

Pricing Power in Advertising Markets: Theory and Evidence[†]

By MATTHEW GENTZKOW, JESSE M. SHAPIRO, FRANK YANG,
AND ALI YURUKOGLU*

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 the model is a good match, qualitatively and quantitatively, to variation in advertising prices across demographic groups, outlets, platforms, and over time. We use the model to quantify the effects of competition within and across platforms. (JEL G34, K21, L13, L82, M37)

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

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

[†]Go to <https://doi.org/10.1257/aer.20220943> to visit the article page for additional materials and author disclosure statements.

Wilbur 2008; Veiga and Weyl 2016). Pricing in advertising markets has become an important issue in antitrust policy (e.g., Competition and Markets Authority 2019).

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

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

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

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

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

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

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

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

between 0.38 and 2.7 log points across television owners, with owners whose audience overlaps more with Netflix tending to experience larger declines in price per impression.

In the final part of the paper, we extend our analysis to social media. We specify a model in which ads may be sold and targeted at the viewer level on social media. The model allows diminishing returns to operate more strongly within television or within social media than between the two, such that ads on the two formats may be imperfect substitutes from the perspective of advertisers. When cross-format substitutability is limited, the model predicts that the age-price relationship on social media should be reversed relative to television, with the young—the heaviest users of social media—commanding the lowest prices and the old—the lightest users—being the “coveted” group. Alternative explanations for the youth premium, such as the young having more malleable preferences (Surowiecki 2002), would not all share this prediction. Using data on prices of Facebook advertisements collected as part of a series of experiments including our own, we show that ads targeting the oldest users indeed command the highest prices. We compare model fit as we vary a parameter governing cross-format diminishing returns and find that the best-fitting parameter implies that television and social media ads are imperfect substitutes.

The primary contribution of this paper is to show that the predictions of a model of a competitive advertising market with multi-homing are a good match, both qualitatively and quantitatively, to existing and novel facts about important real-world markets. In contrast to many prior studies of advertising markets (e.g., Kaiser and Wright 2006; Wilbur 2008; Bel and Domènech 2009; 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

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 they do not study the cross-sectional variation in advertising prices implied by their model. Shi (2016, sec. 3.2) proposes an inverse demand function for advertising that is motivated by a model with multi-homing consumers and diminishing returns to impressions. Greenwood, Ma, and Yorukoglu (2021) calibrate a macroeconomic model in which consumers can consume multiple media goods but do not receive multiple advertisements from the same advertiser. Prat and Valletti (2022, sec. 4) simulate effects of platform mergers under various assumptions about overlap in their audience, though using a microfoundation different from ours.

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

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

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; Strikesocial 2017), we are not aware of prior evidence in the academic literature showing that transaction prices in the United States are greater for Facebook ads targeted to older users.⁵ Turning to time trends, the fact that television advertising revenues have grown despite a declining audience has been noted as a puzzle but is predicted by our quantitative model.⁶

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

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

The remainder of the paper proceeds as follows. Section I presents our model and its implications. Section II describes our data and variable definitions. Section III presents our key findings about the determinants of advertising prices on television. Section IV presents our quantitative implementation of the model, its fit to the

⁵Lambrecht and Tucker's (2019) 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 (Lambrecht and Tucker 2019, columns 1 and 2, Table 7) but does show evidence of interactions between age and gender (Lambrecht and Tucker 2019, column 3, Table 7). Our analysis differs in using transaction price data from campaigns in the United States rather than suggested bid data from a campaign across 191 countries.

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

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

data, and its applications. Section V presents our extension to social media. Section VI concludes.

I. Model

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

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

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

$$(1) \quad V(\{j : j \in B \in \mathcal{B}_n\}) - \sum_{B \in \mathcal{B}_n} p_B,$$

where $V(\cdot)$ is a nonnegative *value function* that is monotone in the set-inclusion order.⁸ We capture the idea that there are diminishing returns to advertising by assuming that $V(\cdot)$ is submodular: an advertiser derives less incremental value from an outlet when adding it to a larger bundle.⁹ We assume that the value function $V(\cdot)$ is the same across advertisers; i.e., the advertisers are homogeneous.

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

$$(2) \quad v_B = V(\mathcal{J}) - V(\mathcal{J} \setminus B),$$

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

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

$$\mathcal{J}' \subseteq \mathcal{J}'' \Rightarrow v(\mathcal{J}') \leq v(\mathcal{J}'').$$

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

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

value, $v_j > 0$ for all $j \in \mathcal{J}$. We use subgame perfect Nash equilibrium in pure strategies as our solution concept and hereafter refer to it as *equilibrium*.

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

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

$$(3) \quad V(S \cup Z) - V(S) \geq V(\mathcal{J}) - V(\mathcal{J} \setminus Z) = v_Z,$$

where the inequality is due to the submodularity of $V(\cdot)$. Thus, in equilibrium, owner Z must secure a profit of at least $N \cdot v_Z$. Now, suppose toward contradiction that some owner Z earns a profit strictly higher than $N \cdot v_Z$. Then, some advertiser n must pay owner Z more than v_Z . Thus, there exists an outlet $j \in \mathcal{J} \setminus Z$ on which advertiser n does not purchase an ad slot because otherwise, by the monotonicity of $V(\cdot)$, the advertiser could profitably deviate by not trading with owner Z at all. But, also because of the monotonicity of $V(\cdot)$, this implies that the owner Z' of outlet j can profitably deviate by adding the ad slot on outlet j to the existing bundle that advertiser n purchases from owner Z' and slightly increasing the total price. The existence of a profitable deviation for owner Z' contradicts that the conjectured strategy is part of an equilibrium, thus implying that in equilibrium each owner must secure a profit of no more than $N \cdot v_Z$ and hence, that incremental pricing holds.

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

A. Special Cases

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

DEFINITION 1 (Viewer-level model): *In the viewer-level model, each viewer $i \in \mathcal{I}$ is endowed with an intrinsic value a_i to advertisers. Each viewer i sees ads on outlet j with probability $\eta_{ij} \in [0, 1]$, independently across outlets. Owners are partitioned*

into formats F , with \mathcal{F} the set of all formats. An advertiser who reaches a given viewer i exactly $M \in \mathbb{N}^{|\mathcal{F}|}$ times in each format obtains value $a_i \cdot u(M)$, where $u(\cdot)$ is monotone and submodular in M and has decreasing differences in each argument with $u(0) = 0$.¹⁰ The reach-only model is a special case of the viewer-level model in which there is a single format and $u(M) = \mathbf{1}_{M>0}$.

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

In the viewer-level model, each viewer i is endowed with an intrinsic value a_i that might reflect, for example, differences in income across households. Each viewer is also endowed with viewing behavior represented by the probability η_{ij} of seeing ads on outlet j . Outlets are partitioned into formats, for example, television and social media. Advertisers' value for ads exhibits diminishing returns, with initial impressions being more valuable than subsequent ones. Diminishing returns may operate more strongly within than across formats. In the reach-only model, there is a single format and only first impressions matter. Theorem 1 implies the following characterization of revenues in the reach-only model.

