Zero rating: The power in the middle

Kunal Phalak, D. Manjunath and Jayakrishnan Nair

Abstract—**Many flavors of differential data pricing are being practiced in different telecom markets. One popular version is zero-rating, where customers do not pay for consuming a certain basket of 'zero-rated' content. These zero-rated services are in turn sponsored by payments to the internet service provider (ISP) by the corresponding content providers (CPs). In this paper, we provide an analytical treatment of a zero-rating platform, highlighting the effect of zero-rating on the structure of the CP market, and also on the surplus of ISPs, CPs, and users.**

A leader-follower game is assumed with the ISP setting the prices for users (for non-sponsored data) and CPs (for sponsored data), CPs making a binary decision on sponsorship, and users consuming content based on the resulting data charges. User consumption is determined by a utility maximization, the sponsorship decision is determined by a Nash equilibrium between the CPs, and the ISP sets prices to maximize its profit. Several scenarios mimicking real-life practices are analyzed. Our results indicate that zero-rating grants the ISP significant power to determine the mix of content consumption, and the profitability of the CPs. Further, the ISP can also take away a significant portion of the surplus in the system.

I. Introduction

Increasing revenue and investment pressures are forcing internet service providers (ISPs) to consider moving away from a flat access charge to a smart data pricing (SDP) regime, at least for mobile internet services. Most SDP schemes propose some form of usage-based pricing in which the user charges are dertermined by multiple attributes of the content data, e.g., volume, time of day, type of data, and source of the data. An excellent survey of smart data pricing proposals is available in [1]. While ISPs look to increase their revenue, household expenditure on telecom devices and services, both in terms of the actual amount and as a fraction of household income, is increasing in most markets [2]–[5]. This is limiting the ability of the ISPs to raise revenue from the consumer side of their services and is making them seek alternative revenue sources, and content providers (CPs) appear to present themselves as an obvious source. More so because the revenues of online service providers appear to be growing much faster than those of ISPs [6], [7].

An SDP scheme that has gained significant traction among the ISPs is differential data pricing in which user charges depend on the content that is consumed. One particular differential pricing scheme that is becoming popular is zero-rating. In this scheme, the ISP sets up a zero-rating platform and invites CPs to sign up on to the platform. The ISP exempts the user from charges for data traffic originating from the websites/apps that have signed up on the zero-rating platform. One way that this manifests to the consumers is as follows. In data plans where users pay a fixed price and are allocated a certain volume for consumption in the billing cycle, zero-rated data that they consume does not count against the data caps that are set in their data plan.

The authors are with the Department of Electrical Engineering, IIT Bombay, India.The authors acknowledge support from DST and CEFIPRA.

Zero-rating has often been compared to the toll-free telephony services. While this is a useful analogy to describe the nature of the service, there are key commercial and technological differences and hence the market effects are significantly dissimilar. A toll-free service that is set up with one service provider is free to subscribers of other competing service providers too. And the termination charges settlement is transparent to the users. However, zero-rating of a service on one ISP does not make it available to customers of another ISP. See [8] for a more detailed discussion on the differences.

Several ISPs are offering some version of zero-rating. *Sponsored Data* from AT&T and *FreeBee Data* from Verizon allow a content provider to sponsor some (e.g., trailers, app downloads) or all the content from its website/app. In the *BingeOn* scheme, T Mobile allows CPs to provide zero-rated content that conforms to their specifications. Zero-rating platforms are also operated by third party organisations and CPs. `internet.org`, the latter day version of FreeBasics, is a zero-rating platform that ISPs and CPs can join. The ISPs that join the platform zero-rate content of the CPs that are on the platform. Like with BingeOn, zero-rated content has to conform to specifications laid out by `internet.org`.

An argument in favor of zero-rating is that it is a form of product differentiation for the ISPs and in competitive ISP markets, the ability to provide such differentiation is crucial to competing in the market. There is also the additional argument of providing societal good in bridging the digital divide. The argument against these schemes is that they violate net neutrality principles, due to their potential to tilt the CP landscape because content from sponsoring CPs are expected to be preferred by users. Specifically, BingeOn has been critiqued in [9], [10] and `internet.org` has been banned in some countries [11]. Note though that as of February 2017, under FCC regulations, zero-rating does not violate net neutrality stipulations in the US. Several other countries, including Brazil and Colombia also permit zero-rating. On the other hand, the European Union imposes certain restrictions on zero-rating practices, while some countries e.g., India, Norway, and Chile, have made it illegal.

Our interest in this paper is to model a zero-rating platform to analyze the power of a strategic ISP in determining the consumption patterns and hence the structure of the CP market. We consider the case where the ISP offers a sponsoring plan such that, for a fee, content providers can have their content zero-rated to the users of the ISP. The users pay for the non zero-rated content that they consume. We take a game-theoretic approach with three types of players—the users, the content providers and the ISP. The users will consume content from competing CPs to maximize a utility function that includes cost of the data. The content providers in turn make a strategic binary decision on whether to sponsor or to not sponsor their content on the ISP. The ISP determines the prices to maximize its profit which in turn is made under different constraints. We consider three cases each of which is motivated by current practice.

1) User price is exogenously determined (by the market or regulator) and the ISP sets the profit maximizing sponsorship price.
2) Both user and sponsorship prices are set by the ISP.
3) User and sponsorship prices are constrained to be equal.

For each of these three cases we analyze the profit-maximizing strategy of the ISP, and its impact on the CP marketplace. We also analyze the surplus of the ISP, the CPs, and the users. Based on these analyses our findings are as follows.

- Being able to charge both users and CPs grants the ISP considerable market power, enabling it to transfer a significant amount of the surplus from the CPs to itself. This power is further magnified when the ISP can optimize both the user charge and the sponsorship charge.
- In many cases, the ISP effects such a surplus transfer by setting the prices such that the most profitable CP will sponsor; this in turn tilts the user traffic and skews the CP marketplace.
- Sponsorship does not always benefit the user base. Thus, zero-rating platforms can result in a scenario where only the ISP stands to benefit; the CPs as well as the users being worse off compared to the case when there is no zero-rating. More importantly, it must be noted that the user utility gains, insofar as they occur, are 'only in the short run.' In the long run, the less profitable CPs may be hastened off the market because of the reduced demand caused by the zero-rating platform. Our models do not capture this aspect.

The rest of the paper is organized as follows. In the next section, we provide an overview of the relevant literature. In Section III, we describe the model in detail and provide some preliminary results. In Sections IV–VI, we analyze each of the three cases mentioned above. We conclude with a detailed discussion in Section VII.

II. Literature survey

Broadly, there are two kinds of non neutral behavior that are practiced by the ISPs—discriminatory QoS and discriminatory pricing. Of course, an ISP could also choose to simultaneously practice both kinds of discrimination. Discriminatory QoS provides a better user experience of the favoured content and hence improves user preference for the same. Under discriminatory pricing, user preference for the favoured content is increased because the surplus of a rational user is increased by not having to pay for it. There is a significant body of research that analyzes the effect of discriminatory QoS on various performance parameters like social surplus, surplus of users, CPs, and ISPs, and the incentive of the ISP to invest in its infrastructure; see, for example, [12]–[17].

In this paper, our interest is in understanding the effect of discriminatory pricing, like in [7], [18]–[22]. Here too there are two key strands in the literature. In one strand, a single ISP and a single CP interact in a game theoretic setting, typically in a Stackelberg or Nash bargaining framework, e.g., [18], [19]. In [18], a Stackelberg game is defined between a single ISP and a single CP in which the ISP sets the sponsorship price and the CP decides the volume of sponsorship. In [19], a sequential game consisting of single CP and a single ISP is considered. In each epoch, the ISP guarantees a certain QoS and the CP chooses the volume of traffic to sponsor. Further, in each epoch the CP and the ISP observe the actual demand which is used to inform the strategy for subsequent epochs. The second strand (see [7], [20]–[22]) considers multiple CPs and one ISP which is also the setting that we consider here.

In the papers [7], [21], [22], the authors consider a macroscopic model of the internet, where the data usage corresponding to each content provider is influenced by a common congestion signal, which is in turn detemined by the aggregate data usage. Thus, the user model is characterized by the solution of a certain fixed point equation. [7] considers the setting where CPs can subsidize the per-byte user charge subject to an upper limit, leading to a subsidization competition between CPs. The results in [7] suggest that subsidization competition can result in increased welfare of the ISP, the CPs, as well as the users, so long as internet access prices are carefully regulated. We note that while the model of [7] is very different from that of the present paper, its message may be viewed as complementary to ours. While [7] suggests that carefully regulated differential pricing can be beneficial to all parties of the internet ecosystem, the present paper shows that unregulated differential pricing grants considerable market power to the ISP, leaving CPs, and potentially even end users, worse off. [22] considers the setting where the ISP sets user-side access charges as well as CP-side sponsorship charges, similar to the setting considered in the present paper. The main take-away of [22] is that the optimal pricing and capacity provisioning decisions of the ISP are strongly influenced by the nature of user traffic (e.g., text or video). However, [22] does not discuss the impact of the ISP's two-sided pricing on the surplus of the CPs and users. Finally, [21] considers a Stackelberg game between the ISP, the CPs, and the users. Here the ISP sets the sponsorship price and the cap on customer data usage, the CPs make the binary decision on sponsorship, and the users are utility maximizing consumers. The key results are on the existence of equilibrium strategies of the CPs and the properties of these equilibria.

