

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*

July 2022

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 and extend this theoretical result and test it using data from television and social media advertising. 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. Also consistent with the theory, we show that social media advertising markets feature a premium for older audiences. 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, the effect of competition on incentives to produce content, and the effect of Netflix ad carriage on prices for linear television advertising.

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 International Industrial Organization Conference, Helsinki GSE, the Leuven Summer Event, and especially discussant Jay Lee 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 digital and traditional media receive substantial revenues from selling consumer attention to advertisers (e.g., Statista 2021). Prices per unit of attention 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' attention commands 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 attention 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 Other documented price premia include 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. The predictions of the model are also in line with less-well-known facts that we document

¹Gabler (2014 pp. 3-4) writes that “We live in a culture of the young, for the young and by the young, and anyone over 49—the demographic breakpoint of old age for most television advertisers—is tossed onto the trash heap of history, all eighty million of them. In effect, these people, just under one-third of the American population, 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).

for social media advertising. 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.

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 outlet's advertising slots is determined by the overlap of its audience with those of other outlets. As a result, the price per viewer that an outlet can charge in equilibrium is decreasing in the overall activity level of its audience, and increasing in the overall size of its audience. In the special case of perfect diminishing returns, where advertisers value only the first impression to a given viewer, each outlet's price per impression is determined solely by the fraction of its audience that is exclusive to that outlet.

We study the model's 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 then turn to social media advertising, using data on prices of Facebook advertisements col-

lected as part of a series of experiments including our own. In contrast to television, the young are the heaviest users of social media. Treating television and social media as distinct ad markets, our model predicts that older users should be the more “coveted” group on social media. Alternative explanations for the youth premium, such as the young having more malleable preferences (Surowiecki 2002), would not all share this prediction. We show that the age-price relationship is indeed reversed on Facebook, with the oldest users commanding the highest prices.

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. Based on this specification we use the audience survey data to calculate the incremental value of advertising on each outlet, 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 35.1 percent of the variation in price per impression across owners, with a slope close to unity, and exhibit the same qualitative patterns as observed prices with respect to age, gender, and outlet size. 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 apply the quantitative model to three 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 competition on the incentive of television network owners to invest in content to attract different kinds of audience members. We find that impressions from those in the oldest age group are more than 82.8 percent less valuable than those in the youngest, and that this gap would attenuate significantly if television network ownership were more concentrated. Third, we study the effect of Netflix carrying advertising on the price of attention 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.

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 the case of social media, while some industry sources report a premium for older audiences on social media (e.g., Ampush 2014), 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

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. Prat and Valletti (forthcoming, Section 5) simulate effects of platform mergers under various assumptions about overlap in their audience, though using a microfoundation different from ours. 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.

³Papazian (2009) 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” (p. 134).

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

⁵Lambrech 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

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 (forthcoming), 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.

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 and social media. Section 5 presents our quantitative implementation of the model, its fit to the data, and its applications. Section 6 concludes.

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, our analysis of outlets' pricing power emphasizes the importance of individual viewers' allocation of attention.

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. Section 2.2 and Appendix A.3 discuss settings with $N > K$. We let $\mathcal{P}(\cdot)$ denote the power set operator.

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 move sequentially in random order and decide which, if any, bundles to buy. When all advertisers have moved, ads are shown and the game ends.

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 that buys slots in a set of bundles $\mathcal{S} \subseteq \mathcal{P}(\mathcal{J})$ is given by

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

where $V(\cdot)$ is a non-negative value function that is monotonic in set-inclusion order.⁸ We capture the idea that there are diminishing returns to advertising by assuming that $V(\cdot)$ is submodular: an advertiser derives less incremental value from an outlet when adding it to a larger bundle.⁹ Section 2.2 discusses settings with partially increasing returns and with heterogeneity in advertisers' value functions.

The following examples exhibit monotonic and submodular value functions.

Example 1. There is a set of viewers. Each viewer sees any ad slot on outlet j with some probability. Each advertiser gets value $a_i \sum_{m=0}^M \beta_m$ from viewer i that views its ad $M \in \mathbb{N}$ times, where $a_i > 0$, $\beta_0 = 0$, $\beta_1 > 0$, and $\beta_m \geq 0$ is non-increasing in m for $m \geq 1$. The following settings are nested in this one:

⁸That is,

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

⁹Formally,

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

(a) (Anderson, Foros, and Kind 2018.) There is a set of viewers, each of which views all ad slots on a subset of the outlets, and $a_i = a$ for all i .

(b) (Awareness advertising with forgetting.) Each advertiser gets value $a_i > 0$ from viewer i that remembers seeing its ad. Each viewer remembers each ad they have seen with some probability, independently across ads.

Example 2. There is a set of viewers, each of which views all ads on a subset of the outlets. There is a partition \mathcal{C} that groups the outlets \mathcal{J} into categories. For each viewer that views its ad on an outlet in category $C \in \mathcal{C}$, each advertiser gets value $a_C > 0$.

Example 3. 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. 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. 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.

Remark 1. In Examples 1 and 2, viewers' choices of which outlets to view are not affected by the outcome of the game. We may interpret this either 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. In Example 3, viewers' choices of which outlets to view are affected by the outcome of the game. We may interpret this as a scenario in which viewers care about advertising and make viewing decisions knowing the outcome of the advertising game.

Remark 2. 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. Our model and analysis are compatible with the presence of such fees provided they 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. Section 5.2 augments our model to include a content investment stage.

Our main result is that each owner is able to extract the incremental value of the outlets it controls. To state this result, for a 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 take an equilibrium to be a subgame perfect equilibrium in pure strategies.

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. Section 2.2 discusses settings with alternative market institutions, including unbundled pricing, bargaining over or auctioning of ad slots, and viewer-level ad pricing and targeting.

2.1 Comparative Statics

Consider a special case of Example 1 in which every owner owns a single outlet $j \in \mathcal{J}$, and diminishing returns are strict in the sense that $\beta_2 < \beta_1$. Suppose that there is a unit mass of viewers subdivided into a set G of mutually exclusive demographic groups, with group g having mass μ_g , so that $\sum_{g \in G} \mu_g = 1$. Members of group $g \in G$ see ads on outlet j with probability $\eta_{gj} \in (0, 1)$, independently across outlets. We assume that all viewers have the same intrinsic value to advertisers, i.e. that $a_i = a$ for all i . Let

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

denote, respectively, the total mass of outlet j 's audience, and the share of this audience that comes from group g . Then p_j^*/λ_j is the price per viewer charged by the owner for an ad slot.

Applying Theorem 1 with the structure of the value function $V(\cdot)$ in this case, we show two comparative statics results. The first result is that, all else equal, an outlet commands a larger price premium for its viewers if its viewers come from less active groups.

Proposition 1. *Suppose that group $g \in G$ is less active than group $h \in G$ in the sense that $\eta_{gj} \leq \eta_{hj}$ for all $j \in \mathcal{J}$. Suppose that outlet $j \in \mathcal{J}$ draws a larger share of its audience from group g and a smaller share of its audience from 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 , $p_j^*/\lambda_j \geq p_k^*/\lambda_k$.*

The inequality in the conclusion of Proposition 1 is strict if $\eta_{gj'} < \eta_{hj'}$ for some $j' \neq j, k$ and $\sigma_{gj} > \sigma_{gk}$.

Intuitively, Proposition 1 holds because, given diminishing returns, the incremental value of showing an ad to a viewer who watches more outlets is lower than the incremental value of showing an ad to a viewer who watches fewer outlets. Since more active viewers tend to watch more outlets, this force puts competitive pressure on the prices that outlets can charge to show ads to these viewers.

Remark 3. 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 each outlet's price per viewer would be invariant to the group composition of its audience.

Remark 4. Appendix A.3 shows that a statement analogous to Proposition 1 holds for multi-outlet owners when diminishing returns are perfect, i.e., $\beta_m = 0$ for $m \geq 2$. Competition is essential for this result. If a single owner were to own all outlets, then 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 outlet commands a larger price premium for its viewers if the outlet attracts a larger share of the total audience.

Proposition 2. *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 , $p_j^*/\lambda_j \geq p_k^*/\lambda_k$.*

The inequality in the conclusion of Proposition 2 is strict if $\delta > 1$. Intuitively, Proposition 2 holds because viewers of the larger outlet tend to watch fewer other outlets, leading to less competitive pressure on the price that the larger outlet can charge to show ads to its viewers. Appendix A.3 shows that a statement analogous to Proposition 2 holds for multi-outlet owners when diminishing returns are perfect.

2.2 Extensions

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.

Proposition 3. *There exists a subgame perfect equilibrium, possibly in mixed strategies, and in any subgame perfect equilibrium each owner Z earns an expected revenue per slot between $(v_Z - \Delta)/|Z|$ and $\sum_{j \in Z} V(\{j\})/|Z|$.*

Allowing for mixed strategies helps to guarantee existence of an equilibrium in this setting.

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 4. *Suppose $V(\cdot)$ is strictly monotone and that $V(\mathcal{J}' \cup Z) - V(\mathcal{J}') \geq V(\mathcal{J}) - V(\mathcal{J} \setminus Z)$ for all $Z \in \mathcal{Z}, \mathcal{J}' \subseteq \mathcal{J} \setminus Z$. 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 $p_Z^* = v_Z$.*

The decreasing differences condition on $V(\cdot)$ in the hypothesis of Proposition 4 is strictly weaker than submodularity. Suppose, for example, that owners are singletons, each of a set of viewers i views at least L outlets, and an advertiser's value for viewer i seeing its ad M times is $a_i \sum_{m=0}^M \beta_m$ 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 $\beta_L \leq \min_{1 \leq m < L} \beta_m$. This setting satisfies the hypotheses of Proposition 4 but not necessarily those of Theorem 1, and allows increasing returns to advertising for viewers receiving few impressions (as in, e.g., Dubé, Hitsch, and Manchanda 2005). Appendix A.3 establishes that analogues of Propositions 1 and 2 hold in this setting, and that the setting continues to satisfy the hypotheses of Proposition 4 if a small number of viewers view fewer than L outlets.

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 $v_Z = \min_{n \in \mathcal{N}} v_{n,Z}$ 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, $\varphi(Z) = v_Z$.

Proposition 5. *Suppose that heterogeneity in the value functions $V_n(\cdot)$ is small in the sense that $\bar{v}_Z - v_Z \leq \frac{1}{N} \varphi(Z)$ 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 $p_Z^* = v_Z$ for all $Z \in \mathcal{Z}$.*

It is immediate that, in the setting of Section 2.1, any efficient equilibrium obeys the comparative statics in Propositions 1 and 2.

The hypothesis of Proposition 5 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.¹⁰ The restriction on incremental values becomes more demanding as the number of advertisers, N , grows large.

Unbundled pricing. It is useful to be able to characterize the price of an owner’s individual outlets. To do this we can imagine that some owners are not allowed to bundle slots on all of their outlets together. Formally, each owner Z is endowed with a partition \mathcal{F}_Z of Z such that they are only allowed to bundle outlets in the same cell of the partition. Denote by $v_B^S = V(S) - V(S \setminus B)$ the incremental value of bundle B in $S \subseteq \mathcal{J}$. 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 6. *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.*

In general, existence of an equilibrium is not guaranteed. Appendix A.3 establishes the existence of an equilibrium in a special case in which the comparative statics of Proposition 2 hold.

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).

Proposition 7. *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.*

The proof of Proposition 7 shows that, in the more general case where owners have different bargaining weights, there exists an equilibrium in which the payment to each owner Z is given by the product of the owner’s bargaining weight and v_Z . If, further, $V(\cdot)$ is strictly monotone, then this outcome is unique.

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.

¹⁰When 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 $c_n \geq 0$.

Proposition 8. *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 .*

Appendix A.3 further characterizes equilibrium in a model where owners conduct auctions, advertising slots are scarce, and advertisers' valuations are heterogeneous.