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

$$(4) \quad p_Z^* = \sum_{i \in \mathcal{I}} a_i \left[\eta_{iZ} \prod_{Z' \neq Z} (1 - \eta_{iZ'}) \right]$$

is the revenue of owner Z ,

$$(5) \quad \lambda_Z = \sum_{i \in \mathcal{I}} \eta_{iZ}$$

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

$$(6) \quad \eta_{iZ} = 1 - \prod_{j \in \mathcal{Z}} (1 - \eta_{ij})$$

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

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

B. Comparative Statics

We provide two comparative statics results in the reach-only model. For these comparative statics, we suppose that all viewers are intrinsically equally valuable; that is, that $a_i = a > 0$ for all i . We also suppose that viewers can be partitioned into a set G of mutually exclusive demographic groups g , with viewing behavior

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

homogeneous within groups, i.e., with $\eta_{iZ} = \eta_{i'Z} = \eta_{gZ}$ for all Z for any two viewers i, i' in the same group g . We define a group $g \in G$ to be *less active* than group $h \in G$ if $\eta_{gZ} \leq \eta_{hZ}$ for all $Z \in \mathcal{Z}$. We let

$$(7) \quad \sigma_{gZ} = \frac{\sum_{i \in g} \eta_{gZ}}{\lambda_Z}$$

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

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

PROPOSITION 1: *Suppose that owner Y draws a larger share of its audience from the less active group g and a smaller share of its audience from the more active group h than owner Z , in the sense that $\sigma_{gY} \geq \sigma_{gZ}$ and $\sigma_{hY} \leq \sigma_{hZ}$, and that the two owners have equal total audience sizes, $\lambda_Y = \lambda_Z$, and equal shares of audience from groups other than g and h , $\sigma_{g'Y} = \sigma_{g'Z}$ for all $g' \neq g, h$. Then owner Y has a higher equilibrium price per viewer than owner Z , $p_Y^*/\lambda_Y \geq p_Z^*/\lambda_Z$.*

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

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

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

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

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

The inequality in the conclusion of Proposition 2 is strict if $\delta > 1$ and $\eta_{gZ} \in (0, 1)$ for all Z . To see the intuition for Proposition 2, consider the pair of owners Y and Z and suppose that there is only one group of viewers. If a viewer sees ads on an outlet of owner $Z' \notin \{Y, Z\}$, then by Corollary 1, the viewer does not contribute to the revenue of owners Y or Z . Among the remaining viewers, the share

of those watching Y 's outlets who also watch Z 's outlets is smaller than the reverse (since for any viewer the probability of watching Y is greater than the probability of watching Z). This implies that Y has a larger share of exclusive viewers and so commands a higher price per viewer.

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

C. Illustrative Example

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

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

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

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

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

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

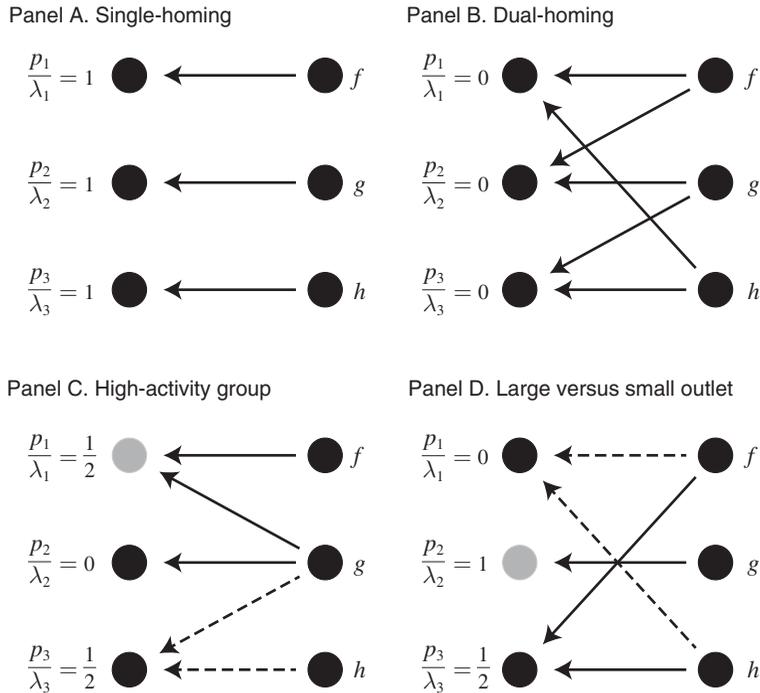


FIGURE 1. ILLUSTRATIVE EXAMPLE

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

D. Discussion

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

Alternative Market Institutions.—In the model, we assume that owners simultaneously post prices. In practice, ad sales can take different forms, including bargaining and auctioning. Online Appendix A.3 shows that incremental pricing holds if

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

owners bargain with advertisers a la Nash-in-Nash (Lee, Whinston, and Yurukoglu 2021) or compete via reserve prices in first-price auctions.

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

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

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

Ad Inventory and Endogenous Response of Viewers.—In the viewer-level model, viewers' choices of which outlets to view are not affected by the outcome of the game. We may interpret this 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. Online Appendix A.3 shows that our model can accommodate settings in which viewers care about advertising and make viewing decisions knowing the outcome of the advertising game provided that the impact of advertising on viewing is additively separable across ad slots. When this assumption is violated, an owner might profitably withhold some ad slots to increase the viewership of the remaining ad slots.

II. Data

We conduct our main analysis of the television market using data from 2015. Here, we describe the concepts we measure for 2015. For sensitivity analysis and extensions, we use data from 2014 through 2019, with concepts defined and

calculated in an analogous manner to those we describe below for 2015. Online Appendix Figure 1 reports sensitivity analysis replacing data from 2015 with data from 2014 or 2016.

A. 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 include the telecast (e.g., NBC Nightly News, June 1), program (e.g., NBC Nightly News), daypart (e.g., early fringe), and network (e.g., NBC). It also includes the duration (e.g., 30 seconds) of the advertising spot, an estimate of its cost, and an estimate of the number of impressions (live viewers) for the associated telecast. We omit from all calculations any advertisements with zero cost or duration. We standardize the cost to a 30-second-spot basis by dividing the cost by the duration of the advertisement (in seconds) and multiplying by 30.

Advertising cost estimates in the AdIntel data are based primarily on information obtained from SQAD at the month-network-daypart level for cable television and from networks at the month-program level for broadcast television (The Nielsen Company 2017).¹³ For consistency, we define our notion of an outlet j to be a network-daypart.¹⁴ Online Appendix Figure 2 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 may think of as an outlet-level analog λ_j of the concept λ_Z defined in Section IB. For each outlet, we also calculate the total (standardized) cost of all advertisements and divide by the number of hours in the corresponding daypart in a 52-week year to get a measure of total cost per hour, which we may think of as an outlet-level analog p_j^* of the concept p_Z^* defined in Section IB. Finally, for each outlet, we divide total cost per hour by total impressions per hour to obtain the average price per impression of a 30-second spot on the outlet, which we may think of as an outlet-level analog p_j^*/λ_j of the concept p_Z^*/λ_Z defined in Section IB.

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

For a subset of advertisements representing 99.9 percent of all impressions, we also have information on the distribution of impressions across household income

¹³Hristakeva and Mortimer (2023) use data from SQAD to analyze price dispersion and discounting in the television advertising market.

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

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

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

bins for the associated program.¹⁷ From this information, we compute the average household income of each outlet's adult impressions (among those for which we measure income) by imputing each bin to its midpoint value and imputing the highest-income bin (\$125,000+) to \$175,000.

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

B. Audience Survey

From GfK MRI's 2015 Survey of the American Consumer, we obtain, for each of 23,978 adult respondents, information on times of day spent watching television in the form of a week-long diary, as well as the implied total weekly television viewing time (GfK Mediamark Research and Intelligence 2017). We compute a measure of each respondent's total viewing time in each daypart by allocating viewing time in each time slot to AdIntel dayparts in proportion to the share of the time slot that is contained within each daypart. We also obtain measures of each respondent's viewership of each of 227 broadcast television programs¹⁸ and time spent watching each of 115 cable television networks in the preceding week. We successfully match 173 broadcast programs and 97 cable television networks to their counterparts in AdIntel.¹⁹

We use the data on viewership by daypart, broadcast program, and cable network to construct a measure of the time that each respondent i viewed each outlet (network-daypart) j . To do this, we first allocate the viewing time of broadcast programs to their respective network-dayparts.²⁰ If in a given daypart there is viewing

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

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

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

²⁰Specifically, we associate each program with a network-daypart following the 2014–2015 United States Network Television Schedule (Wikipedia 2022). We supplement this source with information from the *Sunday News Journal* (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 (Google 2021). 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).

time that cannot be attributed to broadcast programs, we allocate that time to the cable networks in proportion to the respondent's reported viewing time of each network.²¹

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

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

III. Evidence on the Determinants of Advertising Prices

Proposition 1 predicts that more active audiences will command a lower advertising price per impression. Figure 2 shows that this prediction is borne out in a comparison of television outlets. Each panel shows a binned scatterplot of an outlet's $\log(\text{price per impression})$ 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 2 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. Online Appendix Figure 4 shows that older viewers watch more television than younger viewers and that female viewers watch more television than male viewers. The logic of Proposition 1 would lead us to expect that outlets with an older audience would command a lower advertising price than outlets with a younger audience, and likewise for outlets with a more female audience. Figure 3 shows that these predictions are borne out in the data: outlets with older, more female audiences tend to exhibit both lower $\log(\text{price}$

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

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

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

²⁴The household income bins are 0–4,999, 5,000–9,999, 10,000–14,999, 15,000–19,999, 20,000–24,999, 25,000–29,999, 30,000–34,999, 35,000–39,999, 40,000–44,999, 45,000–49,999, 50,000–59,999, 60,000–74,999, 75,000–99,999, 100,000–149,999, 150,000–199,999, 200,000–249,999, and 250,000+ US dollars. We impute each household's income to the midpoint of the corresponding bin, imputing the highest bin to \$300,000.