At this point, it is worth delineating the modeling approach of the present paper with that of [7], [21], [22]. In the latter papers, user data consumption is characterized as an equibrium between the consumption levels and the resulting congestion. This *implicit* characterization of the user behavior makes the three-tier interaction between the users, the CPs, and the ISP hard to analyze mathematically. Indeed, the analytical results in [7], [21], [22] are restricted to existence and monotonity properties of the various equilibria. Most key insights are actually obtained via numerical examples. In contrast, our model does not explicitly capture the effect of congestion—it is assumed that the ISP has made the necessary capacity provisioning. This in turn allows for an *explicit* characterization of user behavior, which further enables an analytical characterization of the impact of zero-rating on the CP marketplace, and the surplus of the ISP, CPs, and users. This modeling approach is based on the observation that (i) market forces drive ISPs to provide a certain minimum QoS to its users, and (ii) the ISP is the leader of our model of the leader-follower interaction between the ISP, CPs, and users, and thus has the ability to ensure that certain pre-defined QoS constraints are satisfied.

In light of the above discussion, the paper closest to ours is [20]. In the model of [20], the ISP sets sponsorship and user prices and each CP controls two variables for each customer—fraction of sponsored traffic and ad volume. Similar to the present paper, the authors do not explicitly capture ISP congestion, and model user behavior as the solution of a utility maximization problem. However, the model of [20] assumes that the data consumption corresponding to each CP is determined independently. In other words, this model ignores competition between CPs. In contrast, the present paper explicitly captures inter-CP competition. To the best of our knowledge, this is the first paper to analyse the interplay between inter-CP competition and discriminatory pricing.

We conclude the discussion on prior work by mentioning that [23] estimates the gains to a CP from sponsoring. This justifies our assumption of the knowledge of CP profitability to the ISP. Also, more recently, there is interest in modeling paid peering between ISPs and CPs, e.g. [24], [25]. Clearly, this line of work is complementary to the body of work on discriminatory pricing.

III. Model and Preliminaries

We consider a single ISP and two competing content providers (CPs).[1] The ISP operates a zero-rating platform, and the CPs have the option of sponsoring their content by joining the zero-rating platform. Specifically, the ISP sets a data charge p for users and a sponsorship charge q for CPs.[2] Users pay the ISP p dollars per byte of non-sponsored data consumed; sponsored content is free for users. A sponsoring CP pays the ISP q dollars per byte of sponsored data consumed.

We capture the interaction between the ISP, the CPs, and the users via the following three-tier leader-follower model. The ISP 'leads' by setting the prices p and q. The CPs respond to these prices by making the (binary) decision of whether or not to sponsor their content. Finally, the users decide how much data to consume from each CP based on both the ISP's data charge p, and sponsorship decisions of the CPs. In the following, we first describe the user behavior model, followed by the behavior model of the CPs and the ISP.

[1]In essence, we are restricting our attention to a single class of internet content providers. For example, social media platforms, or video streaming services, or messaging services. The assumption that there are only two competing CPs in the class under consideration is made primarily for convenience of exposition; several of our results extend easily to a general number of competing CPs.

[2]Note that the sponsorship price is not CP-specific. We revisit this assumption in Section VII.

A. User behavior

This model prescribes the total number of bytes of content that the user base consumes from each CP over a predefined horizon (say a billing cycle), given the data charge p and the sponsorship decisions of both CPs.[3]

Let $\mathcal{N} = \{1, 2\}$ denote the set of CPs, and $\mathcal{S}$ denote the subset of sponsoring CPs. We denote the sponsorship configurations $\mathcal{S} = \phi$, $\mathcal{S} = \{1\}$, $\mathcal{S} = \{2\}$ and $\mathcal{S} = \{1, 2\}$ by NN, SN, NS, and SS, respectively (S denoting the action of sponsoring, and N denoting the action of not sponsoring).

We assume that the user base is partitioned into K classes of users, denoted by $1, 2, \cdots, K$. There are n_k users of class k. Users of Class k derive a utility of $\psi_{i,k}(\theta)$ from consuming θ bytes from CP i. We assume that $\psi_{i,k} : \mathbb{R}_+ \to \mathbb{R}_+$ is continuously differentiable, strictly concave, and strictly increasing, with $\psi'_{i,k}(0) < \infty$. Also, users of Class k have a 'capacity-to-consume' c_k, which is the maximum amount of data (across both CPs) they can consume. Denoting the number of bytes that a user of class k consumes from CP i by $\theta_{i,k}$, we model $\theta_k = (\theta_{1,k}, \theta_{2,k})$ to be the (unique) solution of the following utility maximization.

$$\max_{x=(x_1,x_2)} \sum_{i \in \mathcal{N}} \psi_{i,k}(x_i) - p \sum_{i \in \mathcal{N} \setminus \mathcal{S}} x_i$$
$$\text{s.t. } \sum_{i \in \mathcal{N}} x_i \leq c_k, \quad x \geq 0$$

Note that the objective function above is the surplus of a user of class k, i.e., the total utility from content consumption, minus the amount that the user has to pay the ISP for consuming non-sponsored content.

The total data (in bytes) consumed by the user base from CP i, denoted by θ_i, is given by

$$\theta_i = \sum_{k=1}^{K} n_k \theta_{i,k}.$$

The tuple $\theta = (\theta_1, \theta_2)$ describes the aggregate data usage corresponding to the two CPs by the user base. Note that θ depends on the data charge p as well as the prevailing sponsorship configuration. When we need to make this dependence explicit, we write $\theta^M(p)$, where M is the sponsorship configuration. For example, $\theta^{SN}(p) = (\theta_1^{SN}(p), \theta_1^{SN}(p))$ describes the data usage profile of the user base, when CP 1 sponsors and CP 2 does not.

We make the following assumption throughout the paper.

Assumption 1. *There exists a user class $\hat{k}$ such that*

$$\theta_{1,\hat{k}}^{SS},\ \theta_{2,\hat{k}}^{SS} > 0.$$

In other words, there exists a user class $\hat{k}$ such that if users of that class have free access to the services of both CPs, they consume data from both CPs. The above assumption implies that $\theta_1^{SS}, \theta_2^{SS} > 0$.

We conclude our description of user behavior by collecting some useful consequences of our model. Note that $\theta_1 + \theta_2 \leq c$,

[3]Even though we refer to data consumed 'from' a CP, it should be noted that our results do not depend on the direction of data transfer between the users and content providers.

where $c := \sum_{k=1}^{K} n_k c_k$. Thus, c is an upper bound on the aggregate data consumption seen by both CPs.

Lemma 1. *For the user behavior model described above, the following statements hold.*

1) *If $\mathcal{S} \neq \phi$, then $\theta_1 + \theta_2 = c$.*
2) *For any action m (S or N) of CP 2, $\theta_1^{Sm} > \theta_1^{Nm}$. Moreover, $\theta_1^{SN} > \theta_1^{SS} > 0$.*
3) *$\theta_1^{NN}(p)$, $\theta_2^{NN}(p)$, $\theta_2^{SN}(p)$, and $\theta_1^{NS}(p)$, are non-increasing continuous functions of p. $\theta_1^{SN}(p)$, and $\theta_2^{NS}(p)$, are non-decreasing continuous functions of p.*
4) *For $p \geq \bar{p} := \max_{i,k} \psi'_{i,k}(0)$,*
$$\theta_1^{NN}(p) = \theta_2^{NN}(p) = \theta_2^{SN}(p) = \theta_1^{NS}(p) = 0.$$

Statement 1 above asserts that so long as at least one CP sponsors, the total data usage of the user base equals c, which is the maximum consumption possible. Statement 2 implies that if any CP sponsors, its consumption strictly increases. Moreover, a sponsoring CP attracts a higher consumption when the other CP does not sponsor. Statement 3 implies that the consumption of non-sponsored data is non-increasing in the user price p, while the consumption of sponsored data is non-decreasing in p. Finally, Statement 4 implies that if the user price p is large enough, the consumption of non-sponsored data drops to zero. We give the proof of Lemma 1 in Appendix A.

Next, we describe the behavioral model of the content providers.

B. CP behavior

Recall that in our three-tier leader-follower model, CPs 'follow' the ISP and 'lead' the users, i.e., CPs make the decision of whether or not to sponsor, given the ISP's prices p and q, and knowing *ex-ante* that the user base will respond to the sponsorship configuration as per the model described in Section III-A.

We assume that CP i derives a revenue of a_i dollars per byte of content served.[4] We refer to a_i as the *revenue rate* of CP i. Thus, if CP i decides to sponsor its content on the zero-rating platform, it makes profit $r_i = (a_i - q)\theta_i$. On the other hand, if CP i decides not to sponsor, it makes profit $r_i = a_i\theta_i$. (Note that θ_i itself depends on the decisions of *both* CPs.) Since each CP's decision influences the other's profit, it is natural to model the emerging sponsorship configuration as a *Nash equilibrium* between the CPs; this is the approach we adopt in this paper.

We now characterize the conditions for each of the sponsorship configurations to be a Nash equilibrium. We use the following notation.

$$\alpha(p) := \left(1 - \frac{\theta_1^{NN}(p)}{\theta_1^{SN}(p)}\right), \quad \beta(p) := \left(1 - \frac{\theta_2^{NN}(p)}{\theta_2^{NS}(p)}\right)$$
$$\gamma(p) := \left(1 - \frac{\theta_1^{NS}(p)}{\theta_1^{SS}}\right), \quad \delta(p) := \left(1 - \frac{\theta_2^{SN}(p)}{\theta_2^{SS}}\right)$$

[4]This is a reasonable assumption for ad-supported services, which constitute a major fraction of online services today.

