Capacity Expansion of Neutral ISPs via Content Provider Participation: The Bargaining Edge

Anand Kalvit[a], Saurabh Pinjani[b], Gaurav Kasbekar[b], D. Manjunath[b], Jayakrishnan Nair[b]

[a] *Graduate School of Business, Columbia University.*
[b] *Department of Electrical Engineering, IIT Bombay.*

Abstract

Many Internet service providers (ISPs) operate under network neutrality regulations that forbid differential QoS or differential pricing. This leads to lower profitability for ISPs. However, increasingly bandwidth-hungry content is making the consumers demand significantly improved ISP infrastructure. With the risk of poor consumer experience on them, ISPs have to invest in their infrastructure but with limited scope for monetisation. Hence, they are asking content providers (CPs) to contribute towards ISP capacity expansion. In this paper we explore network neutral capacity expansion sponsored by voluntary contributions from CPs.

We consider the scenario where CPs lead in paying ISPs a 'q-charge,' which has to be used for capacity expansion. Since ISP capacity expansion can benefit all the CPs, and possibly even the ISP, selfish CPs will determine their contributions strategically. We consider three models for CPs to interact in determining the charge—a cooperative model, a non-cooperative model, and a bargaining model. Our analysis reveals a rather surprising result. We show that the bargaining model leads to a higher investment in the ISP infrastructure than even the cooperative model. This leads us to recommend policies that promote transparency in the interconnection agreements between CPs and ISPs.

Keywords:
network neutrality, Internet service providers, capacity expansion, paid peering, Internet economics, interconnection markets

---

[⋆] A preliminary version of this work appeared in Valuetools 2017 [1].

1. Introduction

The Internet has traditionally been a hierarchical network with Tier 1 ISPs at the top and access ISPs and their subscribers (users) at the bottom of the hierarchy. The Tier 1 ISPs interconnected to form a fully connected mesh with settlement-free peering. The Access ISPs bought ‘transit service’ for a fee from one or more intermediate ISPs to provide universal connectivity to their subscribers. This hierarchical topology has started to break down with several networks peering around Tier 1 ISPs and resulting in ‘donut peering’ structures [2]; many of these are paid peering arrangements (as opposed to settlement-free peering) where the ‘source’ ISP pays the ‘target’ so that its traffic can bypass a longer route through a transit ISP. With a small number of content providers dominating Internet traffic[1], such donut peering structures are being strengthened with the CPs also connecting directly to the access ISPs. For example, Netflix Open Connect[2] is a platform for Netflix to directly connect with an access ISP. Google maintains multiple points of presence (POP) in most markets and the access ISPs may connect to them directly.

Direct interconnections between the ISPs and the CPs benefits both the CPs and the end users. The CPs save on transit costs and the users see improved quality of experience (QoE). The improved QoE in turn benefits the CPs with increased customer stickiness. The direct interconnection though imposes a cost to the ISP because the improved QoE can drive up the demand. Hence, as has been pointed out in [3], to avoid the risk of poor user experience, the ISP needs to improve its infrastructure through additional investment. This increased cost to the ISP may not be offset by the possible savings in transit costs.

To recoup such an investment, the ISP can extract the system surplus from the user-side and/or from the CP side. User-side surplus may be extracted through pricing innovations, called smart data pricing; see [4] for

[1] According to the Sandvine report (see `www.sandvine.com`) NetFlix (55.5%), YouTube (17.5%) and seven other streaming applications together contribute about 80% of download traffic on wireline networks. For mobile networks, the top eight CPs contribute more than 80% of the downloads.

[2] https://media.netflix.com/en/company-blog/how-netflix-works-with-isps-around-the-globe-to-deliver-a-great-viewing-experience

several examples. However, the ISPs face several obstacles in this. Firstly, market expectations have been shaped by flat fee regimes and a simple pricing structure is the *de riguer* in many markets. Secondly, in many markets, the pricing regimes can be close to marginal cost because of strong inter-ISP competition. Finally, regulatory requirements like net neutrality stipulations severely limit the ability of the ISP to extract user-side surplus via the many schemes outlined in [4]. There appears to be more manoeuvring room for the ISP to extract CP-side surplus. Also, as has been described in the analysis in [5], with increasing asymmetry between ISP and CP revenues, it is natural for ISPs to demand that CPs contribute to the costs of ISP infrastructure.