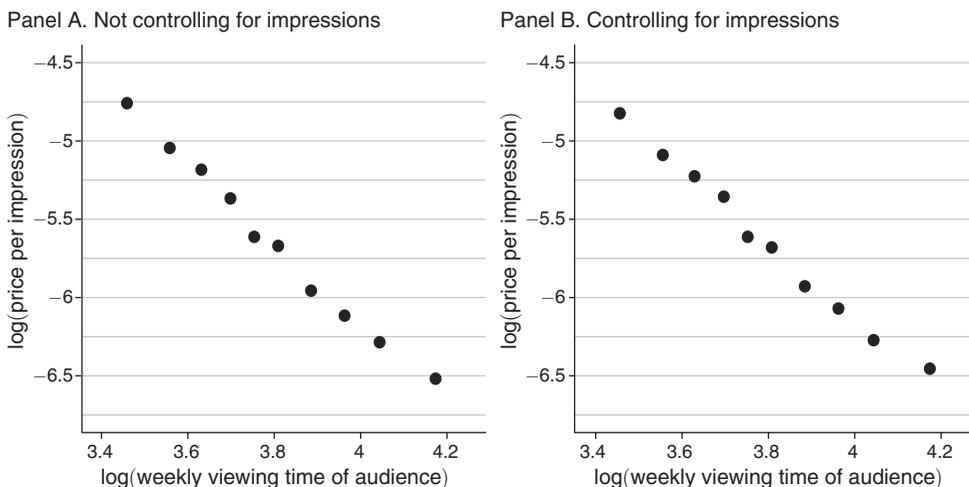


FIGURE 2. ADVERTISING PRICES AND AUDIENCE ACTIVITY LEVELS OF TELEVISION OUTLETS

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

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

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

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

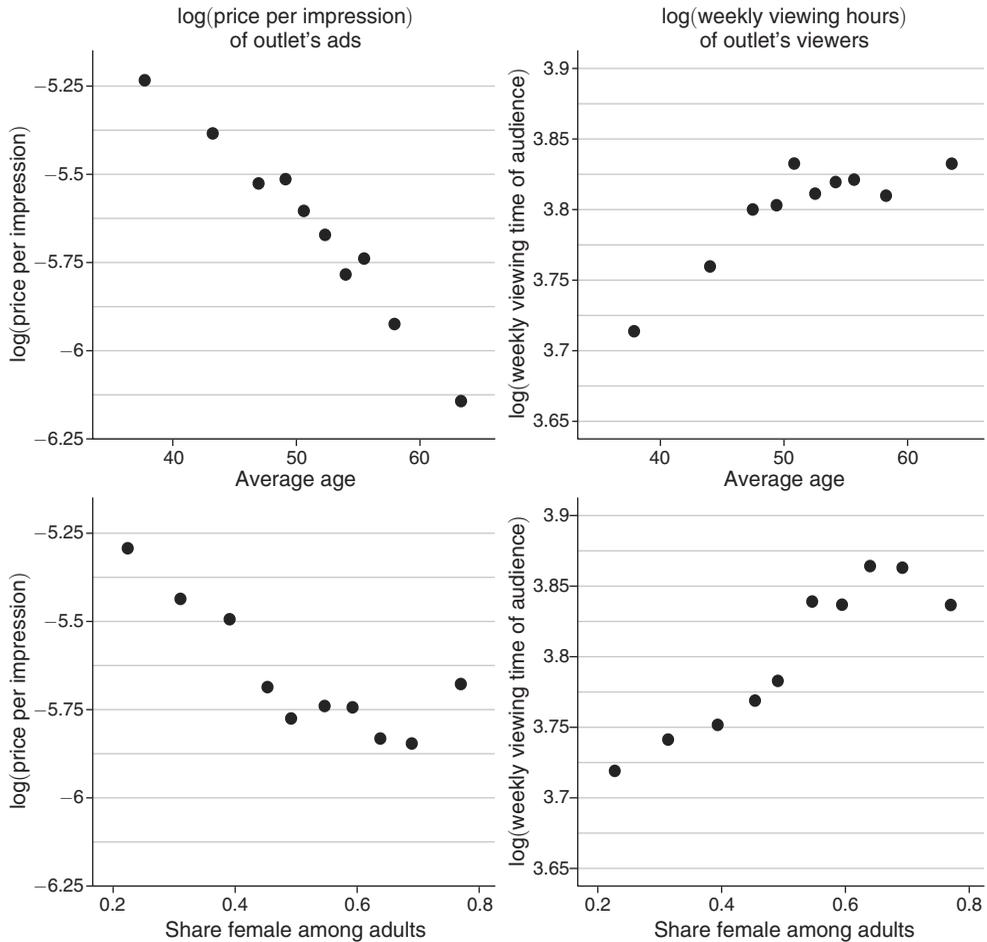


FIGURE 3. ADVERTISING PRICES AND ACTIVITY LEVELS BY AUDIENCE DEMOGRAPHICS OF TELEVISION OUTLETS

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

IV. Quantification and Applications of the Model

A. Implementation and Model Fit

Corollary 1 shows that, in the reach-only model, it is possible to calculate the equilibrium price per viewer p_Z^*/λ_Z given data on the probability η_{ij} that each

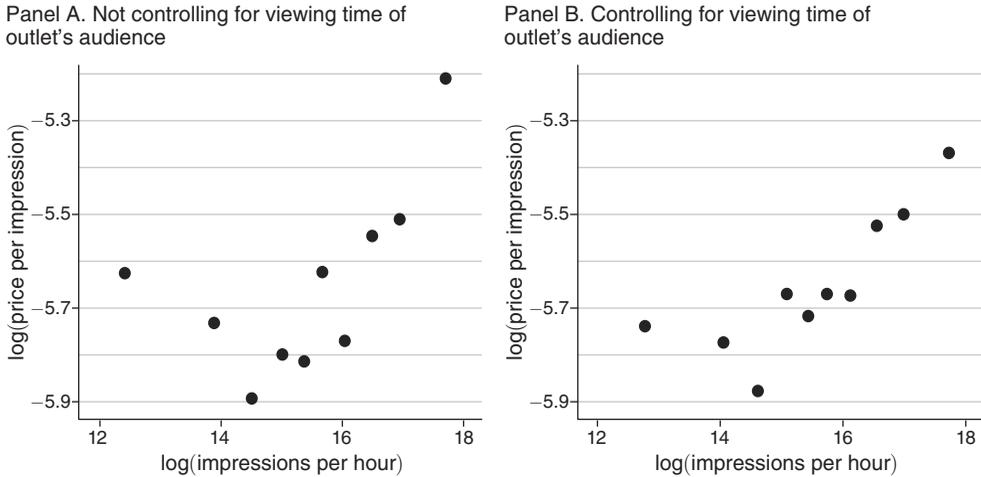


FIGURE 4. ADVERTISING PRICES AND AUDIENCE SIZE OF TELEVISION OUTLETS

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 $\log(\text{price per impression})$ of a 30-second spot on the outlet; the independent variable of interest is the $\log(\text{impressions per hour})$ 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 $\log(\text{weekly viewing time})$ of the outlet's viewers.