Viewer-level ad pricing and targeting. Suppose that each owner can post viewer-specific prices for each bundle. Suppose that $V(\cdot) = \sum_i V_i(\cdot)$, where $V_i(\cdot)$ is the value function if viewer i were the only viewer, and $v_{i,B} = V_i(\mathcal{J}) - V_i(\mathcal{J} \setminus B)$ is the viewer-specific incremental value of bundle B .

Proposition 9. *Suppose that $V_i(\cdot)$ is monotone and submodular and that $v_{i,j} > 0$ for all i and j . Then there exists an equilibrium. In any equilibrium, for any viewer i , all advertisers buy slots on all outlets, and the payment by each advertiser to each owner Z is given by $p_{i,Z}^* = v_{i,Z}$.*

In the setting of Proposition 9, the total payment $\sum_i v_{i,Z} = v_Z$ by each advertiser to each owner Z is equivalent to that under Theorem 1. Note that Example 1 satisfies the hypotheses of Proposition 9 if $\beta_m > 0$ for all $m \geq 1$, and each viewer i views any ad slot on each outlet j with strictly positive probability.

Competition between advertisers. Appendix A.3 provides additional conditions under which a modified incremental value pricing equilibrium exists when an advertiser's value for advertising depends not only on the slots they purchase but also on those purchased by other advertisers.

3 Data

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 on information obtained at the month-network-daypart level for cable television and at the month-program level for broadcast television (The Nielsen Company 2017). For consistency we therefore define our notion of an outlet j to be a network-daypart.¹¹ Appendix Figure 1 reports results when using 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 take as analogous to the concept λ_j defined in Section 2.1. 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 take as analogous to the concept p_j^* defined in Section 2.1. 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 take as analogous to the concept p_j^*/λ_j defined in Section 2.1.

For each advertisement we also have information on the number of impressions by age (in bins) and gender for the associated telecast.¹² 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 75.¹³

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.¹⁴ From this information we compute the average household income of each outlet's adult

¹¹In cases where a telecast spans multiple dayparts, we assign it to the daypart that contains the largest share of broadcast time.

¹²The 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.

¹³Using 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.

¹⁴The 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.

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 and audience demographics of each owner $Z \in \mathcal{Z}$.

We conduct our main analysis of the television market using data from 2015, but for sensitivity analysis and extensions we use data from 2014 through 2019, with concepts defined and calculated in an analogous manner to those we have described for 2015. Appendix Figure 2 reports sensitivity analysis replacing data from 2015 with data from 2014 or 2016.

3.2 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.

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.

3.3 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 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 viewership of each of 227 broadcast television programs,¹⁵ 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.¹⁶

We use the data on viewership by daypart, broadcast program, and cable network to construct a measure of the time that each respondent viewed each outlet (network-daypart) j . To do this, we first allocate the viewing time of broadcast programs to their respective network-dayparts.¹⁷ 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.¹⁸

We thus arrive at a measure of the time each respondent 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 .¹⁹ We treat average log total weekly viewing time as a measure of the

¹⁵The 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.

¹⁶The 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.

¹⁷Specifically, 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).

¹⁸If 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.

¹⁹We exclude from this calculation any respondent with zero total weekly viewing time.

overall activity level of outlet j 's audience.

We also obtain information for each respondent on the total time spent using the internet in an average week (calculated based on reported time spent on three recent days), and on the share of five social media sites (Facebook, Instagram, Reddit, Twitter, and YouTube) visited in the preceding 30 days. In addition, we obtain information on each respondent's gender, age (in bins),²⁰ household income (in bins),²¹ reported attentiveness to different broadcast and cable programs, and attitudes towards television advertising.

As with the television data described in Section 3.1, we conduct our main analysis using data for 2015, but for sensitivity analysis and extensions we use data from 2014 through 2019, with concepts defined and calculated in an analogous manner to those for 2015. For 2019, we additionally obtain data on time spent watching Netflix, which we use in counterfactual analysis in Section 5.

4 Evidence on the Determinants of Advertising Prices

4.1 Television

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

Both panels of Figure 1 show a clear negative relationship between $\log(\text{price per impression})$ and average $\log(\text{weekly viewing time})$. The magnitude of the relationship is large: in Panel B, for example, moving from the bottom to the top decile of average $\log(\text{weekly viewing time})$ corresponds to a decline in $\log(\text{price per impression})$ 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

²⁰The 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.

²¹The 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.

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 2 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 price per impression. Figure 3 shows that this prediction is borne out in the data. Panel A shows a binned scatterplot of log(price per impression) against log(impressions per hour) with baseline controls. Panel B additionally controls for average log(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(impressions per hour) corresponds to an increase in log(price per impression) of roughly 37 log points.

Columns (1) and (2) of Table 1 summarize the patterns in Figures 1 through 3. 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,²² and the industry mix of the outlet's advertisers.

4.2 Social Media

Whereas older people are the heaviest television viewers, younger people are the heaviest users of the internet and social media. Appendix Figure 5 shows that time online (Panel A) and visits to social media sites (Panel B) are decreasing in age. This sets up an interesting test of the predictions of Proposition 1. Treating television and social media advertising as separate markets, Proposition 1 predicts that social media advertising prices should be increasing in audience age, in contrast to the decreasing pattern we see for television advertising prices.

Figure 4 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

²²McGranaghan, Liaukonyte, and Wilbur (forthcoming) 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.

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 consistent with the mechanism underlying Proposition 1, and more difficult to square with some alternative explanations for the differential value of older vs. younger viewers to advertisers, such as intrinsic differences in the malleability of their preferences (Surowiecki 2002).²³ Appendix Figure 6 repeats the analysis in Figure 4 using data on price per click rather than price per impression.

Differences in the age premium between television and social media advertising make sense only if advertisers view ads in the two media as imperfect substitutes. Otherwise, we would expect television and social media to function as a single ad market, with identical relative prices for different groups of viewers. Several facts are consistent with imperfect substitutability. First, industry sources often 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). Second, consistent with this, the firms that advertise the most online are often different from those that advertise the most on television. In 2015 the top 50 television advertisers accounted for 44 percent of television spending but only 17 percent of spending online; the top 50 online advertisers accounted for 36 percent of online spending but only 27 percent of spending on television (see Appendix Table 1). Finally, as we discuss in more detail in Section 5, the dramatic rise in online advertising since the early 2000s has not corresponded with any significant decrease in television advertising spending. Television advertising spending continued to increase through the mid-2010s and has only begun to decline in recent years, in contrast to print advertising spending which fell consistently as online advertising grew (Kitterman 2020). The evidence in Figure 4 supports the joint hypothesis of imperfect substitutability and price determination consistent with our model.

Unlike age, gender is not strongly associated with activity levels online. Appendix Figure 5 shows that males tend to report spending more time on the internet (Panel A), but visiting fewer social media sites (Panel B). Correspondingly, Figure 4 does not show a consistent price premium for advertising to male or female audiences on social media.

Appendix Figure 7 shows the relationship between estimated $\log(\text{price per impression})$ of dis-

²³Smith, 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.

play 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 come from a 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.

5 Quantification and Applications of the Model

Consider the special case of Example 1 with perfect diminishing returns ($\beta_m = 0$ for $m \geq 2$). For viewer $i \in \mathcal{I}$ who spends time T_{ij} viewing outlet j , suppose the probability of seeing an ad placed in one of that outlet’s slots is given by $\eta_{ij} = T_{ij}/T_j$ independently across outlets j , where T_j is outlet j ’s total broadcast time. Given data on viewing times, an ownership partition \mathcal{Z} , and a vector of advertisers’ value of reaching each viewer $(a_1, a_2, \dots, a_{|\mathcal{I}|})$, it is then possible to calculate the equilibrium price $p_Z^* = V(\mathcal{J}) - V(\mathcal{J} \setminus Z)$ implied by Theorem 1, as well as the price per viewer p_Z^*/λ_Z defined analogously to Section 2.1. In particular,

$$p_Z^* = \frac{1}{|\mathcal{I}|} \sum_{i \in \mathcal{I}} a_i \eta_{iZ} \prod_{Z' \neq Z} (1 - \eta_{iZ'}), \quad \lambda_Z = \frac{1}{|\mathcal{I}|} \sum_{i \in \mathcal{I}} \eta_{iZ}, \quad \eta_{iZ} = 1 - \prod_{j \in Z} (1 - \eta_{ij})$$

where η_{iZ} is the probability of viewer i seeing an ad placed in one of owner Z ’s slots.

We perform this calculation in the audience survey data, letting T_{ij} be the number of hours that respondent i spent watching outlet j in the last week. 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 . For each specification, we calculate the predicted price per viewer p_Z^*/λ_Z for each owner. 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 im-

pression). Panel B shows the analogous scatterplot for the specification with values proportional to household income. Predicted $\log(\text{price per viewer})$ explains 34.0 percent of the variation in observed $\log(\text{price per impression})$. In the latter case, the slope of the relationship between observed and predicted prices is economically similar to 1, and statistically indistinguishable from 1. Across both specifications, the model is able to rationalize very large differences in advertising prices between owners.

We can 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.

We can also evaluate the fit of the model to time trends in television advertising revenues. Panel A of Figure 6 shows that annual revenues increased slightly between 2014 and 2019 while 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, in the model as in the data, the price per impression rose substantially 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.²⁴ 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.

²⁴The 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 10 includes an alternative version of Panel B in which impressions are imputed from the audience survey data.

We conclude that the quantitative model provides a reasonable match to variation in advertising prices across owners, demographic groups, and over time. We therefore apply the model to questions of economic and policy interest.

5.1 Application to Mergers of Television Network Owners

Figure 7 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 highlight three mergers that occurred after 2015: Discovery and Scripps (2018), CBS and Viacom (2019), Disney and Fox (2019).

Comparing Panels A and B of Figure 7 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 impact on HHI, but is much closer to the Discovery-Scripps merger in terms of both audience overlap and revenue impact.

5.2 Application to the Incentive to Invest in Content

We can also quantify the effects of competition on the incentives of network owners to invest in content to attract different kinds of audience members. To do this we augment the model in Section 2 to incorporate a content investment game. Specifically, suppose that each viewer i is attracted to each owner Z 's content with probability $\alpha_{iZ} \in [0, 1]$. If the viewer is attracted to owner Z 's content, the viewer again sees ads on outlets $j \in Z$ with probability η_{ij} , independently across j . 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 $\sum_{i \in I} C_{iZ}(\alpha_{iZ})$ where $C_{iZ}(0) = C'_{iZ}(0) = 0$ and $C'_{iZ}(1) > a_i$.

Proposition 18 in Appendix A.3 implies that in any equilibrium we must have that

$$C'_{iZ}(\alpha_{iZ}) = p_{iZ}^*, \quad p_{iZ}^* = a_i \eta_{iZ} \prod_{Z' \neq Z} (1 - \alpha_{iZ'} \eta_{iZ'})$$

where we can think of $C'_{iZ}(\alpha_{iZ})$ as owner Z 's marginal willingness to pay to attract viewer i , p_{iZ}^* as viewer i 's contribution to the value of owner Z 's advertising slots, and

$$m_i = \frac{\sum_{Z \in \mathcal{Z}} C'_{iZ}(\alpha_{iZ})}{\sum_{Z \in \mathcal{Z}} \eta_{iZ}}$$

as the television market's total marginal willingness to pay per impression to attract viewer i . To operationalize the calculation of m_i , for each owner Z we consider the limiting case where $\alpha_{iZ'} = 1$ for all $Z' \neq Z$, so that the resulting value \hat{m}_i represents the total marginal willingness to pay per impression to attract viewer i when each owner believes all other owners will attract the viewer with certainty.

