 2017, Exhibit 23).

the two formats as serving identical functions.³¹ All of these facts suggest that $\phi < 1$.

Importantly, Corollary 2 implies that the cross-format reach-only model inherits the comparative statics of Section 2.2. Specifically, say that an individual i is more active than an individual i' if $\eta_{iZ} \geq \eta_{i'Z}$ for all $Z \in \mathcal{Z}$. We then have the following.

Corollary 3. *In the cross-format reach-only model, if individual i is more active than individual i' , then $p_{iZ}^*/\eta_{iZ} \leq p_{i'Z}^*/\eta_{i'Z}$ for all $Z \in \tilde{\mathcal{Z}}$.*

6.2 Data

Social Media Advertising Prices

We obtain data on the cost of advertising to different audiences on Facebook via an original experiment conducted for this study and a separate advertising campaign conducted for a different study (Allcott et al. 2020a). In both cases advertisements were placed through Facebook’s Ad Manager. In the Facebook advertisement structure, an advertisement set is a group of one or more advertisements with a defined audience target, budget, schedule, bidding, and placement. An advertising campaign is a group of one or more advertisement sets corresponding to a single campaign objective (Facebook 2022). All advertisements targeted English speakers in the United States.

For our experiment, we administered an advertising campaign from July 15, 2017 through July 22, 2017 in partnership with GiveDirectly. The campaign consisted of 14 separate advertisement sets targeting each combination of gender and age group in $\{\text{Men, Women}\} \times \{13-17, 18-24, 25-34, 35-44, 45-54, 55-64, 65+\}$. Each advertisement set included just one advertisement, with fixed budgets of \$20 a day using automated cost-per-click bidding. For each advertisement set, we obtain the price per impression, which we may think of as a group-level average of the concept p_{iZ}^*/η_{iZ} defined in Section 6.1.

From Allcott et al. (2020b), we obtain data from 32 advertisement sets purchased on September 24, 2018: four each targeting each combination of gender and age group in $\{\text{Men, Women}\} \times \{18-24, 25-44, 45-64, 65+\}$. We compute the total cost and total number of impressions for each demographic group, and take the ratio of these to obtain the price per impression.

³¹Appendix Figure 7 shows that a model with $\phi = 1$ would predict declining television advertising revenue in recent years, in contrast to the observed trends, and the predicted trends from a model with $\phi = 0$, depicted in Figure 6.

Audience Survey Data

From the 2015 audience survey described in Section 3.2 we obtain information on whether the respondent visited each of five social media sites—Facebook, Instagram, Reddit, Twitter, and YouTube—in the preceding 30 days. For each respondent i and social media site j , we let η_{ij} be an indicator for whether the respondent visited the site. We treat Facebook and Instagram as jointly owned and the other social media sites as independently owned, and compute η_{iZ} following Corollary 1. We also obtain information on the amount of time each respondent spent using the internet in an average week.³²

6.3 Descriptive Evidence

Whereas older people are the heaviest television viewers, younger people are the heaviest users of social media. Panel A of Appendix Figure 8 shows that older individuals visit fewer social media sites. If ϕ is small, the forces in Corollary 3 would lead us to expect that social media advertising prices will be increasing in audience age, in contrast to the decreasing pattern we see for television advertising prices. Panel B of Appendix Figure 8 shows that a similar qualitative pattern to Panel A obtains for time spent on the internet.


Figure 10 shows that the direction of the age-price gradient is indeed reversed on social media. This is true both according to data we collected in our own experiment (Panel A) and according to data collected as part of Allcott et al.'s (2020a) study (Panel B). These differences are large, with the oldest group commanding a premium of 122 log points (Panel A) or 57 log points (Panel B) relative to the youngest group, on average across genders. We view this evidence as difficult to square with some explanations for the differential value of older vs. younger viewers to advertisers, such as intrinsic differences in the malleability of their preferences (Surowiecki 2002).³³ Figure 10 does not show a consistent price premium for advertising to male or female audiences on social media. Appendix Figure 9 repeats the analysis in Figure 10 using data on price per click rather than price per impression.

³²This is, in turn, calculated based on reported time spent on three recent days.


³³Smith, Moschis, and Moore (1985) find that older consumers rely more on advertising when making purchasing decisions than do younger consumers, though DeLorme, Huh, and Reid (2006) find no evidence of age differences in overall behavioral responses to direct-to-consumer prescription drug advertisements. Lewis and Reiley (2014) find evidence of large effects of digital advertising on purchases by older individuals.

Figure 10: Demographic Premia and Viewing Time on Facebook

Panel A: Data from our experiment



Panel B: Data from Allcott et al. (2020b)



Notes: The plot shows the $\log(\text{price per impression})$ for advertisement sets targeted to a given gender and age group. In Panel A, the data are taken from our own experiment, and the groups are $\{\text{Men, Women}\} \times \{18-24, 25-34, 35-44, 45-54, 55-64, 65+\}$. In Panel B, the data are taken from Allcott et al. (2020b), and the groups are $\{\text{Men, Women}\} \times \{18-24, 25-44, 45-64, 65+\}$. In both panels, the y-axis value is the $\log(\text{price per impression})$ for advertisement sets targeting the given group, and the x-axis value is the midpoint of the age range for the given group, treating 70 as the midpoint for ages 65+.

Appendix Figure 10 shows the relationship between estimated $\log(\text{price per impression})$ of display advertisements and audience demographics across a sample of platforms. The data on prices are imputed from a statistical model that is estimated on data from a range of sources (The Nielsen Company 2017). The data on audience demographics are from the audience survey and do not reflect viewership intensity. The advertisements in the sample are likely more heterogeneous than those in our controlled buying experiment. The plots show no clear relationship between price per impression and the age or gender composition of the platform’s audience.

6.4 Quantification

Figure 11 shows the fit of the model to the patterns in Figure 10 as a function of ϕ . To produce the figure, for each value of ϕ , we calculate the average price per Facebook visitor in each of the demographic cells in Figure 10 following Corollary 2. Panel A of Figure 11 depicts the fit of the model as a function of ϕ , for the values $\phi \in \{0, 0.1, \dots, 1\}$. Appendix Figure 11 depicts corresponding patterns for the fit of the model to the observed $\log(\text{price per impression})$ for television owners.

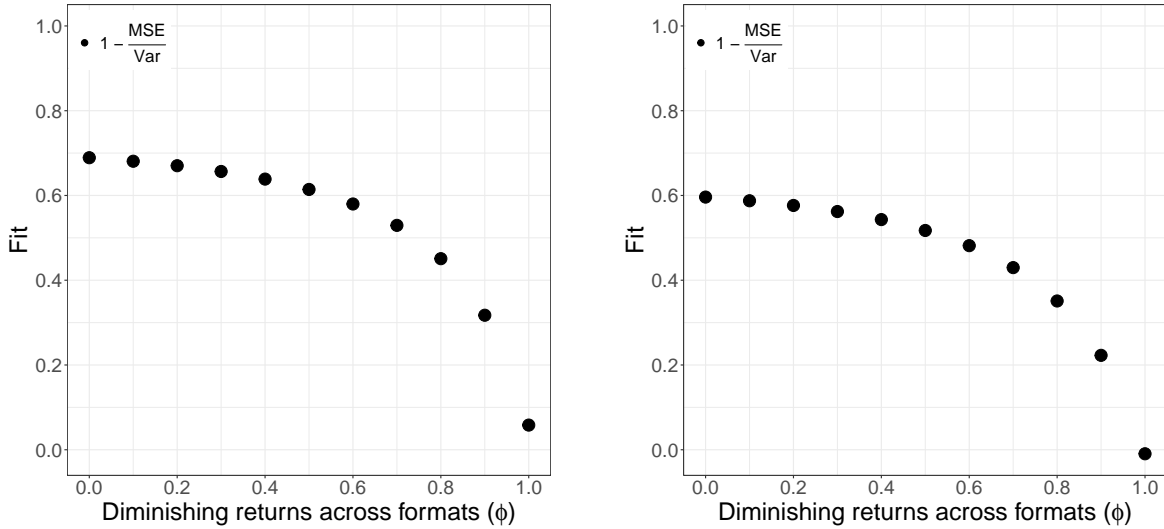
Panel A of Figure 11 shows that the best fit is achieved at $\phi = 0$, but that other low values of ϕ achieve a similar fit. Panel B illustrates the fit of the model at the best-fitting value of ϕ . The model predicts an even lower price for the youngest group of men, and a more systematic gender difference, than what is observed in the data. Although our focus is on the implications of the choice of ϕ for patterns in pricing across demographic groups, we note that our findings regarding the degree of substitutability across platforms, as captured in the parameter ϕ , may be of independent interest, for example in light of questions about appropriate market definition for competition policy in advertising markets (Delrahim 2019).

Figure 11: **Fit of Quantitative Model of Social Media Prices**

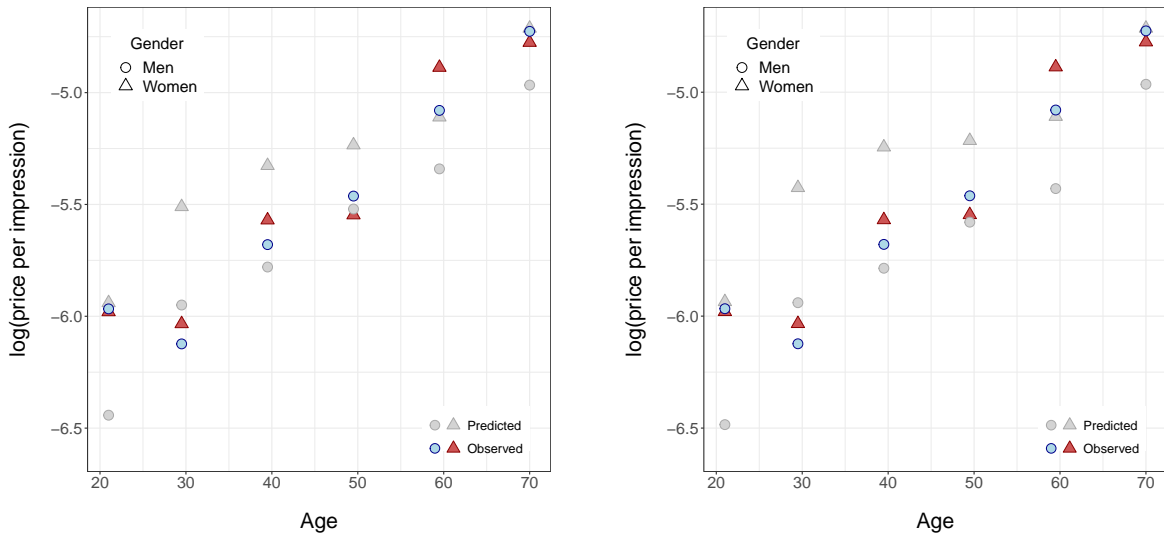
Baseline model with homogeneous value

Model with value proportional to income

Panel A: Fit of model



Panel B: Observed and Predicted Advertising Prices



Notes: Panel A depicts the fit of a model-based prediction of social media advertising prices (y-axis) as a function of the strength of cross-format diminishing returns (x-axis). To produce each plot, for each of the demographic cells depicted in Panel A of Figure 10, and for each value $\phi \in \{0, 0.1, \dots, 1\}$ of the parameter governing the strength of cross-format diminishing returns, we calculate the average price per viewer for visitors to Facebook in the given cell implied by the cross-format reach only model defined in Section 6. For each value of ϕ , we recenter the predicted $\log(\text{average price per viewer})$ by adding a constant so that the predicted $\log(\text{average price per viewer})$ is equal to the observed $\log(\text{price per impression})$ on average across the demographic cells. We then calculate the mean squared error of the model as the mean of squared deviations between the predicted $\log(\text{average price per viewer})$ and the observed $\log(\text{price per impression})$ across the demographic cells. We then plot the fit, given by one minus the ratio of the mean squared error to the sample variance, against the value of ϕ . Panel B depicts, for each demographic cell, the observed $\log(\text{price per impression})$ from Panel A of Figure 10, alongside the predicted $\log(\text{average price per viewer})$ based on the value of ϕ that maximizes the fit. The left column of plots uses $\log(\text{average price per viewer})$ predicted from the baseline model in which advertisers' value of a first impression is homogeneous across viewers. The right column of plots uses $\log(\text{average price per viewer})$ predicted from the model in which advertisers' value of a first impression is proportional to a viewer's income.

7 Conclusions

We extend existing theoretical results on advertising markets with competing outlets and a multi-homing audience. Our model predicts that the equilibrium price per viewer that an outlet charges for its ads is lower the more active is the outlet's audience. We show that this prediction is borne out in data on television advertising. A disciplined, quantitative implementation of the model rationalizes a meaningful portion of the variation in advertising prices across television outlets and owners, the premia associated with specific demographic groups, and recent trends in television advertising revenue.

We conclude that the model captures important competitive forces in the advertising market. We therefore apply the quantitative model to questions of economic and policy interest, including the effects of mergers of television network owners on advertising prices and the effect of Netflix ad carriage on linear television advertising prices.

We extend the analysis to social media, where there is a premium to advertise to older viewers. A model in which diminishing returns are stronger within than between formats can rationalize this fact.

References

- AdStage. 2020. *Paid Media Q1 2020 Benchmark Report*. Report by AdStage (A TopClicks Company).
- Allcott, Hunt, Luca Braghieri, Sarah Eichmeyer, and Matthew Gentzkow. 2020a. “The welfare effects of social media.” *American Economic Review* 110(3): 629-676.
- Allcott, Hunt, Luca Braghieri, Sarah Eichmeyer, and Matthew Gentzkow. 2020b. “Facebook campaign data for: The welfare effects of social media.” (accessed February 2019).
- Alwitt, Linda F. and Paul R. Prabhaker. 1994. “Identifying who dislikes television advertising: Not by demographics alone.” *Journal of Advertising Research* 34(6): 17-29.
- Ambrus, Attila, Emilio Calvano, and Markus Reisinger. 2016. “Either or both competition: A ‘two-sided’ theory of advertising with overlapping viewerships.” *American Economic Journal: Microeconomics* 8(3): 189-222.
- Ampush. 2014. *Mobile Advertising Trends Report: Q3 2014*. Report by Ampush.
- Anderson, Simon and Stephen Coate. 2005. “Market provision of broadcasting: A welfare analysis.” *Review of Economic Studies* 72(4): 947-972.
- Anderson, Simon, Øystein Foros, and Hans Jarle Kind. 2018. “Competition for advertisers and for viewers in media markets.” *Economic Journal* 128(608): 34-54.
- Anderson, Simon and Martin Peitz. 2020. “Media see-saws: Winners and losers in platform markets.” *Journal of Economic Theory* 186: 104990.
- Aral, Sinan. 2021. “What digital advertising gets wrong.” *Harvard Business Review*, February 19. Accessed at <https://web.archive.org/web/20230325071010/https://hbr.org/2021/02/what-digital-advertising-gets-wrong> on April 19 2023.
- Armstrong, Mark and John Vickers. 2022. “Patterns of competitive interaction.” *Econometrica* 90(1): 153-191.
- Athey, Susan, Emilio Calvano, and Joshua S. Gans. 2018. “The impact of consumer multi-homing on advertising markets and media competition.” *Management Science* 64(4): 1574-1590.
- Bel, Germà and Laia Domènech. 2009. “What influences advertising price in television channels? An empirical analysis of the Spanish market.” *Journal of Media Economics* 22: 164-183.
- Berry, Steven, Alon Eizenberg, and Joel Waldfogel. 2016. “Optimal product variety in radio markets.” *RAND Journal of Economics* 47(3): 463-497.
- Chandra, Ambarish and Ulrich Kaiser. 2014. “Targeted advertising in magazine markets and the advent of the internet.” *Management Science* 60(7): 1829-1843.
- Chwe, Michael Suk-Young. 1998. “Believe the hype: solving coordination problems with television advertising.” University of Chicago Working Paper. Accessed at <https://www.ch>

- we.net/michael/h.pdf on July 5 2021.
- Competition and Markets Authority. 2019. *Online Platforms and Digital Advertising: Market Study Interim Report*. Report by UK Competition and Markets Authority. Accessed at https://assets.publishing.service.gov.uk/media/5dfa0580ed915d0933009761/Interim_report.pdf on July 6 2021.
- Crupi, Anthony. 2009. “Cable guy says ratings don’t ad up.” *Hollywood Reporter*, April 14. Accessed at <https://www.hollywoodreporter.com/business/business-news/cable-guy-says-ratings-dont-82305/> on July 13 2021.
- Dee, Jonathan. 2002. “The myth of ‘18 to 34’.” *New York Times Magazine*, October 13. Accessed at <https://www.nytimes.com/2002/10/13/magazine/the-myth-of-18-to-34.html> on July 5 2021.
- DeLorme, Denise E., Jisu Huh, and Leonard N. Reid. 2006. “Age differences in how consumers behave following exposure to DTC advertising.” *Health Communication* 20(3): 255-265.
- Delrahim, Makan. 2019. “Opening remarks for the public workshop on competition in television and digital advertising.” *US Department of Justice Transcript*. Accessed at <https://web.archive.org/web/20210727155932/https://www.justice.gov/atr/page/file/1202076/download> on May 7 2023.
- Dubé, Jean-Pierre, Günter J. Hitsch, and Puneet Manchanda. 2005. “An empirical model of advertising dynamics.” *Quantitative Marketing and Economics* 3: 107–144.
- The Economist. 2021. “The Olympics is a ratings flop. Advertisers don’t care.” *The Economist*, August 11. Accessed at <https://www.economist.com/business/the-olympics-is-a-ratings-flop-advertisers-dont-care/21803502> on June 16 2022.
- Einstein, Mara. 2004. “Broadcast network television, 1955-2003: The pursuit of advertising and the decline of diversity.” *Journal of Media Economics* 17(2): 145-155.
- Facebook. 2018. “Build TV and Facebook campaigns that work better together.” Accessed at <https://web.archive.org/web/20180525071501/https://www.facebook.com/business/news/build-tv-and-facebook-campaigns-that-work-better-together> on February 23 2023.
- Facebook. 2022. “Documentation of marketing API.” Accessed at <https://developers.facebook.com/docs/marketing-api/reference/ad-campaign-group/#Reading> on July 11 2022.
- Fan, Ying. 2013. “Ownership consolidation and product characteristics: A study of the US daily newspaper market.” *American Economic Review* 103(5): 1598-1628.
- Fudenberg, Drew, Jon Kleinberg, Annie Liang, and Sendhil Mullainathan. 2022. “Measuring the completeness of economic models.” *Journal of Political Economy* 130(4): 956-990.