TABLE 1—ADVERTISING PRICES, AUDIENCE DEMOGRAPHICS, AND AUDIENCE ACTIVITY LEVELS OF TELEVISION OUTLETS

Dependent variable:	Observed	
	$\log(\text{price per impression})$	
	(1)	(2)
Average $\log(\text{weekly viewing hours})$ of audience	-1.5556 (0.2913)	
Average age of impressions		-0.0285 (0.0079)
Share female among adult impressions		-0.4690 (0.2599)
$\log(\text{impressions per hour})$	0.0973 (0.0292)	0.1221 (0.0306)
Average household income of impressions (\$1,000)	0.0124 (0.0031)	0.0152 (0.0034)
Number of networks	103	103
Number of network-dayparts	809	809

Notes: Each column reports estimates of a linear regression. The unit of analysis is an outlet (network-daypart). The dependent variable is the $\log(\text{price per impression})$ of a 30-second spot observed in the data, as described in Section IIA. Both 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.

viewer i sees ads on each outlet j , the intrinsic value a_i of each viewer i , and the ownership partition \mathcal{Z} . We operationalize this calculation in the audience survey data, letting the viewers \mathcal{I} be the set of respondents. For each viewer $i \in \mathcal{I}$, we take $\eta_{ij} = T_{ij}/T_j$, where T_j is outlet j 's total broadcast time and we recall that T_{ij} is the time that viewer i spends viewing outlet j . We implement two specifications, a specification in which $a_i = a$ for all i and a specification in which a_i is proportional to household income for all i . We treat the value $(1/|\mathcal{I}|) \sum_{i \in \mathcal{I}} a_i$ 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.

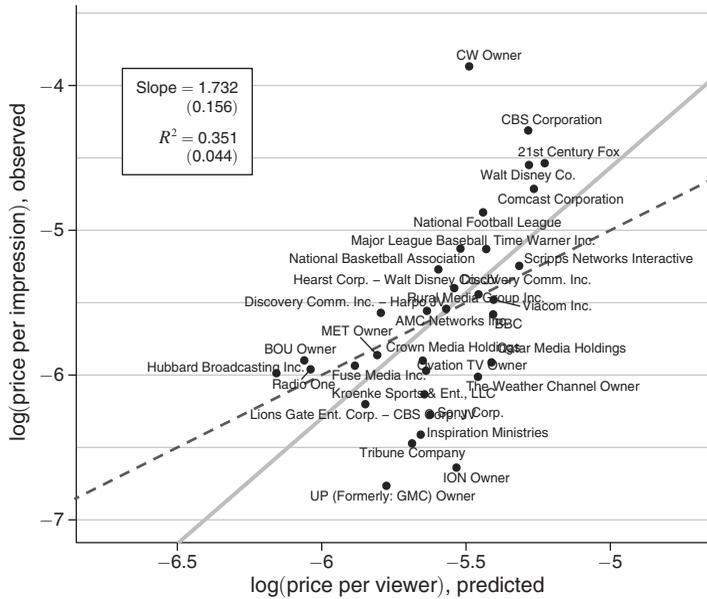
Figure 5 shows that the model does a good job predicting relative advertising prices across owners of television networks. Panel A shows a scatterplot of the observed $\log(\text{price per impression})$ against the predicted $\log(\text{price per viewer})$ for the specification with homogeneous values. Predicted $\log(\text{price per viewer})$ explains 35.1 percent of the variation in observed $\log(\text{price per impression})$. 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 close to 1. Across both specifications, the model is able to rationalize large differences in advertising prices between owners, as illustrated in the x-axes of the scatterplots.

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

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

We can also evaluate the fit of the model to trends in television advertising revenues during recent years in which overall television viewership has fallen. Panel A of Figure 6 shows that annual revenues increased slightly between 2014 and 2019 even as total impressions fell, a pattern that some have regarded as puzzling (*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

Panel A. Baseline model with homogeneous value



Panel B. Model with value proportional to income

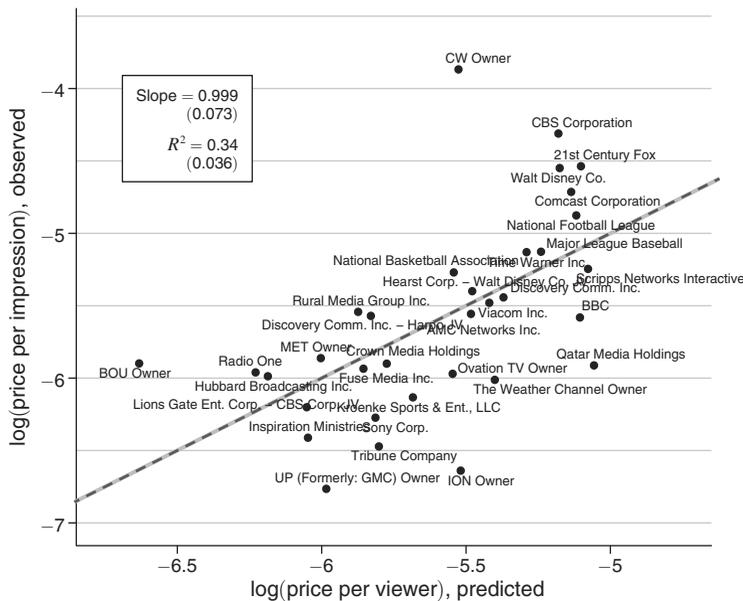


FIGURE 5. OBSERVED AND PREDICTED TELEVISION ADVERTISING PRICES

Notes: Each plot is a scatterplot of the $\log(\text{price per impression})$ of a 30-second spot observed in the data (y-axis), as described in Section IIA, against the $\log(\text{price per viewer})$ predicted by the model (x-axis), as described in Section IV. Panel A uses $\log(\text{price per viewer})$ predicted from the baseline model in which advertisers' value of a first impression is homogeneous across viewers. Panel B uses $\log(\text{price per viewer})$ predicted from the model in which advertisers' value of a first impression is proportional to a viewer's income. The unit of analysis 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 $\log(\text{price per impression})$ observed in the data. The dashed line depicts a 45-degree line. The solid line depicts the line of best fit. The box reports the slope of the line of best fit and the R^2 of the associated linear model, with standard errors in parentheses obtained via a nonparametric bootstrap over survey respondents with 100 replicates.

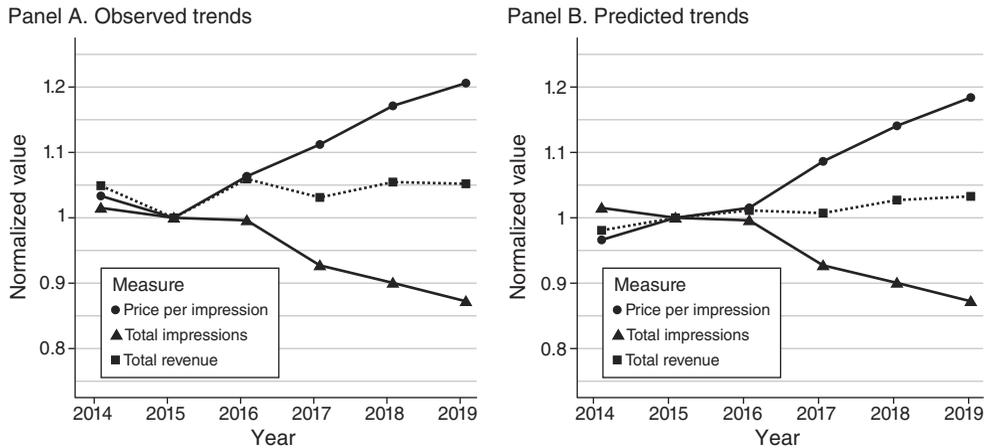


FIGURE 6. OBSERVED AND PREDICTED TELEVISION ADVERTISING REVENUES

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 IIA, 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 IV; the trend in impressions is identical to that in panel A; and the price per impression is the ratio of the two.

the model indeed predicts an increase in the price per impression, and in revenues, over this period. The patterns in panel B provide a reasonable qualitative and quantitative match to those in panel A, even though the revenue calculations underlying panel B are based only on audience survey data, and in particular do not use any information on advertising prices.²⁶ Competition is important for the findings in panel B of Figure 6: with a single monopoly owner of television networks, our model predicts declining, rather than increasing, advertising revenue over the period we study. Diminishing returns also appear to be important for the findings in panel B of Figure 6: online Appendix A.4 shows that two alternative models without diminishing returns can generate increasing prices but not increasing revenues.