Figure 8 depicts the average value of $\ln(\hat{m}_i)$ 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. Under the factual ownership partition, these models imply that willingness to pay per impression is 82.8 (Panel A, SE = 3.18) or 88.8 (Panel B, SE = 4.83) 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 = 3.07) under concentrated ownership. Figure 8 thus illustrates a sense in which the competitive forces that we study can influence the direction of content investment, and suggests that television network owners would have a stronger incentive to target content to older viewers if the television market were less competitive.

Appendix Figure 8 reports the results assuming content investment is made on the network or outlet level, i.e., each network or outlet is an independent player in the content investment game and sets its own α . Appendix Figure 9 reports the results where \hat{m}_i is calculated assuming $\alpha_{iZ'} = \alpha \in \{0.5, 0.75, 1\}$ for all $Z' \neq Z$.

5.3 Application to Netflix Carrying Advertising

Netflix has said that it may introduce a service that carries advertising (Flint and Jacob 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 = 0.04).²⁵ 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.

6 Conclusions

We extend existing theoretical results on competitive advertising markets with 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. The prediction can help us understand why there is a premium for younger viewers on television and a premium for older viewers on social media. 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, the effect of competition on the incentive to invest in content to attract different kinds of viewers, and the effect of Netflix

²⁵Dividing the total daily US Netflix viewing minutes implied by values reported in MoffettNathanson (2022, Exhibit 11) by the 2020 US Population (Census 2021) yields daily viewing of 22.7 minutes per capita, close to the average of 24.4 we calculate from the audience survey data. If, due to coviewing and other factors, Netflix viewing time is larger than what we estimate, then we expect our calculations to understate the effect Netflix ad carriage on advertising prices.

ad carriage on linear television advertising prices.

References

- AdStage. 2020. *Paid Media Q1 2020 Benchmark Report*. San Francisco, CA: 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*. San Francisco, CA: 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.
- 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.chwe.net/michael/h.pdf> on July 5 2021.
- Competition and Markets Authority. 2019. “Online platforms and digital advertising: Market study interim report.” 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 April 14. “Cable guy says ratings don’t ad up.” *Hollywood Reporter*.


- 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.
- 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. 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.
- Flint, Joe and Denny Jacob. 2022. "Netflix explores a version with ads as subscriber base shrinks." *Wall Street Journal*, April 19. Accessed at <https://www.wsj.com/articles/netflix-earnings-q1-2022-11650325682> on June 10 2022.
- Gabler, Neal. 2014. "The tyranny of 18 to 49: American culture held hostage." *The Norman Lear Center*. Accessed at <https://learcenter.org/wp-content/uploads/2014/10/Gabler18to49.pdf> on July 5 2021.
- 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/>.

- mons.com/our-data/national-studies/survey-american-consumer/. (accessed in May 2022).
- 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.
- 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.
- 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.
- Kitterman, Ted. 2020. "Research snapshot: Mobile marketing drives digital ad spending." *PR Daily*. Accessed at <https://www.prdaily.com/mobile-marketing-drives-digital-ad-spending/> on October 5 2021.
- Lambrech, 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.
- 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. Forthcoming. "How viewer tuning, presence, and attention respond to ad content and predict brand search lift." *Marketing Science*. Accessed at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3815349 on February 1 2021.
- Media Intelligence and Kantar Media. 2021. AdSpender. Accessed at <http://edu.adspender.kantarmedia.com> on October 4, 2021.
- MoffettNathanson. 2022. Netflix and Disney: Mad Men to the Rescue? MoffettNathanson (An SVB Company).

- Newman, Andrew Adam. 2012. "In AARP's view, advertisers need to focus." *The 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*. 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 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.
- 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.
- Prat, Andrea. 2018. "Media power." *Journal of Political Economy* 126(4): 1747-1783.
- Prat, Andrea and Tommaso Valletti. Forthcoming. "Attention oligopoly." *American Economic Journal: Microeconomics*.
- Pomerantz, Earl. 2006. "Why do advertisers still covet the 18-49s?" *Television Quarterly* 36:

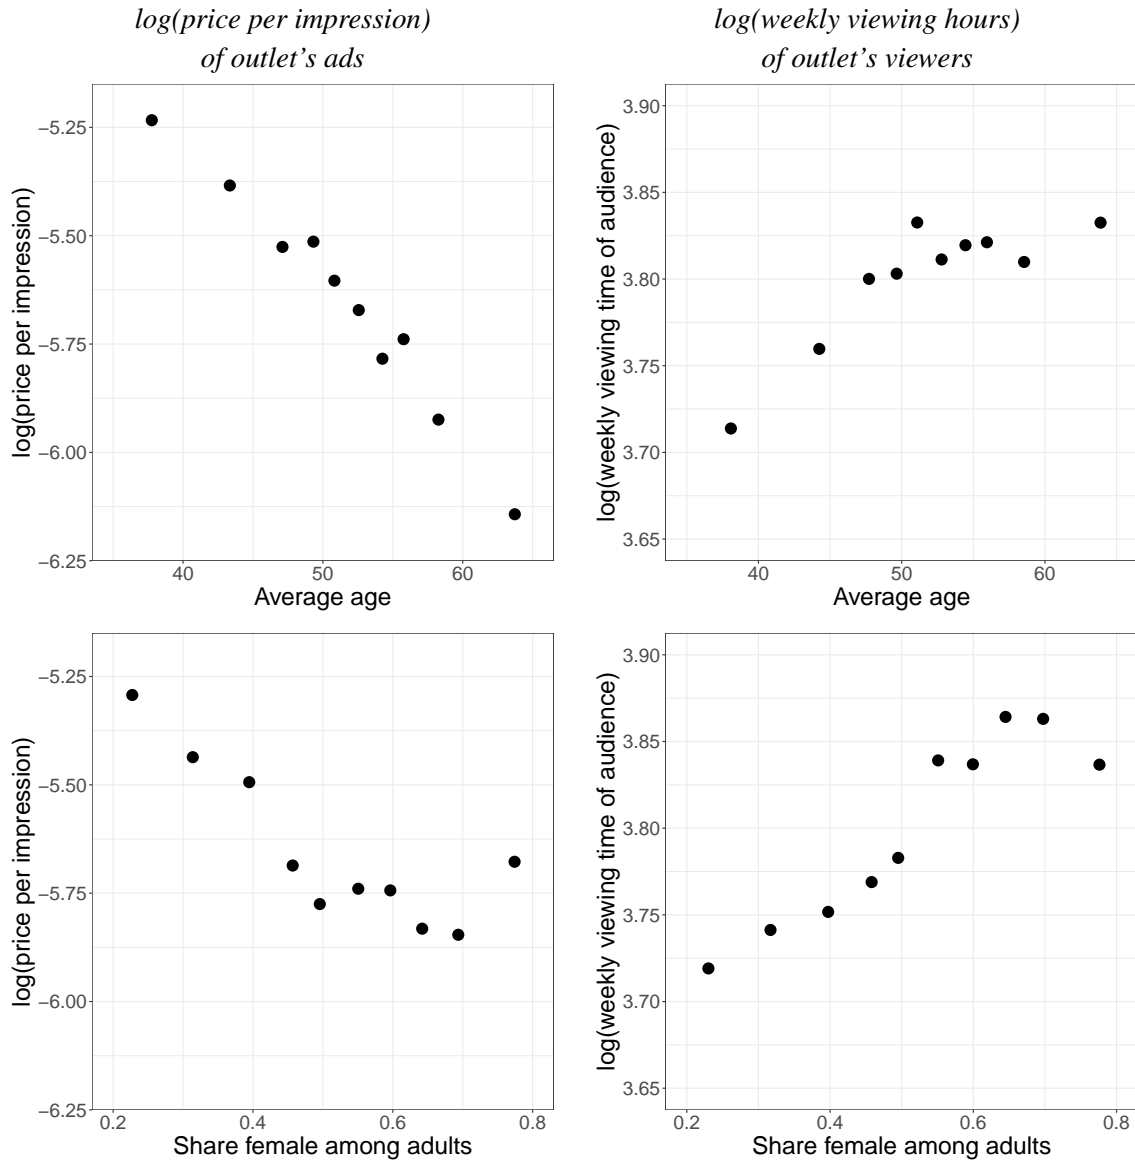
- 40-44. Accessed at https://web.archive.org/web/20200713102219/http://www.tvquarterly.net/tvq_36_3/media/articles/36.3Why_do_Advertisers.pdf on June 3 2021.
- Rochet, Jean-Charles and Jean Tirole. 2003. "Platform competition in two-sided markets." *Journal of the European Economic Association* 1(4): 990–102.
- 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." Statista Research Report. Accessed at <https://www.statista.com/statistics/492908/media-owners-revenue-sector-usa/> on August 16 2021.
- S&P Global Market Intelligence. 2019. "SNL Kagan Report for 2015." S&P Global Market Intelligence. <https://www.spglobal.com/marketintelligence/en/>. Data accessed in November 2019.
- 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.
- 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.
- 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.
- Zubанov, 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.

Figure 1: 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 log of the outlet's impressions per hour.


Figure 2: 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 3: 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



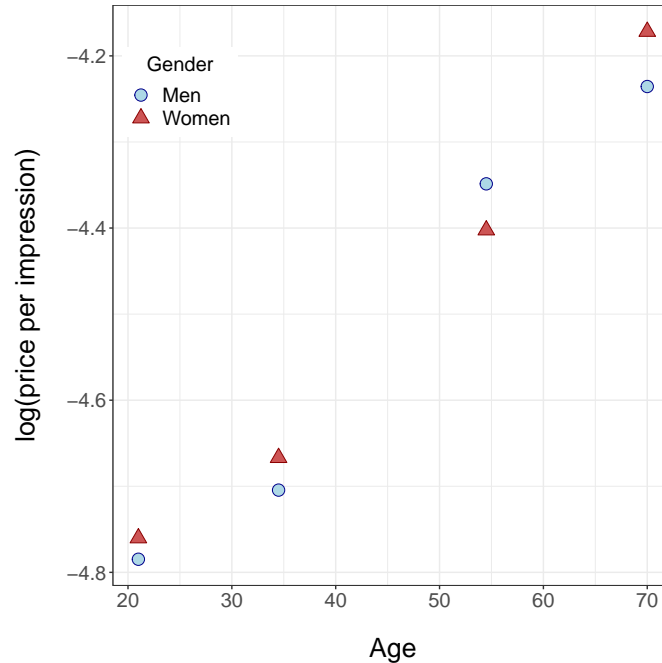
Notes: Each plot is a binned scatterplot of a dependent variable against an independent variable of interest. To construct each plot, we regress the dependent variable on indicators for deciles of the independent variable of interest and a set of controls. The unit of analysis in the regression is an 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 of the 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.

Figure 4: **Demographic Premia and Viewing Time on Facebook**

Panel A: Data from our experiment



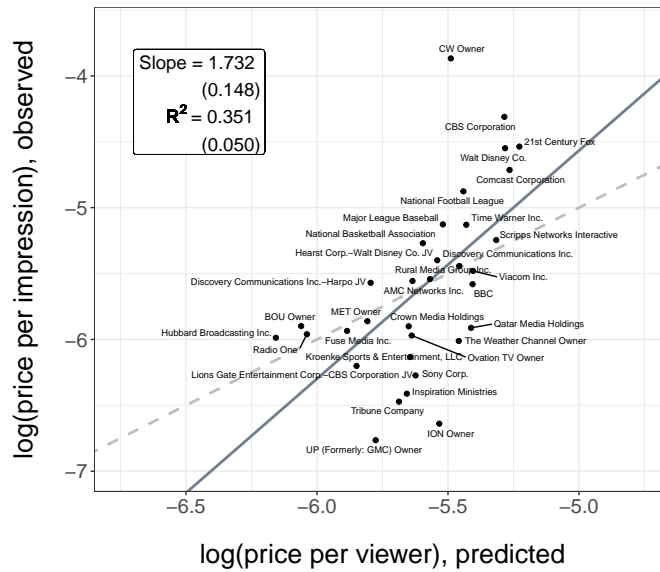
Panel B: Data from Allcott et al. (2020b)




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+.