Currently, the dominant mechanism for surplus transfer from CPs to access ISPs is via bilateral paid peering arrangements. These arrangements, the terms of which are typically private, generally involve the CP paying the ISP in return for improved user QoE, which in turn is effected via higher interconnection bandwidth (the peering arrangement between Netflix and Comcast being a celebrated case; see [6]), or by placing the content of the CP in the data centre of the ISP (and thus very close the end user). Another common practice is for large CPs to establish settlement-free Points-of-Presence (PoPs), like Google's Edge Network[3], that connect them directly to eyeball ISPs. Note that such practices indirectly subsidise the cost of the ISP for providing a high quality of service to its users.

While the above mentioned practices do not at present constitute net neutrality violations, one might argue that they are non-neutral in spirit. After all, end users would have a superior QoE from those CPs that have peering agreements with their ISP, resulting in a non-uniform Internet access. Indeed, the argument has been made that the definition of net neutrality should be expanded to ensure comparable and adequate QoE for all Internet services [6].

The goal of this paper is to propose an alternative, net neutral, mechanism for CPs to contribute towards capacity expansion of ISPs. Specifically, we propose a neutral network with CPs leading in contributing to the ISP infrastructure via 'voluntary' payments—called q-charges in this paper. In our model, the CPs pay the q-charge in their own self interest; this is not unreasonable because, as we saw earlier, only a small number of CPs dominate Internet traffic and they are the ones that would be willing to pay the

[3]`https://peering.google.com`

peering charges. A further motivation for this model is the belief that many CPs wield significant power over the access ISPs, especially in markets where the latter do not enjoy a monopoly.

1.1. Background

The public discourse on net neutrality—arguments for and against differential pricing and differential QoS—has been vigorous for quite some time now and many countries have adopted strong regulations favouring network neutrality. Interconnection agreements, especially between the access ISPs and content provider (CP) networks like those of YouTube, Netflix, and Facebook, or content delivery networks (CDNs) like Akamai and Lightstream can also produce effects similar to neutrality violation. However, not much attention has been paid to the interconnection markets and they have largely remained unregulated.

That interconnection agreements would determine the structure of the Internet was recognized in [7] in the prescient observation that "interconnection agreements do not just route traffic in the Internet, they also route money." It was further recommended in [7] that any policy on interconnections should focus on transparency into the workings of the interconnection markets. However, to the best of our knowledge, the public discourse on interconnection regulation has not yet happened and there are not many details that are public. However, several governments are exploring this matter. This paper develops economic models to inform such regulations.

Specifically, this paper proposes a transparent, neutral mechanism for CPs to contribute towards capacity expansion of ISPs (in effect, bringing about a surplus transfer from CPs to ISPs). Indeed, our results suggest that transparency can actually lead to significantly improved infrastructure investments as compared to an unregulated regime with bilateral arrangements between ISPs and CPs.

1.2. Preview

In Section 2, we describe our model for voluntary contributions towards ISP capacity expansion by CPs. In Section 3, we consider three mechanisms for CPs to determine their q-charges. First, we consider a cooperative model in which the CPs form a single coalition that seeks to maximise the total CP surplus. Next we consider a non-cooperative setup where the CPs strategically decide the q-charge. Finally, we consider a Nash bargaining problem between the CPs to determine the ISP investment. Our key finding in this section is that the bargaining based peering charges leads to a higher investment than even the cooperative model where the CPs maximize their net profit. A second important finding is that in the non-cooperative setup, in most cases, only one CP will contribute to capacity expansion and the others will free-ride. In Section 4, we illustrate the results with numerical examples and characterize the differences in the outcome from the three models.

In Section 5, we provide some policy guidelines for the interconnection market based on our analyses and conclude with a discussion on extensions and future work.

1.3. Previous Work

There is substantial literature on the economics of non-neutral networks. 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 has been studied in, among others, [8, 9, 10]. Analysis of the effect of discriminatory pricing is being increasingly studied, e.g., [11, 12, 13]. Coalitional game theory was used in [14, 5] to analyze the fair sharing of surplus among the access ISPs and CPs; these papers predict that prevalent settlement arrangements are not stable. Early work on paid peering, e.g., [15, 16, 17], implicitly assumed that all the networks were 'similar'. More recently, [18] suggested a Nash-peering model based on the value of peering.