- Fudenberg, Drew, Wayne Gao and Annie Liang. 2023. "How flexible is that functional form? Quantifying the restrictiveness of theories." MIT Working Paper. Accessed at <http://dx.doi.org/10.2139/ssrn.3580408> on March 7 2023.
- Gabler, Neal. 2014. "The tyranny of 18 to 49: American culture held hostage." *The Norman Lear Center*. Accessed at <https://web.archive.org/web/20221006143942/https://learcenter.org/wp-content/uploads/2014/10/Gabler18to49.pdf> on April 17 2023.
- Garland, Eric. 2002. "Disney marshals its parts into a new kind of advertising deal." *New York Times*, June 21. Accessed at <https://web.archive.org/web/20150527211647/https://www.nytimes.com/2002/06/21/business/media-business-advertising-disney-marshals-its-parts-into-new-kind-advertising.html> on April 29 2023.
- Gentzkow, Matthew, Jesse M. Shapiro, and Michael Sinkinson. 2014. "Competition and ideological diversity: Historical evidence from US newspapers." *American Economic Review* 104(10): 3073-3114.
- Geskey, Ronald D. 2016. *Media Planning and Buying in the 21st Century, 4th ed.* CreateSpace Independent Publishing Platform.
- GfK Mediamark Research and Intelligence. 2017. "Survey of the American Consumer, Respondent-Level Data for 2015." GfK Mediamark Research and Intelligence. <https://www.mrisimmons.com/our-data/national-studies/survey-american-consumer/>. (accessed August 2017).
- GfK Mediamark Research and Intelligence. 2019. "Survey of the American Consumer, Respondent-Level Data for 2017." GfK Mediamark Research and Intelligence. <https://www.mrisimmons.com/our-data/national-studies/survey-american-consumer/>. (accessed March 2022).
- GfK Mediamark Research and Intelligence. 2021. "Survey of the American Consumer, Respondent-Level Data for 2019." GfK Mediamark Research and Intelligence. <https://www.mrisimmons.com/our-data/national-studies/survey-american-consumer/>. (accessed in May 2022).
- Goettler, Ronald L. 2012. "Advertising rates, audience composition, and competition in the network television industry." University of Chicago Working Paper. Accessed at <http://goettler.simon.rochester.edu/research/papers/adrates.pdf> on May 31 2022.
- Greenwood, Jeremy, Yueyuan Ma, and Mehmet Yorukoglu. 2021. "'You will:' A macroeconomic analysis of digital advertising." NBER Working Paper No. 28537.
- Hristakeva, Sylvia and Julie H. Mortimer. Forthcoming. "Price dispersion and legacy discounts in

- the national television advertising market.” *Marketing Science*.
- Jeziorski, Przemysław. 2014. “Effects of mergers in two-sided markets: The US radio industry.” *American Economic Journal: Microeconomics* 6(4): 35-73.
- Kaiser, Ulrich and Julian Wright. 2006. “Price structure in two-sided markets: Evidence from the magazine industry.” *International Journal of Industrial Organization* 24(1): 1-28.
- Krouse, Sarah. 2022. “Netflix’s ad-supported plan will launch in November at \$6.99 a month.” *Wall Street Journal*, October 13. Accessed at <https://www.wsj.com/articles/netflix-ads-to-launch-in-november-11665680110> on May 15 2023.
- Lambrecht, Anja and Catherine Tucker. 2019. “Algorithmic bias? An empirical study of apparent gender-based discrimination in the display of STEM career ads.” *Management Science* 65(7): 2647-3448.
- Lee, Robin, Michael Whinston, and Ali Yurukoglu. 2021. “Structural empirical analysis of contracting in vertical markets.” In Kate Ho, Ali Hortacsu, and Alessandro Lizzeri, eds., *Handbook of Industrial Organization* Volume 4, 1st ed. Amsterdam: Elsevier.
- Lewis, Randall A. and David H. Reiley. 2014. “Advertising effectively influences older users: How field experiments can improve measurement and targeting.” *Review of Industrial Organization* 44(2): 147–159.
- Liao, Lu, Alan Sorensen, and Andrey Zubanov. 2020. “Measuring the value of targeted television advertising.” University of Wisconsin Working Paper. Preliminary draft accessed at https://www.ssc.wisc.edu/~sorensen/papers/targeted_advertising_jun2020.pdf on May 31 2022 and cited with author permission.
- McGranaghan, Matthew, Jura Liaukonyte, and Kenneth C. Wilbur. 2022. “How viewer tuning, presence, and attention respond to ad content and predict brand search lift.” *Marketing Science* 41(5): 871-1027.
- MoffettNathanson. 2017. “U. S. internet: From cat videos to the NFL - are we at the tipping point for online video?” MoffettNathanson (An SVB Company). Accessed at <https://web.archive.org/web/20230301231204/https://www.moffettnathanson.com/research.aspx?Section=Media+%2FTelecom> on March 1, 2023.
- MoffettNathanson. 2022. “Netflix and Disney: Mad Men to the rescue?” MoffettNathanson (An SVB Company). Accessed at <https://web.archive.org/web/20230301231204/https://www.moffettnathanson.com/research.aspx?Section=Media+%2FTelecom> on March 1, 2023.
- Newman, Andrew Adam. 2012. “In AARP’s view, advertisers need to focus.” *New York Times*, July 19. Accessed at <https://www.nytimes.com/2012/07/19/business/media/aarp-campaign-tries-to-persuade-advertisers.html> on June 18, 2022.

- The News Journal. 2015. *Sunday News Journal* Sep. 13 - 19, 2015. Wilmington, Delaware. Accessed at <https://www.newspapers.com/image/127413547/> on May 19 2022.
- The Nielsen Company. 2017. *Ad Intel Methodology by Medium*. Report by the Nielsen Company. Accessed at <http://en-us.nielsen.com/sitelets/cls/documents/AdIntel/AdIntel-Methodology-by-Medium-Info-Kit.pdf> on February 24 2021.
- The Nielsen Company. 2018. *CMO Report*. Report by The Nielsen Company. Accessed at <https://www.nielsen.com/wp-content/uploads/sites/3/2019/04/nielsen-cmo-report-2018-2.pdf> on January 7 2022.
- The Nielsen Company. 2019. "Ad Intel dataset." The Nielsen Company [publisher], Kilts Center for Marketing, University of Chicago Booth School of Business [distributor]. <https://www.chicagobooth.edu/research/kilts/datasets/nielsenIQ-nielsen>. (accessed March 2019).
- The Nielsen Company. 2021. *Nielsen Local TV View*. Custom report by the Nielsen Company. Accessed at <https://www.nielsen.com/us/en/client-learning/tv/local/nltv/> in February 2021.
- The Nielsen Company. 2022. "Ad Intel dataset." The Nielsen Company [publisher], Kilts Center for Marketing, University of Chicago Booth School of Business [distributor]. <https://www.chicagobooth.edu/research/kilts/datasets/nielsenIQ-nielsen>. (accessed March 2022).
- Organization for Economic Co-operation and Development. 2022. "Consumer Price Index: Total All Items for the United States [CPALTT01USA661S]." FRED, Federal Reserve Bank of St. Louis [publisher]. <https://fred.stlouisfed.org/series/CPALTT01USA661S>. (accessed June 2022).
- Papazian, Ed, editor. 2009. *TV Dimensions 2009*. New York, NY: Media Dynamics, Inc.
- Patel, Sahil. 2015. "Turner's new ad sales strategy sharpens its focus on clients, data." *Digiday* December 10. Accessed at <https://web.archive.org/web/20230208175713/https://digiday.com/media/turners-new-ad-sales-strategy-sharpens-focus-clients-data/> on April 29 2023.
- Phillips, Robert and Graham Young. 2012. "Television advertisement pricing in the United States." In Özalp Özer and Robert Phillips, eds., *Oxford Handbook of Pricing Management*. Oxford, UK: Oxford University Press.
- Pomerantz, Earl. 2006. "Why do advertisers still covet the 18-49s?" *Television Quarterly* 36: 40-44. Accessed at https://web.archive.org/web/20200713102219/http://www.tvquarterly.net/tvq_36_3/media/articles/36.3Why_do_Adver

- tisers.pdf on June 3 2021.
- Prat, Andrea. 2018. "Media power." *Journal of Political Economy* 126(4): 1747-1783.
- Prat, Andrea and Tommaso Valletti. 2022. "Attention oligopoly." *American Economic Journal: Microeconomics* 14(3): 530-557.
- Rochet, Jean-Charles and Jean Tirole. 2003. "Platform competition in two-sided markets." *Journal of the European Economic Association* 1(4): 990–102.
- Shi, C. Matthew. 2016. "Catching (exclusive) eyeballs: Multi-homing and platform competition in the magazine industry." University of Virginia Working Paper. Accessed at https://web.archive.org/web/20230903201736/https://mattshi.weebly.com/uploads/2/4/8/4/24842967/shi_cmatt_jmp-jan16.pdf in September 2023.
- Smith, Ruth B., George P. Moschis, and Roy L. Moore. 1985. "Some advertising influences on the elderly consumer: Implications for theoretical consideration." *Current Issues and Research in Advertising* 8(1): 187-201.
- Spence, Michael and Bruce Owen. 1977. "Television programming, monopolistic competition, and welfare." *Quarterly Journal of Economics* 91(1): 103-26.
- Statista. 2021. *Media Owners Advertising Revenue in the United States in 2021*, by Medium. Research report by Statista. Accessed at <https://www.statista.com/statistics/492908/media-owners-revenue-sector-usa/> on August 16 2021.
- Strikesocial. 2017. *YouTube Advertising: 2017 Benchmark Report*. Report by Strikesocial. Accessed at https://web.archive.org/web/20211019131344/https://strikesocial.com/wp-content/uploads/2018/08/2017-YouTube_Data_Report_.pdf on March 2 2023.
- Surowiecki, James. 2002. "Ageism in advertising." *The New Yorker*, March 24. Accessed at <https://www.newyorker.com/magazine/2002/04/01/ageism-in-advertising> on July 6 2021.
- S&P Global Market Intelligence. 2019. *SNL Kagan Report for 2015*. Report by S&P Global Market Intelligence. Accessed at <https://www.spglobal.com/marketintelligence/en/> in November 2019.
- U.S. Census Bureau. 2021. "2020 Census apportionment results delivered to the president." U.S. Census Bureau Release Number CB21-CN.30. Accessed at <https://www.census.gov/newsroom/press-releases/2021/2020-census-apportionment-results.html> on July 5 2022.
- Veiga, André and E. Glen Weyl. 2016. "Product design in selection markets." *Quarterly Journal of Economics* 131(2): 1007–1056.
- WARC. 2001. "P&G and Viacom ink \$300m cross-media deal." *WARC* May 31. Accessed at ht

[tps://web.archive.org/web/20230428230056/https://www.warc.com/newsandopinion/news/pg_and_viacom_ink_300m_crossmedia_deal/8911](https://web.archive.org/web/20230428230056/https://www.warc.com/newsandopinion/news/pg_and_viacom_ink_300m_crossmedia_deal/8911) on April 29 2023.

Weprin, Alex. 2015. "Upfronts 2015: NBCU's cable lifestyle channels make their pitch." *Politico* March 31. Accessed at <https://web.archive.org/web/20230420230317/https://www.politico.com/media/story/2015/03/upfronts-2015-nbcus-cable-lifestyle-channels-make-their-pitch-003625/> on April 20 2023.

Wikipedia. 2022. "2014–15 United States network television schedule". Wikimedia Foundation. Accessed at https://en.wikipedia.org/wiki/2014%E2%80%9315_United_States_network_television_schedule on May 19 2022.

Wilbur, Kenneth C. 2008. "A two-sided, empirical model of television advertising and viewing markets." *Marketing Science* 27(3): 356-378.

Zubarov, Andrey. 2020. "The TV advertising market: Demographic segmentation and the impact of viewership decline." University of Wisconsin-Madison Working Paper. Accessed at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3853300 on August 14 2021.

A Proofs and Additional Theoretical Results

A.1 Preliminaries

Lemma 1. (i) For any $V(\cdot)$, not necessarily monotone or submodular, in any equilibrium and for any owner $Z \in \mathcal{Z}$, the payments $p_{nZ}^*, p_{n'Z}^*$ made in equilibrium to owner Z by any two advertisers n, n' must coincide, $p_{nZ}^* = p_{n'Z}^*$.

(ii) If $V(\cdot)$ is either monotone and submodular or strictly monotone and not necessarily submodular, then in any equilibrium each advertiser buys slots on all outlets.

Proof. Fix any equilibrium. First, we claim that because the advertisers are homogeneous, they must have the same equilibrium payoff. Suppose toward contradiction that advertiser n' has a strictly higher equilibrium payoff than advertiser n . Now, let advertiser n imitate the strategy of advertiser n' . Since ad slots are not scarce ($N \leq K$), advertiser n must obtain the equilibrium payoff of advertiser n' , which is a contradiction. So all advertisers have the same equilibrium payoff, which we denote by W .

For part (i), suppose toward contradiction that in the equilibrium there exists some owner $Z \in \mathcal{Z}$ such that not all advertisers make the same total payment to Z . Let n be an advertiser whose equilibrium payment to owner Z is at least as much as any other advertiser's (and therefore strictly greater than some other advertiser's). Let $S \subseteq Z$ be the set of outlets owned by owner Z on which advertiser n buys slots. Let owner Z deviate by offering bundle $B = S$ at price $p_S^* - \varepsilon$ for any $\varepsilon > 0$, where p_S^* denotes the minimum price to buy slots on the set of outlets S at the on-path history in the supposed equilibrium. By subgame perfection, at every possible subgame, each advertiser must purchase a payoff-maximizing set of bundles. We claim that upon the deviation by owner Z , all advertisers buy bundle B offered by owner Z . To see why, note that by buying the bundle B upon the deviation (and imitating advertiser n 's choices at the on-path history), any advertiser can obtain a payoff of $W + \varepsilon$, which is strictly greater than the payoff of W prescribed by the equilibrium, and therefore (because no prices have changed other than the price of bundle B) also strictly greater than the payoff from any other choice of bundles. For sufficiently small ε , the deviation increases the payment to owner Z from at least one advertiser other than n more than it decreases the payment from advertiser n , and therefore strictly increases the payoff to owner Z . This establishes that the deviation is profitable, and hence that there is a contradiction.

For part (ii), suppose toward contradiction that in the equilibrium there exists some advertiser n who does not buy a slot on some outlet $j \in Z$ owned by some owner $Z \in \mathcal{Z}$. Let $T \subset Z$ be the set of outlets owned by owner Z on which advertiser n buys slots, and let $R \subseteq \mathcal{J} \setminus Z$ be the set of outlets owned by other owners on which advertiser n buys slots. By part (i), all advertisers pay the same total amount to owner Z , which then must be p_T^* , the minimum price to buy slots on all

outlets in T . The equilibrium payoff for each advertiser is then given by

$$W = V(T \cup R) - p_R^* - p_T^*,$$

where p_R^* is the minimum price to buy slots on all outlets in R . Let owner Z deviate by offering a single bundle Z with a price $p_T^* + \varepsilon$ for some $\varepsilon > 0$. We show that for ε small enough, every advertiser must purchase this bundle. First, note that $V(Z \cup R) - V(T \cup R) > 0$. This is clearly true if $V(\cdot)$ is strictly monotone. If $V(\cdot)$ is submodular and monotone, then we also have

$$V(Z \cup R) - V(T \cup R) \geq V(\mathcal{J}) - V(\mathcal{J} \setminus (Z \setminus T)) \geq V(\mathcal{J}) - V(\mathcal{J} \setminus \{j\}) = v_j > 0$$

where the strict inequality is due to our assumption that every outlet has positive incremental value. Thus, for ε small enough, we have

$$V(Z \cup R) - p_R^* - (p_T^* + \varepsilon) > V(T \cup R) - p_R^* - p_T^* = W.$$

Now, pick any such sufficiently small ε . Note again that by subgame perfection, at every possible subgame, each advertiser must purchase a payoff-maximizing set of bundles. Thus, upon this deviation, each advertiser buys the bundle Z at the price $p_T^* + \varepsilon$, because any strategy not doing so is a feasible strategy at the on-path history in the supposed equilibrium and therefore generates a payoff less than or equal to W . But this is then a profitable deviation for owner Z , and therefore a contradiction. \square

A.2 Proofs Omitted From the Main Text

Proof of Theorem 1. We first construct an equilibrium. Let each owner Z offer, at a price of $p_Z = v_Z$, a single bundle consisting of all outlets in Z . For any profile p of posted prices (including those off of the equilibrium path), and any set of outlets $S \subseteq \mathcal{J}$, let p_S denote the minimum price to buy slots on all the outlets in S . Now let every advertiser buy the same bundle $S^*(p)$, chosen arbitrarily from the set of solutions to the problem

$$\begin{aligned} \max |S| \\ \text{s.t. } S \in \operatorname{argmax}_{S \subseteq \mathcal{J}} V(S) - p_S \end{aligned}$$

By construction, $S^*(\cdot)$ constitutes an equilibrium strategy profile for the advertisers.

It remains to verify that no owner has a profitable deviation from the proposed profile p^* . Observe that if owner Z plays the proposed strategy, then each advertiser buys the bundle offered by owner Z regardless of the other owners' offerings. This is because for any $S \subseteq \mathcal{J} \setminus Z$, submodularity

of $V(\cdot)$ implies

$$V(S \cup Z) - V(S) \geq V(\mathcal{J}) - V(\mathcal{J} \setminus Z) = v_Z.$$

Fix any owner Z . Suppose all other owners $Z' \neq Z$ follow the proposed strategy. Fix any advertiser. By the above observation we know that the advertiser buys slots on all outlets not in Z regardless of the strategy of owner Z . It follows that the most that owner Z can extract from the advertiser is v_Z because for any $T \supseteq \mathcal{J} \setminus Z$, monotonicity of $V(\cdot)$ implies

$$v_Z = V(\mathcal{J}) - V(\mathcal{J} \setminus Z) \geq V(T) - V(\mathcal{J} \setminus Z)$$

where the right hand side is the most that the advertiser is willing to pay for slots on all the outlets in $T \setminus (\mathcal{J} \setminus Z)$. Therefore, there is no profitable deviation from the proposed strategy for owner Z . Since Z is an arbitrary owner, the proposed profile is an equilibrium.

To prove the second part of the statement, fix any equilibrium of the game, and denote its price profile by p^* . By Lemma 1(ii), all advertisers buy slots on all outlets in \mathcal{J} . Therefore, each advertiser pays p_Z^* to each owner Z . If $p_Z^* > v_Z$ for any owner Z , then any advertiser can profitably deviate by only buying slots on all outlets in $\mathcal{J} \setminus Z$. If $p_Z^* < v_Z$ for any owner Z , then, by the submodularity of $V(\cdot)$, owner Z can profitably deviate by offering a single bundle Z with a price $v_Z - \varepsilon$ for $\varepsilon > 0$ sufficiently small to extract $v_Z - \varepsilon > p_Z^*$ from each advertiser. Thus $p_Z^* = v_Z$ for all $Z \in \mathcal{Z}$.

Monotonicity and Submodularity of $V(\cdot)$ in the Viewer-level Model. Let \mathcal{J}_F be the set of outlets whose owner is of format F . For each viewer $i \in \mathcal{I}$ and each set of outlets $S \subseteq \mathcal{J}$, let X_S^i count the number of outlets in S on which viewer i sees ads. For each $S \subseteq \mathcal{J}$ we can write

$$V(S) = |\mathcal{I}| \cdot \mathbb{E} \left[a_i u \left(X_{S \cap \mathcal{J}_1}^i, \dots, X_{S \cap \mathcal{J}_{|\mathcal{F}|}}^i \right) \right] = |\mathcal{I}| \cdot \mathbb{E} [V_i(S)]$$

where the expectation is taken with respect to the population of viewers i and their random viewing behavior, and the value $V_i(\cdot)$ is based on the realized behavior of viewer i . To show that $V(\cdot)$ is monotone and submodular, it suffices to show that $V_i(S)$ is monotone and submodular for all i and realized X^i , since averaging preserves monotonicity and submodularity.

To show monotonicity of $V_i(\cdot)$, pick some i and realized X^i and fix any $S \subseteq \mathcal{J}$ and $j \in \mathcal{J} \setminus S$. Let F be the format of the owner of outlet j . Then,

$$V_i(S \cup \{j\}) - V_i(S) = a_i \mathbf{1}_{i \rightarrow j} \left(u \left(X_{S \cap \mathcal{J}_1}^i, \dots, X_{S \cap \mathcal{J}_F}^i + 1, \dots, X_{S \cap \mathcal{J}_{|\mathcal{F}|}}^i \right) - u \left(X_{S \cap \mathcal{J}_1}^i, \dots, X_{S \cap \mathcal{J}_F}^i, \dots, X_{S \cap \mathcal{J}_{|\mathcal{F}|}}^i \right) \right)$$

where $i \rightarrow j$ denotes the event that viewer i sees ads on outlet j . This shows the monotonicity of $V_i(\cdot)$ as $u(\cdot)$ is monotone.