It follows from Statement 2 of Lemma 1 that
$$0 < \alpha(p), \gamma(p), \beta(p), \delta(p) \leq 1.$$

Moreover, $\alpha(p)$, $\gamma(p)$, $\beta(p)$, and $\delta(p)$ are non-decreasing in p (from Statement 3 of Lemma 1), and are equal to 1 for $p \geq \bar{p}$ (from Statement 4 of Lemma 1). Using the above notation, the conditions for each sponsorship configuration to be a Nash equilibrium are the following.

Lemma 2.
1) *NN is Nash equilibrium if and only if*
$$q \geq \max\left(a_1\alpha(p), a_2\beta(p)\right).$$
Thus, a sufficient condition for NN to be an equilibrium is $q \geq \max(a_1, a_2)$.
2) *SN is a Nash equilibrium if and only if*
$$a_2\delta(p) \leq q \leq a_1\alpha(p).$$
3) *NS is a Nash equilibrium if and only if*
$$a_1\gamma(p) \leq q \leq a_2\beta(p).$$
4) *SS is a Nash equilibrium if and only if*
$$q \leq \min\left(a_1\gamma(p), a_2\delta(p)\right).$$

The proof of Lemma 2 is provided in Appendix B.

We note that an SN/NS sponsorship configuration can lead to a significantly *skewed* marketplace, with the sponsoring CP commanding a much higher usage compared to the non-sponsoring CP. Such a skew can be a matter of concern, particularly if the CPs themselves provide a comparable service quality, and the skew is primarily a consequence of the asymmetric sponsorship configuration. In Sections IV–VI, we explore the conditions under which an ISP-operated zero-rating platform can *induce* such a skew in the CP marketplace.

C. ISP behavior

Since the ISP is the 'leader' for our three-tier leader-follower model, it sets the user data price p and the sponsorship price q to induce the most profitable Nash equilibrium between the CPs.[5] Note that the ISP's profit r_I is composed of payments from the user base for consumption of non-sponsored content and payments from the CPs corresponding to consumption of sponsored content:
$$r_I = p \sum_{i \in \mathcal{N} \setminus \mathcal{S}} \theta_i + q \sum_{i \in \mathcal{S}} \theta_i.$$

In the following sections, we explore the profit-maximizing strategy of the ISP and its consequences under different constraints on the tuple (p, q). Each of these constraints is motivated by real-life practices of ISPs worldwide.

We will find it instructive to analyze the ISP's profit maximizing strategy (and its consequences) in the regime

[5]If more than one sponsorship configuration is a Nash equilibrium for a given (p, q), we assume the ISP can 'steer' the CPs to the most profitable equilibrium. In other words, we assume that the *leader* (the ISP) can select the most desirable equilibrium between the *followers* (the CPs). This is a standard approach for handling non-unique follower equilibria in leader-follower interactions [26]. Moreover, this assumption is natural in the present setting, where it is the ISP that operates the zero-rating platform.

of *increasing CP revenue rates*. Indeed, one would expect that as (a_1, a_2) become larger, the ISP has a greater incentive to get one or both of the CPs to sponsor. Throughout this paper, for simplicity, we consider the following one-dimensional parameterization for scaling the revenue rates: We take $(a_1, a_2) = (a, \rho a)$, where $a \geq 0$ is the scaling parameter, and $\rho \in (0, 1)$ is fixed.[6] Thus, higher values of a correspond to higher CP revenue rates. This parameterization allows us to analyze how the surplus of the ISP, the CPs, and the users scale with increasing CP revenue rates.

IV. ISP Optimizes q

In this section, we consider the case where the ISP holds the user charge p fixed, and only optimizes the sponsorship price q to maximize its profit. This captures scenarios where the user charge is constrained by market forces or regulation. Even in the absence of such constraints, ISP may prefer to keep p unchanged in order maintain a uniform user charge for all internet access. Note that we are restricting attention to a single class of online services; varying p thus amounts to charging users differently for different classes of content, which may be inconvenient and/or unpopular. Also, there is empirical evidence that an ISP can earn more dollars/byte by selling to a sponsor than from a retail user [27].

Recall that our goal is to analyze the impact of the ISP-operated zero-rating platform on:

(i) The structure of the CP market,
(ii) The surplus of the ISP, the CPs, and the users.

Specifically, we explore the above in the regime of increasing CP revenue rates (scaled via the parameter a).

We first discuss the profit-maximizing strategy of the ISP. It follows from Lemma 2 that the ISP can always enforce an NN equilibrium by setting $q \geq a$. Similarly, the ISP can always enforce an SS equilibrium by setting $q \leq a \min(\gamma, \rho\delta)$.[7] However, an SN configuration is feasible (i.e., can constitute a Nash equilibrium between the CPs) if and only if $\rho\delta \leq \alpha$. Similarly, an NS confguration is feasible if any only if $\gamma \leq \rho\beta$. Clearly, under an NN equilibrium, the ISP's profit does not depend on q. On the other hand, under an SS/SN/NS equilibrium, the ISP maximizes its profit by setting q to the largest feasible value (for that configuration to a Nash equilibrium). The following theorem describes the profit-maximizing strategy of the ISP as a function of the scaling parameter a.

Theorem 1. [ISP's profit maximizing strategy] *As a function of a (for fixed ρ), the profit maximizing strategy for the ISP is the following. There exists a threshold $a_S > 0$ such that:*

1) *For $a < a_S$, the ISP enforces an NN equilibrium by setting $q \geq a$.*
2) *For $a > a_S$, the ISP enforces an SS/SN/NS equilibrium, whichever yields the maximum profit.*

The above result reveals that if CP revenue rates are small, then the ISP enforces an NN equilibrium (or equivalently, does not operate the zero-rating platform), since it can make more money from the user base than from the CPs. However, once CP revenue rates cross a certain threshold, the ISP has the incentive to induce one or both the CPs to sponsor.[8] It is important to note that depending on the system parameters, an SS, SN, or NS configuration might be optimal for the ISP; interestingly, this is in contrast with what occurs if the ISP also optimizes the user charge p (as we see in Sections V and VI).

Next, we consider the profit made by the ISP and the CPs under the profit maximizing strategy of the ISP. Recall that we denote the profit of the ISP by r_I and the profit of CP j by r_j. As the following lemma reveals, the ISP benefits considerably from the zero-rating platform.

Lemma 3. [ISP profit] *Under the ISP's profit maximizing strategy, its profit varies with a (with ρ fixed) as follows. For $a < a_S$, $r_I(a) = p(\theta_1^{NN} + \theta_2^{NN})$. For $a > a_S$, $r_I(a)$ is strictly increasing in a with $r_I(a) \geq ac \min(\rho\delta,\ \gamma)$.*

To interpret the above lemma, we note that in the absence of the zero-rating platform, the ISP's profit would be insensitive to the CP revenue rates; it would simply be equal to $p(\theta_1^{NN} + \theta_2^{NN})$. However, Lemma 3 shows that with the zero-rating platform, the ISP's profit grows (at least) linearly in a for $a > a_S$. This means that once CP revenue rates exceed a certain threshold, the ISP is able to extract a fraction of the CP revenues (which also grow linearly in a) by operating the zero-rating platform. The ISP achieves this by increasing the sponsorship price q in proportion to the CP revenue rates.

In contrast, as the following lemmas show, the CPs do not necessarily stand to benefit from the zero-rating platform.

Lemma 4. [CP profit under SN/NS] *Under the ISP's profit maximizing strategy, the sponsorship price q is set such that under an NS/SN equilibrium:*

1) *The sponsoring CP makes the same profit as it would without the zero-rating platform.*
2) *The non-sponsoring CP makes a profit less than or equal to that it would make without the zero-rating platform.*

Lemma 5. [CP profit under SS] *Under the ISP's profit maximizing strategy, the sponsorship price q is set such that under an SS equilibrium, at least one of the CPs makes a profit less than or equal to that it would make without the zero-rating platform.*

The above lemmas reveal that at least one of the CPs is worse off with the introduction of the zero-rating platform. Indeed, it is possible that the zero-rating platform leaves both CPs worse off; we illustrate this via a numerical example later. It is also important to note that Lemmas 4 and 5 do not rely on the scaling regime of revenue rates assumed in the preceding results. In other words, these results imply that if it is optimal for the ISP to induce one or both the CPs to sponsor, then one or both of the CPs is necessarily worse off.

Finally, when the ISP leaves the user charge p unchanged, the users are better off with the introduction of the zero-rating platform.

[6] The case $\rho > 1$ is subsumed in the above, by simply switching the labels of the two CPs.

[7] Since p is considered fixed in this section, the dependence of $\alpha,\ \beta,\ \gamma,\ \delta$ on p is suppressed throughout this section.

[8] Note that the threshold a_S in general depends on ρ.

Lemma 6. [User surplus] *Under an SS/SN/NS configuration, each user has a strictly greater surplus than she would without the zero-rating platform.*

In conclusion, when the ISP optimizes only the data sponsorship price q, our results show that the ISP induces one or both CPs to sponsor if their revenue rates are large enough. Moreover, while the ISP and the user base benefits from this flavor of zero-rating, one of both of the CPs end up being worse off.