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 dot-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.

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.


Figure 7: **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. For the mergers that took place, the log of the simulated change in revenue is -6.34 for Disney-Scripps (SE = 0.03), -6.35 for CBS-Viacom (SE = 0.02), and -5.31 for Disney-FOX (SE = 0.02), where standard errors in parentheses are obtained via a nonparametric bootstrap over survey respondents with 100 replicates.

Figure 8: 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 average log(total marginal willingness to pay per impression), $\ln(\hat{m}_i)$, 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.

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 3.3 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).

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.

A Proofs and Additional Theoretical Results

A.1 Preliminaries

Lemma 1. *Write SPEPS to abbreviate subgame perfect equilibrium in pure strategies.*

- (i) *For any $V(\cdot)$, not necessarily monotone or submodular, in any SPEPS all advertisers make the same total payment to any given owner. This holds even if each owner Z is endowed with an arbitrary partition \mathcal{F}_Z of Z such that they are only allowed to bundle outlets in the same cell of the partition.*
- (ii) *If $V(\cdot)$ is either monotone and submodular or strictly monotone and not necessarily submodular, then in any SPEPS each advertiser buys slots on all outlets.*

Proof. Fix any SPEPS. First, observe that because the advertisers are homogeneous, they must have the same equilibrium payoff (say W).

For part (i), suppose for 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 who pays the most to owner Z . 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 any $\varepsilon > 0$, where p_S^* denotes the minimum price to buy slots on the outlets in a given set S in the SPEPS. 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 in the original SPEPS), any advertiser can obtain a payoff of $W + \varepsilon|\mathcal{B}|$. Note that any set of outlets an advertiser wants to buy slots on after this deviation is also a valid choice in the original equilibrium. Therefore, 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 . But then this is a profitable deviation for owner Z when ε is small enough. Contradiction.

For part (ii), suppose for contradiction that there exist some owner $Z \in \mathcal{Z}$, some outlet $j \in Z$, and some advertiser n who does not buy a slot on outlet j . By part (i), all advertisers pay the same total amount to owner Z (say t). Let $T \subset Z$ be the set of outlets that advertiser n buys slots on from owner Z . Let R be the set of outlets in $\mathcal{J} \setminus Z$ that advertiser n buys slots on. The equilibrium payoff for each advertiser is thus given by

$$W = V(T \cup R) - p_R^* - t.$$

Let owner Z offer a single bundle Z with a price $\tilde{p}_Z = t + \varepsilon$, for some $\varepsilon > 0$ (recall that here we assume the owner can bundle anything). If $V(\cdot)$ is strictly monotone, then we clearly have

$V(Z \cup R) - V(T \cup R) > 0$. 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. Therefore, for ε small enough, we have

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

Pick any such ε . Every advertiser would buy the bundle Z at the price $t + \varepsilon$, because any strategy not doing so is a feasible strategy in the equilibrium and generates a payoff less than or equal to W . But this is then a profitable deviation for owner Z . Contradiction. \square

A.2 Proofs Omitted From the Main Text

Example 1. Let i denote a viewer uniformly drawn from the set of viewers. Let X_S denote the random number of outlets in $S \subseteq \mathcal{J}$ watched by i . We can write

$$V(S) = \mathbb{E} \left[a_i \sum_{m=0}^{X_S} \beta_m \right].$$

To show V is monotone and submodular, it suffices to fix a realization of viewer i 's decision, and show the realized value function

$$\tilde{V}(S) := a_i \sum_{m=0}^{X_S} \beta_m$$

is monotone and submodular, since averaging preserves monotonicity and submodularity. For any $S \subseteq \mathcal{J}$ and $j \in \mathcal{J} \setminus S$, we have

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

where $i \rightarrow j$ denotes the event that viewer i views outlet j . This shows monotonicity as $a_i > 0$ and $\beta_m \geq 0$ for all m . For submodularity, fix any $S \subseteq T \subseteq \mathcal{J}$, and $j \in \mathcal{J} \setminus T$. Since $S \subseteq T$, we have $1 \leq X_S + 1 \leq X_T + 1$. Since β_m is non-increasing in m for $m \geq 1$, it follows immediately that

$$\tilde{V}(S \cup \{j\}) - \tilde{V}(S) = a_i \mathbf{1}_{i \rightarrow j} \beta_{X_S+1} \geq a_i \mathbf{1}_{i \rightarrow j} \beta_{X_T+1} = \tilde{V}(T \cup \{j\}) - \tilde{V}(T),$$

which shows submodularity.

Example 2. We can write

$$V(S) = \sum_{C \in \mathcal{C}} V(S \cap C; a_C)$$

where $V(\cdot; a)$ is the value function given in the proof for Example 1 with $\beta_m = 0$ for $m \geq 2$. Since $V(\cdot; a)$ is monotone and submodular for any $a > 0$, and both monotonicity and submodularity are preserved under restriction and addition, we have that $V(\cdot)$ is monotone and submodular.

Example 3. Let \mathcal{K} be the set of programs with a generic element denoted by k , and let $\mathcal{K}_j \subseteq \mathcal{K}$ be the programs associated with outlet j . Let i denote a viewer uniformly drawn from the set of viewers. Let $i \rightarrow k; A$ denote the event that viewer i watches program k and program k carries an ad. For a set of programs $\mathcal{K}' \subseteq \mathcal{K}$, let

$$X_{\mathcal{K}'} = \sum_{k \in \mathcal{K}'} \mathbf{1}_{i \rightarrow k; A}$$

be the number of programs in \mathcal{K}' watched by i when each program carries an ad. Let

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

consist of sets of representative programs (i.e., one program for each outlet). Let $\mathcal{K}_S = \cup_{j \in S} \mathcal{K}_j$. Note that for each advertiser, the value of a set of outlets $S \subseteq \mathcal{J}$ can be written as

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

where the expectation is taken over a fixed probability distribution (regardless of the choice of S) which specifies that every program carries an ad. As in Example 1, it suffices to show that for any realization and any $\mathcal{K}' \in \mathcal{R}$ fixed, we have

$$\tilde{V}(S) := u(X_{\mathcal{K}' \cap \mathcal{K}_S})$$

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

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

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

Proof of Theorem 1. We first construct a SPEPS. Let each owner Z offer a single bundle consisting of all outlets in Z with a price $p_Z = v_Z$. For any profile of posted prices (including off-the-equilibrium-path histories), let every advertiser solve in the second stage the problem,

$$\max_{S \subseteq \mathcal{J}} V(S) - p_S^*,$$

where p_S^* denotes the minimum price to buy slots on all of the outlets in S for a given profile of prices p . The problem may have multiple solutions. Pick a solution S^* such that $|S^*|$ is the largest. Let advertisers buy slots on all outlets in S^* .

It remains to verify that no owner has a profitable deviation. Observe that if $p_Z = v_Z$ 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 p_{-Z} of other owners' bundles. 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 players follow the proposed strategy. Fix any advertiser. By the above observation we know that the advertiser would always buy slots on all outlets not in Z . Then the maximal amount that owner Z can extract from this advertiser is v_Z because for any $S \supseteq \mathcal{J} \setminus Z$, monotonicity of $V(\cdot)$ implies

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

where the right hand side is the maximal price that the advertiser is willing to pay for the slots on outlets in $S \setminus (\mathcal{J} \setminus Z)$. Therefore, following the proposed strategy is optimal for owner Z . Since Z is an arbitrary owner, our construction is a SPEPS.

To prove the second part of the statement, fix any SPEPS of the game. 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 earlier observation, 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 $Z \in \mathcal{Z}$.

Proof of Proposition 1. 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

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

where $g \rightarrow j$ denotes the event that a randomly sampled viewer g views outlet j and X_j^g counts the random number of outlets viewed by g that are not j . 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,

$$\mathbb{E}[\mathbf{1}_{g' \rightarrow j} \beta_{X_j^{g'}+1}] = \mathbb{E}[\mathbf{1}_{g' \rightarrow k} \beta_{X_k^{g'}+1}].$$

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

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

Using independence, we can write the above as

$$\mu_g [\eta_{gj}(1 - \eta_{gk}) - \eta_{gk}(1 - \eta_{gj})] \mathbb{E}[\beta_{X^g+1}] \geq \mu_h [\eta_{hk}(1 - \eta_{hj}) - \eta_{hj}(1 - \eta_{hk})] \mathbb{E}[\beta_{X^h+1}]$$

where X^g counts the random number of outlets viewed by viewer g that are not in $\{j, k\}$. Since $\lambda_j = \lambda_k$, this reduces to

$$(\sigma_{gj} - \sigma_{gk}) \mathbb{E}[\beta_{X^g+1}] \geq (\sigma_{hk} - \sigma_{hj}) \mathbb{E}[\beta_{X^h+1}].$$

It follows easily from our assumptions that $\sigma_{gj} - \sigma_{gk} = \sigma_{hk} - \sigma_{hj} \geq 0$. So it suffices to show $\mathbb{E}[\beta_{X^g+1}] \geq \mathbb{E}[\beta_{X^h+1}]$. 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 for all $j \in \mathcal{J}$,

$$\mathbf{1}_{g \rightarrow j} \leq \mathbf{1}_{h \rightarrow j}.$$

Under this coupling, we have $X^g \leq X^h$ pointwise. The claim then follows directly by noting that β_m is non-increasing in m for $m \geq 1$.

Now suppose $\sigma_{gj} > \sigma_{gk}$ and $\eta_{gj'} < \eta_{hj'}$ for some $j' \neq j, k$. Using integration by parts, we have

$$\mathbb{E}[\beta_{X^g+1}] - \mathbb{E}[\beta_{X^h+1}] = \int_0^\infty \mathbb{P}(\beta_{X^g+1} > s) ds - \int_0^\infty \mathbb{P}(\beta_{X^h+1} > s) ds$$

$$= \sum_{m=1}^{\infty} (\beta_m - \beta_{m+1}) \left(\mathbb{P}(X^g + 1 \leq m) - \mathbb{P}(X^h + 1 \leq m) \right) > 0$$

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

Proof of Proposition 2. We follow the same notation as in the proof of Proposition 1. By Theorem 1, we can write

$$\begin{aligned} p_j^* &= a \sum_g \mu_g \mathbb{E}[\mathbf{1}_{g \rightarrow j} \beta_{X_j^g+1}] \\ &= a \sum_g \mu_g (\eta_{gj} \eta_{gk} \mathbb{E}[\beta_{X^g+2}] + \eta_{gj} (1 - \eta_{gk}) \mathbb{E}[\beta_{X^g+1}]) \\ &= a \sum_g \mu_g (\eta_{gk} \eta_{gj} \mathbb{E}[\beta_{X^g+2}] + \delta \eta_{gk} (1 - \frac{1}{\delta} \eta_{gj}) \mathbb{E}[\beta_{X^g+1}]) \\ &= a \sum_g \mu_g (\eta_{gk} \eta_{gj} \mathbb{E}[\beta_{X^g+2}] + \eta_{gk} (1 - \eta_{gj}) \mathbb{E}[\beta_{X^g+1}] + \eta_{gk} (\delta - 1) \mathbb{E}[\beta_{X^g+1}]) \\ &= a \sum_g \mu_g \eta_{gk} (\mathbb{E}[\beta_{X_k^g+1}] + (\delta - 1) \mathbb{E}[\beta_{X^g+1}]) \\ &\geq a \sum_g \mu_g \eta_{gk} (\mathbb{E}[\beta_{X_k^g+1}] + (\delta - 1) \mathbb{E}[\beta_{X_k^g+1}]) = \delta p_k^* = \frac{\lambda_j}{\lambda_k} p_k^* \end{aligned}$$

where we have used independence, $\delta \geq 1$, $X_k^g \geq X^g$, and β_m is non-increasing for $m \geq 1$. Now suppose $\delta > 1$. For any group g , using integration by parts, we have

$$\begin{aligned} \mathbb{E}[\beta_{X^g+1}] - \mathbb{E}[\beta_{X_k^g+1}] &= \int_0^{\infty} \mathbb{P}(\beta_{X^g+1} > s) ds - \int_0^{\infty} \mathbb{P}(\beta_{X_k^g+1} > s) ds \\ &= \sum_{m=1}^{\infty} (\beta_m - \beta_{m+1}) \left(\mathbb{P}(X^g + 1 \leq m) - \mathbb{P}(X_k^g + 1 \leq m) \right) > 0 \end{aligned}$$

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

Proof of Proposition 3. We prove the second part of the statement first. Fix any subgame perfect equilibrium allowing for mixed strategies (SPEMS) and any owner Z . Suppose, for contradiction, the expected revenue per slot r_Z is strictly higher than $\sum_{j \in Z} V(\{j\})/|Z|$. Then the 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 the advertiser buys slots on. Since any non-negative submodular function is also sub-additive, we have

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

So simply not buying anything from Z is a profitable deviation for the advertiser. Contradiction.