Our interest in this paper is more along the lines of the work of [19, 20, 21]. In [19], a Nash bargaining model is used to determine peering prices when there is a churn in the system. In [20], using a transactional model for demand, a Stackelberg game with the ISP as the leader and the CPs as the followers is formulated to study peering prices. In [21], a choice model is used to determine the value of direct peering and the peering dynamics are analyzed. We differ from these in that we consider the contribution of the CPs to improving the quality of service seen by the customers. This contribution in turn directly affects their revenues. Hence there is a strategic element to their decision. In this sense, our work has similarities to the public good game (see [22, 23] for different variants) where the public good here is ISP infrastructure. In the classic public good game, all the players receive a uniform benefit from the public good. In our problem of interest, the benefit of each CP is a (CP-specific) function of the ISP service capacity.

2. Model and preliminaries

We consider a system with a single access ISP and K CPs, indexed by $i = 1, 2, \ldots, K$, serving a fixed user population. The CPs can pay the ISP a voluntary q-charge that is to be used for capacity expansion by the ISP, thus increasing the quality of service (QoS) that the users see and hence their consumption of content from the CPs.

The content consumption of the users from the CPs depends on the inherent interest in the corresponding content and also on the quality of service (QoS) that the network provides. This QoS in turn depends on the investment made by the ISP towards its infrastructure. Let μ_0 be the baseline investment without contribution of q-charge by the CPs and let μ be the additional investment enabled by the q-charge on the CP side. The increased QoS is seen by all the users. Further, since the ISP is neutral, the increased consumption cannot be tilted in favour (or against) any CP by the ISP.

The consumption of content from CP i by users is given by $x_{i0} + x_i(\mu)$, where x_{i0} is the 'baseline' consumption, and $x_i(\mu)$ is the extra consumption that is enabled by the additional investment of μ (derived from the q-charges) in the ISP infrastructure. We make the reasonable assumption that $x_i(\mu)$ is continuously differentiable, increasing, and concave in μ with $x_i(0) = 0$ and $\lim_{\mu \to \infty} x'(\mu) = 0$. Note that this model allows consumptions of content from different CPs to be dependent. We further assume that the profitability of a CP is linear in the consumption of its content, i.e., CP i has a revenue of v_i per unit of traffic, which is a natural model for ad-supported services.[4]

Let Q_i be the q-charge paid by CP i, and $Q = (Q_i, 1 \leq i \leq K)$. The q-charges that are paid by the CPs are used by the ISP towards enhancing its infrastructure; for simplicity, we assume that $\mu = \mu(Q) = \sum_{i=1}^{K} Q_i$. This capacity expansion in turn increases consumption which is profitable to each of the CPs. Thus the surplus of each CP will be a function of Q; denote this by $f_i(Q)$. We can now write

$$f_i(Q) \quad = \quad v_i(x_{i0} + x_i\ (\mu(Q))) - Q_i \tag{1}$$

Since our interest is in the extra investment afforded by the CP payments, we will let $\mu(0) = 0$. Also note that $f_i(0) = v_i x_{i0}$. Without loss of generality,

[4]Our model can also be applied to subscription based services as follows. Interpret $x_i(\mu)$ as the growth in the subscriber base of CP i due to a capacity expansion of μ, and v_i as the per-user subscription charge.

we assume $f_i(0) = 0$ since our interest is in the analysis of the incremental benefits of contributions towards capacity expansion by the CPs. Throughout, we assume that users pay a fixed charge p for connection to the ISP. Thus there is no incentive for the ISP to contribute to μ.

We will analyse the complete information game, i.e., the v_i, $x_i(\mu)$ are known to all the CPs. We will also assume that the user charge of p is exogenously determined and is not part of the strategy space of either the ISP or the CPs. Our interest is to analyze the incentives for CPs to pay for ISP capacity expansion.

We will consider the following kinds of interactions between the CPs and the ISP.

- We first consider a cooperative game between the CPs with the objective of maximizing the total profit, determining μ cooperatively. This is a 'benchmark' objective and hence the actual sharing of the μ among the CPs is not relevant. The μ obtained from this analysis is useful for purposes of comparison with the more 'realistic' schemes.

- Next we consider the case when the CPs are non-cooperative in determining Q.

- The third interaction model that we consider is for the CPs to determine the Q using a bargaining framework.

Our interest is to analyze the ISP investment μ under the different interaction models described above. We will use the superscript C to indicate these quantities under the cooperative framework, B for the bargaining framework, and N for the non-cooperative framework. For example μ^C would be the ISP investment with cooperation, and f_i^B the surplus of CP i when they bargain on the q-charge.