To show the submodularity of $V_i(\cdot)$, pick some i and realized X^i and fix any $S \subseteq T \subseteq \mathcal{J}$, and $j \in \mathcal{J} \setminus T$. Let F be the format of the owner of outlet j . Since $X_{S \cap \mathcal{J}_{F'}}^i \leq X_{T \cap \mathcal{J}_{F'}}^i$ for all formats F' , we have

$$\begin{aligned}
V_i(S \cup \{j\}) - V_i(S) &= a_i \mathbf{1}_{i \rightarrow j} \left(u \left(X_{S \cap \mathcal{J}_1}^i, \dots, X_{S \cap \mathcal{J}_F}^i + 1, \dots, X_{S \cap \mathcal{J}_{|\mathcal{J}|}}^i \right) - u \left(X_{S \cap \mathcal{J}_1}^i, \dots, X_{S \cap \mathcal{J}_F}^i, \dots, X_{S \cap \mathcal{J}_{|\mathcal{J}|}}^i \right) \right) \\
&\geq a_i \mathbf{1}_{i \rightarrow j} \left(u \left(X_{S \cap \mathcal{J}_1}^i, \dots, X_{T \cap \mathcal{J}_F}^i + 1, \dots, X_{S \cap \mathcal{J}_{|\mathcal{J}|}}^i \right) - u \left(X_{S \cap \mathcal{J}_1}^i, \dots, X_{T \cap \mathcal{J}_F}^i, \dots, X_{S \cap \mathcal{J}_{|\mathcal{J}|}}^i \right) \right) \\
&\geq a_i \mathbf{1}_{i \rightarrow j} \left(u \left(X_{T \cap \mathcal{J}_1}^i, \dots, X_{T \cap \mathcal{J}_F}^i + 1, \dots, X_{T \cap \mathcal{J}_{|\mathcal{J}|}}^i \right) - u \left(X_{T \cap \mathcal{J}_1}^i, \dots, X_{T \cap \mathcal{J}_F}^i, \dots, X_{T \cap \mathcal{J}_{|\mathcal{J}|}}^i \right) \right) \\
&= V_i(T \cup \{j\}) - V_i(T)
\end{aligned}$$

where the second line follows because $u(\cdot)$ has decreasing differences in each argument, and the third line follows because $u(\cdot)$ is submodular. This shows the submodularity of $V_i(\cdot)$.

Proof of Corollary 1. Following the notation in the proof for the viewer-level model, we have that

$$\begin{aligned}
v_Z &= |\mathcal{I}| \cdot \mathbb{E} \left[a_i \left(u \left(X_{\mathcal{J}}^i \right) - u \left(X_{\mathcal{J} \setminus Z}^i \right) \right) \right] \\
&= |\mathcal{I}| \cdot \mathbb{E} \left[a_i \mathbf{1}_{i \rightarrow Z} \mathbf{1}_{i \not\rightarrow \mathcal{J} \setminus Z} \right] \\
&= \sum_i a_i \mathbb{E} \left[\mathbf{1}_{i \rightarrow Z} \right] \mathbb{E} \left[\mathbf{1}_{i \not\rightarrow \mathcal{J} \setminus Z} \right] \\
&= \sum_i a_i \eta_{iZ} \prod_{j \notin Z} (1 - \eta_{ij}) \\
&= \sum_i a_i \eta_{iZ} \prod_{Z' \neq Z} (1 - \eta_{iZ'})
\end{aligned}$$

where we use the notation $\mathbf{1}_{i \rightarrow S}$ to denote the event that viewer i watches at least one outlet in the set S and the notation $\mathbf{1}_{i \not\rightarrow S}$ to denote the opposite. The second line follows because $u(M) = \mathbf{1}_{M > 0}$, the third because of the independence of ad viewing across outlets, and the fourth and fifth by the definition of η_{iZ} and the structure of viewing behavior in the viewer-level model. The corollary then follows by Theorem 1.

Proof of Proposition 1. The setting of Proposition 1 is formally equivalent to a special case of the setting of Proposition 4 in Appendix A.3.1, in which $\beta_1 = u(1) - u(0) = 1 > 0 = u(2) - u(1) = \beta_2$ and, for each owner Z , we consider that it owns a single outlet with viewing probability $\eta_{gj} = \eta_{gZ}$. The desired result then follows by Proposition 4.

Proof of Proposition 2. The setting of Proposition 2 is formally equivalent to a special case of the setting of Proposition 5 in Appendix A.3.1, in which $\beta_1 = u(1) - u(0) = 1 > 0 = u(2) - u(1) = \beta_2$ and, for each owner Z , we consider that it owns a single outlet with viewing probability $\eta_{gj} = \eta_{gZ}$. The desired result then follows by Proposition 5.

Proof of Proposition 3. We proceed by showing that the setting of Proposition 3 is equivalent to a special case of the setting of Theorem 1. Let $\tilde{\mathcal{J}}$ be the set of outlets owned by owners in $\tilde{\mathcal{Z}}$. Let

$$\mathcal{J}^* = (\mathcal{J} \setminus \tilde{\mathcal{J}}) \cup (\tilde{\mathcal{J}} \times \mathcal{I})$$

be an augmented set of outlets that includes viewer-specific instances of all outlets in $\tilde{\mathcal{J}}$. Let $\mathcal{Z}^* = (\mathcal{Z} \setminus \tilde{\mathcal{Z}}) \cup (\tilde{\mathcal{Z}} \times \mathcal{I})$ be the corresponding augmented set of owners. Let $\mathcal{J}_i^* = (\mathcal{J} \setminus \tilde{\mathcal{J}}) \cup (\tilde{\mathcal{J}} \times \{i\}) \subseteq \mathcal{J}^*$ denote the set of outlets on which it is possible to show ads to viewer i . Let $V_i^*(\cdot)$ be the restriction of $V_i(\cdot)$ to \mathcal{J}_i^* , i.e.

$$V_i^*(S) = V_i(S \cap \mathcal{J}_i^*)$$

for all $S \subseteq \mathcal{J}^*$, where with slight abuse of notation we take $\{(j, i)\}$ as equivalent to $\{j\}$ when evaluating $V_i(\cdot)$. Let $V^*(\cdot) = \sum_{i \in \mathcal{I}} V_i^*(\cdot)$. Recall that both monotonicity and submodularity are preserved under restriction and addition. Observe that an instance of the game considered in Theorem 1 with primitives $\mathcal{J}^*, V^*(\cdot), \mathcal{Z}^*$ is equivalent to the game considered in Proposition 3. The conclusions of Proposition 3 therefore follow from Theorem 1 and the structure of the value function $V_i^*(\cdot)$.

Proof of Corollary 2. Following the notation in the proof of Corollary 1, for any owner Z in format 1 we have that

$$\begin{aligned} v_Z &= |\mathcal{I}| \cdot \mathbb{E} \left[a_i \left(u(X_{\mathcal{J}_1}^i, X_{\mathcal{J}_2}^i) - u(X_{\mathcal{J}_1 \setminus Z}^i, X_{\mathcal{J}_2}^i) \right) \right] \\ &= |\mathcal{I}| \cdot \mathbb{E} \left[a_i \left(\mathbf{1}_{X_{\mathcal{J}_1}^i > 0} - \mathbf{1}_{X_{\mathcal{J}_1 \setminus Z}^i > 0} - \phi \left(\mathbf{1}_{X_{\mathcal{J}_1}^i X_{\mathcal{J}_2}^i > 0} - \mathbf{1}_{X_{\mathcal{J}_1 \setminus Z}^i X_{\mathcal{J}_2}^i > 0} \right) \right) \right] \\ &= |\mathcal{I}| \cdot \mathbb{E} \left[a_i \left(\mathbf{1}_{X_{\mathcal{J}_1}^i > 0} - \mathbf{1}_{X_{\mathcal{J}_1 \setminus Z}^i > 0} - \phi \mathbf{1}_{X_{\mathcal{J}_2}^i > 0} \left(\mathbf{1}_{X_{\mathcal{J}_1}^i > 0} - \mathbf{1}_{X_{\mathcal{J}_1 \setminus Z}^i > 0} \right) \right) \right] \\ &= |\mathcal{I}| \cdot \mathbb{E} \left[a_i \left(\mathbf{1}_{X_{\mathcal{J}_1}^i > 0} - \mathbf{1}_{X_{\mathcal{J}_1 \setminus Z}^i > 0} \right) \left(1 - \phi \mathbf{1}_{X_{\mathcal{J}_2}^i > 0} \right) \right] \\ &= \sum_i a_i \mathbb{E}[\mathbf{1}_{i \rightarrow Z}] \mathbb{E}[\mathbf{1}_{i \not\rightarrow \mathcal{J}_1 \setminus Z}] (1 - \phi \mathbb{E}[\mathbf{1}_{i \rightarrow \mathcal{J}_2}]) \end{aligned}$$

$$= \sum_i a_i \eta_{iZ} \left(1 - \phi \left(1 - \prod_{Z' \in \mathcal{Z} \setminus F(Z)} (1 - \eta_{iZ'}) \right) \right) \prod_{Z' \in F(Z) \setminus \{Z\}} (1 - \eta_{iZ'})$$

and

$$v_{iZ} = a_i \eta_{iZ} \left(1 - \phi \left(1 - \prod_{Z' \in \mathcal{Z} \setminus F(Z)} (1 - \eta_{iZ'}) \right) \right) \prod_{Z' \in F(Z) \setminus \{Z\}} (1 - \eta_{iZ'}).$$

The second line follows because $u(M_1, M_2) = \mathbf{1}_{M_1 > 0} + \mathbf{1}_{M_2 > 0} - \phi \mathbf{1}_{M_1 M_2 > 0}$ and the remaining lines follow from the independence of ad viewing across outlets, the structure of viewing behavior, and the definition of η_{iZ} . An analogous construction applies for any owner Z in format 2. The corollary then follows by Proposition 3.

Proof of Corollary 3. This is an immediate consequence of Corollary 2.

A.3 Extensions

A.3.1 Comparative Statics with General Diminishing Returns

Consider a special case of the viewer-level model in which there is a single format, every owner owns a single outlet, $a_i = a > 0$ for all $i \in \mathcal{I}$, and $\eta_{ij} = \eta_{i'j}$ for any $i, i' \in g$ for $g \in G$ and G a partition of \mathcal{I} . To ease notation, let $\beta_0 = 0$, and let

$$\beta_m := u(m) - u(m-1)$$

for $m \geq 1$, so that

$$u(M) = \sum_{m=0}^M \beta_m.$$

Let μ_g denote the size of viewers from group g . Let

$$\lambda_j = \sum_{g \in G} \mu_g \eta_{gj}, \quad \sigma_{gj} = \frac{\mu_g \eta_{gj}}{\lambda_j}$$

denote, respectively, the expected number of viewers seeing an ad on outlet j , and the share of this audience that comes from group g . Then p_j^*/λ_j is the equilibrium price per viewer charged by the owner for an ad slot. In this setting, we define a group $g \in G$ to be *less active* than group $h \in G$ if $\eta_{gj} \leq \eta_{hj}$ for all $j \in \mathcal{J}$.

Proposition 4. *Suppose that outlet $j \in \mathcal{J}$ draws a larger share of its audience from a less active group g and a smaller share of its audience from a more active group h than outlet $k \in \mathcal{J}$, in the sense that $\sigma_{gj} \geq \sigma_{gk}$ and $\sigma_{hj} \leq \sigma_{hk}$, and that the two outlets have equal total audience sizes,*

$\lambda_j = \lambda_k$, and equal shares of audience from groups other than g and h , $\sigma_{g'j} = \sigma_{g'k}$ for all $g' \neq g, h$. Then outlet j has a higher equilibrium price per viewer than outlet k , $p_j^*/\lambda_j \geq p_k^*/\lambda_k$, with strict inequality whenever $\prod_{l \neq j, k} (1 - \eta_{gl}) > \prod_{l \neq j, k} (1 - \eta_{hl})$, $\sigma_{gj} > \sigma_{gk}$, and $\beta_1 > \beta_2$.

Proof. With a slight abuse of notation, for any group g , let g denote both the group and a randomly sampled viewer from the group. By Theorem 1, we can write

$$p_j^* = a \sum_g \mu_g \mathbb{E}[\mathbf{1}_{g \rightarrow j} \beta_{X_{\mathcal{J} \setminus \{j\}}^g + 1}]$$

where $g \rightarrow j$ denotes the event that a randomly sampled viewer g views ads on outlet j and $X_{\mathcal{J} \setminus \{j\}}^g$ counts the random number of outlets other than j on which viewer g views ads. For any $g' \neq g, h$, we have

$$\eta_{g'j} = \frac{\lambda_j \sigma_{g'j}}{\mu_{g'}} = \frac{\lambda_k \sigma_{g'k}}{\mu_{g'}} = \eta_{g'k}.$$

Therefore, for any $g' \neq g, h$, by independence and symmetry,

$$\mathbb{E}[\mathbf{1}_{g' \rightarrow j} \beta_{X_{\mathcal{J} \setminus \{j\}}^{g'} + 1}] = \mathbb{E}[\mathbf{1}_{g' \rightarrow k} \beta_{X_{\mathcal{J} \setminus \{k\}}^{g'} + 1}].$$

To prove that $p_j^*/\lambda_j \geq p_k^*/\lambda_k$, it then suffices to show

$$\mu_g \mathbb{E}[\mathbf{1}_{g \rightarrow j} \beta_{X_{\mathcal{J} \setminus \{j\}}^g + 1} - \mathbf{1}_{g \rightarrow k} \beta_{X_{\mathcal{J} \setminus \{k\}}^g + 1}] \geq \mu_h \mathbb{E}[\mathbf{1}_{h \rightarrow k} \beta_{X_{\mathcal{J} \setminus \{k\}}^h + 1} - \mathbf{1}_{h \rightarrow j} \beta_{X_{\mathcal{J} \setminus \{j\}}^h + 1}].$$

Because ad viewing is independent across outlets for any given viewer, we can write the above as

$$\mu_g [\eta_{gj}(1 - \eta_{gk}) - \eta_{gk}(1 - \eta_{gj})] \mathbb{E}[\beta_{X_{\mathcal{J} \setminus \{j, k\}}^g + 1}] \geq \mu_h [\eta_{hk}(1 - \eta_{hj}) - \eta_{hj}(1 - \eta_{hk})] \mathbb{E}[\beta_{X_{\mathcal{J} \setminus \{j, k\}}^h + 1}]$$

where $X_{\mathcal{J} \setminus \{j, k\}}^g$ counts the random number of outlets not in $\{j, k\}$ on which viewer g views ads. Since $\lambda_j = \lambda_k$, this reduces to

$$(\sigma_{gj} - \sigma_{gk}) \mathbb{E}[\beta_{X_{\mathcal{J} \setminus \{j, k\}}^g + 1}] \geq (\sigma_{hk} - \sigma_{hj}) \mathbb{E}[\beta_{X_{\mathcal{J} \setminus \{j, k\}}^h + 1}].$$

It follows easily from our assumptions that $\sigma_{gj} - \sigma_{gk} = \sigma_{hk} - \sigma_{hj} \geq 0$. So it suffices to show $\mathbb{E}[\beta_{X_{\mathcal{J} \setminus \{j, k\}}^g + 1}] \geq \mathbb{E}[\beta_{X_{\mathcal{J} \setminus \{j, k\}}^h + 1}]$. Since $\eta_{gj} \leq \eta_{hj}$ for all $j \in \mathcal{J}$ and viewing decisions are independent across outlets for both g and h , there exists a monotone coupling of the viewing decisions by g and h in the sense that $\mathbf{1}_{g \rightarrow j} \leq \mathbf{1}_{h \rightarrow j}$ for all $j \in \mathcal{J}$. Under this coupling, we have $X_{\mathcal{J} \setminus \{j, k\}}^g \leq X_{\mathcal{J} \setminus \{j, k\}}^h$ pointwise. The claim then follows directly by noting that β_m is non-increasing in m for $m \geq 1$ and by the definition of the viewer-level model.

To show strict inequality, suppose that $\prod_{l \neq j, k} (1 - \eta_{gl}) > \prod_{l \neq j, k} (1 - \eta_{hl})$, $\sigma_{gj} > \sigma_{gk}$, and

$\beta_1 > \beta_2$. Using integration by parts and the definition of β_m , we have

$$\begin{aligned}\mathbb{E}[\beta_{X_{\mathcal{J}\setminus\{j,k\}}^g+1}] - \mathbb{E}[\beta_{X_{\mathcal{J}\setminus\{j,k\}}^h+1}] &= \int_0^\infty \mathbb{P}(\beta_{X_{\mathcal{J}\setminus\{j,k\}}^g+1} > s) ds - \int_0^\infty \mathbb{P}(\beta_{X_{\mathcal{J}\setminus\{j,k\}}^h+1} > s) ds \\ &= \sum_{m=1}^\infty (\beta_m - \beta_{m+1}) \left(\mathbb{P}(X_{\mathcal{J}\setminus\{j,k\}}^g + 1 \leq m) - \mathbb{P}(X_{\mathcal{J}\setminus\{j,k\}}^h + 1 \leq m) \right) > 0\end{aligned}$$

where the strict inequality follows from the fact that each term in the summation is nonnegative, $\beta_1 > \beta_2$, and $\mathbb{P}(X_{\mathcal{J}\setminus\{j,k\}}^g = 0) = \prod_{l \neq j,k} (1 - \eta_{gl}) > \prod_{l \neq j,k} (1 - \eta_{hl}) = \mathbb{P}(X_{\mathcal{J}\setminus\{j,k\}}^h = 0)$. Since $\sigma_{gj} > \sigma_{gk}$, we then have $(\sigma_{gj} - \sigma_{gk})\mathbb{E}[\beta_{X_{\mathcal{J}\setminus\{j,k\}}^g+1}] > (\sigma_{hk} - \sigma_{hj})\mathbb{E}[\beta_{X_{\mathcal{J}\setminus\{j,k\}}^h+1}]$ and hence $p_j^*/\lambda_j > p_k^*/\lambda_k$. \square

Proposition 5. *Suppose that outlet j has a larger audience than outlet k in the sense that for some $\delta \geq 1$, $\eta_{gj} = \delta\eta_{gk}$ for all $g \in G$. Then outlet j has a higher price per viewer than outlet k , $p_j^*/\lambda_j \geq p_k^*/\lambda_k$ with strict inequality whenever $\eta_{gj} \prod_{l \neq j,k} (1 - \eta_{gl}) > 0$, $\delta > 1$, and $\beta_1 > \beta_2$.*

Proof. We follow the same notation as in the proof of Proposition 4. By Theorem 1, we can write

$$\begin{aligned}p_j^* &= a \sum_g \mu_g \mathbb{E} \left[\mathbf{1}_{g \rightarrow j} \beta_{X_{\mathcal{J}\setminus\{j\}}^g+1} \right] \\ &= a \sum_g \mu_g \left(\eta_{gj} \eta_{gk} \mathbb{E} \left[\beta_{X_{\mathcal{J}\setminus\{j,k\}}^g+2} \right] + \eta_{gj} (1 - \eta_{gk}) \mathbb{E} \left[\beta_{X_{\mathcal{J}\setminus\{j,k\}}^g+1} \right] \right) \\ &= a \sum_g \mu_g \left(\eta_{gk} \eta_{gj} \mathbb{E} \left[\beta_{X_{\mathcal{J}\setminus\{j,k\}}^g+2} \right] + \delta \eta_{gk} \left(1 - \frac{1}{\delta} \eta_{gj} \right) \mathbb{E} \left[\beta_{X_{\mathcal{J}\setminus\{j,k\}}^g+1} \right] \right) \\ &= a \sum_g \mu_g \left(\eta_{gk} \eta_{gj} \mathbb{E} \left[\beta_{X_{\mathcal{J}\setminus\{j,k\}}^g+2} \right] + \eta_{gk} (1 - \eta_{gj}) \mathbb{E} \left[\beta_{X_{\mathcal{J}\setminus\{j,k\}}^g+1} \right] + \eta_{gk} (\delta - 1) \mathbb{E} \left[\beta_{X_{\mathcal{J}\setminus\{j,k\}}^g+1} \right] \right) \\ &= a \sum_g \mu_g \eta_{gk} \left(\mathbb{E} \left[\beta_{X_{\mathcal{J}\setminus\{k\}}^g+1} \right] + (\delta - 1) \mathbb{E} \left[\beta_{X_{\mathcal{J}\setminus\{j,k\}}^g+1} \right] \right) \\ &\geq a \sum_g \mu_g \eta_{gk} \left(\mathbb{E} \left[\beta_{X_{\mathcal{J}\setminus\{k\}}^g+1} \right] + (\delta - 1) \mathbb{E} \left[\beta_{X_{\mathcal{J}\setminus\{j,k\}}^g+1} \right] \right) = \delta p_k^* = \frac{\lambda_j}{\lambda_k} p_k^*\end{aligned}$$

where we have used independence of ad viewing across outlets, $\eta_{gj} = \delta\eta_{gk}$, $\delta \geq 1$, $X_{\mathcal{J}\setminus\{k\}}^g \geq X_{\mathcal{J}\setminus\{j,k\}}^g$, and the fact that β_m is nonincreasing in m for $m \geq 1$.