Now, for contradiction, suppose $r_Z < (v_Z - \Delta)/|Z|$. 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 the argument 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 $K\tilde{p}_Z$. But this is then a profitable deviation. 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 , 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 assign the resulting expected revenue (averaged over different orders) for owner Z as the payoff to owner Z in the auxiliary game given the pure strategy profile p . This constructs a finite normal-form game among the owners, and thus a Nash equilibrium (possibly in mixed strategies) exists (say \mathcal{E}). Now let each owner play the strategy prescribed by \mathcal{E} in the original game, followed by advertisers choosing which slots to buy in the same way as before. Evidently, this constructs a SPEMS for the original game.

Proof of Proposition 4. We first construct a SPEPS. We use the same construction as in the proof of Theorem 1. Note that 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 SPEPS 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 is the same as in the proof of Theorem 1.

Proof of Proposition 5. As in the proof of Theorem 1, we first construct a SPEPS. Let each owner Z offer a single bundle consisting of all outlets in Z with a price $p_Z = \underline{v}_Z$. For any profile of posted prices (including off-the-equilibrium-path histories), let each advertiser n solve the following problem in the second stage

$$\max_{S \subseteq \mathcal{J}} V_n(S) - p_S^*$$

where p_S^* denotes the minimum price to buy slots on all of the outlets in S for a given profile of prices p . Pick a solution S_n^* such that $|S_n^*|$ is the largest. Let advertiser n buy slots on outlets in S_n^* .

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 p_{-Z} . This is because for any $S \subseteq \mathcal{J} \setminus Z$, submodularity of $V_n(\cdot)$ implies

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

Fix an owner Z . 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 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 $n \in \mathcal{N}$ who does not buy a slot on outlet j . We claim that the total revenue that owner Z extracts is no more than

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

Indeed, this is the maximal revenue that owner Z can possibly get, even if the owner price dis-

criminate using the identities of the advertisers but is subject to the constraint that at least one advertiser does not buy on some outlet $j \in Z$. Now note 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 a SPEPS. The outcome is efficient because all advertisers buy slots on all outlets.

To prove the second part of the statement, fix any efficient SPEPS. Note that 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 $Z \in \mathcal{Z}$.

Proof of Proposition 6. 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 .

Fix any SPEPS. Fix any owner Z and any advertiser n . Let S_c be the set of outlets that advertiser n buys slots on in the cell $c \in \mathcal{F}_Z$. Consider owner Z offering the bundles $\mathcal{B} := \{S_c : S_c \neq \emptyset, c \in \mathcal{F}_Z\}$ with prices $\{p_{S_c}^* : S_c \neq \emptyset, c \in \mathcal{F}_Z\}$, where p_S^* denotes the minimum price to buy slots on the outlets for any set $S \subseteq \mathcal{J}$ in the SPEPS. We claim that owner Z weakly increases the payoff with this strategy. Note that 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 in the original SPEPS). Because in the original SPEPS all advertisers pay the same total amount to Z by Lemma 1(i), 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 in the original SPEPS by advertiser n' due to the tie-breaking rule. But then advertiser n' pays strictly less than advertiser n to owner Z in the original SPEPS, contradicting Lemma 1(i).

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 is a partition of S . In particular, bundles not contained in S are not offered by the owners. For any bundle B offered by any owner Z , by rationality of the advertisers,

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

Now, for contradiction, suppose 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 $B' \in \mathcal{B}_Z$. Contradiction.

Proof of Proposition 7. 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 .

We first show that $\hat{\mathbb{C}} := \{(Z, \xi_Z(V(\mathcal{J}) - V(\mathcal{J} \setminus Z)))\}$ 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 $\tilde{\mathbb{C}}$. Note that 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.

Proof of Proposition 8. Fix any profile of announced reserve prices $p := \{p_B : B \subseteq Z, Z \in \mathcal{Z}\}$. Note that, 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, every 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.

Proof of Proposition 9. This follows directly from Theorem 1, treating each viewer i as a separate market.

A.3 Additional Results

Comparative statics for multi-outlet owners. Consider a setting identical to that of Section 2.1 except that each owner Z may own multiple outlets. Suppose diminishing returns are perfect in the sense that $\beta_m = 0$ for $m \geq 2$. Members of group $g \in G$ see ads on outlet j with probability η_{gj} , independently across outlets. For a given owner Z , let

$$\eta_{gZ} = 1 - \prod_{j \in Z} (1 - \eta_{gj}), \quad \lambda_Z = \sum_{g \in G} \mu_g \eta_{gZ}, \quad \sigma_{gZ} = \frac{\mu_g \eta_{gZ}}{\lambda_Z}$$

denote the share of group g that is in the owner's audience, the total mass of the owner's audience, and the share of this audience that comes from group g , respectively. Then p_Z^*/λ_Z is the price per viewer collected by owner Z for an ad slot on each of its outlets. By Theorem 1, we know that

$$p_Z^* = a\beta_1 \sum_g \mu_g \eta_{gZ} \prod_{j \notin Z} (1 - \eta_{gj}) = a\beta_1 \sum_g \mu_g \eta_{gZ} \prod_{Z' \neq Z} (1 - \eta_{gZ'}).$$

Note that the equilibrium prices above are identical to those in the setting of Section 2.1, replacing outlets with owners, if we specify perfect diminishing returns. Therefore our results on comparative statics apply immediately.

Proposition 10. *Suppose that group $g \in G$ is less active than group $h \in G$ in the sense that $\eta_{gZ} \leq \eta_{hZ}$ for all $Z \in \mathcal{Z}$. Suppose that owner $Y \in \mathcal{Z}$ draws a larger share of its audience from group g and a smaller share of its audience from group h than owner $Z \in \mathcal{Z}$, 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 , $p_Y^*/\lambda_Y \geq p_Z^*/\lambda_Z$.*

Proposition 11. *Suppose that owner Y has a larger audience than owner Z in the sense that for*

some $\delta \geq 1$, $\eta_{gY} = \delta \eta_{gZ}$ for all $g \in G$. Then owner Y has a higher price per viewer than owner Z , $p_Y^*/\lambda_Y \geq p_Z^*/\lambda_Z$.

Partially increasing returns.

Example 4. Owners are singletons, each of a set of viewers i views at least L outlets, each outlet has a strictly positive mass of viewers, and an advertiser's value for viewer i seeing its ad M times is $a_i \sum_{m=0}^M \beta_m$ 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 $\beta_L \leq \min_{1 \leq m < L} \beta_m$.

Proposition 12. *The value function $V(\cdot)$ in Example 4 satisfies the hypotheses of Proposition 4.*

Proof. As in the proof of Example 1, let $i \rightarrow j$ denote the event that a random viewer i watches outlet j , X_S count the number of outlets viewed in set S , and \tilde{V} denote the realized value function. Strict monotonicity follows because

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

is strictly positive with positive probability. For the decreasing differences condition, consider any j and any $S \subseteq \mathcal{J} \setminus \{j\}$, and note that

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

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

Example 5. Owners are singletons. Let \mathcal{I} be the set of viewers who view at least L number of outlets, and $\varepsilon = \mathbb{P}(i \notin \mathcal{I})$. For every pair of outlets $\{j, k\}$, $\mathbb{P}(i \text{ views } j \text{ and } k | i \in \mathcal{I}) > 0$. An advertiser's value for viewer i seeing its ad M times is $a_i \sum_{m=0}^M \beta_m$ 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 13. *There exists $\bar{\varepsilon} > 0$ such that for all $\varepsilon \in [0, \bar{\varepsilon}]$ the value function $V(\cdot)$ in Example 5 satisfies the hypotheses of Proposition 4.*

Proof. Strict monotonicity follows by the same argument as in the proof of Proposition 12. Now, fix any j and any $S \subset \mathcal{J} \setminus \{j\}$, and any $k \in \mathcal{J} \setminus (S \cup \{j\})$. Note that because $\mathbb{P}(i \text{ views } j \text{ and } k | i \in \mathcal{I}) > 0$

$\mathcal{I}) > 0$, we have

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

Thus, since β_m is strictly decreasing for all $m \geq L$, and $\beta_L < \min_{1 \leq m < L} \beta_m$, by the argument in the proof of Example 1, we have

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

Let

$$\tau_{S,j} := \mathbb{E}[a_i \mathbf{1}_{i \rightarrow j} \beta_{X_S+1} | i \in \mathcal{I}] - \mathbb{E}[a_i \mathbf{1}_{i \rightarrow j} \beta_{X_{\mathcal{J} \setminus \{j\}}+1} | i \in \mathcal{I}] > 0.$$

Let

$$\tau := \min_{j, S \subset \mathcal{J} \setminus \{j\}} \tau_{S,j} > 0.$$

Let $\bar{\beta} = \max_{1 \leq m < L} \beta_m$. Now, we claim that for any ε such that

$$0 \leq \varepsilon \leq \frac{\tau}{\tau + \bar{\beta}}$$

we have that $V(\cdot)$ satisfies the decreasing differences condition. To see this, consider any j and any $S \subset \mathcal{J} \setminus \{j\}$. Note that

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

where the second inequality follows from the construction of τ and the third inequality follows from $0 \leq \varepsilon \leq \frac{\tau}{\tau + \bar{\beta}}$. \square

Proposition 14. *Consider a special case of Example 4 with $a_i = a$ for all i , and further impose the structure in Section 2.1, where for each group g , there are at least L many outlets (denoted by set \mathcal{L}_g) such that $\eta_{gj} = 1$ 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 1 and 2. By Propositions 4 and 12, we have the equilibrium prices equal to the incremental values.

For the conclusion of Proposition 1, recall that we use a monotone coupling. Under that coupling we have $X^g \leq X^h$ pointwise. Recall that X^g counts the random number of outlets viewed by viewer g that are not in $\{j, k\}$ and hence $X^g, X^h \geq L - 2$. Hence $X^g + 1, X^h + 1 \geq L - 1$. 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 $X^g + 1 \leq X^h + 1$ pointwise, we have $\beta_{X^g+1} \geq \beta_{X^h+1}$ pointwise and so $\mathbb{E}[\beta_{X^g+1}] \geq \mathbb{E}[\beta_{X^h+1}]$, which concludes the proof as before.

For the conclusion of Proposition 2, recall that X_k^g counts the random number of outlets viewed by viewer g that are not k . So $X_k^g + 1 \geq X^g + 1 \geq L-1$. Because β_m is non-increasing in m for $m \geq L-1$, we have $\beta_{X^g+1} \geq \beta_{X_k^g+1}$ pointwise and so $\mathbb{E}[\beta_{X^g+1}] \geq \mathbb{E}[\beta_{X_k^g+1}]$, which concludes the proof as before. \square

Existence of unbundled pricing equilibrium.