B. Implications for Demographic Premia

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

$$(8) \quad p_{iZ}^* = a_i \eta_{iZ} \prod_{Z' \neq Z} (1 - \eta_{iZ'}).$$

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

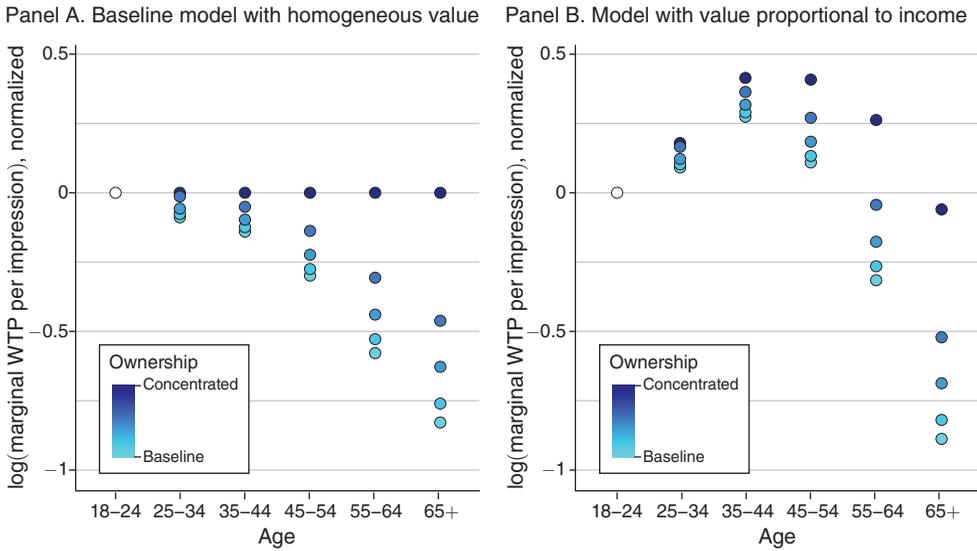


FIGURE 7. MARGINAL WILLINGNESS TO PAY TO ATTRACT OLDER VERSUS YOUNGER VIEWERS

Notes: In each plot, the y-axis value corresponds to the estimated average $\log(\text{total marginal willingness to pay per impression})$, $\ln(m_i)$, as described in Section IVB, 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.

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

$$(9) \quad m_i = \frac{\sum_{Z \in \mathcal{Z}} p_{iZ}^*}{\sum_{Z \in \mathcal{Z}} \eta_{iZ}}$$

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

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

TABLE 2—PREDICTED ADVERTISING PRICES, AUDIENCE DEMOGRAPHICS, AND AUDIENCE ACTIVITY LEVELS OF TELEVISION OUTLETS

Dependent variable:	Predicted $\log(\text{price per viewer})$			
	Homogeneous value		Value proportional to income	
	(1)	(2)	(3)	(4)
<i>Average log(weekly viewing hours) of audience</i>	-1.6799 (0.0607)		-1.8388 (0.1027)	
<i>Average age of impressions</i>		-0.0028 (0.0024)		-0.0020 (0.0029)
<i>Share female among adult impressions</i>		-0.3056 (0.0933)		-0.5230 (0.1228)
$\log(\text{impressions per hour})$	0.0082 (0.0044)	0.0418 (0.0109)	0.0198 (0.0075)	0.0628 (0.0125)
<i>Average household income of impressions (\$1,000)</i>	0.0002 (0.0004)	0.0057 (0.0016)	0.0102 (0.0008)	0.0152 (0.0018)
Number of networks	103	103	103	103
Number of network-dayparts	809	809	809	809

Notes: Each column reports estimates of a linear regression. The unit of analysis is an outlet (network-daypart). The dependent variable is the $\log(\text{price per viewer})$ predicted by the model, as described in Section IV. Columns 1 and 2 use $\log(\text{price per viewer})$ predicted from the baseline model in which advertisers' value of a first impression is homogeneous across viewers. Columns 3 and 4 use $\log(\text{price per viewer})$ 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.

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

We can also evaluate the fit of the model to the patterns we documented in Section III. 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 1 and 2 of Table 2 report estimates of the same regression models as in columns 1 and 2 of Table 1, replacing the observed $\log(\text{price per impression})$ with $\log(\text{price per viewer})$ $\ln(p_j^*/\lambda_j)$ 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 3 and 4 of Table 2 repeat the exercise in columns 1 and 2, with $\ln(p_j^*/\lambda_j)$ 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.

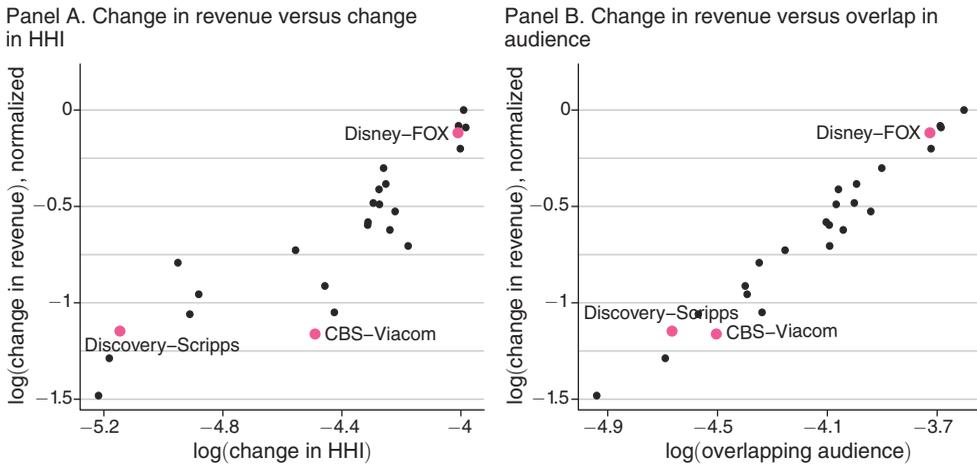


FIGURE 8. PREDICTED EFFECTS OF MERGERS ON ADVERTISING REVENUE

Notes: For each of a set of simulated mergers, panel A plots the log of the simulated change in revenue (y-axis) against the log of the simulated change in HHI (x-axis), and panel B plots the log of the simulated change in revenue (y-axis) against the log of the overlapping audience (x-axis). Larger, more lightly shaded circles indicate mergers that occurred after 2015; smaller, more darkly shaded circles indicate hypothetical mergers that have not occurred. We construct the plots as follows. For each owner, we compute the audience size as the probability of an average viewer seeing an ad on at least one of the owner's outlets, as described in Section IV. 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 IV. 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 premerger audience. We normalize the log of the simulated change in revenue by subtracting the value for the merger with the largest simulated change in revenue. For the mergers that took place, the normalized log of the simulated change in revenue is -1.15 for Discovery-Scripps ($SE = 0.04$), -1.16 for CBS-Viacom ($SE = 0.03$), and -0.12 for Disney-FOX ($SE = 0$), where standard errors in parentheses are obtained via a non-parametric bootstrap over survey respondents with 100 replicates.