To show strict inequality, suppose that $\eta_{gj} \prod_{l \neq j,k} (1 - \eta_{gl}) > 0$, $\delta > 1$, and $\beta_1 > \beta_2$. For any group g , using integration by parts and the definition of β_m , we have

$$\begin{aligned}\mathbb{E} \left[\beta_{X_{\mathcal{J}\setminus\{j,k\}}^g+1} \right] - \mathbb{E} \left[\beta_{X_{\mathcal{J}\setminus\{k\}}^g+1} \right] &= \int_0^\infty \mathbb{P}(\beta_{X_{\mathcal{J}\setminus\{j,k\}}^g+1} > s) ds - \int_0^\infty \mathbb{P}(\beta_{X_{\mathcal{J}\setminus\{k\}}^g+1} > s) ds \\ &= \sum_{m=1}^\infty (\beta_m - \beta_{m+1}) \left(\mathbb{P}(X_{\mathcal{J}\setminus\{j,k\}}^g + 1 \leq m) - \mathbb{P}(X_{\mathcal{J}\setminus\{k\}}^g + 1 \leq m) \right) > 0\end{aligned}$$

where the strict inequality follows from that each term in the summation is nonnegative, $\beta_1 > \beta_2$, and $\mathbb{P}\left(X_{\mathcal{J}\setminus\{j,k\}}^g = 0\right) - \mathbb{P}\left(X_{\mathcal{J}\setminus\{k\}}^g = 0\right) = \eta_{gj} \prod_{l \neq j,k} (1 - \eta_{gl}) > 0$. Since $\delta > 1$, we then have $(\delta - 1) \mathbb{E}\left[\beta_{X_{\mathcal{J}\setminus\{j,k\}}^g + 1}\right] > (\delta - 1) \mathbb{E}\left[\beta_{X_{\mathcal{J}\setminus\{k\}}^g + 1}\right]$ and hence $p_j^*/\lambda_j > p_k^*/\lambda_k$. \square

A.3.2 Endogenous Response of Viewers

There is a set of programs \mathcal{K} . Each outlet consists of K programs. Each program has one ad slot. Each advertiser that purchases an ad slot on a given outlet is randomly assigned to the slot in one of the outlet's programs. There is a set of viewers \mathcal{I} . Each viewer views a subset of programs. Whether a given viewer views a given program depends on whether that program carries an ad, but not on whether other programs do. Thus, as in Ambrus, Calvano, and Reisinger (2016, Section II), advertising may be a nuisance (or attraction) that affects a given outlet's audience, but advertising on one outlet does not drive audience to (or attract audience from) another. For each viewer that views its ad on $M \in \mathbb{N}$ distinct outlets, each advertiser gets value $u(M) \geq 0$ where $u(\cdot)$ is nondecreasing and exhibits decreasing differences.

Proposition 6. *There exists an equilibrium. In any equilibrium, all advertisers buy slots on all outlets, and the payment by each advertiser to each owner Z is given by $p_Z^* = v_Z$.*

Proof. By Theorem 1, it suffices to show that the value function $V(\cdot)$, in this setting, satisfies monotonicity and submodularity. Let $k \in \mathcal{K}$ denote a generic program and let $\mathcal{K}_j \subseteq \mathcal{K}$ be the programs associated with outlet j . For each viewer $i \in \mathcal{I}$, let $i \dashrightarrow k$ denote the event that viewer i watches program k in the scenario that program k carries an ad. This event may be random if viewing behavior is probabilistic. For a set of programs $\mathcal{K}' \subseteq \mathcal{K}$, let

$$X_{\mathcal{K}'}^i = \sum_{k \in \mathcal{K}'} \mathbf{1}_{i \dashrightarrow k}$$

be the (possibly random) number of programs in \mathcal{K}' watched by i in the scenario that each program $k \in \mathcal{K}'$ carries an ad. Let

$$\mathcal{R} = \{\mathcal{K}' \subseteq \mathcal{K} : |\mathcal{K}' \cap \mathcal{K}_j| = 1 \text{ for all } j \in \mathcal{J}\}$$

be the set of all sets of representative programs, such that within each set there is one program for each outlet. For any set of outlets $S \subseteq \mathcal{J}$, let $\mathcal{K}_S = \cup_{j \in S} \mathcal{K}_j$. For each advertiser, the value of a set of outlets $S \subseteq \mathcal{J}$ can be written as

$$V(S) = |\mathcal{I}| \cdot \mathbb{E} \left[\frac{1}{|\mathcal{R}|} \sum_{\mathcal{K}' \in \mathcal{R}} u \left(X_{\mathcal{K}' \cap \mathcal{K}_S}^i \right) \right]$$

where the expectation is taken with respect to the population of viewers and the (possibly random) viewing behavior of each viewer. Because averaging preserves monotonicity and submodularity, it suffices to show that for any realized viewing behavior and representative programs \mathcal{R} , and any $\mathcal{K}' \in \mathcal{R}$, we have that the function

$$\tilde{V}_i(S) := u\left(X_{\mathcal{K}' \cap \mathcal{K}_S}^i\right)$$

is monotone and submodular. It is clear that $\tilde{V}_i(S)$ is monotone since $u(\cdot)$ is nondecreasing. For submodularity, note that for any $S \subseteq T \subseteq \mathcal{J}$ and any $j \in \mathcal{J} \setminus T$,

$$\begin{aligned} \tilde{V}_i(S \cup \{j\}) - \tilde{V}_i(S) &= u\left(X_{\mathcal{K}' \cap \mathcal{K}_{S \cup \{j\}}}^i\right) - u\left(X_{\mathcal{K}' \cap \mathcal{K}_S}^i\right) \\ &\geq u\left(X_{\mathcal{K}' \cap \mathcal{K}_{T \cup \{j\}}}^i\right) - u\left(X_{\mathcal{K}' \cap \mathcal{K}_T}^i\right) \\ &= \tilde{V}_i(T \cup \{j\}) - \tilde{V}_i(T) \end{aligned}$$

where the second line follows from the assumption that $u(\cdot)$ has decreasing differences. \square

A.3.3 Rationing of Ad Slots

Suppose that we may have $N > K$ and assume that bundle prices can only take on values in the set $\{0, \Delta, 2\Delta, \dots\}$ where $\Delta > 0$ is some fixed increment. For this extension, we allow for mixed strategies and assume that the advertisers make purchasing decisions sequentially (in a random order) rather than simultaneously.

Proposition 7. *There exists a subgame perfect Nash equilibrium, possibly in mixed strategies, and in any subgame perfect Nash equilibrium each owner Z earns an expected revenue per ad slot between $(v_Z - \Delta) / |Z|$ and $\sum_{j \in Z} V(\{j\}) / |Z|$.*

Proof. We prove the second part of the statement first. Fix any subgame perfect equilibrium allowing for mixed strategies (SPEMS). Suppose, toward contradiction, that expected revenue per slot is strictly higher than $\sum_{j \in Z} V(\{j\}) / |Z|$ for some owner $Z \in \mathcal{Z}$. Then the owner's expected total revenue is strictly higher than $K \sum_{j \in Z} V(\{j\})$. Thus with positive probability, the owner earns a realized revenue strictly higher than $K \sum_{j \in Z} V(\{j\})$. In any such event, there is at least one advertiser who buys slots on a set of outlets $B \subseteq Z$ and pays strictly more than $\sum_{j \in B} V(\{j\})$ to the owner. Let $S \subseteq \mathcal{J}$ be the set of outlets that this advertiser buys slots on. Since any non-negative submodular function is also sub-additive, we have that

$$V(S) - V(S \setminus B) \leq V(B) - V(\emptyset) \leq \sum_{j \in B} V(\{j\}).$$

So not buying anything from Z is a profitable deviation for the advertiser, implying a contradiction.

Next, toward contradiction suppose that the expected revenue per slot is strictly lower than $(v_Z - \Delta)/|Z|$ for some owner $Z \in \mathcal{Z}$. Then the expected total revenue is strictly lower than $K(v_Z - \Delta)$. Let the owner deviate by offering a single bundle Z with a price $\tilde{p}_Z = \lceil v_Z - \Delta \rceil$, where $\lceil x \rceil$ denotes the operator that rounds x up to the closest value in $\{0, \Delta, 2\Delta, \dots\}$. Note that

$$v_Z - \Delta \leq \lceil v_Z - \Delta \rceil < v_Z.$$

Since $\tilde{p}_Z < v_Z$, by an argument analogous to the one in the proof of Theorem 1, submodularity of $V(\cdot)$ implies that, in any realization, the owner would be able to sell all the slots and secure revenue $K\tilde{p}_Z$. Because this is a profitable deviation for the owner, we have a contradiction.

To show the existence of a SPEMS, we construct an auxiliary finite game in normal form, apply the standard existence result, and then recover a SPEMS in the original game. Consider a simultaneous-move game between all the owners. Let

$$\mathcal{A}(Z) = \{0, \Delta, 2\Delta, \dots, \lceil V(\mathcal{J}) \rceil, \infty\}^{|\mathcal{P}(Z)|}$$

be the set of pure strategies that an owner can choose from. Clearly, $\mathcal{A}(Z)$ is finite for any Z . For each pure strategy profile $p \in \mathcal{A}(Z)$, draw a random order for the advertisers and then let the advertisers, in that order, choose which slots to buy given the posted prices specified in p . Then, for each $p \in \mathcal{A}(Z)$, assign the resulting expected revenue (averaged over different orders) for owner Z as the payoff to owner Z in the auxiliary game. This constructs a finite normal-form game among the owners. Standard results then imply the existence of a Nash equilibrium, possibly in mixed strategies. Call this equilibrium \mathcal{E} . Now let each owner play the strategy prescribed by \mathcal{E} in the original game. Evidently, this constructs a SPEMS for the original game. \square

A.3.4 Partially Increasing Returns

Theorem 1 relies on submodularity of $V(\cdot)$. Under strict monotonicity the conclusion of Theorem 1 obtains under a weakening of submodularity.

Proposition 8. *Suppose that $V(\cdot)$ is strictly monotone and that $V(S \cup Z) - V(S) \geq V(\mathcal{J}) - V(\mathcal{J} \setminus Z)$ for all $Z \in \mathcal{Z}$ and $S \subseteq \mathcal{J} \setminus Z$. Then there exists an equilibrium, and in any equilibrium, all advertisers buy slots on all outlets, and the payment by each advertiser to each owner Z is given by $p_Z^* = v_Z$.*

Proof. We first construct an equilibrium. We use the same construction as in the proof of Theorem 1. When verifying the construction, the only properties of $V(\cdot)$ used in the proof of Theorem 1 are

that $V(\cdot)$ is monotone and that for any $S \subseteq \mathcal{J} \setminus Z$,

$$V(S \cup Z) - V(S) \geq V(\mathcal{J}) - V(\mathcal{J} \setminus Z)$$

which we assume.

To prove the second part of the statement, fix any equilibrium of the game. By Lemma 1(ii) and strict monotonicity of $V(\cdot)$, all advertisers buy slots on all outlets in \mathcal{J} . The rest follows analogously to the proof of Theorem 1. \square

The decreasing differences condition on $V(\cdot)$ in the hypothesis of Proposition 8 is strictly weaker than submodularity. In particular, consider the following example.

Example 1. Owners are singletons, each of a set \mathcal{I} of viewers $i \in \mathcal{I}$ sees ads on at least L outlets, each outlet has a strictly positive number of viewers, and an advertiser's value for viewer i seeing its ad M times is $a_i u(M) = a_i \sum_{m=0}^M \beta_m$ where $a_i > 0$ for all i , $\beta_0 = 0, \beta_m > 0$ for all m , β_m is non-increasing for all $m \geq L$, and $\beta_L \leq \min_{1 \leq m < L} \beta_m$.

Example 1 allows increasing returns to advertising for viewers receiving few impressions (as in, e.g., Dubé, Hitsch, and Manchanda 2005). Although Example 1 need not satisfy the hypotheses of Theorem 1, Example 1 does satisfy the hypotheses of Proposition 8.

Proposition 9. *The value function $V(\cdot)$ in Example 1 satisfies the hypotheses of Proposition 8.*

Proof. For any viewer $i \in \mathcal{I}$ and outlet $j \in \mathcal{J}$, let $i \rightarrow j$ denote the event that viewer i watches outlet j . For any viewer $i \in \mathcal{I}$ and set of outlets $S \subseteq \mathcal{J}$, let X_S^i denote the number of outlets in set S on which viewer i views ads. We have that

$$V(S) = |\mathcal{I}| \cdot \mathbb{E} [a_i u(X_S^i)] = |\mathcal{I}| \cdot \mathbb{E} [V_i(S)]$$

where the expectation is taken with respect to the population of viewers and their (possibly random) viewing behavior. To show strict monotonicity, observe that for any $j \in \mathcal{J}$ and any $S \subseteq \mathcal{J} \setminus \{j\}$,

$$V_i(S \cup \{j\}) - V_i(S) = a_i \mathbf{1}_{i \rightarrow j} \beta_{X_S+1}$$

is weakly positive, and strictly so with positive probability. To show that

$$V(S \cup Z) - V(S) \geq V(\mathcal{J}) - V(\mathcal{J} \setminus Z)$$

for all $Z \in \mathcal{Z}, S \subseteq \mathcal{J} \setminus Z$, because owners are singletons, it suffices to show that for any j and any

$S \subseteq \mathcal{J} \setminus \{j\}$, we have

$$V_i(S \cup \{j\}) - V_i(S) = a_i \mathbf{1}_{i \rightarrow j} \beta_{X_S+1} \geq a_i \mathbf{1}_{i \rightarrow j} \beta_{X_{\mathcal{J} \setminus \{j\}}+1} = V_i(\mathcal{J}) - V_i(\mathcal{J} \setminus \{j\}).$$

To see the above, consider the event $i \rightarrow j$. Then, $X_{\mathcal{J} \setminus \{j\}} + 1 = X_{\mathcal{J}} \geq L$. Note that $X_S + 1 = X_{S \cup \{j\}} \leq X_{\mathcal{J}}$. If $X_{S \cup \{j\}} \geq L$, then $\beta_{X_{S \cup \{j\}}} \geq \beta_{X_{\mathcal{J}}}$ because β_m is non-increasing for $m \geq L$. If $X_{S \cup \{j\}} < L$, then we have

$$\beta_{X_{\mathcal{J}}} \leq \beta_L \leq \min_{1 \leq m < L} \beta_m \leq \beta_{X_{S \cup \{j\}}}.$$

So in either case, the claimed inequality holds. \square

The setting of Example 1 continues to satisfy the hypotheses of Proposition 8 if a small number of viewers view fewer than L outlets.

Example 2. Owners are singletons. There is a set of viewers $\bar{\mathcal{I}}$. There is a subset of viewers $\mathcal{I} \subseteq \bar{\mathcal{I}}$ such that (i) each viewer $i \in \mathcal{I}$ sees ads on at least L outlets and (ii) each viewer $i \in \mathcal{I}$ sees ads on both of any pair of outlets $\{j, k\} \in \mathcal{J}^2$ with strictly positive probability. An advertiser's value for viewer i seeing its ad M times is $a_i u(M) = a_i \sum_{m=0}^M \beta_m$ where $a_i \in (0, \bar{a})$ for all i , $\beta_0 = 0, \beta_m > 0$ for all m , β_m is strictly decreasing for all $m \geq L$, and $\beta_L < \min_{1 \leq m < L} \beta_m$.

Proposition 10. *There exists $\bar{\epsilon} > 0$ such that the value function $V(\cdot)$ in Example 2 satisfies the hypotheses of Proposition 8 as long as the share of viewers not in \mathcal{I} is no more than $\bar{\epsilon}$.*

Proof. Strict monotonicity follows by the same argument as in the proof of Proposition 9. Now, fix any $j \in \mathcal{J}$, any $S \subset \mathcal{J} \setminus \{j\}$, and any $k \in \mathcal{J} \setminus (S \cup \{j\})$. Because each viewer $i \in \mathcal{I}$ sees ads on both of any pair of outlets $\{j, k\} \in \mathcal{J}^2$ with strictly positive probability, following the notation in the proof of Proposition 9 we have that

$$\mathbb{P}(i \rightarrow j, X_S^i < X_{\mathcal{J} \setminus \{j\}}^i | i \in \mathcal{I}) > 0.$$

Thus, since β_m is strictly decreasing for all $m \geq L$, and $\beta_L < \min_{1 \leq m < L} \beta_m$, we have

$$\mathbb{E}[a_i \mathbf{1}_{i \rightarrow j} \beta_{X_S^i+1} | i \in \mathcal{I}] > \mathbb{E}[a_i \mathbf{1}_{i \rightarrow j} \beta_{X_{\mathcal{J} \setminus \{j\}}^i+1} | i \in \mathcal{I}].$$

Now let

$$\tau := \min_{j \in \mathcal{J}, S \subset \mathcal{J} \setminus \{j\}} \left(\mathbb{E}[a_i \mathbf{1}_{i \rightarrow j} \beta_{X_S^i+1} | i \in \mathcal{I}] - \mathbb{E}[a_i \mathbf{1}_{i \rightarrow j} \beta_{X_{\mathcal{J} \setminus \{j\}}^i+1} | i \in \mathcal{I}] \right) > 0.$$

Let $\bar{\beta} = \max_{1 \leq m < L} \beta_m$. Let

$$\bar{\epsilon} = \frac{\tau}{\tau + \bar{a}\bar{\beta}} > 0.$$

We claim that if $\mathbb{P}(i \notin \mathcal{I}) \leq \bar{\varepsilon}$ then $V(\cdot)$ satisfies that $V(S \cup Z) - V(S) \geq V(\mathcal{J}) - V(\mathcal{J} \setminus Z)$ for all $Z \in \mathcal{Z}, S \subseteq \mathcal{J} \setminus Z$. Because owners are singletons, it suffices to consider any $j \in \mathcal{J}$ and any $S \subseteq \mathcal{J} \setminus \{j\}$. Note that

$$\begin{aligned}
V(S \cup \{j\}) - V(S) &\geq |\bar{\mathcal{I}}| \cdot \mathbb{P}(i \in \mathcal{I}) \mathbb{E}[a_i \mathbf{1}_{i \rightarrow j} \beta_{X_{S+1}} | i \in \mathcal{I}] \\
&\geq |\bar{\mathcal{I}}| \cdot \mathbb{P}(i \in \mathcal{I}) \left(\mathbb{E}[a_i \mathbf{1}_{i \rightarrow j} \beta_{X_{\mathcal{J} \setminus \{j\} + 1}} | i \in \mathcal{I}] + \tau \right) \\
&\geq |\bar{\mathcal{I}}| \cdot \left(\mathbb{P}(i \in \mathcal{I}) \mathbb{E}[a_i \mathbf{1}_{i \rightarrow j} \beta_{X_{\mathcal{J} \setminus \{j\} + 1}} | i \in \mathcal{I}] + \mathbb{P}(i \notin \mathcal{I}) \bar{a} \bar{\beta} \right) \\
&\geq |\bar{\mathcal{I}}| \cdot \left(\mathbb{P}(i \in \mathcal{I}) \mathbb{E}[a_i \mathbf{1}_{i \rightarrow j} \beta_{X_{\mathcal{J} \setminus \{j\} + 1}} | i \in \mathcal{I}] + \mathbb{P}(i \notin \mathcal{I}) \mathbb{E}[a_i \mathbf{1}_{i \rightarrow j} \beta_{X_{\mathcal{J} \setminus \{j\} + 1}} | i \notin \mathcal{I}] \right) \\
&= V(\mathcal{J}) - V(\mathcal{J} \setminus \{j\}),
\end{aligned}$$

where the second inequality follows from the construction of τ and the third inequality follows from the fact that $\mathbb{P}(i \notin \mathcal{I}) \leq \bar{\varepsilon}$. \square

Lastly, we show that analogues of Propositions 1 and 2 hold in a special case of Example 1 that is analogous to the setting of Section 2.2.