3. Determining the q-charges: Analytical results

In this section, we study an idealized cooperative model between CPs, as well as more practical models based on the Nash equilibrium and the Nash bargaining framework. The cooperative model corresponds to the scenario where the CPs form a single coalition seeking to maximize its aggregate profit. At the other extreme, the Nash equilibrium based model captures a non-cooperative setting where each CP seeks to maximize its own profit, in the absence of any signalling between the CPs. Finally, the bargaining framework corresponds to a scenario where the CPs can communicate and 'bargain' to arrive at an agreement on the q-charges paid by each CP. The above models enable us to analyze the impact of non-cooperation as well as bargaining between the CPs.

Our analysis reveals that from the standpoint of capacity expansion, the bargaining based capacity expansion arrangement is more efficient than the cooperative arrangement, which in turn is more efficient than the non-cooperative arrangement. In other words, the strategic interaction that enables a bargaining solution between the CPs is actually beneficial to the user base since it leads to the maximum capacity expansion. Interestingly, except when the CPs are perfectly symmetric, any Nash equilibrium has only one CP contributing towards capacity expansion and the others free-riding. This is undesirable not just because it leads to lower capacity expansion, but also because it can result in a push towards (non-neutral) preferential treatment to traffic of the 'sponsoring' CP. (The latter is not explicitly modeled in this paper.)

From a regulatory standpoint, our analysis suggests that it is socially beneficial to have a transparent platform for CPs to commit on their contributions towards network neutral capacity expansion. Indeed, this might be preferable to the present practice where CPs enter into bilateral (and seemingly non-neutral [6]) peering arragements with ISPs, the terms of which are kept private.

Since we have assumed $\mu = \sum_{i=1}^{K} Q_i$, the CP surplus functions are given by $f_i = v_i x_i(\sum_{i=1}^{K} Q_i) - Q_i \quad (i \in \{1, 2, ...K\})$.

Remark 1. *Under our 'lumpsum' model for q-charges, it is not meaningful to consider the ISP as being strategic. Indeed, since the ISP surplus is not tied to the capacity expansion (q-charges as well as user payments being flat and not volume based), the ISP has no incentive to increase capacity. If the q-charges are volume-based, it is indeed meaningful to consider models with a strategic ISP. This is discussed in Section 5.*

3.1. Cooperative peering

We begin our analyses by first considering the cooperative regime—the idealised setting where the CPs act as a single coalition seeking to maximize its net surplus,

$$\sum_{i=1}^{K} v_i x_i(\mu) - \mu.$$

The first term above is the revenue from the increase in user data consumption due to capacity expansion μ, and the second is the q-charge paid by the coalition. This is an 'idealisation' in the sense that we disregard the strategic interactions between the CPs. Thus, it is natural to benchmark the capacity expansion under this model to that under more practical settings where the CPs act strategically.

The concavity assumptions on the x_i lead to the following elementary characterization of the cooperative capacity expansion, denoted by μ^C.

Lemma 1.

$$\mu^C = \begin{cases} 0 & \textit{if } \left(\sum_{i=1}^{K} v_i x_i'(0)\right) \leq 1 \\ h^{-1}(1) & \textit{otherwise} \end{cases},$$

where $h(\mu) := \sum_{i=1}^{K} v_i x_i'(\mu)$, *and* $h^{-1}(\cdot)$ *denotes the inverse of the (monotone) function* $h(\cdot)$.

Note that if $\left(\sum_{i=1}^{K} v_i x_i'(0)\right) \leq 1$ then the CP coalition does not have an incentive to invest in capacity expansion, since the resultant usage increase does not generate sufficient revenue.

While the above cooperative model does prescribe the capacity expansion uniquely, it does not prescribe the q-charges paid by the respective CPs. Cooperative game theory provides several solution concepts to capture this settlement between the CPs, e.g., the Shapely value mechanism [24]. However, since our focus is primarily on the capacity expansion corresponding to the cooperative regime (which forms the benchmark against which we compare the capacity expansion under the non-cooperative and bargaining models), we do not address the issue of cooperative settlement in this paper.