Numerical results: We now present some sample numerical results to illustrate the conclusions of this section. While a comprehensive case study is beyond the scope of the present paper, our purpose is simply to give the reader a visual interpretation of our analytical results.

We consider a single class of users, associated with the utility function $\psi_i(\theta) = \log(1+\theta)$ for both CPs. Figures 1 and 2 show how the profits of the ISP and the CPs, and the surplus of the user base, scale with a, for two different parameter settings.

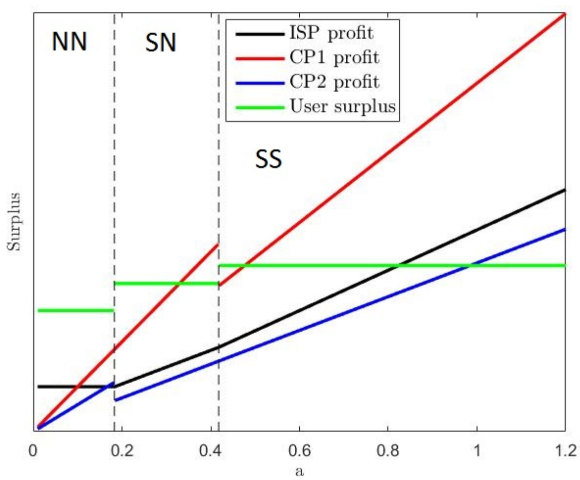

Fig. 1. Surplus of ISP, CPs, and users as a function of a, when ISP optimizes only q ($C = 30,\ \rho = 0.6,\ p = 0.05$)

Note that in the setting depicted in Figure 1, the ISP enforces an NN equilibrium for small a, an SN equilibrium for moderate a, and an SS equilibrium for large a. Also note that the ISP and the users clearly benefit from sponsorship. Interestingly, it can be verified that the SS equilibrium is a *prisoner's dilemma* between the CPs. That is, starting from an NN configuration, CP 1 has an incentive to sponsor. However, once CP 1 sponsors, CP 2 sees a sharp reduction in usage, and thus is also induced to sponsor. However, the resulting SS equilibrium has both CPs worse off compared to the NN configuration.

In the setting depicted in Figure 2, the ISP induces an NN equilibrium for small values of a and an SN equilibrium for larger values of a. Note that the sponsoring CP does not benefit from sponsorship, while the non-sponsoring CP suffers.

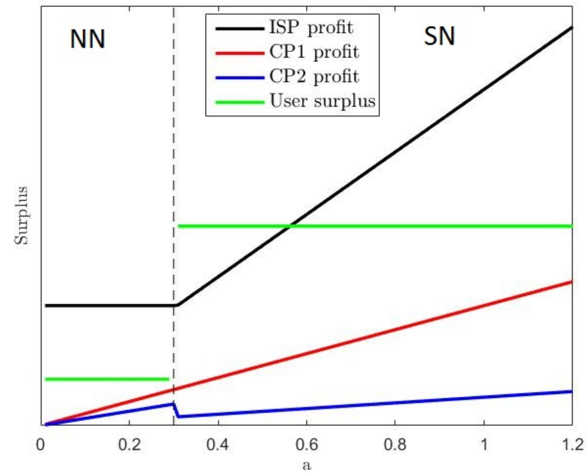

Fig. 2. Surplus of ISP, CPs, and users as a function of a, when ISP optimizes only q ($C = 4,\ \rho = 0.6,\ p = 0.5$)

V. ISP Jointly Optimizes p and q

In this section, we consider the case where the ISP jointly optimizes the user price p as well as the sponsorship price q to maximize its profit. This means the ISP has the power to set prices freely on both sides of the two-sided market (linking users and CPs) it serves. Indeed, we show that this grants considerable market power to the ISP.

Recall that throughout this paper, we focus our attention on a single class of internet services. Thus, varying p freely (for this class) amounts to charging users different prices for different classes of online services. Indeed, ISPs have advertised for such app-specific differential pricing [28], [29].

As in Section IV, we explore the impact of the zero-rating platform in the regime of increasing CP revenue rates. We begin by considering the profit-maximizing strategy for the ISP. It is easy to see that an NN equilibrium is always feasible for the ISP, i.e., it can set p and q to satisfy the condition in Lemma 2 for an NN equilibrium (for example, $p > 0,\ q = a$). Moreover, in this case, it is easy to see that the ISP maximizes its profit by setting the user price as

$$p = p^*_{NN} := \arg\max_{p>0} [p(\theta_1^{NN}(p) + \theta_2^{NN}(p))].$$

Similarly, it can be shown that an SS equilibrium is always feasible (e.g., $p > 0,\ q = a \min(\gamma(p), \rho\delta(p))$, as is an SN equilibrium (e.g., $p = \bar{p},\ q = a$). However, an NS equilibrium is not necessarily feasible. The following theorem describes the profit-maximizing strategy of the ISP as a function of the CP revenue rates.

Theorem 2. [ISP's profit maximizing strategy] *As a function of a (for fixed ρ), the profit maximizing strategy for the ISP is the following. There exist thresholds a_S and a_M, where $a_M \geq a_S > 0$, such that:*

1) *For $a < a_S$, the ISP enforces an NN equilibrium, by setting $p = p^*_{NN}$, $q \geq a$.*

2) *For $a_S < a < a_M$, the ISP enforces an SN/NS equilibrium, whichever yields the maximum profit.*[9]
3) *For $a > a_M$, the ISP enforces an SN equilibrium, by setting $q = a$, $p \geq \bar{p}$, so that CP 2 gets zero usage and there is a complete monopoly for CP 1.*[10]

There are several important take-aways from Theorem 2:

1) If CP revenue rates are small, the ISP enforces an NN equilibrium (or equivalently, does not operate the zero-rating platform). This is to be expected, and is analogous to the conclusion of Theorem 1 in Section IV.[11]
2) If CP revenue rates exceed a certain threshold a_S, the ISP will induce either an SN or an NS equilibrium. In other words, the ISP has an incentive to *skew* the CP marketplace. This is in contrast to our conclusion for the setting where the ISP only optimizes the sponsorship price (see Theorem 1); in that case, it is also possible for the profit-maximizing sponsorship configuration to be SS.
3) If CP revenue rates are high enough, then the ISP induces an SN equilibrium with an extreme skew, such that CP 1 (which has the higher revenue rate) gains a *complete monopoly*.[12]

Compared to the setting where the ISP only optimizes q (see Theorem 1), note that the profit-maximizing strategy of the ISP is structurally different when the ISP optimizes prices on both the user side and the CP side. Next, we turn to the implications of this strategy on the surplus of the ISP, the CPs, and the users. As before, we use r_I to denote the profit of the ISP, and r_j to denote the profit of CP j.

Lemma 7. [ISP profit] *Under the ISP's profit maximizing strategy, its profit varies with a (with ρ fixed) as follows. For $a < a_S$, $r_I(a) = p(\theta_1^{NN}(p^*_{NN}) + \theta_2^{NN}(p^*_{NN}))$. For $a > a_S$, $r_I(a)$ is strictly increasing in a with $r_I(a) \geq ac$.*

As before, note that in the absence of the zero-rating platform, the ISP's profit would be insensitive to the CP revenue rates. However, Lemma 7 shows that once CP revenue rates exceed the threshold that makes sponsorship attractive for the ISP, its profit grows linearly in a. It is also important to note that the lower bound ac on the ISP's profit matches the maximum possible revenue that the CPs can make combined. This means that by optimizing both p and q, the ISP can extract nearly all of the CP revenue. Intuitively, this is achieved as follows. The ISP raises both the user price p as well as the sponsorship price q, such that the sponsoring CP does not have an incentive to stop sponsoring (this would cut usage drastically thanks to the high user price), but also ends up passing on most of its revenue to the ISP (thanks to the high sponsorship price). This is further highlighted in the following result.

Lemma 8. [CP profit] *Under the ISP's profit maximizing strategy, $r_1(a) = r_2(a) = 0$ for $a > a_M$. Moreover, if the CPs are identical from the user standpoint (i.e., $\psi_{1,k}(\cdot) = \psi_{2,k}(\cdot)$ for all user classes k), then both $\frac{r_1(a)}{a}$ and $\frac{r_2(a)}{a}$ are non-decreasing over $a > a_S$.*

Note that in the absence of the zero-rating platform, CP revenues would grow linearly in a. However, with the ISP-operated zero-rating platform, Lemma 8 shows that if the CPs have identical utility functions, then CP profits are *sub-linear* in a for $a > a_S$. Moreover, even with non-symmetric CPs, the ISP is able to extract *all* the CP revenue if $a > a_M$.[13]

Finally, we turn to the surplus of the users under the profit-maximizing strategy of the ISP. As the following lemma shows, even the user base can be worse off when the ISP optimizes both p and q.

Lemma 9. [User surplus] *Suppose that for some user class k,*

$$\psi_{1,k}(\theta_1^{NN}(p))+\psi_{2,k}(\theta_2^{NN}(p))-p(\theta_1^{NN}(p)+\theta_2^{NN}(p)) > \psi_{1,k}(c),$$

where p denotes the prevailing user charge before the introduction of the zero-rating platform. Then for $a > a_M$, users of class k have a lower surplus under the zero-rating platform than they did without.