Proposition 11. *Consider a special case of Example 1 with $a_i = a$ for all i , and further impose the structure in Section 2.2, where for each group $g \in G$, there is a set $\mathcal{L}_g \subseteq \mathcal{J}$ such that $|\mathcal{L}_g| \geq L$ and $\eta_{gj} = 1$ for all $j \in \mathcal{L}_g$. Then the conclusions of Propositions 1 and 2 hold.*

Proof. We follow the same arguments and notation as in the proofs of Propositions 4 and 5. By Propositions 8 and 9, we have only to characterize the incremental value for each owner Z .

For the conclusion of Proposition 1, recall that we use a monotone coupling in the proof of Proposition 4. Under that coupling we have $X_{\mathcal{J} \setminus \{j,k\}}^g \leq X_{\mathcal{J} \setminus \{j,k\}}^h$ pointwise, where we recall that $X_{\mathcal{J} \setminus \{j,k\}}^g$ counts the random number of outlets not in $\{j,k\}$ on which viewer g sees ads. It follows that $X_{\mathcal{J} \setminus \{j,k\}}^g, X_{\mathcal{J} \setminus \{j,k\}}^h \geq L - 2$ and so $X_{\mathcal{J} \setminus \{j,k\}}^g + 1, X_{\mathcal{J} \setminus \{j,k\}}^h + 1 \geq L - 1$. Since $\beta_L \leq \min_{1 \leq m < L} \beta_m \leq \beta_{L-1}$ and β_m is non-increasing in m for $m \geq L$, β_m is non-increasing in m for $m \geq L - 1$. Since $X_{\mathcal{J} \setminus \{j,k\}}^g + 1 \leq X_{\mathcal{J} \setminus \{j,k\}}^h + 1$ pointwise, we have $\beta_{X_{\mathcal{J} \setminus \{j,k\}}^g + 1} \geq \beta_{X_{\mathcal{J} \setminus \{j,k\}}^h + 1}$ pointwise and so $\mathbb{E} \left[\beta_{X_{\mathcal{J} \setminus \{j,k\}}^g + 1} \right] \geq \mathbb{E} \left[\beta_{X_{\mathcal{J} \setminus \{j,k\}}^h + 1} \right]$, which concludes the proof as before.

For the conclusion of Proposition 2, recall that $X_{\mathcal{J} \setminus \{k\}}^g$ counts the random number of outlets other than k on which viewer g sees ads. So $X_{\mathcal{J} \setminus \{k\}}^g + 1 \geq X_{\mathcal{J} \setminus \{j,k\}}^g + 1 \geq L - 1$. Because β_m is non-increasing in m for $m \geq L - 1$, we have $\beta_{X_{\mathcal{J} \setminus \{j,k\}}^g + 1} \geq \beta_{X_{\mathcal{J} \setminus \{k\}}^g + 1}$ pointwise and so $\mathbb{E} \left[\beta_{X_{\mathcal{J} \setminus \{j,k\}}^g + 1} \right] \geq \mathbb{E} \left[\beta_{X_{\mathcal{J} \setminus \{k\}}^g + 1} \right]$, which concludes the proof as before. \square

A.3.5 Heterogeneous Advertisers

Suppose now that each advertiser $n \in \mathcal{N}$ has a monotone and submodular value function $V_n(\cdot)$. If outlets can post advertiser-specific prices, then the result is parallel to that in Theorem 1, in the sense that the equilibrium price of owner Z 's bundle to advertiser n is given by $v_{n,Z} = V_n(\mathcal{J}) - V_n(\mathcal{J} \setminus Z)$. If outlets cannot post advertiser-specific prices, then incremental pricing holds if heterogeneity among the advertisers is sufficiently small compared to the incremental value of a single outlet. Let $\underline{v}_Z = \min_{n \in \mathcal{N}} v_{n,Z}$ and $\bar{v}_Z = \max_{n \in \mathcal{N}} v_{n,Z}$ denote the minimum and maximum values of $v_{n,Z}$, respectively, with respect to n . Let $\varphi(Z) = \min_{n \in \mathcal{N}, j \in Z} V_n((\mathcal{J} \setminus Z) \cup \{j\}) - V_n(\mathcal{J} \setminus Z)$ denote the minimal incremental value of any one of owner Z 's outlets. In the special case where Z is a single-outlet owner, $\varphi(Z) = \underline{v}_Z$.

Proposition 12. *Suppose that heterogeneity in the value functions $V_n(\cdot)$ is small in the sense that $\bar{v}_Z - \underline{v}_Z \leq \frac{1}{N} \varphi(Z)$ for all $Z \in \mathcal{Z}$. Then there exists an efficient equilibrium, and in any efficient equilibrium, all advertisers buy slots on all outlets, and $p_Z^* = \underline{v}_Z$ for all $Z \in \mathcal{Z}$.*

Proof. As in the proof of Theorem 1, we first construct an equilibrium. For any profile p of posted prices (including those off of the equilibrium path), let p_S denote the minimum price to buy slots on all the outlets in S . Now let every advertiser n buy the bundle $S_n^*(p)$, chosen arbitrarily from the set of solutions to the problem

$$\begin{aligned} & \max |S| \\ & \text{s.t. } S \in \operatorname{argmax}_{S \subseteq \mathcal{J}} V_n(S) - p_S \end{aligned}$$

By construction, $(S_1^*(\cdot), \dots, S_N^*(\cdot))$ constitutes an equilibrium strategy profile for the advertisers.

Now, let each owner Z offer a single bundle consisting of all outlets in Z with a price $p_Z = \underline{v}_Z$. We only need to check that each owner has no profitable deviation. Observe that if $p_Z = \underline{v}_Z$ is offered by some owner Z and there is no proper subset $W \subset Z$ being offered, then any advertiser will buy the bundle Z regardless of the prices posted by owners other than Z . This is because for any $S \subseteq \mathcal{J} \setminus Z$, submodularity of $V_n(\cdot)$ implies that

$$V_n(S \cup Z) - V_n(S) \geq V_n(\mathcal{J}) - V_n(\mathcal{J} \setminus Z) \geq \min_{n' \in \mathcal{N}} V_{n'}(\mathcal{J}) - V_{n'}(\mathcal{J} \setminus Z) = \underline{v}_Z.$$

Fix an owner $Z \in \mathcal{Z}$ and suppose all other players follow the proposed strategy. We claim that offering a single bundle Z with a price \underline{v}_Z is an optimal strategy for owner Z . To see this, consider two cases.

Case 1: Suppose Z offers some set of bundles \mathcal{B}_Z such that every advertiser buys a slot on every outlet in Z . Then the minimal price to buy all outlets in Z must be no more than \underline{v}_Z because

otherwise there is one advertiser who can profitably deviate by simply not buying anything in \mathcal{B}_Z . Hence in this case the owner cannot do better than simply offering the bundle Z with a price \underline{v}_Z .

Case 2: Suppose Z offers some set of bundles \mathcal{B}_Z such that there exist at least one outlet $j \in Z$ and one advertiser $\tilde{n} \in \mathcal{N}$ such that advertiser \tilde{n} does not buy a slot on outlet j . The total revenue that owner Z extracts is no more than

$$\max_{n,j} \left\{ \sum_{n' \neq n} v_{n',Z} + v_{n,Z \setminus \{j\}} \right\}$$

where $v_{n,Z \setminus \{j\}} = V_n(\mathcal{J}) - V_n(\mathcal{J} \setminus (Z \setminus \{j\}))$ is the incremental value to advertiser n of the owner's outlets excluding outlet j . We also have that

$$\begin{aligned} \max_{n,j} \left\{ \sum_{n' \neq n} v_{n',Z} + v_{n,Z \setminus \{j\}} \right\} &\leq \max_{n,j} \{N\bar{v}_Z - (v_{n,Z} - v_{n,Z \setminus \{j\}})\} \\ &= N\bar{v}_Z - \min_{n,j} \{V_n((\mathcal{J} \setminus Z) \cup \{j\}) - V_n(\mathcal{J} \setminus Z)\} \\ &= N\bar{v}_Z - \varphi(Z) \\ &\leq N\bar{v}_Z - N(\bar{v}_Z - \underline{v}_Z) = N\underline{v}_Z \end{aligned}$$

where we have used the assumption that $\bar{v}_Z - \underline{v}_Z \leq \frac{1}{N}\varphi(Z)$. Hence the owner also cannot do better than simply offering the bundle Z with a price \underline{v}_Z .

Thus the construction is an equilibrium. The outcome is efficient because all advertisers buy slots on all outlets.

To prove the second part of the statement, fix any efficient equilibrium. All outlets must sell N slots, because the preferences for each player are quasilinear in money and thus the total surplus is maximized only if all potential trades are realized (recall $K \geq N$). Then by the argument in Case 1, we know that $p_Z^* \leq \underline{v}_Z$ for all $Z \in \mathcal{Z}$. Moreover, p_Z^* cannot be strictly lower than \underline{v}_Z for any owner Z , because if this were the case then it would be a profitable deviation for owner Z to offer a single bundle Z with a price $\underline{v}_Z - \varepsilon$ for $\varepsilon > 0$ small enough. Hence $p_Z^* = \underline{v}_Z$ for all $Z \in \mathcal{Z}$. \square

Note that the hypothesis of Proposition 12 restricts the incremental values rather than the level of $V_n(\cdot)$, in the sense that it allows for $V_n(\cdot) = V(\cdot) + c_n$ for any c_n that preserves non-negativity.³⁴ The restriction on incremental values becomes more demanding as the number of advertisers, N , grows large.

³⁴When there are two or more owners, it also allows for $V_n(\emptyset) = V(\emptyset)$ and $V_n(\mathcal{J}') = V(\mathcal{J}') + c_n$ for $\emptyset \neq \mathcal{J}' \subseteq \mathcal{J}$, where $c_n \geq 0$.

A.3.6 Unbundled Pricing

Suppose that each owner $Z \in \mathcal{Z}$ is endowed with a partition \mathcal{F}_Z of Z such that owner Z can offer a bundle $B \subseteq Z$ if and only if there exists $C \in \mathcal{F}_Z$ such that $B \subseteq C$. For example, if \mathcal{F}_Z partitions owner Z 's outlets into singletons, then the owner cannot bundle together ads on different outlets. Denote by $v_B^S = V(S) - V(S \setminus B)$ the incremental value of bundle $B \subseteq S$ relative to set $S \subseteq \mathcal{J}$. We refine the notion of equilibrium by assuming that, when indifferent, owners break ties in favor of offering fewer bundles, and each advertiser breaks ties by favoring owners according to a prespecified ordering.

Proposition 13. *In any equilibrium satisfying the tie-breaking rule, each bundle sold has a price of $p_B^* = v_B^S$, where $S \subseteq \mathcal{J}$ is the set of all outlets sold.*

Proof. Let \mathcal{O}_n denote advertiser n 's tie-breaking ordering over owners; that is, if indifferent among one or more sets of bundles, advertiser n chooses in a manner that maximizes the payoffs of the owners according to a lexicographic preference over owners defined by \mathcal{O}_n .

If in some equilibrium advertiser n' obtains a strictly greater payoff than advertiser n , advertiser n can improve their payoff by mimicking the strategy of advertiser n' . It follows that in any equilibrium all advertisers must obtain the same payoff.

Now, fix some equilibrium with advertiser payoff W and pricing profile p . For any set $S \subseteq \mathcal{J}$, let p_S be the minimum price to buy slots on all the outlets in S in the equilibrium.

We first prove that, analogous to Lemma 1(i), in the equilibrium all advertisers pay the same total amount to any given owner Z . Suppose toward contradiction that there exists some owner $Z \in \mathcal{Z}$ such that not all advertisers make the same total payment to Z . Let n be an advertiser whose payment to owner Z in the supposed equilibrium is at least as large as any other advertiser's, and strictly larger than some other advertiser's. Let S_C be the set of outlets that advertiser n buys slots on in the cell $C \in \mathcal{F}_Z$. Let owner Z deviate by offering the bundles $\mathcal{B} := \{S_C : S_C \neq \emptyset, C \in \mathcal{F}_Z\}$ with prices $\{p_{S_C} - \varepsilon : S_C \neq \emptyset, C \in \mathcal{F}_Z\}$ for some $\varepsilon > 0$. Note that $|\mathcal{B}| \geq 1$ since advertiser n pays a positive amount to owner Z . By buying all the bundles offered in this deviation of Z (and imitating advertiser n 's choices at the on-path history in the supposed equilibrium), any advertiser can obtain a payoff of $W + \varepsilon|\mathcal{B}|$. Because any set of outlets an advertiser wants to buy slots on after this deviation is also a valid choice at the on-path history, if an advertiser does not buy all the bundles in \mathcal{B} , then the advertiser gets at most $W + \varepsilon(|\mathcal{B}| - 1)$. Hence after this deviation, all advertisers buy the bundles in \mathcal{B} offered by Z . Because this is a profitable deviation for owner Z when ε is small enough, we have a contradiction.

We next prove that every advertiser must buy the same set of bundles from any given owner Z and that owner Z offers at most one bundle from each cell in \mathcal{F}_Z . Fix any owner $Z \in \mathcal{Z}$ and any advertiser $n \in \mathcal{N}$. Let S_C be the set of outlets that advertiser n buys slots on in the

cell $C \in \mathcal{F}_Z$. Suppose that owner Z offers the bundles $\mathcal{B} := \{S_C : S_C \neq \emptyset, C \in \mathcal{F}_Z\}$ with prices $\{p_{S_C} : S_C \neq \emptyset, C \in \mathcal{F}_Z\}$. We claim that owner Z weakly increases their payoff with this strategy compared to any equilibrium strategy. For each advertiser, this change restricts the set of possible choices while keeping at least one choice that maintains the equilibrium payoff (imitating the choice of advertiser n at the on-path history in the equilibrium). Because at the on-path history all advertisers pay the same total amount to Z , this change can only decrease owner Z 's payoff if there is some advertiser n' (not necessarily different from n) who now breaks ties in favor of some owner that ranks higher than Z in $\mathcal{O}_{n'}$. However, that choice must also be made at the on-path history by advertiser n' due to the tie-breaking rule. But then advertiser n' pays strictly less than advertiser n to owner Z in the equilibrium, which contradicts what we previously showed.

Because an owner chooses to offer fewer bundles when indifferent, the above observation implies that every advertiser must buy the same set of bundles from any given owner Z and that owner Z offers at most one bundle from each cell in \mathcal{F}_Z . (Otherwise, owner Z may simply pick an advertiser n who buys the smallest number of bundles from Z and offer the set of bundles \mathcal{B} as defined above to strictly decrease the total number of bundles offered without decreasing payoff.) Then all advertisers buy slots on the same set of outlets (say S) and any owner Z offers $\mathcal{B}_Z := \{S \cap B : S \cap B \neq \emptyset, B \in \mathcal{F}_Z\}$ as the available bundles.

Therefore, in the second stage, the set of feasible bundles that advertisers can choose

$$\{S \cap B : S \cap B \neq \emptyset, B \in \mathcal{F}_Z\}_{Z \in \mathcal{Z}}$$

must be a partition of the set S . For any bundle B offered by any owner Z , by rationality of the advertisers,

$$p_B \leq V(S) - V(S \setminus B) = v_B^S.$$

Next we show that $p_B \geq v_B^S$. Suppose toward contradiction that there exist some owner Z and some bundle $B' \in \mathcal{F}_Z$, $B' \subseteq S$ such that $p_{B'} < v_{B'}^S$. Consider the following deviation. Let owner Z offer all bundles in \mathcal{B}_Z as in the equilibrium but change the price for each bundle B to $\tilde{p}_B = v_B^S - \varepsilon$ for some $\varepsilon > 0$. We claim that after this deviation, all advertisers continue buying slots on the same outlets from owner Z as in the equilibrium. Indeed, if an advertiser stops buying some bundle $B \in \mathcal{B}_Z$, then the advertiser can only choose $S' \subseteq S \setminus B$ since the set of available bundles is a partition of S . But submodularity of $V(\cdot)$ implies

$$V(S' \cup B) - V(S') \geq V(S) - V(S \setminus B) = v_B^S > \tilde{p}_B.$$

Therefore owner Z can extract $v_B^S - \varepsilon$ for each bundle $B \in \mathcal{B}_Z$ from each advertiser. For ε sufficiently small, this is then a profitable deviation for owner Z since in the equilibrium we have

$p_B \leq v_B^S$ for all $B \in \mathcal{B}_Z$ and $p_{B'} < v_{B'}^S$ for some bundle $B' \in \mathcal{B}_Z$. We therefore have contradiction. \square

A.3.7 Bargaining between Owners and Advertisers

Suppose that rather than simultaneously posting prices, owners bargain with advertisers a la Nash-in-Nash (Lee, Whinston, and Yurukoglu 2021). We follow the notation in Lee, Whinston, and Yurukoglu (2021). For each owner Z and each advertiser n , let

$$\mathcal{C}_{Zn} := \{(B, p) : B \subseteq Z, p \in \mathbb{R}_+\}$$

be the contract space, with an element denoted by \mathbb{C}_{Zn} . For a contract \mathbb{C}_{Zn} , let $B(\mathbb{C}_{Zn})$ and $p(\mathbb{C}_{Zn})$ denote the associated bundle and price. Let $\mathbb{C}_0 = \{(\emptyset, 0)\}$ denote the null contract. For a given set of contracts $\mathbb{C} := \{\mathbb{C}_{Zn}\}_{Z \in \mathcal{Z}, n=1, \dots, N}$, owner Z 's payoff is given by

$$\Pi_Z(\mathbb{C}) = \sum_n p(\mathbb{C}_{Zn})$$

and advertiser n 's payoff is given by

$$\Pi_n(\mathbb{C}) = V\left(\bigcup_{Z \in \mathcal{Z}} B(\mathbb{C}_{Zn})\right) - \sum_{Z \in \mathcal{Z}} p(\mathbb{C}_{Zn}).$$

Given the set of contracts \mathbb{C}_{-Zn} excluding pair (Z, n) , let

$$\mathcal{C}_{Zn}^+(\mathbb{C}_{-Zn}) = \{\mathbb{C}_{Zn} \in \mathcal{C}_{Zn} : \Pi_n(\{\mathbb{C}_{Zn}, \mathbb{C}_{-Zn}\}) - \Pi_n(\{\mathbb{C}_0, \mathbb{C}_{-Zn}\}) \geq 0\}$$

be the set of contracts between Z and n that give non-negative gains from trade to owner Z and advertiser n (note that only the constraint for the advertiser is relevant as any contract would give non-negative gains from trade to owner Z). Recall that a set of contracts $\hat{\mathbb{C}}$ is a *Nash-in-Nash equilibrium* if:

(i) For all Z, n such that $\hat{\mathbb{C}}_{Zn} \neq \mathbb{C}_0$,

$$\hat{\mathbb{C}}_{Zn} \in \operatorname{argmax}_{\mathbb{C}_{Zn} \in \mathcal{C}_{Zn}^+(\hat{\mathbb{C}}_{-Zn})} [\Pi_Z(\{\mathbb{C}_{Zn}, \hat{\mathbb{C}}_{-Zn}\}) - \Pi_Z(\{\mathbb{C}_0, \hat{\mathbb{C}}_{-Zn}\})]^{\xi_Z} [\Pi_n(\{\mathbb{C}_{Zn}, \hat{\mathbb{C}}_{-Zn}\}) - \Pi_n(\{\mathbb{C}_0, \hat{\mathbb{C}}_{-Zn}\})]^{1-\xi_Z},$$

where $\xi_Z \in [0, 1]$ denotes the bargaining weight for owner Z .