3.2. Non-cooperative peering

Next, we consider the non-cooperative setting, where CPs decide on their contributions towards capacity expansion selfishly and without coordination. In this case, it is natural to consider Nash equilibria between the CPs, with Q_i being the 'action' of CP i. Note that $Q^N = (Q_1^N, Q_2^N, ..Q_K^N)$ is a (pure) Nash equilibrium if

$$Q_i^N \in \arg\max_{Q_i} f_i(Q_1^N, Q_2^N, ..Q_i, ..., Q_K^N)$$

One key issue with the non-cooperative setting is 'free-riding' where some of the CPs do not contribute to μ and free ride on the contributions of the others. Indeed, as we will see, except when the CPs are perfectly symmetric, a Nash equilibrium involves at most one of the CPs making a positive contribution towards capacity expansion. To formalise this, we define the following types of Nash equilibria. A Nash equilibrium Q^N is said to be

- Type 0 if none of the CPs pay the q-charge, i.e., $Q_i^N = 0 \; \forall \; i$.
- Type k, $1 \leq k \leq K$, if exactly k CPs pay the q-charge, i.e., $|\{i : \; Q_i^N > 0\}| = k$.

A second issue, not unrelated to the first, is the 'tragedy of commons', where the CPs under-invest in capacity expansion (often to their own disadvantage) by acting selfishly.

The following theorem characterizes the conditions for existence of different types of Nash equilibria.

Theorem 1. *A (pure) Nash equilibrium always exists.*

1. *If* $\max_i[v_i x_i'(0)] \leq 1$, *then the only Nash equilibrium is of Type 0.*
2. *If* $\max_i[v_i x_i'(0)] > 1$ *and* $i_{\max} = \arg\max_i \left[(x_i')^{-1}(\frac{1}{v_i}) \right]$, *then the only Nash equilibrium is of Type 1 and is of the form* $(0, 0, ..Q_{i_{\max}}^N, 0.., 0))$, *where* $Q_{i_{\max}}^N = \max_i \left[(x_i')^{-1}(\frac{1}{v_i}) \right]$.
3. *If* $\max_i[v_i x_i'(0)] > 1$ *and* $\{i_1, i_2, ...i_k\} = \arg\max_i \left[(x_i')^{-1}(\frac{1}{v_i}) \right]$, *then there is a continuum of Type j Nash equilibria, $j \leq k$, satisfying*

$$\sum_{j=1}^{k} Q_{i_j}^N = \mu^N, \;\; Q_{i_1}^N, Q_{i_2}^N, ..Q_{i_k} \geq 0,$$

where $\mu^N = \max_i \left[(x_i')^{-1}(\frac{1}{v_i}) \right]$.

Theorem 1 determines that there is a non zero q-charge if the marginal gain at $\mu = 0$ for at least one of the CPs is greater than 1. Further, only the CP with the maximum value of $(x_i')^{-1}(\frac{1}{v_i})$ contributes, except when there is more than one maximizer. The exception requires a certain symmetry among a subset of the CPs, which is unlikely in practice.

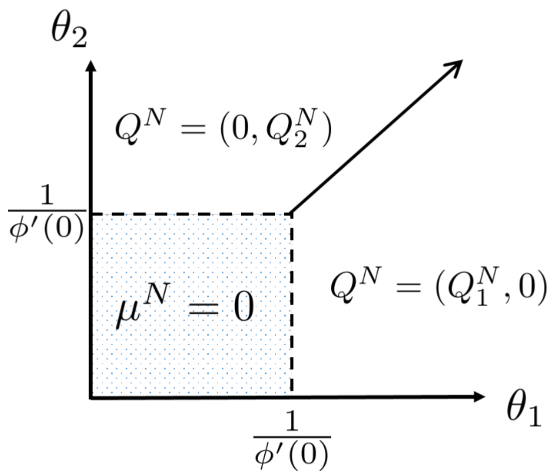

Figure 1: Conditions for existence of different types of Nash equilibria with $K = 2$, $x_i(\mu) = a_i\phi(\mu)$, $\theta_i := a_i v_i$. The ray $\theta_1 = \theta_2 > 1/\phi'(0)$ corresponds to a continuum of equilibria.

It is illustrative to consider a special case with two CPs and $x_i(\mu) = a_i\phi(\mu)$, where a_i captures user's relative preferences for the two CPs. Defining $\theta_i = a_i v_i$, Theorem 1 partitions the $\theta_1 \times \theta_2$ space into four regions, and provides a precise characterization of the Nash equilibria in each region. Figure 1 provides a pictorial depiction of these regions. Interpreting θ_i to be the size of CP i, we see that when the sizes of both CPs are small (precisely, less than or equal to $\frac{1}{\phi'(0)}$), the only Nash equilibrium is $(0, 0)$. Else, except when the sizes are exactly matched, the unique Nash equilibrium is Type 1, with the 'larger' CP being the only contributor towards capacity expansion. This means that in general, non-cooperative peering results in highly asymmetric contributions by the CPs.