Proposition 15. *Consider the setting of Section 2.1 with $G = 1$, \mathcal{F}_Z being the finest partition (i.e., the owners are not allowed to bundle), and $\beta_m = \beta^{m-1}$ for some constant $\beta \geq 0$. Then, there exists a SPEPS satisfying the tie-breaking rule, and the conclusion of Proposition 2 holds.*

Proof. We construct the set of outlets S that all advertisers buy slots on. Let i denote a viewer uniformly drawn from the population. Let X_T denote the random number of outlets in T viewed by viewer i . Since $\beta_m = \beta^{m-1}$, for any $T \subseteq \mathcal{J}$ we can write

$$V(T) = a \mathbb{E} \left[\sum_{m=1}^{X_T} \beta^{m-1} \right]$$

where the sum is interpreted as zero when $X_T = 0$. Fix an owner Z . Let $F \subseteq Z$ denote the menu Z offers. Since we assume no owner can bundle any outlets, every outlet j in F is sold individually at some price. Let Z solve

$$\max_{F \subseteq Z} \sum_{j \in F} V(F) - V(F \setminus \{j\}).$$

Let F_Z^* be a maximizer of the above problem. For any $T \subseteq \mathcal{J} \setminus Z$, we claim that F_Z^* also solves

$$\max_{F \subseteq Z} \sum_{j \in F} V(T \cup F) - V((T \cup F) \setminus \{j\}).$$

Indeed, for any $F \subseteq Z$ and any $j \in F$,

$$\begin{aligned} V(T \cup F) - V((T \cup F) \setminus \{j\}) &= a \mathbb{E} \left[\sum_{m=1}^{X_{T \cup F}} \beta^{m-1} - \sum_{m=1}^{X_{(T \cup F) \setminus \{j\}}} \beta^{m-1} \right] \\ &= a \mathbb{E} [\mathbf{1}_{i \rightarrow j} \beta^{X_{T \cup F} - 1}] \\ &= a \mathbb{E} [\mathbf{1}_{i \rightarrow j} \beta^{X_T + X_F - 1}] \\ &= a \mathbb{E} [\beta^{X_T}] \mathbb{E} [\mathbf{1}_{i \rightarrow j} \beta^{X_F - 1}] \\ &= \mathbb{E} [\beta^{X_T}] (V(F) - V(F \setminus \{j\})) \end{aligned}$$

where we have used the fact that viewing probabilities are independent across outlets. Therefore,

$$\max_{F \subseteq Z} \sum_{j \in F} V(T \cup F) - V((T \cup F) \setminus \{j\}) = \mathbb{E}[\beta^{X_T}] \max_{F \subseteq Z} \sum_{j \in F} V(F) - V(F \setminus \{j\})$$

and thus is also solved by F_Z^* . Now let $S = \bigcup_{Z \in \mathcal{Z}} F_Z^*$. We can then construct a SPEPS using the set S and equipping outlets in S with the prices identified in Proposition 6: $p_j = v_j^S$ for all $j \in S$ and $p_j = \infty$ for all $j \notin S$. By the proof of Proposition 6, we know that with the proposed bundles and prices, it is optimal for any advertiser to buy slots on all outlets in S . We also know that given these prices, for any owner Z , regardless of what owner Z does, the advertisers buy slots on outlets in $S \setminus Z$ (in this construction, for any advertiser n , we may let the tie-breaking ordering \mathcal{O}_n be any complete ordering over the owners). Therefore, owner Z simply solves the problem

$$\max_{F \subseteq Z} \sum_{j \in F} V((S \setminus Z) \cup F) - V((S \setminus Z) \cup F \setminus \{j\})$$

which has a maximizer F_Z^* as shown earlier. So it is optimal for owner Z to offer menu F_Z^* in which each outlet $j \in F_Z^*$ has a price v_j^S . Since this holds for any owner, the construction is a SPEPS.

To see how the conclusion of Proposition 2 holds, note that

$$p_j^* = V(S) - V(S \setminus \{j\}) = a \sum_g \mu_g \mathbb{E}[\mathbf{1}_{g \rightarrow j} \beta_{X_j^g+1}]$$

where we use the notations in the proof of Proposition 2, except that X_j^g now counts the number of outlets in $S \setminus \{j\}$ that are viewed by g . The rest follows verbatim. \square

Auctioning of scarce advertising slots to heterogeneous advertisers. Each owner owns one outlet, and each outlet has K slots, where $K < N$ (so ad slots are scarce). The advertisers are heterogeneous, with value functions given by $a_n V(\cdot)$ (with $V(\emptyset)$ normalized to 0). We order the advertisers so that $a_1 > a_2 > \dots > a_N > 0$. We assume that a_{K+1} is sufficiently smaller than a_K in the sense that for all j ,

$$a_{K+1} V(\{j\}) < a_K (V(\mathcal{J}) - V(\mathcal{J} \setminus \{j\})). \quad (\text{A1})$$

Each owner runs a uniform price auction (i.e., the K slots are sold at the $(K+1)$ -th highest bid) with ties broken in favor of advertisers with higher a_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 payoffs 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 (i) $V(\cdot)$ is monotone and submodular; and (ii) 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 $a_K(V(\mathcal{J}) - V(\mathcal{J} \setminus \{j\}))$.*

Proof. Consider the following strategy profile: in every auction j , each advertiser n with $n \leq K$ bids $a_n(V(\mathcal{J}) - V(\mathcal{J} \setminus \{j\}))$; advertiser $K + 1$ bids $a_K(V(\mathcal{J}) - V(\mathcal{J} \setminus \{j\}))$; and advertiser n with $n > K + 1$ 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. Note that regardless of the choices the advertiser makes on other auctions, the advertiser weakly prefers to win auction j at the price of $a_K(V(\mathcal{J}) - V(\mathcal{J} \setminus \{j\}))$ because for any $S \subseteq \mathcal{J} \setminus \{j\}$, we have

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

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

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 $a_K(V(\mathcal{J}) - V(\mathcal{J} \setminus \{j\}))$. 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

$$a_n(V(S \cup \{j\}) - V(S)) \leq a_n V(\{j\}) \leq a_{K+1} V(\{j\}) < a_K(V(\mathcal{J}) - V(\mathcal{J} \setminus \{j\})),$$

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

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 $a_K(V(\mathcal{J}) - V(\mathcal{J} \setminus \{j\}))$. Therefore, the K winning bidders in auction j must be advertisers $1, \dots, K$. However, at a price strictly higher than $a_K(V(\mathcal{J}) - V(\mathcal{J} \setminus \{j\}))$ 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 $a_K(V(\mathcal{J}) - V(\mathcal{J} \setminus \{j'\}))$. But then advertiser n' can

profitably deviate to losing auction j' by the argument above. 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 . Suppose toward contradiction that x is not the equilibrium allocation (i.e., the ad slots are not all allocated to advertisers $1, \dots, K$). 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,

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

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

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

and hence \tilde{x} gives a strictly higher total surplus than x . But x is assumed to be an efficient allocation. 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 $a_K(V(\mathcal{J}) - V(\mathcal{J} \setminus \{j\}))$, then advertiser K can profitably deviate to losing auction j . Therefore, in every auction j , the clearing price must be weakly lower than $a_K(V(\mathcal{J}) - V(\mathcal{J} \setminus \{j\}))$. Now, if there is any auction j' in which the clearing price is strictly lower than $a_K(V(\mathcal{J}) - V(\mathcal{J} \setminus \{j'\}))$, the equilibrium cannot be owner-optimal, because we have just shown an equilibrium that has clearing prices equal to $a_K(V(\mathcal{J}) - V(\mathcal{J} \setminus \{j\}))$ for all j . Thus, in every efficient owner-optimal equilibrium, the clearing price in every auction j must be exactly $a_K(V(\mathcal{J}) - V(\mathcal{J} \setminus \{j\}))$. \square

Competitors' ad effect. We consider a setting in which each owner owns a single outlet, and 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} is the vector of outlets 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:

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

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

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

Proof. We construct a SPEPS 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 a SPEPS 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 if $p_j = \tilde{v}_j$ is offered by an owner, 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 holding everything else fixed (including other advertisers' decisions).

Now fix any owner j . Suppose all other players follow the proposed strategy. Note that 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), the maximal amount owner j can extract from an advertiser is at most $V(\mathcal{J}, \overrightarrow{\mathcal{J} \setminus \{j\}}) - V(\mathcal{J} \setminus \{j\}, \overrightarrow{\mathcal{J} \setminus \{j\}})$. Further, we claim that at least one advertiser would stop buying the slot on outlet j after the price increase. Suppose not. Then all advertisers buy slots on all outlets. But the last advertiser moving in sequence has a profitable deviation of buying only the slots on 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 a SPEPS. \square

Incentive to invest in content. 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,*

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

which, under perfect diminishing returns, is equivalent to

$$C'_{iZ}(\alpha_{iZ}) = a_i \eta_{iZ} \prod_{Z' \neq Z} (1 - \alpha_{iZ'} \eta_{iZ'}).$$

Proof. For a given investment profile α , by Theorem 1, the equilibrium prices 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 the payoff to owner Z when making investment choices is given by

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

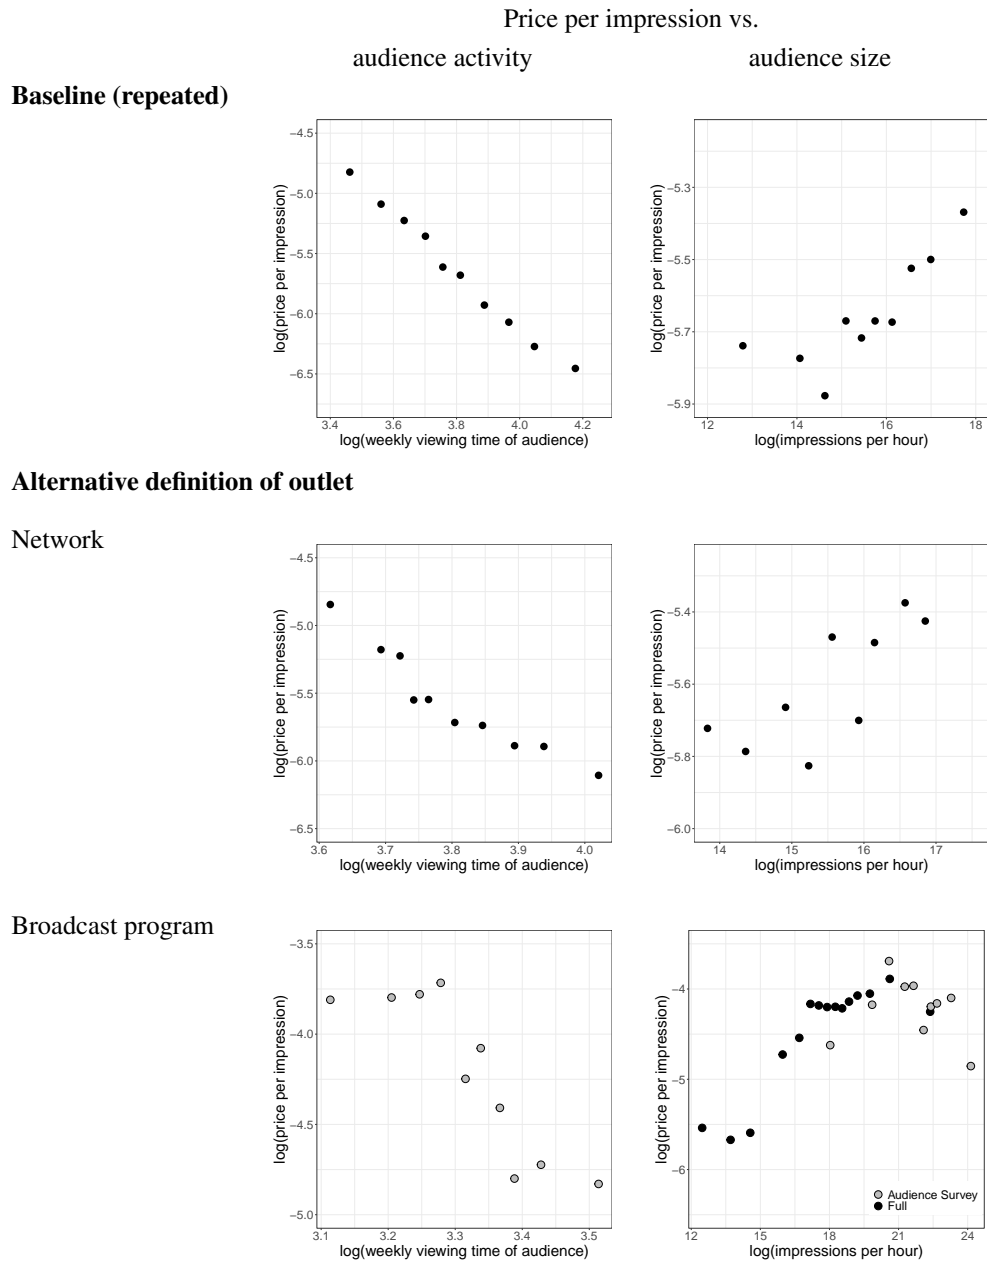
which is a continuous and strictly concave function in $(\alpha_{iZ})_{i \in \mathcal{I}}$ (note that $V_i^Z(\mathcal{J}; \alpha) - V_i^Z(\mathcal{J} \setminus Z; \alpha)$ does not depend on α_{iZ}). For owner Z , the first order condition for α_{iZ} is given by