(ii) For all Z, n such that $\hat{\mathbb{C}}_{Zn} = \mathbb{C}_0$, there is no contract in $\mathcal{C}_{Zn}^+(\hat{\mathbb{C}}_{-Zn})$ that gives strictly positive gains from trade to both Z and n .

Proposition 14. *If all owners have identical bargaining weights, there exists a Nash-in-Nash equilibrium in which all advertisers buy slots on all outlets, and the payment by each advertiser to each owner Z is proportional to v_Z . If $V(\cdot)$ is strictly monotone, then this outcome is unique.*

Proof. We first show that $\hat{\mathbb{C}} := \{(Z, \xi_Z v_Z)\}$ is a Nash-in-Nash equilibrium. Condition (ii) clearly holds. For (i), note that

$$\Pi_Z(\{\mathbb{C}_{Zn}, \hat{\mathbb{C}}_{-Zn}\}) - \Pi_Z(\{\mathbb{C}_0, \hat{\mathbb{C}}_{-Zn}\}) = p(\mathbb{C}_{Zn})$$

and

$$\Pi_n(\{\mathbb{C}_{Zn}, \hat{\mathbb{C}}_{-Zn}\}) - \Pi_n(\{\mathbb{C}_0, \hat{\mathbb{C}}_{-Zn}\}) = V(B(\mathbb{C}_{Zn}) \cup (\mathcal{J} \setminus Z)) - V(\mathcal{J} \setminus Z) - p(\mathbb{C}_{Zn}).$$

Because $V(\cdot)$ is monotone, a solution to the Nash bargaining problem is given by $B(\mathbb{C}_{Zn}) = Z$ and $p(\mathbb{C}_{Zn}) = \xi_Z(V(\mathcal{J}) - V(\mathcal{J} \setminus Z))$. This proves that $\hat{\mathbb{C}}$ is a Nash-in-Nash equilibrium.

For uniqueness, suppose that $V(\cdot)$ is strictly monotone and fix any Nash-in-Nash equilibrium $\tilde{\mathbb{C}}$. For any Z and n , regardless of $\tilde{\mathbb{C}}_{-Zn}$, given that $V(\cdot)$ is strictly monotone, any solution to the Nash bargaining problem must have $B(\mathbb{C}_{Zn}) = Z$. Therefore, for any Z and n , any solution to the Nash bargaining problem must have $p(\mathbb{C}_{Zn}) = \xi_Z(V(\mathcal{J}) - V(\mathcal{J} \setminus Z))$, proving the claim.

Lastly, observe that when $\xi_Z = \xi$ for all $Z \in \mathcal{Z}$, the payments to each owner under $\hat{\mathbb{C}}$ are proportional to v_Z . □

A.3.8 Auctioning of Advertising Slots

Suppose that rather than simultaneously posting prices, owners simultaneously set reserve prices for each of their bundles, and then conduct simultaneous first-price auctions.

Proposition 15. *If owners simultaneously set reserve prices and then conduct simultaneous first-price auctions, there exists an equilibrium, and in any equilibrium all advertisers buy slots on all outlets, and the payment by each advertiser to each owner Z is given by v_Z .*

Proof. Fix any profile of announced reserve prices $p := \{p_B : B \subseteq Z, Z \in \mathcal{Z}\}$. For any advertiser, bidding strictly above the reserve price for any bundle B is strictly dominated by bidding at the reserve price p_B , because in both cases the advertiser is guaranteed to win the bundle (as $K \geq N$). Thus, for any bundle B , each advertiser either bids at the reserve price for that bundle, or bids below the reserve price and loses the auction. Therefore, after eliminating the strictly dominated strategies for the advertisers, this game is strategically equivalent to the pricing game of our main model. Hence, the claim follows directly from Theorem 1. □

Now, in addition, we further characterize equilibrium in a model where owners conduct auctions, advertising slots are scarce, and advertisers' valuations are heterogeneous. Suppose that each owner owns one outlet, and identify each owner with the outlet j that the owner owns. Suppose further that each outlet has K slots, where $K < N$ (so ad slots are scarce). The advertisers are heterogeneous, with value functions given by $\alpha_n V(\cdot)$ where $V(\cdot)$ is monotone and submodular with $V(\emptyset) = 0$. We order the advertisers so that $\alpha_1 > \alpha_2 > \dots > \alpha_N > 0$. We assume that

$$\alpha_{K+1}V(\{j\}) < \alpha_K(V(\mathcal{J}) - V(\mathcal{J} \setminus \{j\})) \quad \forall j \in \mathcal{J}. \quad (\text{A1})$$

Each owner runs a uniform price auction in which the K slots are allocated to the K highest bidders at a price equal to the $(K + 1)^{\text{th}}$ highest bid, with ties broken in favor of advertisers with higher α_n . The auctions happen simultaneously. Each advertiser simultaneously submits bids to every auction. We take an equilibrium to be a Nash equilibrium in pure strategies. We say an equilibrium is *owner-optimal* if there is no other equilibrium that gives weakly higher payoffs to all owners and strictly higher payoff to at least one owner. We say an equilibrium is *efficient* if the equilibrium allocation maximizes total surplus among all possible allocations.

Proposition 16. *Suppose that Assumption (A1) holds. Then, there exists an efficient owner-optimal equilibrium, and in every efficient owner-optimal equilibrium, for every owner j , the clearing price of auction j is $\alpha_K(V(\mathcal{J}) - V(\mathcal{J} \setminus \{j\}))$.*

Proof. Consider the following strategy profile: in every auction j , each advertiser $n \leq K$ bids $\alpha_n(V(\mathcal{J}) - V(\mathcal{J} \setminus \{j\}))$; advertiser $K + 1$ bids $\alpha_K(V(\mathcal{J}) - V(\mathcal{J} \setminus \{j\}))$; and each advertiser $n > K + 1$ bids 0.

We show that this is an equilibrium. Fix any j and any advertiser n with $n \leq K$. Because this is a $(K + 1)$ -th price auction, the advertiser cannot influence the price it pays conditional on winning. Regardless of the choices the advertiser makes on other auctions, the advertiser weakly prefers to win auction j at the price of $\alpha_K(V(\mathcal{J}) - V(\mathcal{J} \setminus \{j\}))$ rather than to lose the auction at that price because for any $S \subseteq \mathcal{J} \setminus \{j\}$, we have

$$\alpha_n(V(S \cup \{j\}) - V(S)) \geq \alpha_K(V(S \cup \{j\}) - V(S)) \geq \alpha_K(V(\mathcal{J}) - V(\mathcal{J} \setminus \{j\})),$$

where we have used submodularity of $V(\cdot)$. Therefore, advertiser n has no profitable deviation.

Next, fix any j and any advertiser n with $n > K$. Note that advertiser n loses every auction under the proposed strategy profile. Also note that to win auction j , advertiser n has to pay $\alpha_K(V(\mathcal{J}) - V(\mathcal{J} \setminus \{j\}))$. However, regardless of the choices the advertiser makes on other auctions, the advertiser strictly prefers not to win auction j at this price, because for any $S \subseteq \mathcal{J} \setminus \{j\}$,

we have

$$\alpha_n(V(S \cup \{j\}) - V(S)) \leq \alpha_n V(\{j\}) \leq \alpha_{K+1} V(\{j\}) < \alpha_K (V(\mathcal{J}) - V(\mathcal{J} \setminus \{j\})),$$

where we have used submodularity of $V(\cdot)$ and Assumption (A1). Therefore, advertiser n has no profitable deviation.

Now, we show that this equilibrium is owner-optimal. Suppose toward contradiction that there is another equilibrium that gives some owner j a strictly higher payoff and all other owners weakly higher payoffs. Fix any such equilibrium \mathcal{E}' . By the argument above, no advertiser n with $n > K$ would want to win auction j at a price strictly higher than $\alpha_K (V(\mathcal{J}) - V(\mathcal{J} \setminus \{j\}))$. Therefore, the K winning bidders in auction j must be advertisers $1, \dots, K$. However, at a price strictly higher than $\alpha_K (V(\mathcal{J}) - V(\mathcal{J} \setminus \{j\}))$ for auction j , advertiser K must lose some auction $j' \neq j$, because otherwise the advertiser can profitably deviate to losing auction j . Then, since there are K winners in auction j' , there must be an advertiser n' with $n' > K$ who wins auction j' . For owner j' to have a weakly higher payoff in equilibrium \mathcal{E}' than in the original equilibrium, the clearing price in auction j' must be weakly higher than $\alpha_K (V(\mathcal{J}) - V(\mathcal{J} \setminus \{j'\}))$. But then advertiser n' can profitably deviate to losing auction j' by the argument above, which is a contradiction.

We claim that the allocation of ad slots to advertisers under this equilibrium is the unique efficient allocation. To see this, fix any efficient allocation x and suppose toward contradiction that x differs from the allocation under the equilibrium. Then, it must be that some advertiser $n \leq K$ is not allocated to an ad slot on some outlet j , which means that some advertiser $n' > K$ is allocated to an ad slot on outlet j . Consider an allocation \tilde{x} that is the same as x except that it allocates the ad slot on outlet j to n instead of n' . We claim that this change strictly increases the total surplus. Indeed, let S be the set of outlets whose slots are assigned to advertiser n under allocation x , and similarly define S' for advertiser n' . Then,

$$\begin{aligned} \alpha_n (V(S \cup \{j\}) - V(S)) &\geq \alpha_K (V(\mathcal{J}) - V(\mathcal{J} \setminus \{j\})) \\ &> \alpha_{K+1} V(\{j\}) \geq \alpha_{n'} (V(S') - V(S' \setminus \{j\})), \end{aligned}$$

where we have used submodularity of $V(\cdot)$ and Assumption (A1). Therefore,

$$\alpha_n V(S \cup \{j\}) + \alpha_{n'} V(S' \setminus \{j\}) > \alpha_n V(S) + \alpha_{n'} V(S'),$$

and hence \tilde{x} gives a strictly higher total surplus than x , which is a contradiction.

Finally, fix any efficient owner-optimal equilibrium. By efficiency and the argument above, the winning bidders in every auction must be advertisers $1, \dots, K$. If there is any auction j in which the clearing price is strictly higher than $\alpha_K (V(\mathcal{J}) - V(\mathcal{J} \setminus \{j\}))$, then advertiser K can profitably

deviate to losing auction j . Therefore, in every auction j , the clearing price must be weakly lower than $\alpha_K(V(\mathcal{J}) - V(\mathcal{J} \setminus \{j\}))$. Now, if there is any auction j' in which the clearing price is strictly lower than $\alpha_K(V(\mathcal{J}) - V(\mathcal{J} \setminus \{j'\}))$, the equilibrium cannot be owner-optimal, because we have just shown an equilibrium that has clearing prices equal to $\alpha_K(V(\mathcal{J}) - V(\mathcal{J} \setminus \{j\}))$ for all j . Thus, in every efficient owner-optimal equilibrium, the clearing price in every auction j must be exactly $\alpha_K(V(\mathcal{J}) - V(\mathcal{J} \setminus \{j\}))$. \square

A.3.9 Competition between Advertisers

We consider a setting in which each owner owns a single outlet and advertisers make purchasing decisions sequentially in random order. We modify the value function $V(\cdot)$ as follows. Let advertiser n 's value for buying ads on the set of outlets S_n be $V(S_n, \vec{S}_{-n})$, where $\vec{S}_{-n} \in \mathcal{J}^{N-1}$ is the vector of sets bought by other advertisers. We say $\vec{S}_{-n} \leq \vec{S}'_{-n}$ if each entry of the vector is smaller in the set-inclusion order. Since all owners are single-outlet owners, we use j to denote both an outlet and the owner associated with the outlet. Let $\vec{\mathcal{J}}$ be the vector of length $N-1$ with \mathcal{J} in each entry, and $\overrightarrow{\mathcal{J} \setminus \{j\}}$ be the vector of length $N-1$ with $\mathcal{J} \setminus \{j\}$ in each entry. We impose two assumptions:

$$V(\mathcal{J}, \vec{S}) - V(\mathcal{J} \setminus \{j\}, \vec{S}) \geq V(\mathcal{J}, \vec{S}') - V(\mathcal{J} \setminus \{j\}, \vec{S}') \text{ for any } \vec{S} \leq \vec{S}' \text{ and } j; \quad (\text{A2})$$

$$V(\mathcal{J}, \overrightarrow{\mathcal{J} \setminus \{j\}}) - V(\mathcal{J} \setminus \{j\}, \overrightarrow{\mathcal{J} \setminus \{j\}}) \leq \left(1 + \frac{1}{N}\right) \left(V(\mathcal{J}, \vec{\mathcal{J}}) - V(\mathcal{J} \setminus \{j\}, \vec{\mathcal{J}})\right) \text{ for any } j. \quad (\text{A3})$$

Let $\tilde{v}_j = V(\mathcal{J}, \vec{\mathcal{J}}) - V(\mathcal{J} \setminus \{j\}, \vec{\mathcal{J}})$ denote the modified incremental value of outlet j in this setting.

Proposition 17. *Suppose that $V(\cdot, \vec{S})$ is monotone and submodular for any \vec{S} , and that $V(\cdot, \cdot)$ satisfies (A2) and (A3). Then there exists an equilibrium in which all advertisers buy slots on all outlets, and the price for outlet j is $p_j^* = \tilde{v}_j$.*

Proof. We construct an equilibrium as follows. Let each owner j announce price \tilde{v}_j . For each profile of prices p announced (including off-the-equilibrium-path histories), the subgame in the second stage is a finite extensive-form game and hence admits an equilibrium by backward induction. When doing the backward induction, if an advertiser is indifferent between different sets of outlets to buy slots on, we pick one with the largest cardinality. Now we verify that no owner has a profitable deviation.

Observe that at any history, if $p_j = \tilde{v}_j$ is offered by an owner $j \in \mathcal{J}$, then any advertiser will buy a slot on outlet j regardless of p_{-j} and what other advertisers do. This is because for any $S \subseteq \mathcal{J} \setminus \{j\}$ and any $\vec{S}_{-n} \leq \vec{\mathcal{J}}$,

$$V(S \cup \{j\}, \vec{S}_{-n}) - V(S, \vec{S}_{-n}) \geq V(\mathcal{J}, \vec{S}_{-n}) - V(\mathcal{J} \setminus \{j\}, \vec{S}_{-n}) \geq V(\mathcal{J}, \vec{\mathcal{J}}) - V(\mathcal{J} \setminus \{j\}, \vec{\mathcal{J}})$$

where we have used submodularity of $V(\cdot, \vec{S}_{-n})$ and Assumption (A2). Further, when all other advertisers buy slots on all outlets, the incremental value for an advertiser to buy a slot on some outlet j is exactly $V(\mathcal{J}, \vec{\mathcal{J}}) - V(\mathcal{J} \setminus \{j\}, \vec{\mathcal{J}})$. Therefore, at the proposed price profile, for any outlet j , each advertiser is indifferent between buying and not buying a slot on outlet j in the proposed equilibrium.

Now fix any owner $j \in \mathcal{J}$. Suppose all other players follow the proposed strategy. Owner j is selling N slots by announcing price \tilde{v}_j and clearly has no incentive to decrease the price. Consider the deviation of raising the price. By the earlier observation, all advertisers would continue buying slots on outlets in $\mathcal{J} \setminus \{j\}$. Therefore, by (A2), owner j cannot extract more than $V(\mathcal{J}, \overrightarrow{\mathcal{J} \setminus \{j\}}) - V(\mathcal{J} \setminus \{j\}, \overrightarrow{\mathcal{J} \setminus \{j\}})$ from any advertiser. Further, we claim that at least one advertiser would stop buying the slot on outlet j after the price increase, because if not, the last advertiser to move can profitably deviate by buying only the slots on the outlets in $\mathcal{J} \setminus \{j\}$. Thus there are at most $N - 1$ advertisers buying a slot on outlet j . Hence owner j 's revenue is at most

$$(N-1) \left(V(\mathcal{J}, \overrightarrow{\mathcal{J} \setminus \{j\}}) - V(\mathcal{J} \setminus \{j\}, \overrightarrow{\mathcal{J} \setminus \{j\}}) \right) \leq (N-1) \left(1 + \frac{1}{N} \right) \left(V(\mathcal{J}, \vec{\mathcal{J}}) - V(\mathcal{J} \setminus \{j\}, \vec{\mathcal{J}}) \right) \leq N\tilde{v}_j$$

where the first inequality is due to (A3). So there is no profitable deviation for owner j . Since this holds for any owner, the construction is an equilibrium. \square

A.3.10 Incentive to Invest in Content

Suppose that there is a set of viewers \mathcal{I} . Each viewer $i \in \mathcal{I}$ is attracted to each owner Z 's content with probability $\alpha_{iZ} \in [0, 1]$. If a viewer i is attracted to owner Z 's content, the viewer sees ads on outlets $j \in \mathcal{Z}$ with probability $\eta_{ij} \in (0, 1)$, independently across j , and other details follow the reach-only model. Prior to the game specified in Section 2, each owner simultaneously announces a choice of α_{iZ} for all viewers i , paying a content cost $\sum_{i \in \mathcal{I}} C_{iZ}(\alpha_{iZ})$ where $C_{iZ}(0) = C'_{iZ}(0) = 0$ and $C'_{iZ}(1) > a_i$, for $C'_{iZ}(\cdot)$ the first derivative of $C_{iZ}(\cdot)$. For a given investment profile $\{(\alpha_{iZ})_{i \in \mathcal{I}}\}_{Z \in \mathcal{Z}}$, a viewer i , and an owner Z , let $V_i^Z(\cdot; \alpha)$ denote the value function induced by the viewing probabilities of viewer i conditional on viewer i being attracted to owner Z .

Proposition 18. *Suppose the investment profile $\{(\alpha_{iZ})_{i \in \mathcal{I}}\}_{Z \in \mathcal{Z}}$ is an equilibrium. Then,*

$$C'_{iZ}(\alpha_{iZ}) = V_i^Z(\mathcal{J}; \alpha) - V_i^Z(\mathcal{J} \setminus Z; \alpha) = a_i \eta_{iZ} \prod_{Z' \neq Z} (1 - \alpha_{iZ'} \eta_{iZ'}).$$

Proof. For a given investment profile α , by Theorem 1, the equilibrium payments in the subgame are given by

$$p_Z^*(\alpha) = \sum_i \alpha_{iZ} (V_i^Z(\mathcal{J}; \alpha) - V_i^Z(\mathcal{J} \setminus Z; \alpha)).$$

So when making investment choices, owner Z maximizes the objective

$$\sum_i \alpha_{iZ} (V_i^Z(\mathcal{J}; \alpha) - V_i^Z(\mathcal{J} \setminus Z; \alpha)) - \sum_i C_{iZ}(\alpha_{iZ}),$$

which is differentiable in $(\alpha_{iZ})_{i \in \mathcal{I}}$, and where $V_i^Z(\mathcal{J}; \alpha) - V_i^Z(\mathcal{J} \setminus Z; \alpha)$ does not depend on α_{iZ} . For owner Z , the first order condition for α_{iZ} is given by

$$C'_{iZ}(\alpha_{iZ}) = V_i^Z(\mathcal{J}; \alpha) - V_i^Z(\mathcal{J} \setminus Z; \alpha).$$

Since $C'_{iZ}(1) > a_i \geq V_i^Z(\mathcal{J}; \alpha) - V_i^Z(\mathcal{J} \setminus Z; \alpha)$ for all α , i , and Z , in any equilibrium no owner Z will choose $\alpha_{iZ} = 1$ for any viewer i . Then, since $C'_{iZ}(0) = 0$ and $\eta_{iZ} > 0$, in any equilibrium no owner Z will choose $\alpha_{iZ} = 0$ for any viewer i . Hence, in any equilibrium, the above first order condition must hold for all i and all Z . The form of $V_i^Z(\mathcal{J}; \alpha) - V_i^Z(\mathcal{J} \setminus Z; \alpha)$ follows by Corollary 1. □

In the (unattainable) limiting case where $\alpha_{iZ} = 1$ for all $Z \in \mathcal{Z}$, we have that $C'_{iZ}(1) = p_{iZ}^*$ for all Z .