In contrast, recall that users provably benefit from the zero-rating platform if the ISP only optimizes the sponsorship price (see Lemma 6).

In conclusion, we see that optimizing both the user price as well as the sponsorship price grants the ISP considerably more power as compared to the case where the ISP only optimizes the sponsorship price. In particular, if the CP revenue rates are large enough, the ISP skews the CP marketplace and extracts most of the CP revenue, leaving both CPs, and potentially even some users, worse off.

Numerical results: We now present some numerical results (see Figures 3 and 4) to visualize the conclusions of this section. For these results, we use the same settings as described in Section IV.

In Figure 3, note that the profit-maximizing configuration switches from NN for small values of a to SN for larger values of a. Note that CP revenues drop sharply after sponshorship kicks in. Moreover, even user surplus tends to descrease with increasing a.

In Figure 4, we see as before that there is a single threshold in a beyond which the ISP induces an SN equilibrium. However, note that the CP revenues drop to zero immediately after this threshold (i.e., $a_S = a_M$ in this example). Interestingly, users derive a higher surplus with sponsorship.

[9] This case is only relevant if $a_S < a_M$.

[10] Note that Case 3 is just a special case of Case 2, where we are more specific about the ISP's action.

[11] Note that the *value* of the threshold a_S defined in Theorem 2 is in general different from the one defined in Theorem 1. Indeed, it is easy to show that the former is less than or equal to the latter.

[12] It is not hard to see that for a large enough, this extreme skew actually maximizes the *social welfare*, which is the sum of the surplus of the ISP, the CPs, and the users. However, it is important to note that this welfare maximization comes at the expense of considerable inequity in the *distribution* of that welfare: The ISP corners most of the surplus, leaving the CPs and even potentially users worse off (see Lemmas 8 and 9).

[13] The conclusion of *zero* profit needs to be interpreted carefully. Mathematically, we define a Nash equilibrium to be an action profile where either player cannot obtain a *strictly higher* payoff from unilaterally switching her action. Thus, the SN configuration enforced by the ISP for $a > a_M$ is a Nash equilibrium inspite of zero profits for both CPs, since neither CP can make a positive profit from deviating from their sponsorship decision. In practical terms, what our result implies is that the ISP can set prices such that is is able to extract *most* of the CP revenue.

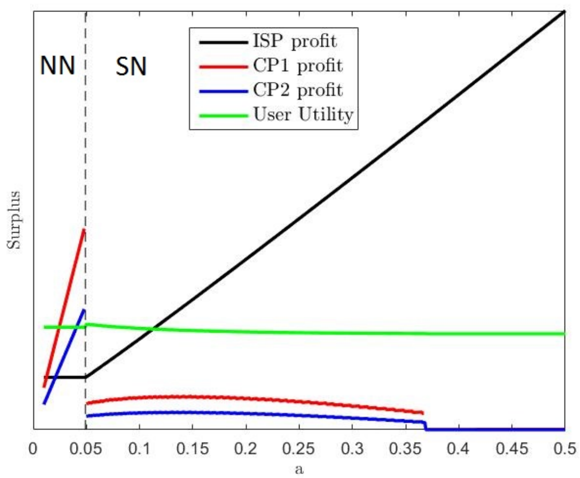

Fig. 3. Surplus of ISP, CPs, and users as a function of a, when ISP optimizes p and q ($C = 30,\ \rho = 0.6$)

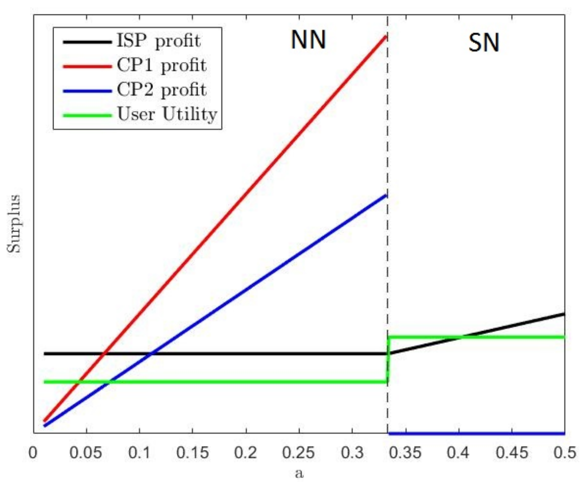

Fig. 4. Surplus of ISP, CPs, and users as a function of a, when ISP optimizes p and q ($C = 30,\ \rho = 0.6$)

VI. ISP Optimizes Subject to $p = q$

In this section, we consider the scenario where the ISP maximises its profit maintaining the user price and the sponsorship price equal. This has the advantage of appearing fair to CPs, since they pay the same price per byte for sponsoring as users do. The present setting is also motivated by the emergence of third party platforms like mCent (`mcent.com`) and Gigato (`gigato.co`) that enable CPs to provide a 'data-back' to their users corresponding to the data consumed on their apps. Platforms like these enable CPs to essentially sponsor their content even without an ISP-operated zero-rating platform.

We remind the reader that since the setting under consideration involves the ISP optimizing the user price p, this amounts to differential data pricing by the ISP (as discussed in Section V). Compared to the setting considered in Section V, the restriction that $p = q$ does at first glance seem to constrain the ISP's profit maximization. Surprisingly however, our results in this section reveal that this restriction does not impact the ISP notably, and the broad conclusions of Section V continue to hold.

As before, we explore the impact of the zero-rating platform in the regime of increasing CP revenue rates. The following theorem sheds light on the ISP's profit-maximizing strategy.

Theorem 3. [ISP's profit maximizing strategy] *As a function of a (for fixed ρ), the profit maximizing strategy for the ISP is the following. There exist thresholds $\tilde{a}_S,\ a_S$, and a_M, where $a_M \geq a_S \geq \tilde{a}_S > 0$, such that:*

1) *For $a < \tilde{a}_S$, the ISP does not operate the zero-rating platform, and sets $p = p^*_{NN}$.*[14]
2) *For $a_S < a < a_M$, the ISP enforces an SS/SN/NS equilibrium, whichever yields the maximum profit.*
3) *For $a > a_M$, the ISP enforces an SN equilibrium, by setting $q = p = a$ such that CP 2 gets zero usage and there is a complete monopoly for CP 1.*

It is instructive to interpret the above strategy by comparing it with the ISP's profit-maximizing strategy when it optimizes p and q freely (see Theorem 2). In both cases, when CP revenue rates are small, the ISP does not benefit from the zero-rating platform. Similarly, in both cases, when CP revenue rates are sufficiently large ($a > a_M$), the ISP enforces an SN equilibrium with a complete monopoly for the firm with the greater revenue rate. However, for intermediate values of a, the optimal strategy of the ISP can differ across the two settings. Moreover, the characterization of the ISP strategy is less precise for intermediate values of a. (For instance, under the constraint that $p = q$, there is the possibility of an interval $a \in (\tilde{a}_S, a_S)$ where the ISP alternates between SN, SS, NS, and not operating the zero-rating platform at all.) This is because the constraint that $p = q$ makes the feasibility regions of different sponsorship configurations more complex (see Lemma 2).

In conclusion, we note that the profit-maximizing strategy of the ISP under the constraint that $p = q$ is structurally similar to that when p and q may be set freely when the CP revenue rates are small, or sufficiently large. Next, we turn to the profit of the ISP and the CPs.

Lemma 10. [ISP profit] *Suppose we take $(a_1, a_2) = (a, \rho a)$, where $\rho \in (0, 1)$ and $a > 0$. Under the ISP's profit maximizing strategy, its profit varies with a (with ρ fixed) as follows. For $a < \tilde{a}_S$, $r_I(a) = p(\theta_1^{NN}(p^*_{NN}) + \theta_2^{NN}(p^*_{NN}))$. For $a > a_S$, $r_I(a)$ is strictly increasing in a. For $a > a_M$, $r_I(a) = ac$.*

Lemma 11. [CP profit] *Suppose we take $(a_1, a_2) = (a, \rho a)$, where $\rho \in (0, 1)$ and $a \geq 0$. Under the ISP's profit maximizing strategy, $r_1(a) = r_2(a) = 0$ for $a > a_M$.*

We see that similar to our conclusions from Section V, the ISP is able to extract all the CP revenue for large enough revenue rates. As before, this is achieved by setting the p and

[14] Unlike in Sections IV and V, there is a subtle distinction in this section between an NN equilibrium (with $p = q$) and the ISP not operating the zero-rating platform. This is because under the settings considered in the previous sections, the ISP could enforce an NN equilibrium by setting the desired user charge and simply setting $q \geq a$, to achieve the same effect as not operating the zero-rating platform. The same is not possible in the present setting.

q to match the higher of the CP revenue rates, such that the sponsoring CP does not have an incentive to stop sponsoring, but also ends up passing on all its revenue to the ISP.[13].

Finally, we note that users are not guaranteed to benefit from the zero-rating platform when the ISP optimizes p (in spite of the the constraint that $p = q$).

Lemma 12. [User surplus] *Suppose that some user class k,*

$$\psi_{1,k}(\theta_1^{NN}(p))+\psi_{2,k}(\theta_2^{NN}(p))-p(\theta_1^{NN}(p)+\theta_2^{NN}(p)) > \psi_{1,k}(c),$$

where p denotes the prevailing user charge before the introduction of the zero-rating platform. Then for $a > a_M$, users of class k have a lower surplus with the zero-rating platform than they did without.