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

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 $C'_{iZ}(0) = 0$, 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 . \square

B Additional Empirical Results

Appendix Figure 1: 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 1 and the plot “Price per impression vs. audience size” corresponding to Panel B of Figure 3. 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 (3060 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.

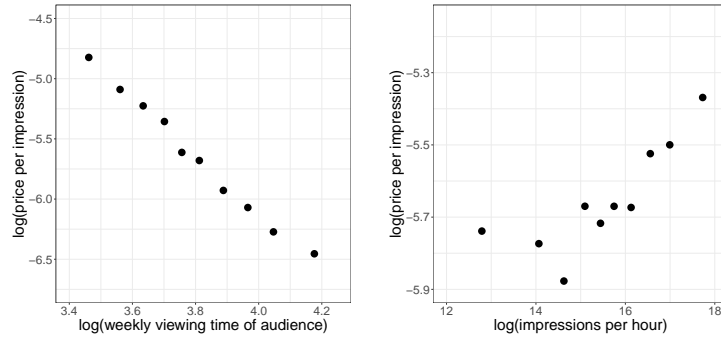
Appendix Figure 2: Sensitivity to Alternative Samples

Price per impression vs.

audience activity


audience size

Baseline (repeated)




Alternative sample years

2014

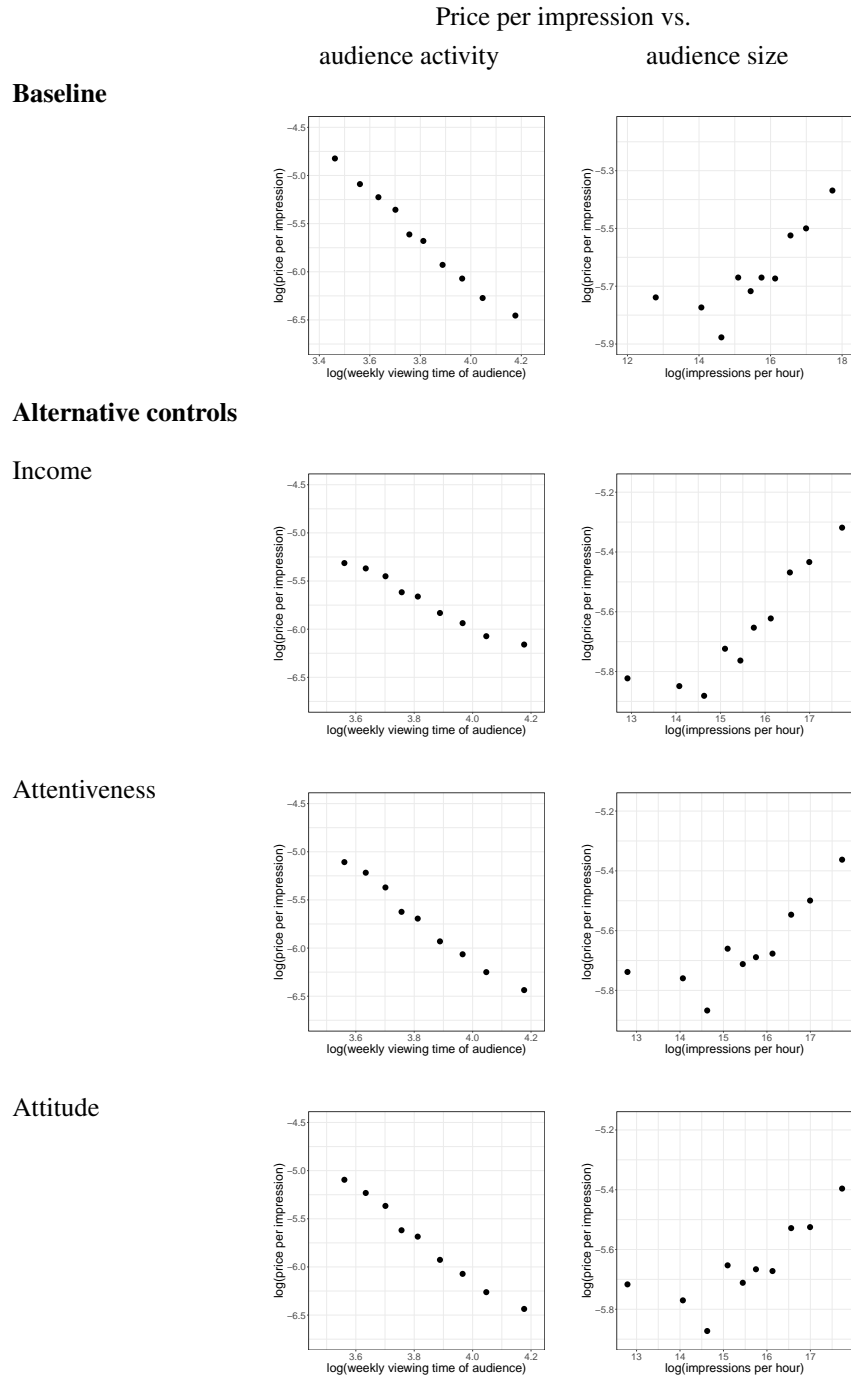


2016



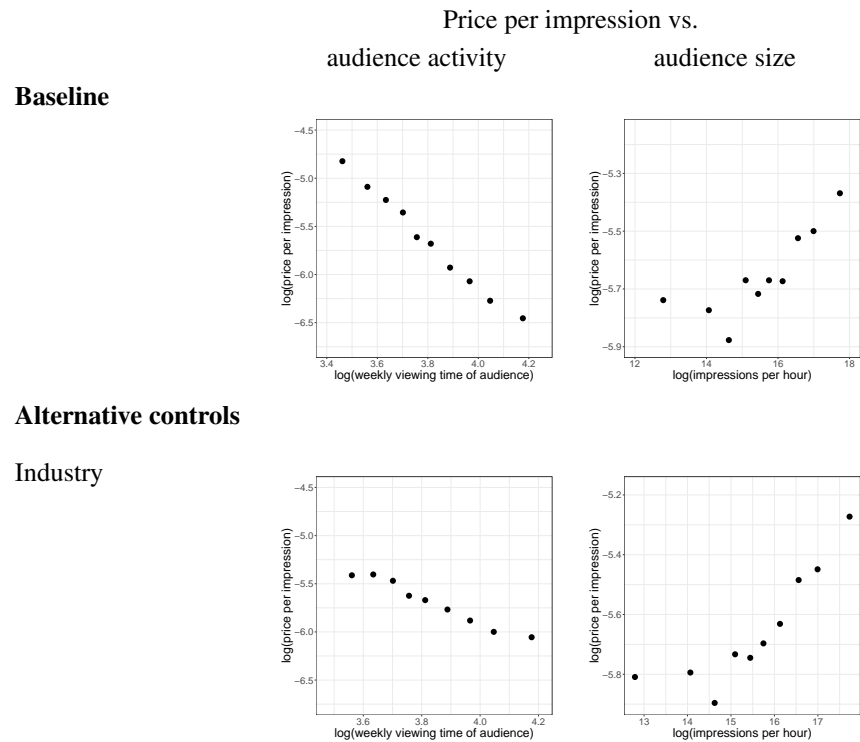
Notes: Within a given row, both plots are based on the same regression specification. The row labeled “Baseline” corresponds to the main specification in the paper, with the plot “Price per impression vs. audience activity” corresponding to Panel B of Figure 1 and the plot “Price per impression vs. audience size” corresponding to Panel B of Figure 3. The rows under the header “Alternative Sample Years” present results for alternative sample years.

Appendix Figure 3: Sensitivity to Alternative Controls



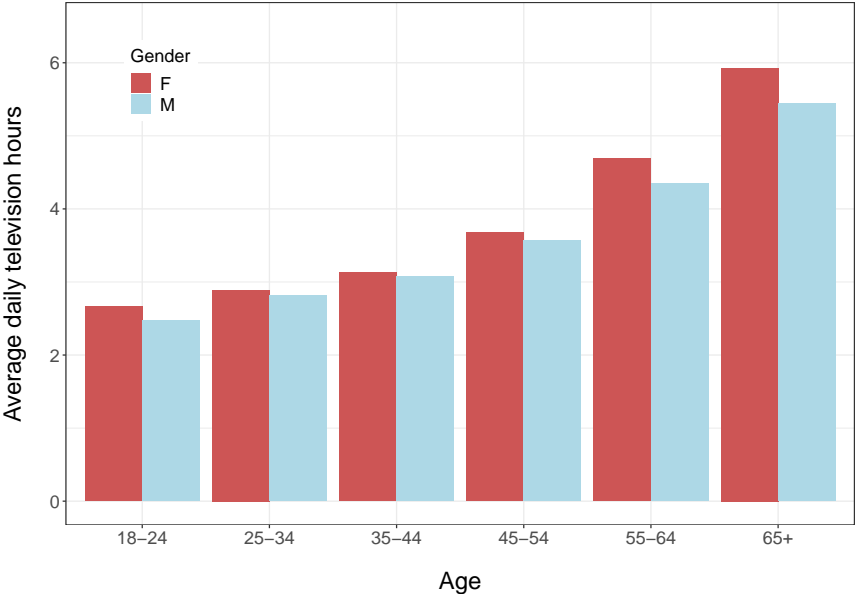
Notes: Within a given row, both plots are based on the same regression specification. The row labeled “Baseline” corresponds to the main specification in the paper, with the plot “Price per impression vs. audience activity” corresponding to Panel B of Figure 1 and the plot “Price per impression vs. audience size” corresponding to Panel B of Figure 3. 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 1 and the plot “Price per impression vs. audience size” corresponding to Panel B of Figure 3. 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: Measures of Online Activity by Age and Gender

Panel A: Average internet hours per day




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




Notes: Panel A shows average daily hours spent on the internet across age groups by gender. Panel B 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.