A.4 Alternative Models with Declining Audience

In this section, we present two alternative models for ad markets, and show that these models cannot generate rising ad revenues with declining audience, even though they can generate rising prices.

A.4.1 Heterogeneous Advertisers with Additively Separable Values

Here we interpret declining audience as a decline in advertisers' separable valuations for outlets. Consider a special case of the extension in A.3.5 with heterogeneous advertisers in which each owner owns a single outlet $j \in \mathcal{J}$ and each advertiser $n \in \mathcal{N}$ has value $V_n(\cdot)$ such that

$$V_n(S) = \sum_{j \in S} V_n(\{j\})$$

for any $S \subseteq \mathcal{J}$.

Proposition 19. *Consider two markets \mathcal{M} and $\tilde{\mathcal{M}}$, one in which the advertisers' values are V_n and the other in which the advertisers' values are \tilde{V}_n . Suppose that $\tilde{V}_n \leq V_n$ for all advertisers n . Then the total ad revenue in any equilibrium of market $\tilde{\mathcal{M}}$ must be weakly lower than the total ad revenue in any equilibrium of market \mathcal{M} .*

Proof. We show that for any outlet j , its ad revenue must be lower in market $\tilde{\mathcal{M}}$. Note that because each advertiser has an additively separable value function, the pricing problem for outlet j , regardless of how other outlets price, is equivalent to the pricing problem where the N advertisers have values $V_n(\{j\})$. Now, note that for any price p ,

$$p \cdot \sum_n \mathbf{1}_{V_n(\{j\}) \geq p} \geq p \cdot \sum_n \mathbf{1}_{\tilde{V}_n(\{j\}) \geq p}.$$

Thus, for any price p ,

$$\sup_{p'} \left\{ p' \cdot \sum_n \mathbf{1}_{V_n(\{j\}) \geq p'} \right\} \geq p \cdot \sum_n \mathbf{1}_{\tilde{V}_n(\{j\}) \geq p}.$$

Therefore, we have

$$\sup_{p'} \left\{ p' \cdot \sum_n \mathbf{1}_{V_n(\{j\}) \geq p'} \right\} \geq \sup_p \left\{ p \cdot \sum_n \mathbf{1}_{\tilde{V}_n(\{j\}) \geq p} \right\}.$$

The left-hand side is the ad revenue for outlet j in any equilibrium of market \mathcal{M} , and the right-hand side is the ad revenue for outlet j in any equilibrium of market $\tilde{\mathcal{M}}$. The claim follows. \square

Proposition 19 does not preclude that the equilibrium price for any outlet j can be higher in market $\tilde{\mathcal{M}}$ than in market \mathcal{M} . For example, suppose that there are two advertisers such that $V_1 = V_2 - \varepsilon > 0$ in market \mathcal{M} , and $\tilde{V}_1 = 0$ and $\tilde{V}_2 = V_2$ in market $\tilde{\mathcal{M}}$. For $\varepsilon > 0$ small enough, the equilibrium price for each outlet is greater in market $\tilde{\mathcal{M}}$ than in market \mathcal{M} .

A.4.2 Falling Supply of Ad Slots

Here we interpret declining audience as a decrease in the total supply of ad slots K . Suppose that the advertisers are heterogeneous. Specifically, each advertiser n values the number of total impressions at b_n per impression, and has a budget constraint c_n . We use perfect competition as a solution concept, so that a price per impression p^* is an equilibrium if and only if

$$D(p^*) = K$$

where $D(\cdot)$ is the aggregate market demand for ad slots, with $D(p) = \sum_{n \in \mathcal{N}} D_n(p) = \frac{1}{p} \sum_{n \in \mathcal{N}} c_n \mathbf{1}_{b_n \geq p}$ for any price $p \geq 0$.

Proposition 20. *Consider two markets \mathcal{M} and $\tilde{\mathcal{M}}$, one in which the supply of ad slots is K and the other in which the supply of ad slots is \tilde{K} . Suppose that $\tilde{K} < K$. Then the total ad revenue in any equilibrium of market $\tilde{\mathcal{M}}$ must be weakly lower than the total ad revenue in any equilibrium of market \mathcal{M} , although the price per impression in any equilibrium of market $\tilde{\mathcal{M}}$ must be weakly greater than the price per impression in any equilibrium of market \mathcal{M} .*

Proof. At any equilibrium price p^* for market \mathcal{M} , we have that

$$\frac{1}{p^*} \sum_n c_n \mathbf{1}_{b_n \geq p^*} = K.$$

At any equilibrium price \tilde{p}^* for market $\tilde{\mathcal{M}}$, we have that

$$\frac{1}{\tilde{p}^*} \sum_n c_n \mathbf{1}_{b_n \geq \tilde{p}^*} = \tilde{K}.$$

It follows by inspection that $\tilde{p}^* \geq p^*$ because $\tilde{K} < K$. We also observe that revenues are weakly lower in market $\tilde{\mathcal{M}}$ because

$$\tilde{p}^* \tilde{K} = \sum_n c_n \mathbf{1}_{b_n \geq \tilde{p}^*} \leq \sum_n c_n \mathbf{1}_{b_n \geq p^*} = p^* K$$

and $\tilde{p}^* \geq p^*$. □

B Additional Empirical Results


Appendix Figure 1: Sensitivity to Alternative Samples

Price per impression vs.

audience activity


audience size

Baseline




Alternative sample years

2014

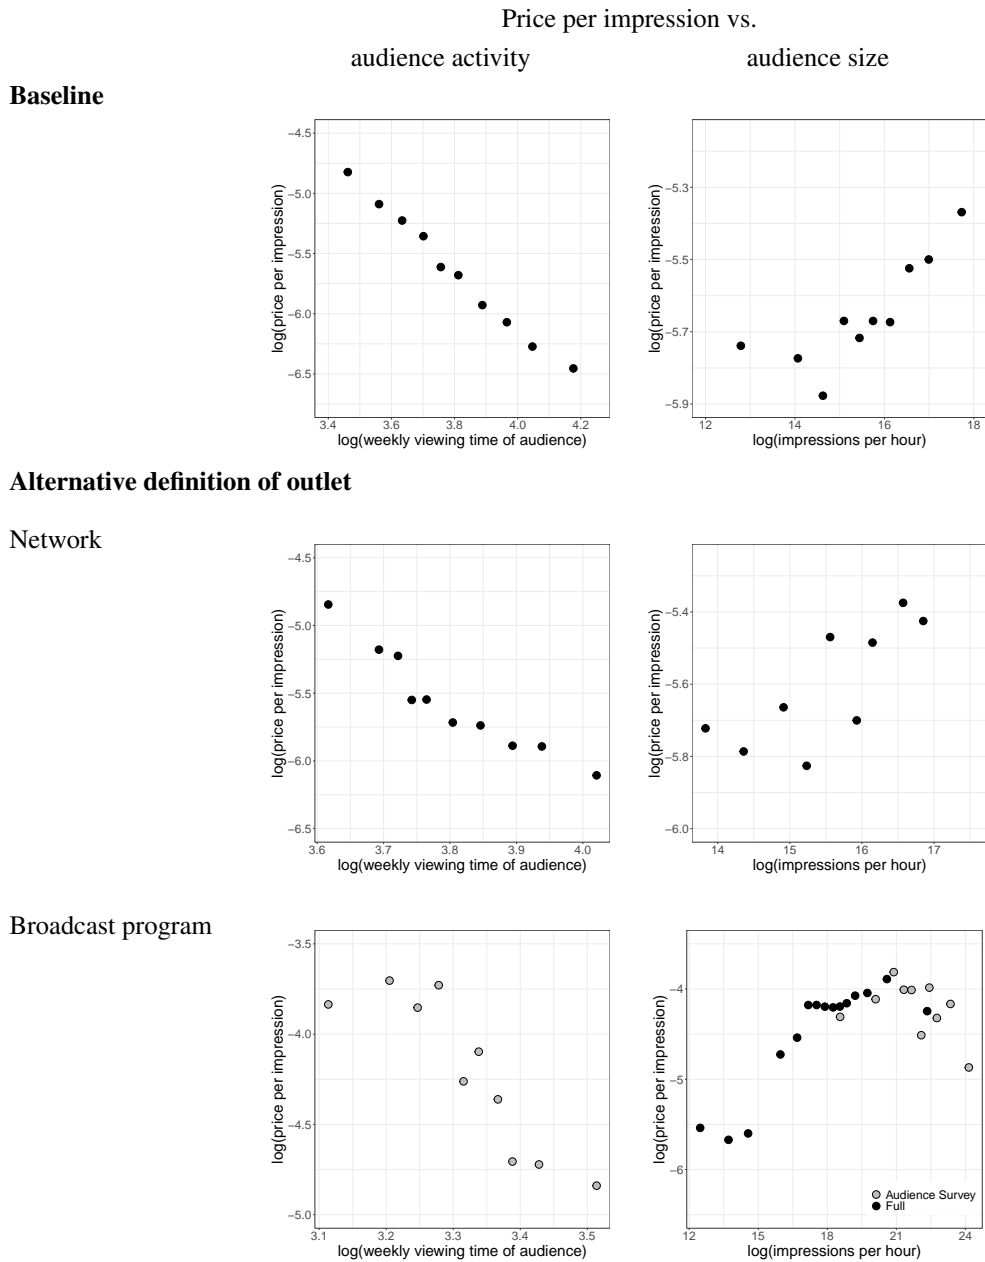


2016



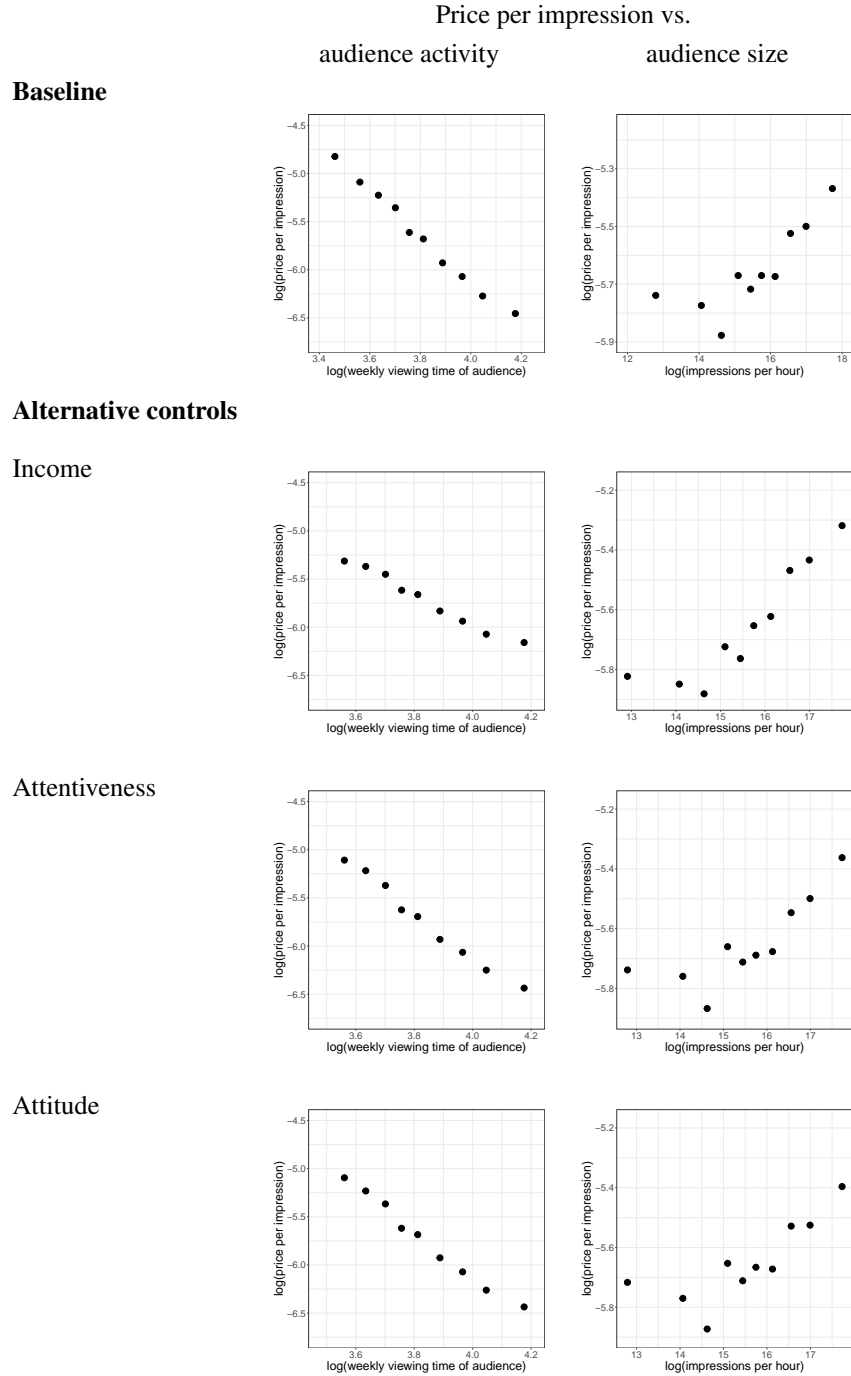
Notes: Within a given row, both plots are based on the same regression specification. The row labeled “Baseline” corresponds to the main specification in the paper, with the plot “Price per impression vs. audience activity” corresponding to Panel B of Figure 2 and the plot “Price per impression vs. audience size” corresponding to Panel B of Figure 4. The rows under the header “Alternative sample years” present results for alternative sample years.

Appendix Figure 2: Sensitivity to Alternative Outlet Definitions



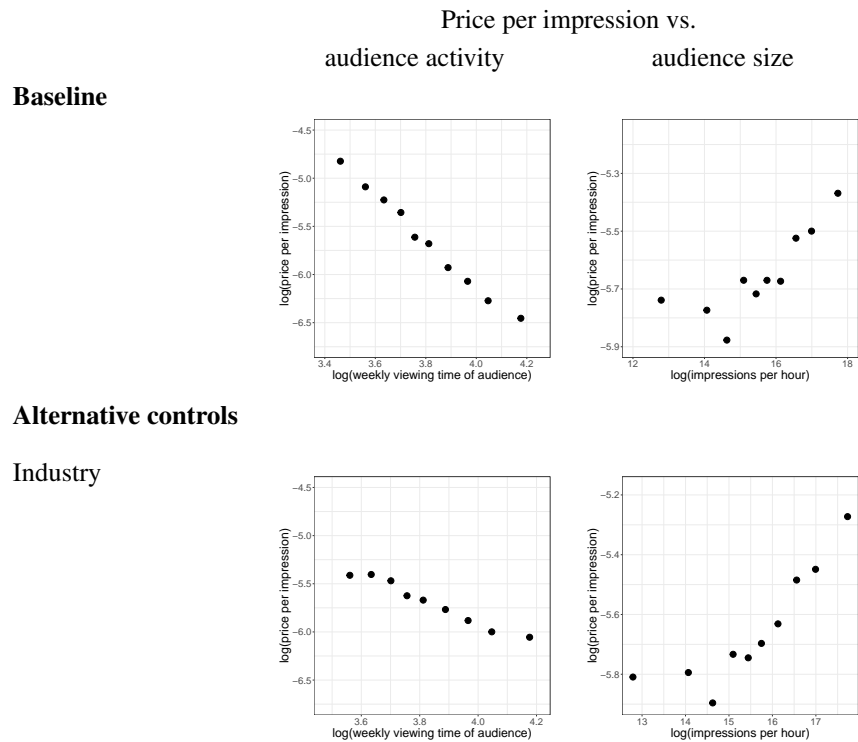
Notes: Within a given row, both plots are based on the same regression specification. The row labeled “Baseline” corresponds to the main specification in the paper, with the plot “Price per impression vs. audience activity” corresponding to Panel B of Figure 2 and the plot “Price per impression vs. audience size” corresponding to Panel B of Figure 4. The rows under the header “Alternative definition of outlet” consider different outlet definitions. In the row labeled “Network” an outlet j is a network. In the “Network” row, the “Price per impression vs. audience activity” specification includes controls for the share of total impressions that are to adults and for indicators of deciles of audience size, and the “Price per impression vs. audience size” specification includes controls for the share of total impressions that are to adults and for indicators of deciles of audience activity. In the row labeled “Broadcast program” an outlet j is a broadcast program, with bins corresponding to 15 quantiles of the full sample of broadcast programs (3055 programs) colored black and bins corresponding to deciles of the subsample of broadcast programs included in the audience survey (173 programs) colored gray. In the “Broadcast program” row, the “Price per impression vs. audience activity” specification includes controls for the share of total impressions that are to adults and for indicators of deciles of audience size, and the “Price per impression vs. audience size” specification includes a control for the share of total impressions that are to adults.

Appendix Figure 3: Sensitivity to Alternative Controls




Notes: Within a given row, both plots are based on the same regression specification. The row labeled “Baseline” corresponds to the main specification in the paper, with the plot “Price per impression vs. audience activity” corresponding to Panel B of Figure 2 and the plot “Price per impression vs. audience size” corresponding to Panel B of Figure 4. The rows under the header “Alternative controls” consider different sets of control variables. The row labeled “Income” adds controls for indicators for deciles of the average household income of adult impressions. The row labeled “Attentiveness” adds controls for indicators for deciles of the time-weighted average attentiveness of the outlet’s viewers, where a viewer’s attentiveness is the viewer’s average self-reported attentiveness across broadcast and cable programs, coded as some (0.5), most (0.75), or full (1), and measured for each program relative to the mean among all respondents who rate the program. The row labeled “Attitude” adds controls for indicators for deciles of the time-weighted average of viewers’ attitudes toward television advertising, where a viewer’s attitude toward advertising is measured as the first principal component of the viewer’s responses (on a five-point scale) to a series of eight questions about TV advertising.