In conclusion, we see that the constraint that $p = q$ does not inhibit the ISP significantly (compared to the case where both p and q are can be set freely), particularly when the CP revenue rates are large. When the CP revenue rates are large, the ISP can skew the CP marketplace to extract most of the CP revenue. Moreover, even the user base may suffer as a consequence of the zero-rating platform.

Numerical results: Finally, we present some numerical results (see Figures 5 and 6) to visualize the conclusions of this section. Again, we use the same settings as described in Section IV.

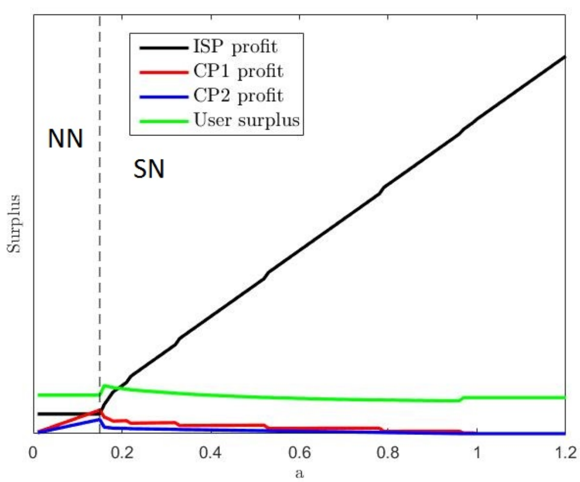

Fig. 5. Surplus of ISP, CPs, and users as a function of a, when ISP optimizes with $p = q$ $(C = 30,\ \rho = 0.6)$

In both Figures 5 and 6, note that there is a single thresold in a, beyond which the ISP induces an SN equilibrium. We observe that CP profits drop sharply beyond the threshold. However, while users suffer due to sponsorship in the setting corresponding to Figure 5, they benefit from sponsorship in the setting corresponding to Figure 6.

VII. Discussion and Conclusion

In this paper, we analyze an ISP-operated zero-rating platform, where CPs can zero-rate their content by paying the ISP a sponsorship fee. We show that being able to charge on both sides of the two-sided market it serves grants the ISP substantial market power. Specifically, we show that the ability to set the sponsorship price allows the ISP to extract a sizeable fraction of the CP surplus. This power is even more magnified when the ISP can also optimize the user-side price. Moreover, the optimal strategy of the ISP often results in a skewing of the CP marketplace, with the CPs as well as the user base being worse off. From a policy standpoint, our work makes a case against unregulated discriminatory pricing by ISPs. This complements the conclusion of [7], which suggests that carefully regulated discriminatory priciing can be socially desirable.

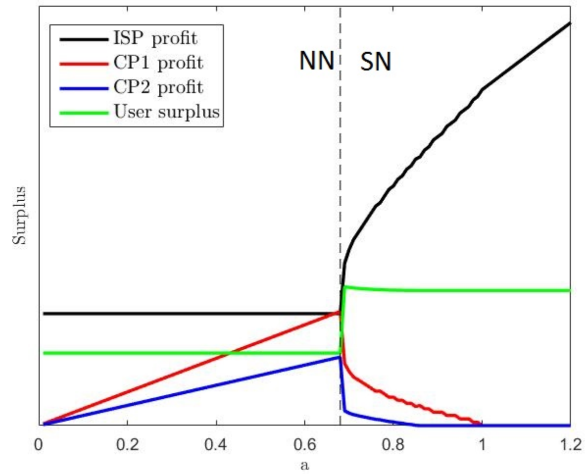

Fig. 6. Surplus of ISP, CPs, and users as a function of a, when ISP optimizes with $p = q$ $(C = 4,\ \rho = 0.6)$

Our theoretical conclusion that zero-rating will drive consumption away from non sponsored content is now supported by empirical evidence. `dflmonitor.eu` has reported that the ISPs that provide zero rated content actually sell significantly less bandwdith to end users than those that do not zero-rate. Specifically, they have found that average and the median value of the number of bytes sold is halved in networks that allow zero rating.

It is important at this point to clarify the scope of our model and our conclusions. Our leader-follower interaction model assumes the ISP as the leader and the CPs as followers. This is natural when a 'large' ISP operates a zero-rating platform for 'smaller' CPs. Examples of such platforms include *Sponsored Data* from AT&T and *FreeBee Data* from Verizon in the US, and the erstwhile *Airtel Zero* from Airtel in India. However, it should be noted that there are also situations where the dominance is reversed, e.g., the interaction between small ISPs and large CPs like Google and Facebook. Capturing such interactions (which typically involve peering arrangements; see [30], [31]), would require very different models, and is beyond the scope of this paper.

Our analysis is based on two key assumptions, which we discuss next.

1) *Single ISP:* Throughout this paper, we considered a single ISP for analytical tractability. This assumption, which essentially boils down to ignoring the user churn

across ISPs, is standard in the literature (see, for example, [7], [18]–[22]). However, it would be interesting to understand the impact of inter-ISP competition in discriminatory pricing regimes. This is the focus on ongoing work.

2) *Congestion modeling:* Our model does not capture the impact of congestion on user behavior; this is a departure from the modeling approach adopted in [7], [20]–[22]. The tractability afforded by this modeling simplification allows us make explicit conclusions on the impact of zero-rating on the CP marketplace, as well as ISP, CP, and user surplus. One justification for this assumption is that market forces automatically drive the ISP to provide a certain level of QoS to its users.

Next, we make some remarks on other natural extensions of our model.

1) We have assumed that the sponsorship price is the same for both CPs. It is easy to see that allowing the ISP to set separate prices to the CPs only gives it more power and does not change the conclusions.
2) Some zero-rated content is of lower quality (e.g., lower bit rate video, websites with lower resolution images) than those that are not zero-rated. It would be interesting to capture the effect of this explicitly in the user utility functions.
3) We have assumed an all-or-nothing model of sponsorship. Allowing the CP to sponsor only a fraction of the content, or a fixed volume, should make the model richer.

It is also instructive to compare our conclusions with those from the analysis of [32] and [21] in which the CPs act independently of each other but the ISP action is constrained by its capacity. Both these papers conclude that zero-rating will be preferred by the 'larger' CPs. From our model of the strategic behavior of the CPs, and the ISP's knowledge, we conclude that the ISP will in fact *force* the larger CP to sponsor through its pricing policy. Indeed, this sponsoring is not necessarily beneficial to either CP.

Finally, we note that the utility of models like the one in this paper is in the insights they provide (in this case, into behavioral tendencies of the parties involved) rather than in the exact quantitative characterisations. Specifically, we conclude that if the ISP is allowed to set prices on both sides of the two-sided market it serves, it will tend to extract the CP surplus by raising both the user price and the sponsorship price; implications like zero profit for the CPs are artefacts of the model that must be interpreted with the appropriate grain of salt.

Appendix A
Proof of Lemma 1

Statement 1 is a trivial consequence of the monotonicity of the utility functions $\psi_{i,k}(\cdot)$.

Proof of Statement 2: We first prove that $\theta_1^{SN} > \theta_1^{SS} > 0$. It is easy to see that $\theta_{1,k}^{SN}$ is the maximizer of the strictly concave function

$$f_k(x) = \psi_{1,k}(x) + \psi_{2,k}(c_k - x) - p(c_k - x)$$

over $[0, c_k]$. Similarly, $\theta_{1,k}^{SS}$ is the maximizer of the strictly concave function

$$g_k(x) = \psi_{1,k}(x) + \psi_{2,k}(c_k - x)$$

over $[0, c_k]$. Since $f_k(x) = g_k(x) + px - pc_k$, it follows that $\theta_{1,k}^{SN} \geq \theta_{1,k}^{SS}$, which implies $\theta_1^{SN} \geq \theta_1^{SS}$. To show that the preceding inequality is strict, it suffices to show that

$$\theta_{1,\hat{k}}^{SN} > \theta_{1,\hat{k}}^{SS}. \tag{1}$$

To show this, note that by Assumption 1, $\theta_{1,\hat{k}}^{SS} \in (0, c_{\hat{k}})$ and therefore satisfies

$$g'_{\hat{k}}(\theta_{1,\hat{k}}^{SS}) = 0 \iff \psi'_{1,\hat{k}}(\theta_{1,\hat{k}}^{SS}) = \psi'_{2,\hat{k}}(c_{\hat{k}} - \theta_{1,\hat{k}}^{SS}). \tag{2}$$

Thus,

$$f'_{\hat{k}}(\theta_{1,\hat{k}}^{SS}) = p > 0,$$

which implies (1).

Next, we prove that for any action m of CP 2, $\theta_1^{Sm} > \theta_1^{Nm}$.

Case 1: $m = S$. Note that $\theta_{1,k}^{NS}$ is the maximizer of the strictly concave function

$$h_k(x) = \psi_{1,k}(x) + \psi_{2,k}(c_k - x) - px = g_k(x) - px,$$

over $[0, c_k]$. If follows then that $\theta_{1,k}^{SS} \geq \theta_{1,k}^{NS}$, which implies that $\theta_1^{SS} \geq \theta_2^{NS}$. To prove that the preceding inequality is strict, we now show that

$$\theta_{1,\hat{k}}^{SS} > \theta_{1,\hat{k}}^{NS}. \tag{3}$$

As before, note that by Assumption 1, $\theta_{1,\hat{k}}^{SS} \in (0, c_{\hat{k}})$. It then follows from (2) that

$$h'_k(\theta_{1,\hat{k}}^{SS}) = -p < 0,$$

which implies (3).