Appendix Figure 6: **Demographic Premia (Per Click) and Viewing Time on Facebook**

Panel A: Data from our experiment

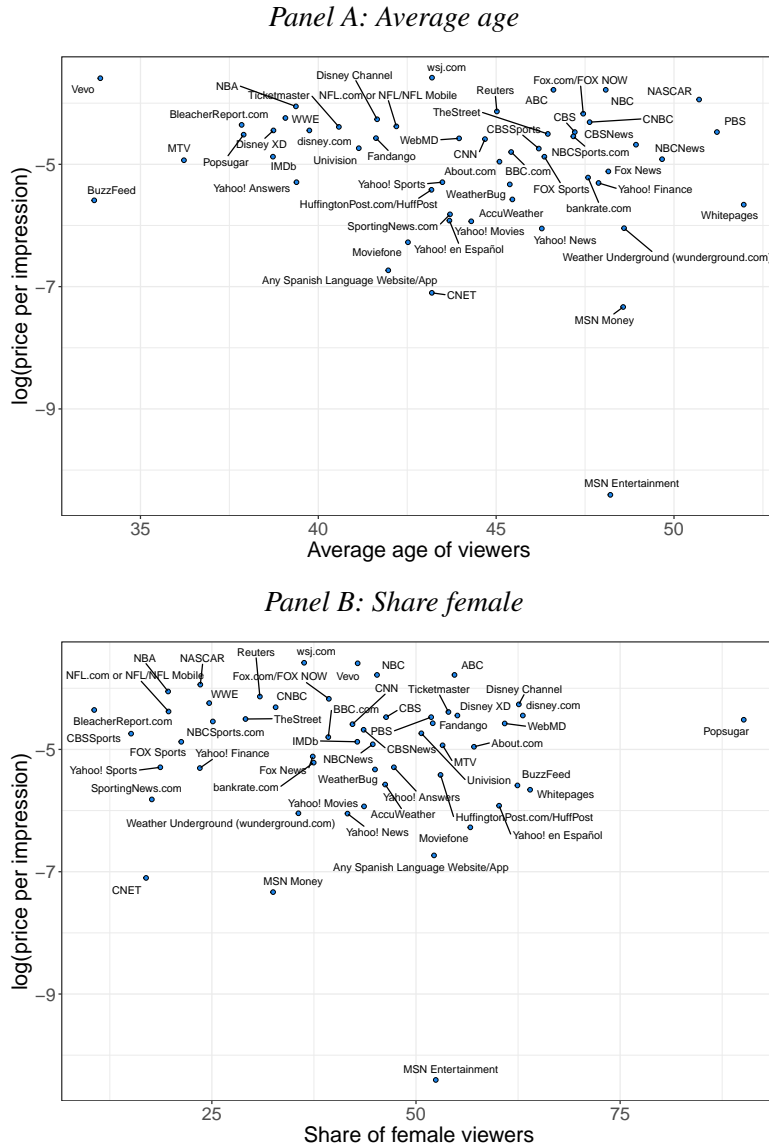


Panel B: Data from Allcott et al. (2020b)



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 7: Advertising Prices and Demographics of Digital Platforms



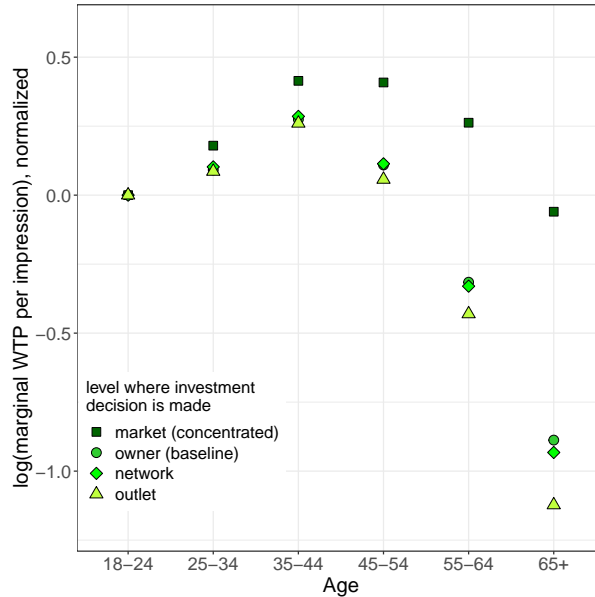
Notes: Each plot is a scatterplot of the $\log(\text{price per impression})$ of display advertising on a platform against the demographic characteristics of the platform's viewers. We construct the price per impression by computing the ratio of total revenue to total impressions across all display ads 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 8: **Marginal Willingness to Pay to Attract Older vs. Younger Viewer Under Alternative Levels of Investment Decision-Making**

Panel A: Baseline model with homogeneous value

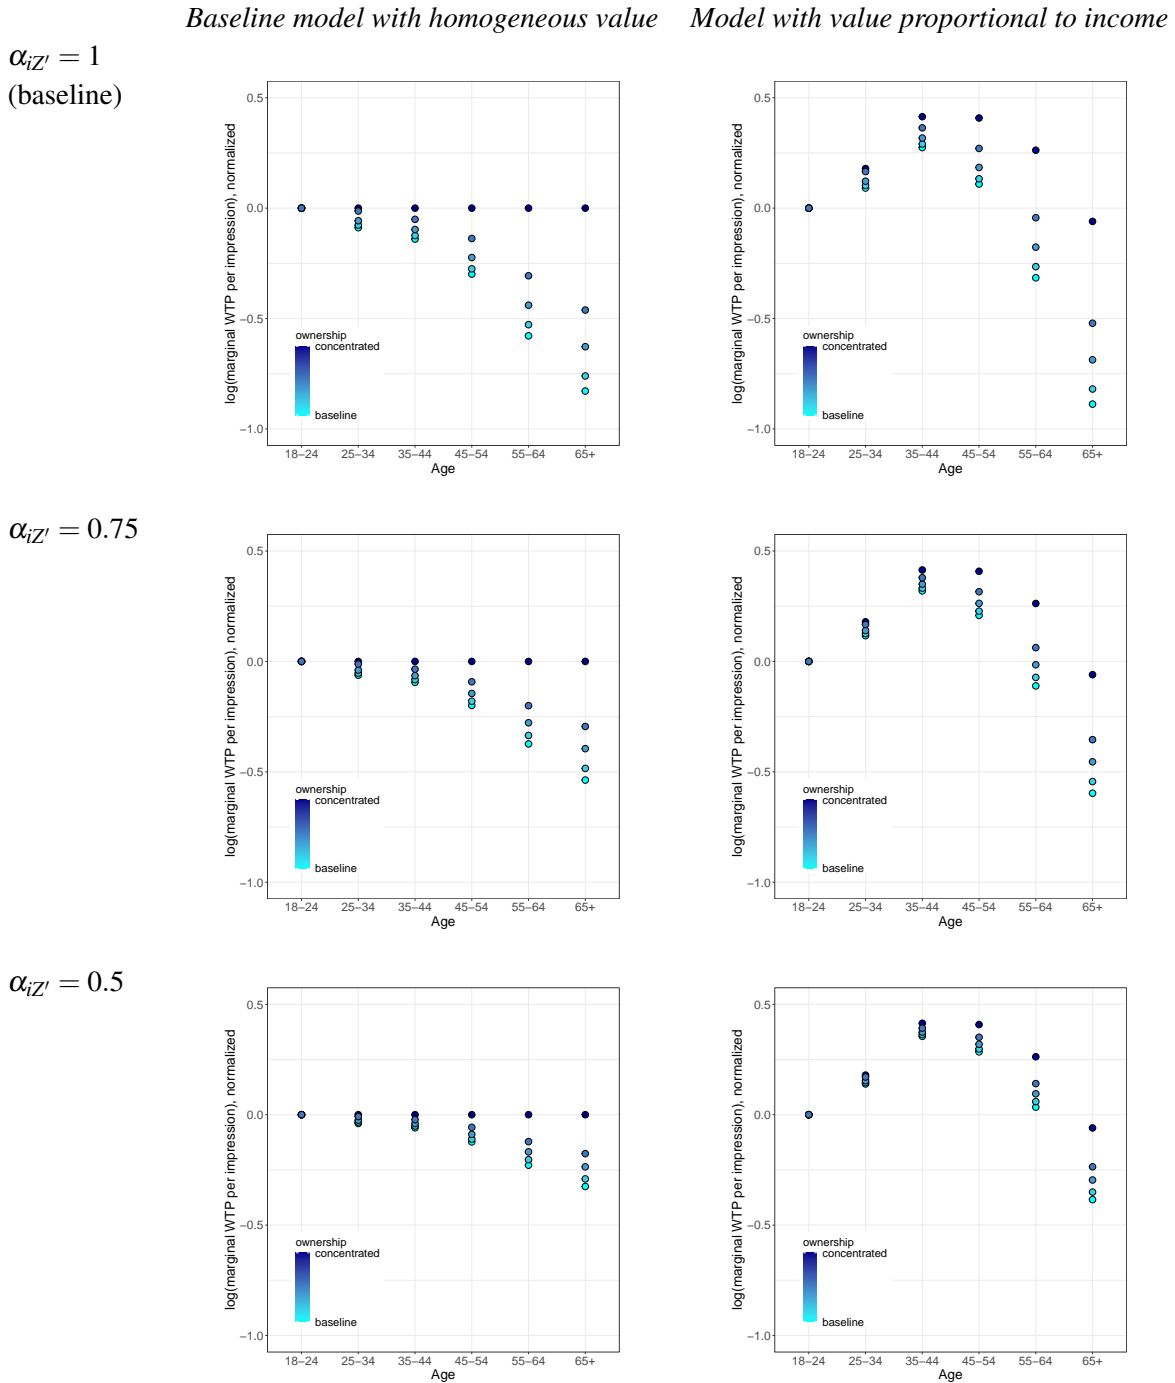


Panel B: Model with value proportional to income



Notes: In each plot, the y-axis value corresponds to the average $\log(\text{total marginal willingness to pay per impression})$, $\ln(\hat{m}_i)$, as described in Section 5.2, for viewers in the age bin listed on the x-axis, under models in which the content investment decision is made at different levels. 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, "market (concentrated)" corresponds to the model in which investment is made by a single entity owning all television networks, "owner (baseline)" corresponds to the model in which investment is made by each owner of the observed partition \mathcal{Z} , "network" corresponds to the model in which investment is made by each individual network, and "outlet" corresponds to the model in which investment is made by each individual outlet (network-daypart). In both plots, and for each level of investment, the y-axis value is normalized by adding a scalar so that its average value is zero in the youngest age group.

Appendix Figure 9: **Marginal Willingness to Pay to Attract Older vs. Younger Viewers Under Alternative Values of $\alpha_{iZ'}$**




Notes: The plots show the calculations described in Section 5.2 when using different values of $\alpha_{iZ'}$ for $Z' \neq Z$ when calculating \hat{m}_i . Each row corresponds to a different value of $\alpha_{iZ'}$, with the first row corresponding to the baseline case of $\alpha_{iZ'} = 1$ depicted in in Figure 8. The left column uses the baseline model in which advertisers' value of a first impression a_i is homogeneous across viewers, whereas the right column uses the model in which a_i is proportional to a viewer i 's income. In each plot, the y-axis value corresponds to the average log(total marginal willingness to pay per impression), $\ln(\hat{m}_i)$, 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. In all plots, and for each ownership scenario, the y-axis value is normalized by adding a scalar so that its average value is zero in the youngest age group.

Appendix Table 1: **Overlap in Advertising Spending, TV vs. Online**

	Share of TV advertising	Share of online advertising
Top 50 TV advertisers	44	17
Top 50 online advertisers	27	36

Notes: Each column shows the share of advertising spending on the given medium (TV or online) in 2015 coming from the top 50 television advertisers by spending and the top 50 online advertisers by spending. Each advertiser is a parent company that can advertise for multiple brands. Online advertising includes non-mobile display advertising. Television advertising includes cable and broadcast. These values were computed using data from Media Intelligence and Kantar Media (2021).

Appendix Figure 10: **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.