Appendix Figure 3: Sensitivity to Alternative Controls (continued)



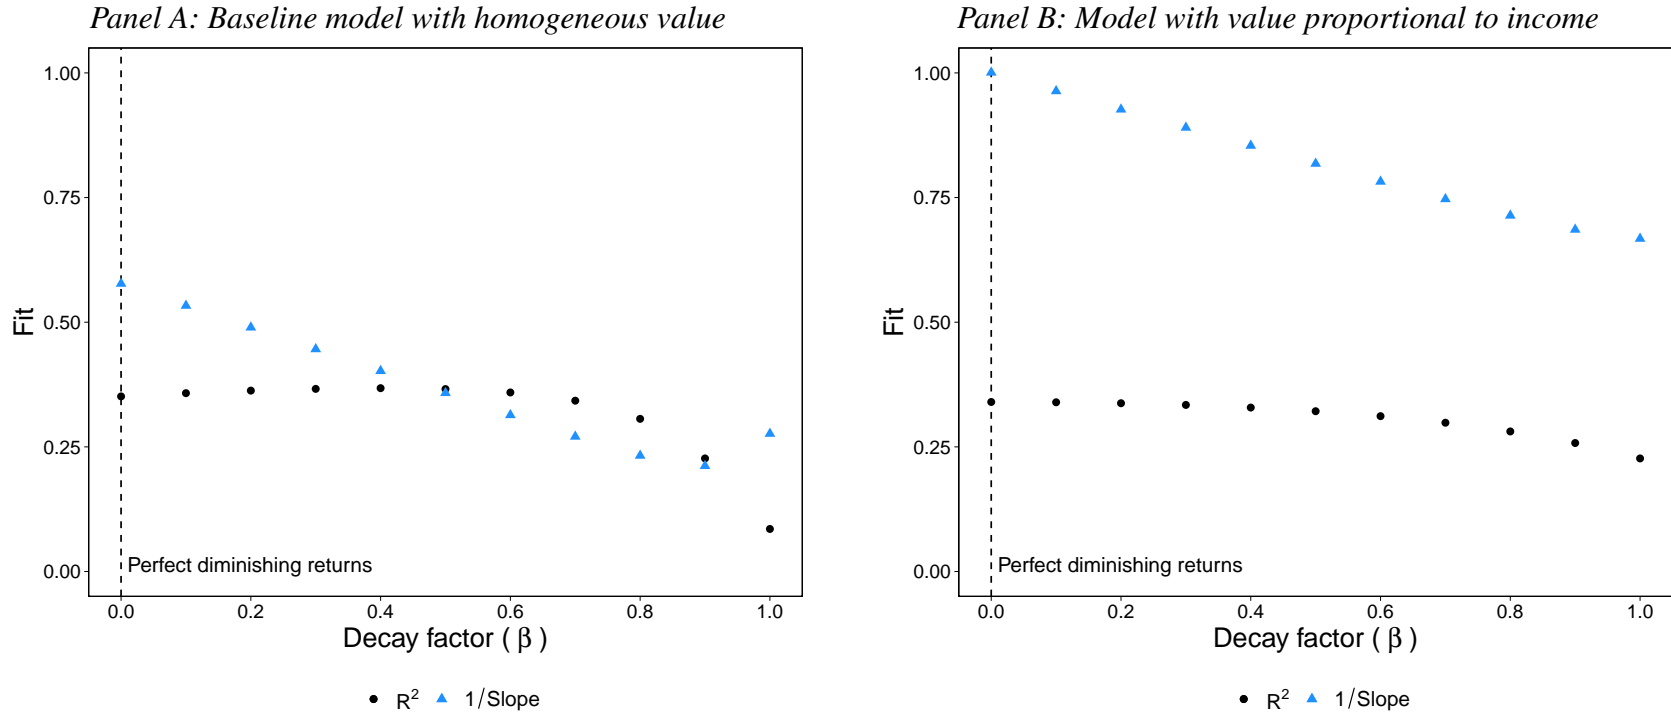
Notes: Within a given row, both plots are based on the same regression specification. The row labeled “Baseline” corresponds to the main specification in the paper, with the plot “Price per impression vs. audience activity” corresponding to Panel B of Figure 2 and the plot “Price per impression vs. audience size” corresponding to Panel B of Figure 4. The rows under the header “Alternative controls” consider different sets of control variables. The row labeled “Industry” adds controls for the share of the outlet’s adult impressions that are to ads whose advertisers are in each of 11 industry categories: automotive; business and consumer services; business supplies; drugs and remedies; entertainment; food and drink; home and garden; insurance and real estate; retail; travel; and other.

Appendix Figure 4: Average Television Viewing Hours Per Day by Age and Gender



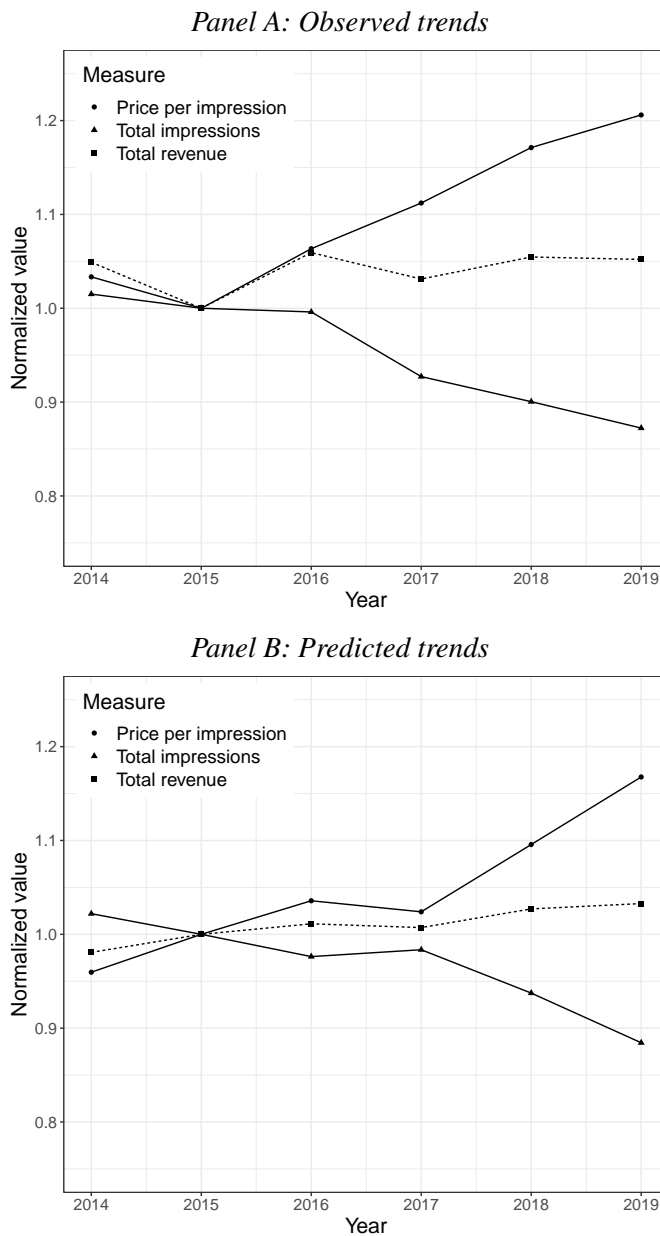
Notes: The figure shows the average daily viewing hours spent on television across age groups by gender.

Appendix Figure 5: **Fit of Quantitative Model Under Imperfect Diminishing Returns**




Notes: In each plot we depict the fit of a model-based prediction of advertising prices (y-axis) as a function of the extent of diminishing returns from advertising assumed in the model (x-axis). To produce the plot, we parameterize the viewer-level model with a single format such that $u(M) = \sum_{m=1}^M \beta_m$ where $\beta_m = \beta^{m-1}$ for $m \in \{1, \dots, \bar{M}\}$, $\beta_m = 0$ for $m > \bar{M}$, and $\beta \in [0, 1]$. Here, β describes the extent of diminishing returns, with $\beta = 0$ denoting the special case of the reach-only model and $\beta = 1$ denoting the special case of no diminishing returns up to the \bar{M}^{th} impression. We assume that $\bar{M} = 10$ and calculate η_{ij} as described in Section 5. We calculate the log(price per viewer) implied by the model for each value $\beta \in \{0, 0.1, \dots, 1\}$ depicted on the x-axis. We then regress the log(price per impression) of a 30-second spot observed in the data, as described in Section 3.1, against the log(price per viewer) predicted by the model and depict on the y-axis the R^2 of the regression and the multiplicative inverse of the estimated slope. Panel A uses log(price per viewer) predicted from the baseline model in which advertisers' value of a first impression is homogeneous across viewers. Panel B uses log(price per viewer) predicted from the model in which advertisers' value of a first impression is proportional to a viewer's income. The unit of analysis for the regression is an owner Z , and all variables in the regression are residualized with respect to the share of the owner's impressions that are to adults.

Appendix Figure 6: **Observed and Predicted Television Advertising Revenues, Alternate Estimates of Impressions**



Notes: Each plot depicts trends in the television advertising market over the sample period. We plot trends in total revenue, total impressions, and price per impression (total revenues divided by total impressions), all normalized relative to their 2015 value. In Panel A, all series are as observed in the data, as described in Section 3.1, and revenue is deflated to 2015 dollars using the US Consumer Price Index (Organization for Economic Co-operation and Development 2022). In Panel B, the trends in revenue and impressions are predicted by the baseline model in which advertisers' value of a first impression is homogeneous across viewers, as described in Section 5.


Appendix Figure 7: **Predicted Television Advertising Revenues, Strong Cross-Format Diminishing Returns**




Notes: The plot depicts trends in total revenue, total impressions, and price per impression (total revenues divided by total impressions), all normalized relative to their 2015 value. Total revenue is predicted by the cross-format reach-only model defined in Section 6 where $\phi = 1$ such that diminishing returns operate just as strongly between as within formats. Total impressions are as observed in the data, as described in Section 3.1. Price per impression is calculated as the ratio of the predicted revenue to total impressions.

Appendix Figure 8: Measures of Online Activity by Age and Gender

Panel A: Share of social media sites visited in the past 30 days




Panel B: Average internet hours per day




Notes: Panel A shows the average share of five social media sites (Facebook, Instagram, Reddit, Twitter and YouTube) visited in the past 30 days across age groups by gender. Panel B shows average daily hours spent on the internet across age groups by gender.

Appendix Figure 9: **Demographic Premia (Per Click) and Viewing Time on Facebook**

Panel A: Data from our experiment



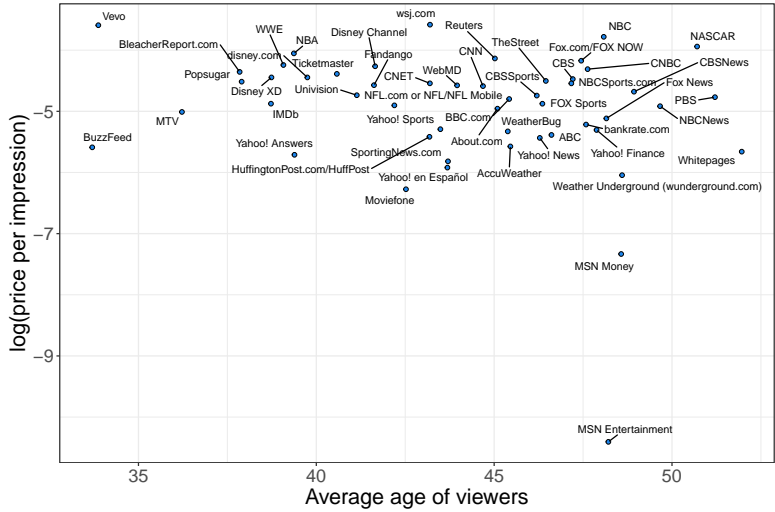
Panel B: Data from Allcott et al. (2020b)



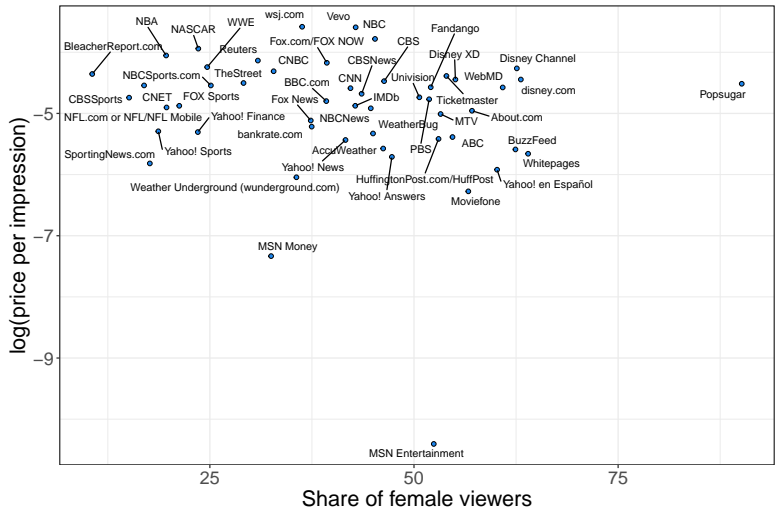
Notes: The plot shows the log(price per click) for advertisement sets targeted to a given gender and age group. In Panel A, the data are taken from our own experiment, and the groups are $\{\text{Men, Women}\} \times \{18-24, 25-34, 35-44, 45-54, 55-64, 65+\}$. In Panel B, the the data are taken from Allcott et al. (2020b), and the groups are $\{\text{Men, Women}\} \times \{18-24, 25-44, 45-64, 65+\}$. In both panels, the y-axis value is the log(price per click) for advertisement sets targeting the given group, and the x-axis value is the midpoint of the age range for the given group, treating 70 as the midpoint for ages 65+.

Appendix Figure 10: Advertising Prices and Demographics of Digital Platforms

Panel A: Average age



Panel B: Share female

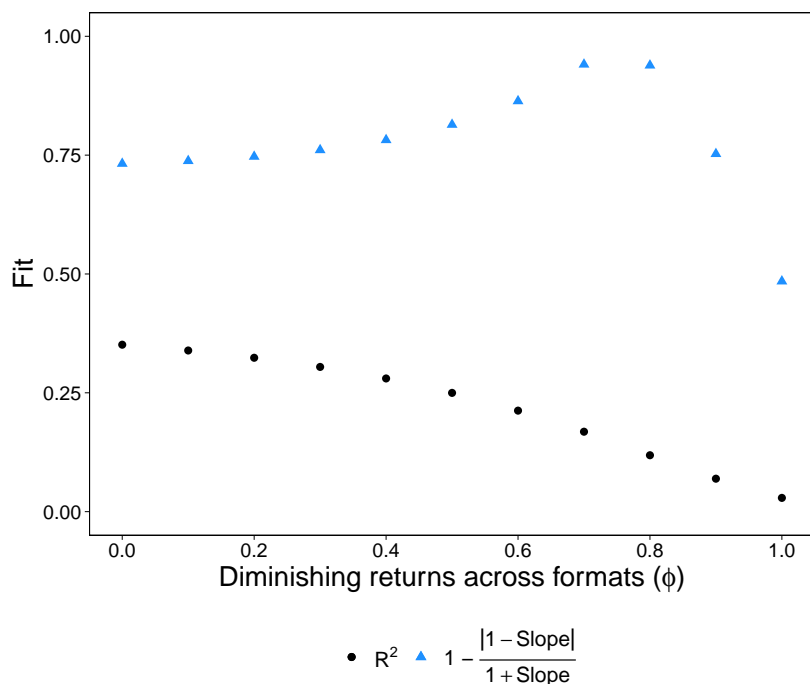


Notes: Each plot is a scatterplot of the $\log(\text{price per impression})$ of display advertising on a platform against the demographic characteristics of the platform's viewers. We construct the price per impression by computing the ratio of total revenue to total impressions across all display ads on the platform reported in AdIntel 2017 (The Nielsen Company 2022). The sample of platforms is the set of platforms that AdIntel 2017 (The Nielsen Company 2022) classifies as Entertainment, Finance, Information/Reference, News/Commentary, Spanish, Sports, Technology, or Weather, excluding some platforms such as those that focus primarily on direct sales of products or services. The x-axis shows the average age (Panel A) or share female (Panel B) of those who report visiting the platform in the previous 30 days in GfK MRI's 2017 Survey of the American Consumer (GfK Mediamark Research and Intelligence 2019).


Appendix Figure 11: **Fit of Quantitative Model of Television Prices with Social Media Competition**

68

Panel A: Baseline model with homogeneous value



Panel B: Model with value proportional to income



Notes: In each plot we depict the fit of a model-based prediction of advertising prices (y-axis) as a function of the strength of cross-format diminishing returns (x-axis). To produce the plot, for each television owner Z and for each value $\phi \in \{0, 0.1, \dots, 1\}$ of the parameter governing the strength of cross-format diminishing returns, we calculate the predicted $\log(\text{price per viewer})$ implied by the cross-format reach-only model defined in Section 6. For each value of ϕ , we then regress the $\log(\text{price per impression})$ of a 30-second spot observed in the data, as described in Section 3.1, against the $\log(\text{price per viewer})$ predicted by the model and depict on the y-axis the R^2 of the regression and one minus the percent deviation of the slope from one. Panel A uses $\log(\text{price per viewer})$ predicted from the baseline model in which advertisers' value of a first impression is homogeneous across viewers. Panel B uses $\log(\text{price per viewer})$ predicted from the model in which advertisers' value of a first impression is proportional to a viewer's income. The unit of analysis for the regression is an owner Z , and all variables in the regression are residualized with respect to the share of the owner's impressions that are to adults.

Appendix Table 1: **Restrictiveness and Completeness of Quantitative Economic Model**

<i>Panel A: Restrictiveness</i>		
<i>Economic model:</i>	<i>Homogeneous values</i>	<i>Value proportional to income</i>
	(1)	(2)
MSE of log(simulated price per impression) w.r.t.		
....constant model	0.6650	0.6650
....economic model	0.7190	0.8200
Restrictiveness of economic model	1.0800	1.2330
Number of owners	33	33
Number of viewers	21506	21506

Notes: The table evaluates the restrictiveness (Fudenberg, Gao, and Liang 2023) and completeness (Fudenberg et al. 2022) of the quantitative model of log(price per viewer) described in Section 5. Column (1) uses log(price per viewer) predicted from the baseline model in which advertisers' value of a first impression is homogeneous across viewers. Column (2) uses log(price per viewer) predicted from the model in which advertisers' value of a first impression is proportional to a viewer's income. The constant model predicts log(price per impression) with its mean across owners. In calculating the mean squared error, we represent both the observed and predicted values in terms of deviation from the mean across owners.

To evaluate restrictiveness (Panel A), in each of 10,000 replicates, we randomly draw values of each owner's log(price per impression), independently uniform over the support of the observed log(price per impression). Restrictiveness is the ratio of the mean, across replicates, of the mean squared error of the economic model of log(price per viewer) described in Section 5, and the mean squared error of the constant model, with respect to the random draws.

Appendix Table 1: (continued): **Restrictiveness and Completeness of Quantitative Economic Model***Panel B: Completeness*

<i>Economic model:</i>	<i>Homogeneous values</i>	<i>Value proportional to income</i>
	(1)	(2)
MSE of log(observed price per impression) w.r.t.		
....constant model	0.4570	0.4570
....owner-level regression model	0.4360	0.4240
....viewer-level lasso model	0.4910	0.4910
....economic model	0.3290	0.3090
Completeness of economic model w.r.t.		
....owner-level regression model	6.0480	4.4490
....viewer-level lasso model	—	—
Number of owners	33	33
Number of viewers	21506	21506

Notes (continued): The table evaluates the restrictiveness (Fudenberg, Gao, and Liang 2023) and completeness (Fudenberg et al. 2022) of the quantitative model of log(price per viewer) described in Section 5.

To evaluate completeness (Panel B), we consider two comparison models. In the first comparison model, we estimate a linear regression of the observed log(price per impression) of each owner Z on the log value, $\ln(\lambda_Z)$, of total impressions, the log value, $\ln\left(\frac{\sum_{i \in \mathcal{I}} a_i \eta_{iZ}}{\sum_{i \in \mathcal{I}} a_i}\right)$, of total weighted impressions, and the log value, $\ln\left(\frac{\sum_{i \in \mathcal{I}} a_i \prod_{Z' \neq Z} (1 - \eta_{iZ'})}{\sum_{i \in \mathcal{I}} a_i}\right)$, of the weighted fraction of viewers not seeing an ad on other owners' networks.

In the second comparison model, for each viewer i , we estimate a linear regression of the observed log(price per impression) of each owner Z on the log value, $\ln(\eta_{iZ})$, of the viewer's probability of seeing an ad on that owner's networks, as well as the log values, $\ln(1 - \eta_{i(Z)})$, of the viewer's probability of not seeing an ad on each of the other owners' networks, indexed in descending order. When a value inside a logarithm is zero we replace it with its minimum across all owners Z for the given viewer i , and we include indicators for imputed values in the regression. We estimate the model via lasso using 10-fold cross-validation to choose the penalty, and for each viewer i we hold out one randomly chosen target owner whose data is excluded from the estimation sample and for which we predict the log(price per impression) from the final lasso fit. In the 1.9 percent of cases where there is insufficient variation in the regressors to estimate the model, we use the mean of the dependent variable as the lasso fit. For each owner Z , we take the weighted mean predicted log(price per impression) across all viewers i for which the given owner is the target, and treat this mean as the lasso-predicted log(price per impression) for the given owner. When comparing to the model with homogeneous values we use uniform weights; when comparing to the model with value proportional to household income we use household income as the weight.

Completeness is the ratio of the improvement in mean squared error between the economic model of log(price per viewer) described in Section 5 and the comparison model, each evaluated relative to the constant model. We treat completeness as undefined when a given model has higher mean squared error than the constant model.