Case 2: $m = N$. We first show that

$$\theta_{1,k}^{SN} \geq \theta_{1,k}^{NN} \quad \forall\ k. \tag{4}$$

If $\theta_{1,k}^{NN} = 0$, (4) is trivially true. If $\theta_{1,k}^{NN} > 0$, then it is easy to show that

$$\psi'_{1,k}(\theta_{1,k}^{NN}) - p \geq \max\left[0, \psi'_{2,k}(\theta_{2,k}^{NN}) - p\right].$$

It follows from the above inequality that

$$\begin{aligned}\psi'_{1,k}(\theta_{1,k}^{NN}) &\geq \psi'_{2,k}(\theta_{2,k}^{NN}) \\ &\geq \psi'_{2,k}(c_k - \theta_{1,k}^{NN}).\end{aligned}$$

The last inequality above follows from the concavity of $\psi_{2,k}(\cdot)$. We therefore have

$$f'_k(\theta_{1,k}^{NN}) \geq p > 0,$$

which implies (4). It follows from (4) that

$$\theta_1^{SN} \geq \theta_k^{NN}.$$

To prove that the above inequality is strict, it suffices to show that

$$\theta_{1,\hat{k}}^{SN} > \theta_{1,\hat{k}}^{NN}. \tag{5}$$

Since we have already shown that $f'_{\hat{k}}(\theta_{1,\hat{k}}^{NN}) > 0$, (5) holds so long as $\theta_{1,\hat{k}}^{NN} < c_{\hat{k}}$. For the purpose of obtaining a contradiction, let us assume that $\theta_{1,\hat{k}}^{NN} = c_{\hat{k}}$. This in turn would imply that $\psi'_{1,\hat{k}}(c_{\hat{k}}) \geq \psi'_{2,\hat{k}}(0)$, which implies that $g'_{\hat{k}}(c_{\hat{k}}) \geq 0$, which implies $\theta_{1,\hat{k}}^{SS} = c_{\hat{k}}$. But this contradicts Assumption 1. Thus, (5) is proved.

Proof of Statement 3 and 4: These proofs are elementary and are omitted.

Appendix B
Proof of Lemma 1

Proof of Statement 1: For NN to be an equilibrium, neither CP should have a unilateral incentive to sponsor their content. For CP 1, this condition is

$$a_1\theta_1^{NN}(p) \geq (a_1 - q)\theta_1^{SN}(p)$$
$$\iff q \geq a_1\left(1 - \frac{\theta_1^{NN}(p)}{\theta_1^{SN}(p)}\right) = a_1\alpha(p).$$

Similarly, CP 2 has no unilateral incentive to sponsor under NN if and only if $q \geq a_2\beta(p)$. Thus, we conclude that NN is a Nash equilibrium between the CPs if and only if $q \geq \max(a_1\alpha(p), a_2\beta(p))$.

Proofs of Statements (2), (3), and (4) follow along similar lines, and are omitted.

Appendix C
Proofs of results in Section IV

This section is devoted to proving the results stated in Section IV.

Proof of Theorem 1:

Note that an SS equilibrium is always feasible, and the maximum profit of the ISP under an SS equilibrium equals $ac\min(\gamma, \rho\delta)$. An SN equilibrium, if feasible, yields a maximum profit of $a\alpha\theta_1^{SN} + p\theta_2^{SN}$. Similarly, an NS equilibrium, if feasible, yields a maximum profit of $a\rho\beta\theta_2^{NS} + p\theta_1^{NS}$.

Let $g(a)$ denote the maximum profit the ISP makes from an SS/SN/NS equilibrium. Clearly, $g(a)$ is strictly increasing in a, and grows unboundedly as $a \to \infty$. The statement of the theorem now follows, with

$$a_S = g^{-1}(p(\theta_1^{NN} + \theta_2^{NN})).$$

Proof of Lemma 3:

The statement of the lemma follows from the arguments in the proof of Theorem 1: Since an SS equilibrium is always feasible, $g(a) \geq ac\min(\gamma, \rho\delta)$.

Proof of Lemma 4:

It suffices to verify the claim of the lemma for an SN equilibrium; a symmetric argument applies for NS as well. Under an SN equilibrium, the ISP sets $q = a\alpha$.

Thus, the sponsoring CP (CP 1) has a profit equal to

$$\left(a - a\left(1 - \frac{\theta_1^{NN}}{\theta_1^{SN}}\right)\right)\theta_1^{SN} = a\theta_1^{NN},$$

which matches the profit of CP 1 in the absence of the zero-rating platform.

On the other hand, the non-sponsoring CP (CP 2) has a profit equal to

$$a_2\theta_2^{SN} \leq a_2\theta_2^{NN}.$$

Proof of Lemma 5:

Under an SS equilibrium, the ISP sets $q = a\min(\gamma, \rho\delta)$. This implies that either $q = a\gamma$, or $q = a\rho\delta$.

If $q = a\gamma$, note that CP 1's profit equals

$$\left(a\left(1 - \frac{\theta_1^{NS}}{\theta_1^{SS}}\right) - a\right)\theta_1^{SS} = a\theta_1^{NS} \leq a\theta_1^{NN}.$$

Alternatively, if $q = a\rho\delta$, then CP 2's profit equals

$$\left(a\rho\left(1 - \frac{\theta_2^{SN}}{\theta_2^{SS}}\right) - a\rho\right) = a\rho\theta_2^{SN} \leq a\rho\theta_2^{NN}.$$

Proof of Lemma 6:

This proof is elementary and is omitted.

Appendix D
Proofs of results in Section V

This section is dedicated to proving the results stated in Section V.

Proof of Theorem 2:

The profit of the ISP under an SS equilibrium is given by

$$r_I^{SS}(a) = \max_{p\in(0,\bar{p}]}\left[ac\ \min(\gamma(p), \rho\delta(p))\right]$$
$$= ac\rho.$$

The above equality holds because $\gamma(p)$ and $\delta(p)$ and non-decreasing, implying that $\bar{p}$ is a maximizer of the above optimization.

The ISP's profit under an SN equilibrium equals

$$r_I^{SN}(a) = \max_{p\in(0,\bar{p}]}\left[p\theta_2^{SN}(p) + a\alpha(p)\theta_1^{SN}(p)\right] \quad (6)$$
$$= \max_{p\in(0,\bar{p}]}\left[p\theta_2^{SN}(p) + a(\theta_1^{SN}(p) - \theta_1^{NN}(p))\right]. \quad (7)$$

Setting $p = \bar{p}$ in the above maximization yields a profit of ac, implying that $r_I^{SN}(a) > r_I^{SS}(a)$. Finally, the profit of the ISP under an NS equilibrium equals

$$r_I^{NS}(a) = \max_{p\in(0,\bar{p}]}\left[p\theta_1^{NS}(p) + a\rho(\theta_2^{NS}(p) - \theta_2^{NN}(p))\right]. \quad (8)$$

Let $g(a)$ denote the maximum profit the ISP can derive out of a sponsorship configuration, i.e., $g(a) = \max(r_I^{SN}(a), r_I^{NS}(a))$. Clearly, $g(a)$ is strictly increasing in a, and grows unboundedly as $a \to \infty$. Statements 1 and 2 of the theorem now follow, with

$$a_S = g^{-1}(r_I^{NN}).$$

Focusing now on the maximization (6), note that $f(p) := (\theta_1^{SN}(p) - \theta_1^{NN}(p))$ is non-decreasing in p, with $f(\bar{p}) = c$. Thus, for large enough a, $\bar{p}$ is a maximizer of the optimization in (6), and $r_I^{SN}(a) = ac$. (Note that in this case, CP 2 gets zero usage and the ISP sets $q = a$.) Similarly, for large enough a, $\bar{p}$ is a maximizer of the optimization in (8), and $r_I^{NS}(a) = \rho ac$. This implies the existence of a threshold a_M beyond which the ISP enforces an SN equilibrium wherein CP 2 gets zero usage, and the ISP extracts the entire revenue of CP 1.

Proof of Lemma 7:

The statement of the lemma follows easily from the arguments in the proof of Theorem 2.

Proof of Lemma 8:

The claim that both CPs make zero profit for $a > a_M$ follows directly from the arguments in the proof of Theorem 2. If the CPs are identical from the user perspective, then it is easy to see that $r_I^{SN}(a) > r_I^{NS}(a)$, meaning that the ISP prefers an SN configuration to an NS configuration. Under the SN configuration, the profit of CP 1 is given by $r_1(a) = a\theta_1^{NN}(p^*_{SN})$, where p^*_{SN} is the maximizer of (6). Since p^*_{SN} is non-decreasing in a, the claim regarding the profit of CP 1 follows. The monotonicity of p^*_{SN} with respect to a also implies the claim regarding the profit of CP 2 since $r_2(a) = \rho a \theta_2^{SN}(p^*_{SN})$.

Proof of Lemma 9:

This proof is elementary and is omitted.

APPENDIX E
PROOFS OF RESULTS IN SECTION VI

In this section, we provide the proof of Theorem 3. We omit the proofs of Lemma 10 (on ISP profit), Lemma 11 (on CP profit), and Lemma 12 (on user surplus), since these can be proved along the lines of the corresponding results in Section V.