C. Application to Mergers of Television Network Owners

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

Comparing panels A and B of Figure 8 shows that, according to the model, the revenue effects of a given merger are more strongly related to the overlap in

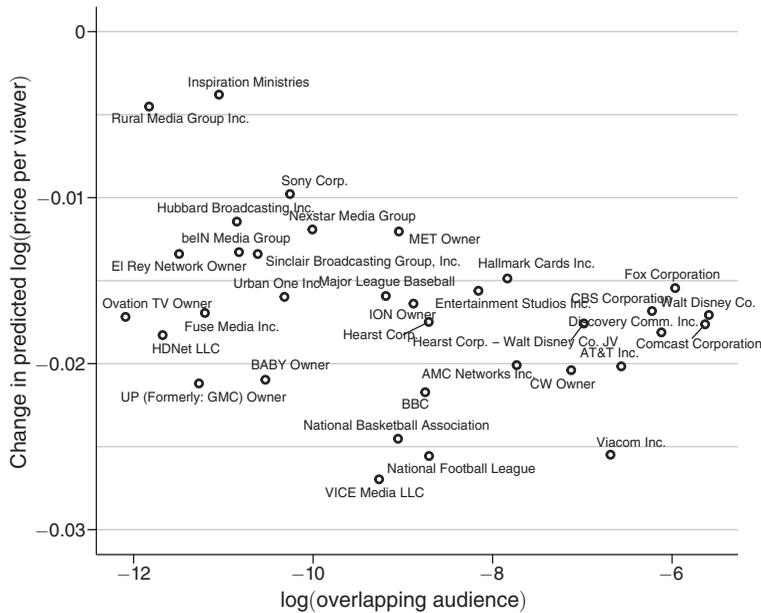


FIGURE 9. PREDICTED EFFECT OF NETFLIX ADVERTISING ON ADVERTISING PRICES

Notes: The plot shows a scatterplot, across owners of television networks, of the change in the owner's $\log(\text{price per viewer})$ 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 IIB 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 $\log(\text{price per viewer})$ implied by the baseline model in which advertisers' value of a first impression is homogeneous across viewers, as described in Section IV, 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).

audience between the merging entities than to the change in HHI induced by the merger. Among the three mergers that occurred, for example, the CBS-Viacom merger is roughly midway between the Discovery-Scripps merger and the Disney-Fox merger in terms of its log impact on HHI but is much closer to the Discovery-Scripps merger in terms of both the log of audience overlap and the log of the predicted revenue impact.²⁷

D. Application to Netflix Carrying Advertising

In 2022, Netflix launched a new service tier that carries advertising (Krouse 2022). Figure 9 visualizes the implications, under the model with homogeneous values, for television network owners of Netflix counterfactually adding advertising to its platform in 2019, assuming no change in audience behavior and that Netflix

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

shows advertisements across all of its content and subscribers. Across the owners, we estimate that Netflix ad carriage would reduce the price per viewer by between 0.38 and 2.7 log points, with a mean reduction of 1.68 (SE = 0.05).²⁸ As the plot illustrates, owners whose outlets have greater audience overlap with Netflix tend to experience greater proportional declines in price per viewer in this counterfactual, though there is substantial variation in the effect of Netflix for a given level of audience overlap, owing to variation across owners in the overlap of their outlets' audience with that of other owners. Our estimates imply that Netflix itself would have a relatively high price per impression—about 24.4 log points larger than the average of the 5 largest TV owners—consistent with its relatively young audience.

V. Extension to Social Media

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

A. Model

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

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

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

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

DEFINITION 2 (Cross-format reach-only model): *In the cross-format reach-only model, owners are partitioned into two formats, 1 and 2, corresponding to the sets $\tilde{\mathcal{Z}}$ and $\mathcal{Z} \setminus \tilde{\mathcal{Z}}$, respectively. Advertisers' valuations follow the viewer-level model where*

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

$u(M)$ depends on M only via $(\mathbf{1}_{M_1 > 0}, \mathbf{1}_{M_2 > 0})$ and $u(M)$ is invariant to exchanging the elements of M .

In the cross-format reach-only model, there are two formats, with owners in one format (e.g., social media) able to post viewer-specific prices and one format (e.g., television) unable to do so. An advertiser's value from a given viewer is determined by whether that viewer sees its ad at least once on each format, with no incremental return to multiple impressions on a given format. In this case, it is without loss of generality to write that

$$(10) \quad u(M_1, M_2) = \mathbf{1}_{M_1 > 0} + \mathbf{1}_{M_2 > 0} - \phi \mathbf{1}_{M_1 M_2 > 0},$$

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

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

$$(11) \quad p_{iZ}^* = a_i \eta_{iZ} \left\{ \left[1 - \phi \left(1 - \prod_{Z \in \tilde{\mathcal{Z}} \setminus F(Z)} (1 - \eta_{iZ}) \right) \right] \prod_{Z \in F(Z) \setminus \{Z\}} (1 - \eta_{iZ}) \right\},$$

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

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

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

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

advertisers do not view the two formats as serving identical functions.³¹ All of these facts suggest that $\phi < 1$.

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

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

B. Data

Social Media Advertising Prices.—We obtain data on the cost of advertising to different audiences on Facebook via an original experiment conducted for this study (Gentzkow et al. 2024) 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 $\{Men, Women\} \times \{13-17, 18-24, 25-34, 35-44, 45-54, 55-64, 65+\}$. Each advertisement set included just one advertisement, with fixed budgets of \$20 a day using automated cost-per-click bidding. For each advertisement set, we obtain the price per impression, which we may think of as a group-level average of the concept p_{iZ}^*/η_{iZ} defined in Section VA.

From Allcott et al. (2020b), we obtain data from 32 advertisement sets purchased on September 24, 2018: 4 each targeting each combination of gender and age group in $\{Men, Women\} \times \{18-24, 25-44, 45-64, 65+\}$. 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.

Audience Survey Data.—From the 2015 audience survey described in Section IIB, we obtain information on whether the respondent visited each of 5 social media sites—Facebook, Instagram, Reddit, Twitter, and YouTube—in the preceding 30 days. For each respondent i and social media site j , we let η_{ij} be an indicator for whether the respondent visited the site. We treat Facebook and Instagram as jointly owned and the other social media sites as independently owned and compute η_{iZ} following Corollary 1. We also obtain information on the amount of time each respondent spent using the Internet in an average week.³²

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

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

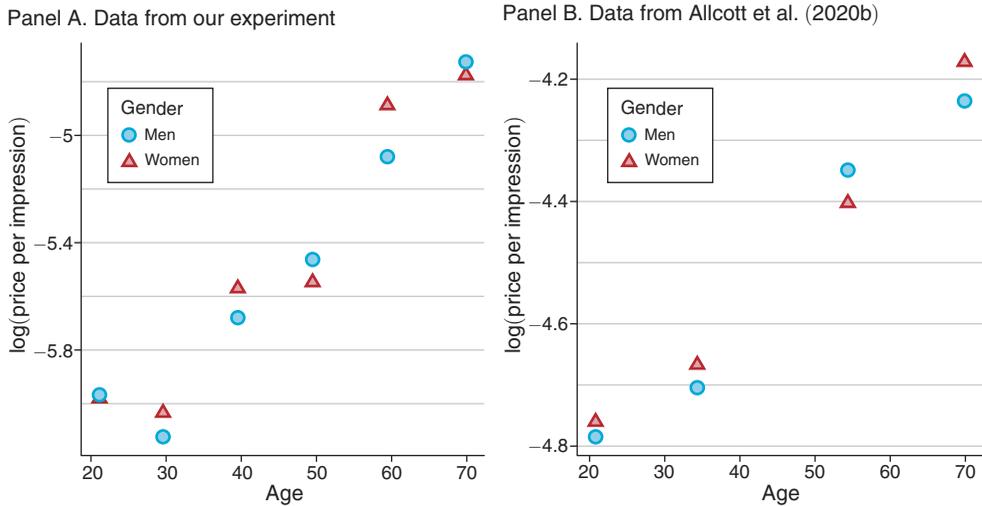


FIGURE 10. DEMOGRAPHIC PREMIUMS AND VIEWING TIME ON FACEBOOK

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 $\{\text{Men}, \text{Women}\} \times \{18\text{--}24, 25\text{--}34, 35\text{--}44, 45\text{--}54, 55\text{--}64, 65+\}$. In panel B, the data are taken from Allcott et al. (2020b), and the groups are $\{\text{Men}, \text{Women}\} \times \{18\text{--}24, 25\text{--}44, 45\text{--}64, 65+\}$. 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+.

C. Descriptive Evidence

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

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

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

or female audiences on social media. Online Appendix Figure 9 repeats the analysis in Figure 10 using data on price per click rather than price per impression.