Proof of Theorem 3:

The first statement of Theorem 3 follows from the fact that under any (feasible) sponsorship configuration M, the profit of the ISP equals $q(\theta_1^M(q) + \theta_2^M(q)) \leq ac$. Thus, not operating the zero-rating platform is optimal for the ISP when a is small enough.

The second and third statement follow from the following obvervations:

1) For $a \geq \bar{p}$, an SN equilibrium is feasible, the optimal strategy of the ISP (under SN) being to set $p = q = a$, resulting in profit ac. Note that in this case, there would be a complete monopoly for CP 1, and both CPs would receive zero profit.
2) Moreover for $a \geq \bar{p}$, the ISP profit under NS/SS (if feasible) is at most $\rho ac < ac$.

We conclude that there exists a threshold a_s such that it is optimal for the ISP to induce some sponsorship configuration for $a > a_S$, and another threshold a_M such that for $a > a_M$, it is optimal for the ISP to induce an SN equilibrium setting $p = q = a$. Clearly, $a_S \leq a_M$. This completes the proof.

REFERENCES

[1] S. Sen, C. Joe-Wong, S. Ha, and M. Chiang, "A survey of smart data pricing: Past proposals, current plans, and future trends," *ACM Computing Surveys*, vol. 46, no. 2, pp. 15:1–15:37, November 2013.

[2] A. Troianovski, "Cellphones are eating the family budget," http://www.wsj.com/articles/SB10000872396390444083304578018731890309450, 28 October 2012.

[3] P. Marceux, "Special report: The consumer," http://blog.euromonitor.com/2013/08/special-report-the-telecom-consumer-in-2020.html, 27 August 2013.

[4] ," http://www.huffingtonpost.ca/2014/10/16/telecom-costs-canada_n_5999242.html, 16 October 2014.

[5] J. W. Baek, "The effects of the internet and mobile services on urban household expenditures: The case of south korea," *Telecommunications Policy*, vol. 40, no. 1, pp. 22–38, February 2016.

[6] M. Page, L. Rossi, and C. Rand, "A viable model for the future internet," Tech. Rep., A. T. Kearney, 2010.

[7] R. T. B. Ma, "Subsidization competition: Vitalizing the neutral internet," in *Proc. of ACM CoNEXT*, 2014.

[8] V. Misra, "Zero rating: Slows down innovation, distorts competition & fractures the internet," http://www.medianama.com/2015/04/223-zero-rating-slows-down-innovation-distorts-competition-fractures-the-internet, 17 April 2015.

[9] B. van Schewick, "T-mobile's Binge On violates key net neutrality principles," 2016.

[10] B. van Schewick, "Comments on public consultation CRTC 2016-192 – examination of differential pricing of Internet data plans," Stanford Law School, 28 June 2016.

[11] A. Soni, "India deals blow to facebook in people-powered 'net neutrality' row," 08 February 2016.

[12] R. Hahn and Scott Wallsten, "The economics of net neutrality," Tech. Rep. Related Publication 06-13, AEI-Brookings Joint Center for Regulatory Studies, April 2006.

[13] J. Musacchio, G. Schwartz, and J. Walrand, "A two-sided market analysis of provider investment incentives with an application to the net-neutrality issue," *Review of Network Economics*, vol. 8, no. 1, March 2009.

[14] J. P. Choi and B. C. Kim, "Net neutrality and investment incentives," *RAND Journal of Economics*, vol. 41, no. 3, pp. 446–471, Autumn 2010.

[15] N. Economides and J. Tag, "Network neutrality on the internet: A two-sided market analysis," *Information Economics and Policy*, vol. 24, no. 2, pp. 91–104, June 2012.

[16] N. Economides and E. Hermalin, "The economics of network neutrality," *RAND Journal of Economics*, vol. 43, no. 4, pp. 602–629, 2012.

[17] R. ElDelgawy and R. J. La, "Interaction between a content provider and a service provider and its efficiency," in *Proc. of IEEE ICC*, 2015.

[18] M. Andrews, U. Ozen, M. I. Reiman, and Q. Wang, "Economic models of sponsored content in wireless networks with uncertain demand," in *Proc. of IEEE INFOCOM*, 2013, pp. 3213–3218.

[19] M. H. Lotfi, K. Sundaresan, S. Sarkar, and M. A. Khojastepour, "The economics of quality sponsored data in non-neutral networks," http://arxiv.org/abs/1501.02785, September 2015.

[20] C. Joe-Wong, S. Ha, and M. Chiang, "Sponsoring mobile data: An economic analysis of the impact on users and content providers," in *Proc. of IEEE INFOCOM*, 2015.

[21] L. Zhang, W. Wu, and D. Wang, "Sponsored data plan: A two-class service model in wireless data networks," in *Proceedings of ACM SIGMETRICS*, June 2015, pp. 85–96.

[22] Xin Wang, Richard TB Ma, and Yinlong Xu, "On optimal two-sided pricing of congested networks," *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, vol. 1, no. 1, pp. 7, 2017.

[23] M. Andrews, G. Burns, and H. Lee, "Calculating the benefits of sponsored data for an individual content provider," in *Proceedings of Conference on Information Sciences and Systems (CISS)*, 2014.

[24] C. Courcoubetis, L. Gyarmati, N. Laoutaris, P. Rodriguez, and K. Sdrolias, "Negotiating premium peering prices: A quantitative model with applications," *ACM Transactions on Internet Technology*, vol. 16, no. 2, pp. 1–22, April 2016.

[25] C. Courcoubetis, K. Sdrolias, and R. Weber, "Paid peering: pricing and adoption incentives," *Journal of Communications and Networks*, vol. 18, no. 6, December 2016.

[26] Ankur A Kulkarni and Uday V Shanbhag, "An existence result for hierarchical stackelberg v/s Stackelberg games," *IEEE Transactions on Automatic Control*, vol. 60, no. 12, pp. 3379–3384, 2015.

[27] N. Pahwa, "Some thoughts on T-Mobile zero-rating select video services," http://www.medianama.com/2015/11/223-tmobile-zero-rating-video/, 12 November 2015.

[28] S. N. Vikas, "Whatsapp signs telco deal with airtel; claims to have 50m active users in india," http://www.medianama.com/2014/05/223-whatsapp-airtel-50m-india/, 12 May 2014.

[29] R. Mukherjee, "Tata Docomo users can stream YouTube videos for Rs. 9 per day," http://www.medianama.com/2014/07/223-tata-docomo-youtube/, 08 July 2014 (Retrieved on 10 July 2016).

[30] Arpit Gupta, Matt Calder, Nick Feamster, Marshini Chetty, Enrico Calandro, and Ethan Katz-Bassett, "Peering at the internet's frontier: A first look at isp interconnectivity in africa," in *International Conference on Passive and Active Network Measurement*, 2014, pp. 204–213.

[31] Craig Labovitz, Scott Iekel-Johnson, Danny McPherson, Jon Oberheide, and Farnam Jahanian, "Internet inter-domain traffic," in *ACM SIGCOMM Computer Communication Review*. ACM, 2010, vol. 40, pp. 75–86.

[32] C. Joe-Wong, S. Ha, and M. Chiang, "Sponsoring mobile data: An economic analysis of the impact on users and content providers," in *Proc. of IEEE INFOCOM*, 2015.

Kunal Phalak obtained his Bachelors and Masters in Electrical Engineering from IIT Bombay with specialization in Communications and Signal Processing in 2016. His interests are in game theory and understanding the new age economy built on the Internet and networks. After graduating from IIT Bombay, Kunal spent about two years with the Advanced Analytics Practice in McKinsey & Company. He is currently working as an Investment Associate with Vertex Venture Capital investing in technology & internet businesses in India and South East Asia.

D. Manjunath received his BE from Mysore University, MS from IIT Madras and PhD from Rensselaer Polytechnic Institute in 1986, 1989 and 1993 respectively. He has been with the Electrical Engineering Department of IIT Bombay since July 1998 where he is now an Institute Chair Professor. He has previously worked in the Corporate R & D Center of General Electric in Scehenectady NY (1990), Computer and Information Sciences Department of the University of Delaware (1992–93), Computer Science Department, University of Toronto (1993–94) and the Department of Electrical Engineering of IIT Kanpur (1994–98). At IIT Bombay, he was Head of the Computer Centre during 2011-15. His research interests are in the general areas of communication networks and performance analysis. His recent research has concentrated on random networks with applications in wireless and sensor networks, network pricing and queue control. He is a recipient of the best paper award at ACM SIGMETRICS 2010. He is an associate editor of *IEEE Transactions on Networking*, *Queueing Systems: Theory and Applications,* and of *Sadhana: The Proceedings of the Indian Academy of Sciences.* He was TPC chair for COMSNETS 2011 and NCC 2015 and general chair for ACM MobiHoc 2013 and COMSNETS 2015. He is a coauthor of two textbooks, *Communication Networking: An Analytical Approach* (May 2004) and *Wireless Networking* (Apr 2008), both of which are published by Morgan-Kaufman Publishers.

Jayakrishnan Nair received his BTech and MTech in Electrical Engg. (EE) from IIT Bombay (2007) and Ph.D. in EE from California Inst. of Tech. (2012). He has held post-doctoral positions at California Inst. of Tech. and Centrum Wiskunde & Informatica. He is currently an Assistant Professor in EE at IIT Bombay. His research focuses on modeling, performance evaluation, and design issues in queueing systems and communication networks.