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

D. Quantification

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

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

VI. Conclusions

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

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

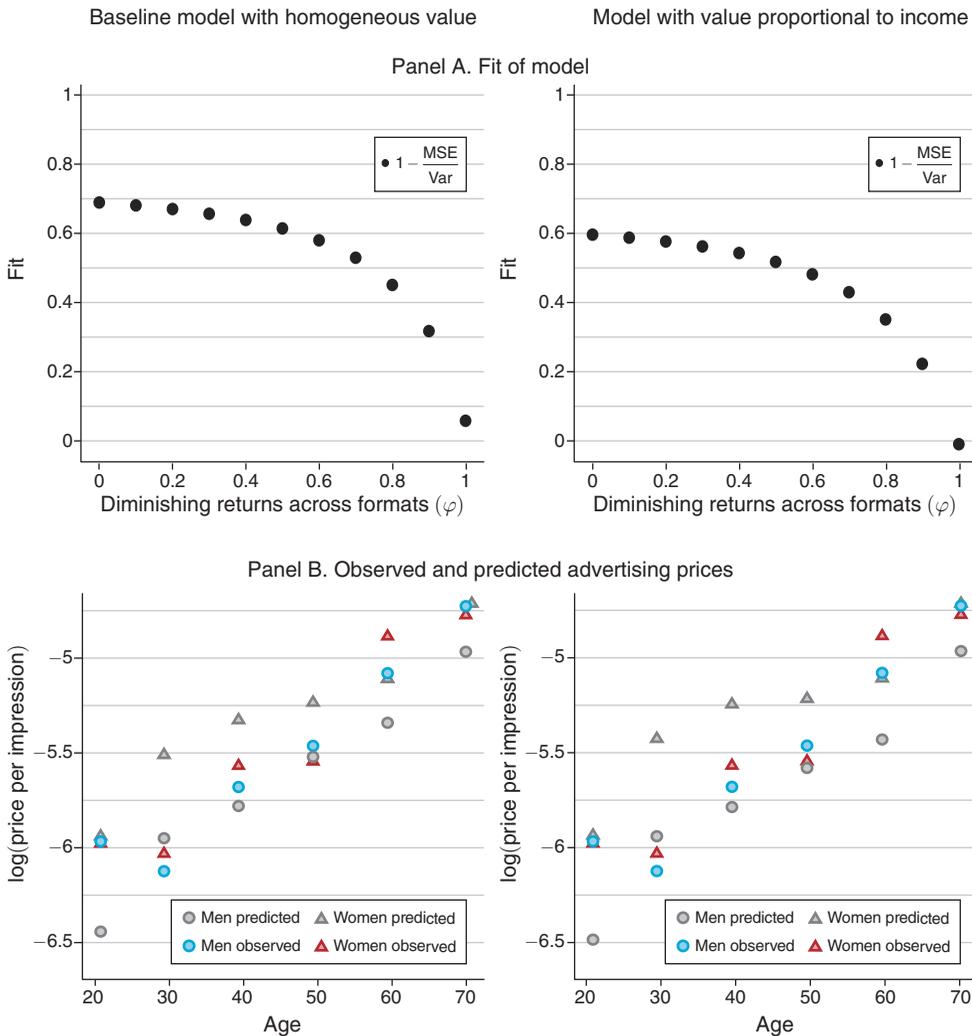


FIGURE 11. FIT OF QUANTITATIVE MODEL OF SOCIAL MEDIA PRICES

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

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

REFERENCES

- AdStage.** 2020. *Paid Media Q1 2020 Benchmark Report*. San Francisco, CA: AdStage.
- Allcott, Hunt, Luca Braghieri, Sarah Eichmeyer, and Matthew Gentzkow.** 2020a. "The Welfare Effects of Social Media." *American Economic Review* 110 (3): 629–76.
- 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*. New York, NY: Ampush.
- Anderson, Simon, and Stephen Coate.** 2005. "Market Provision of Broadcasting: A Welfare Analysis." *Review of Economic Studies* 72 (4): 947–72.
- Anderson, Simon, Øystein Foros, and Hans Jarle Kind.** 2018. "Competition for Advertisers and for Viewers in Media Markets." *Economic Journal* 128 (608): 34–54.
- Anderson, Simon, and Martin Peitz.** 2020. "Media See-Saws: Winners and Losers in Platform Markets." *Journal of Economic Theory* 186: 104990.
- Aral, Sinan.** 2021. "What Digital Advertising Gets Wrong." *Harvard Business Review*, February 19. <https://web.archive.org/web/20230325071010/https://hbr.org/2021/02/what-digital-advertising-gets-wrong>.
- Armstrong, Mark, and John Vickers.** 2022. "Patterns of Competitive Interaction." *Econometrica* 90 (1): 153–91.
- 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): 1574–90.
- 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: 164–83.
- Berry, Steven, Alon Eizenberg, and Joel Waldfogel.** 2016. "Optimal Product Variety in Radio Markets." *RAND Journal of Economics* 47 (3): 463–97.
- Chandra, Ambarish, and Ulrich Kaiser.** 2014. "Targeted Advertising in Magazine Markets and the Advent of the Internet." *Management Science* 60 (7): 1829–43.
- Chwe, Michael Suk-Young.** 1998. "Believe the Hype: Solving Coordination Problems with Television Advertising." Unpublished.
- Competition and Markets Authority.** 2019. *Online Platforms and Digital Advertising: Market Study Interim Report*. London: UK Competition and Markets Authority.
- Crupi, Anthony.** 2009. "Cable Guy Says Ratings Don't Ad Up." *Hollywood Reporter*, April 14. <https://www.hollywoodreporter.com/business/business-news/cable-guy-says-ratings-dont-82305/>.
- Dee, Jonathan.** 2002. "The Myth of '18 to 34'." *New York Times Magazine*, October 13. <https://www.nytimes.com/2002/10/13/magazine/the-myth-of-18-to-34.html>.
- 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): 255–65.
- Delrahim, Makan.** 2019. *Opening remarks for the public workshop on competition in television and digital advertising*. Washington, DC: US Department of Justice.
- Dubé, Jean-Pierre, Günter J. Hitsch, and Puneet Manchanda.** 2005. "An Empirical Model of Advertising Dynamics." *Quantitative Marketing and Economics* 3 (2): 107–44.
- Economist.** 2021. "The Olympics is a ratings flop. Advertisers don't care." August 11. <https://www.economist.com/business/the-olympics-is-a-ratings-flop-advertisers-dont-care/21803502>.
- Einstein, Mara.** 2004. "Broadcast Network Television, 1955-2003: The Pursuit of Advertising and the Decline of Diversity." *Journal of Media Economics* 17 (2): 145–55.
- Facebook.** 2018. "Build TV and Facebook campaigns that work better together." Facebook. <https://web.archive.org/web/20180525071501/https://www.facebook.com/business/news/build-tv-and-facebook-campaigns-that-work-better-together> (accessed February 23, 2023).
- Facebook.** 2022. "Documentation of marketing API." Facebook. <https://developers.facebook.com/docs/marketing-api/reference/ad-campaign-group/#Reading> (accessed July 11, 2022).

- Fan, Ying.** 2013. "Ownership Consolidation and Product Characteristics: A Study of the US Daily Newspaper Market." *American Economic Review* 103 (5): 1598–1628.
- Fudenberg, Drew, Jon Kleinberg, Annie Liang, and Sendhil Mullainathan.** 2022. "Measuring the Completeness of Economic Models." *Journal of Political Economy* 130 (4): 956–90.
- Fudenberg, Drew, Wayne Gao, and Annie Liang.** 2023. "How Flexible is That Functional Form? Quantifying the Restrictiveness of Theories." Unpublished.
- Gabler, Neal.** 2014. "The tyranny of 18 to 49: American culture held hostage." Unpublished.
- Garland, Eric.** 2002. "Disney marshals its parts into a new kind of advertising deal." *New York Times*, June 21. <https://web.archive.org/web/20150527211647/https://www.nytimes.com/2002/06/21/business/media-business-advertising-disney-marshals-its-parts-into-new-kind-advertising.html>.
- Gentzkow, Matthew, Jesse M. Shapiro, and Michael Sinkinson.** 2014. "Competition and Ideological Diversity: Historical Evidence from US Newspapers." *American Economics Review* 104 (10): 3073–3114.
- Gentzkow, Matthew, Jesse M. Shapiro, Frank Yang, and Ali Yurukoglu.** 2024. "Replication data for: Pricing Power in Advertising Markets: Theory and Evidence." American Economic Association [Publisher], Inter-university Consortium for Political and Social Research [Distributor]. <https://doi.org/10.3886/E193824V1>.
- Geskey, Ronald D.** 2016. *Media Planning and Buying in the 21st Century*. 4th ed. Scotts Valley, CA: CreateSpace Independent Publishing Platform.
- GfK Mediamark Research and Intelligence.** 2017. "Survey of the American Consumer, Respondent-Level Data for 2015." GfK Mediamark Research and Intelligence. <https://www.mrisimmons.com/our-data/national-studies/survey-american-consumer/> (accessed August 2017).
- GfK Mediamark Research and Intelligence.** 2021. "Survey of the American Consumer, Respondent-Level Data for 2019." GfK Mediamark Research and Intelligence. <https://www.mrisimmons.com/our-data/national-studies/survey-american-consumer/> (accessed May 2022).
- Goettler, Ronald L.** 2012. "Advertising Rates, Audience Composition, and Competition in the Network Television Industry." Unpublished.
- Google.** 2021. Geocoding and Time Zone APIs. <https://developers.google.com/maps/documentation/geocoding/overview> and <https://developers.google.com/maps/documentation/timezone/overview>. (accessed April 2021).
- Greenwood, Jeremy, Yueyuan Ma, and Mehmet Yorukoglu.** 2021. "'You Will:' A Macroeconomic Analysis of Digital Advertising." NBER Working Paper 28537.
- Hristakeva, Sylvia, and Julie H. Mortimer.** 2023. "Price Dispersion and Legacy Discounts in the National Television Advertising Market." *Marketing Science*. <https://doi.org/10.1287/mksc.2023.1442>.
- Jeziorski, Przemysław.** 2014. "Effects of Mergers in Two-Sided Markets: The US Radio Industry." *American Economic Journal: Microeconomics* 6 (4): 35–73.
- Kaiser, Ulrich, and Julian Wright.** 2006. "Price Structure in Two-Sided Markets: Evidence from the Magazine Industry." *International Journal of Industrial Organization* 24 (1): 1–28.
- Krouse, Sarah.** 2022. "Netflix's ad-supported plan will launch in November at \$6.99 a month." *Wall Street Journal*, October 13. <https://www.wsj.com/articles/netflix-ads-to-launch-in-november-11665680110>.
- Lambrecht, Anja, and Catherine Tucker.** 2019. "Algorithmic Bias? An Empirical Study of Apparent Gender-Based Discrimination in the Display of STEM Career Ads." *Management Science* 65 (7): 2966–81.
- Lee, Robin, Michael Whinston, and Ali Yurukoglu.** 2021. "Structural Empirical Analysis of Contracting in Vertical Markets." In *Handbook of Industrial Organization*, Vol. 4, edited by Kate Ho, Ali Hortacsu, and Alessandro Lizzeri, 673–742. Amsterdam: Elsevier.
- Lewis, Randall A., and David H. Reiley.** 2014. "Advertising Effectively Influences Older Users: How Field Experiments Can Improve Measurement and Targeting." *Review of Industrial Organization* 44 (2): 147–59.
- Liao, Lu, Alan Sorensen, and Andrey Zubanov.** 2020. "Measuring the Value of Targeted Television Advertising." Unpublished.
- McGranaghan, Matthew, Jura Liaukonyte, and Kenneth C. Wilbur.** 2022. "How Viewer Tuning, Presence, and Attention Respond to Ad Content and Predict Brand Search Lift." *Marketing Science* 41 (5): 445–67.
- MoffettNathanson.** 2017. "U. S. internet: From cat videos to the NFL - are we at the tipping point for online video?" MoffettNathanson (An SVB Company). <https://web.archive.org/web/20230301231204/https://www.moffettnathanson.com/research.aspx?Section=Media+%2FTelecom> (accessed March 1, 2023).

- MoffettNathanson.** 2022. "Netflix and Disney: Mad Men to the rescue?" MoffettNathanson (An SVB Company). <https://web.archive.org/web/20230301231204/https://www.moffettnathanson.com/research.aspx?Section=Media+%2FTelecom> (accessed March 1, 2023).
- Newman, Andrew Adam.** 2012. "In AARP's view, advertisers need to focus." *New York Times*, July 19. <https://www.nytimes.com/2012/07/19/business/media/aarp-campaign-tries-to-persuade-advertisers.html>.
- News Journal.** 2015. *Sunday News Journal* Sep. 13–19, 2015. Wilmington, Delaware. <https://www.newspapers.com/image/127413547/>.
- The Nielsen Company.** 2017. *Ad Intel Methodology by Medium*. New York, NY: Nielsen Company.
- The Nielsen Company.** 2018. *CMO Report*. New York, NY: Nielsen Company.
- The Nielsen Company.** 2019. "Ad Intel dataset." Nielsen Company. <https://www.chicagobooth.edu/research/kilts/datasets/nielsenIQ-nielsen> (accessed March, 2019).
- The Nielsen Company.** 2021. "Nielsen Local TV View." Nielsen Company. <https://www.nielsen.com/us/en/client-learning/tv/local/nltv/> (accessed February, 2021).
- Organization for Economic Co-operation and Development.** 2022. "Consumer Price Index: Total All Items for the United States [CPALTT01USA661S]." Federal Reserve Bank of St. Louis. <https://fred.stlouisfed.org/series/CPALTT01USA661S> (accessed June 2022).
- Papazian, Ed.** 2009. *TV Dimensions 2009*. New York, NY: Media Dynamics, Inc.
- Patel, Sahil.** 2015. "Turner's new ad sales strategy sharpens its focus on clients, data." *Digiday*, December 10. <https://web.archive.org/web/20230208175713/https://digiday.com/media/turners-new-ad-sales-strategy-sharpens-focus-clients-data/>.
- Phillips, Robert, and Graham Young.** 2012. "Television Advertisement Pricing in the United States." In *Oxford Handbook of Pricing Management*, edited by Özalp Özer and Robert Phillips. Oxford, UK: Oxford University Press.
- Pomerantz, Earl.** 2006. "Why do advertisers still covet the 18-49s?" *Television Quarterly* 36: 40–44.
- Prat, Andrea.** 2018. "Media Power." *Journal of Political Economy* 126 (4): 1747–83.
- Prat, Andrea, and Tommaso Valletti.** 2022. "Attention Oligopoly." *American Economic Journal: Microeconomics* 14 (3): 530–57.
- Rochet, Jean-Charles, and Jean Tirole.** 2003. "Platform Competition in Two-Sided Markets." *Journal of the European Economic Association* 1 (4): 990–1029.
- Shi, C. Matthew.** 2016. "Catching (exclusive) eyeballs: Multi-homing and platform competition in the magazine industry." Unpublished.
- 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. Hamburg, Germany: Statista.
- Strikesocial.** 2017. *YouTube Advertising: 2017 Benchmark Report*. Chicago, IL: Strikesocial.
- Surowiecki, James.** 2002. "Ageism in advertising." *New Yorker*, March 24. <https://www.newyorker.com/magazine/2002/04/01/ageism-in-advertising>.
- S&P Global Market Intelligence.** 2019. *SNL Kagan Report for 2015*. New York, NY: S&P Global Market Intelligence.
- US Census Bureau.** 2021. *2020 Census apportionment results delivered to the president*. Washington, DC: US Census Bureau.
- Veiga, André, and E. Glen Weyl.** 2016. "Product Design in Selection Markets." *Quarterly Journal of Economics* 131 (2): 1007–56.
- WARC.** 2001. "P&G and Viacom ink \$300m cross-media deal," May 31. https://web.archive.org/web/20230428230056/https://www.warc.com/newsandopinion/news/pg_and_viacom_ink_300m_crossmedia_deal/8911.
- Wepriin, Alex.** 2015. "Upfronts 2015: NBCU's cable lifestyle channels make their pitch." *Politico*, March 31. <https://www.politico.com/media/story/2015/03/upfronts-2015-nbcus-cable-lifestyle-channels-make-their-pitch-003625/>.
- Wikipedia, s. v.** "2014–15 United States network television schedule," accessed May 19, 2022. https://en.wikipedia.org/wiki/2014%E2%80%9315_United_States_network_television_schedule.
- Wilbur, Kenneth C.** 2008. "A Two-Sided, Empirical Model of Television Advertising and Viewing Markets." *Marketing Science* 27 (3): 356–78.
- Zubanov, Andrey.** 2020. "The TV Advertising Market: Demographic Segmentation and the Impact of Viewership Decline." Unpublished.