A direct corollary of Theorem 1 is the following.

Corollary 1. *The capacity expansion μ^N under any Nash equilibrium between the CPs is unique.*

Proof of Theorem 1. If $\max_i[v_i x_i'(0)] \leq 1$ (Condition (1)), it is easy to check that the dominant response of CP i given any actions of the other CPs is to set $Q_i = 0$. It therefore follows that Type 0 is the unique equilibrium for this case.

Next, suppose $\max_i[v_i x_i'(0)] > 1$ and $\{i_1, i_2, ... i_k\} = \arg\max_i \left[(x_i')^{-1}(\frac{1}{v_i})\right]$. Define $\hat{\mu} := \max_i \left[(x_i')^{-1}(\frac{1}{v_i})\right]$. Since $\max[v_i x_i'(0)] > 1$, it follows that $\hat{\mu} > 0$.

Moreover, for $1 \leq j \leq k$, $v_{i_j} x'_{i_j}(\hat{\mu}) - 1 = 0$.

Given any action profile Q, it is easy to check that

1. If $\mu = \sum_{i=1}^{K} Q_i < \hat{\mu}$, then Q is not a Nash equilibrium, since for CPs i_j $(1 \leq j \leq k)$,
$$\frac{\partial f_{i_j}}{\partial Q_{i_j}} = v_{i_j} x'_{i_j}(\mu) - 1 > 0$$
i.e., these CPs have an incentive to increase their contribution.
2. If $\mu = \sum_{i=1}^{K} Q_i > \hat{\mu}$, then Q is not a Nash equilibrium, since
$$\frac{\partial f_i}{\partial Q_i} = v_i x'_i(\mu) - 1 < 0$$
for all i, i.e., all CPs that are making a contribution under Q have an incentive to decrease their contributions.
3. If $\mu = \sum_{i=1}^{K} Q_i = \hat{\mu}$, then Q is a Nash equilibrium if and only if $Q_j = 0$ for $j \notin \arg\max_i \left[(x'_i)^{-1}(\frac{1}{v_i}) \right]$. Thus, we have the continuum of equilibria as claimed.

□

Our next result states that a non-cooperative determination of the q-charge always leads to a lower capacity expansion compared to the cooperative case.

Theorem 2. *If* $\sum_{i=1}^{K} v_i x'_i(0) \leq 1$, *then* $\mu^N = \mu^C = 0$. *If* $\sum_{i=1}^{K} v_i x'_i(0) > 1$, *then* $\mu^N < \mu^C$.

Proof. If $\sum_{i=1}^{K} v_i x'_i(0) \leq 1$, then $\mu_C = 0$ as shown in Lemma 1. However, $\max_i[v_i x'_i(0)] < \sum_{i=1}^{K} v_i x'_i(0)$. This implies that $\max_i[v_i x'_i(0)] < 1$ and therefore $\mu^N = 0$ by Theorem 1.

If $\max_i v_i x'_i(0) \leq 1 < \sum_{i=1}^{K} v_i x'_i(0)$, then $\mu_C > 0 = \mu^N$, and therefore $\mu^C > \mu^N$.

If $1 < \max_i v_i x'_i(0) < \sum_{i=1}^{K} v_i x'_i(0)$, then $v_{i_{\max}} x'_{i_{\max}}(\mu^N) = 1 < \sum_{i=1}^{K} v_i x'_i(\mu^N)$. However, $\sum_{i=1}^{K} v_i x'_i(\mu^C) = 1$. From the concavity assumptions on x_i, it now follows that $\mu^N < \mu^C$. □

Finally, we consider the price of anarchy (POA), defined as

$$\text{POA} = \frac{\mu^C}{\mu^N}.$$

The POA captures the inefficiency due to the non-cooperation between the CPs with respect to capacity expansion. Clearly, the POA is well defined when $\sum_{i=1}^{K} v_i x_i'(0) > 1$ (otherwise, we have $\mu^C = \mu^N = 0$). Moreover, if $\sum_{i=1}^{K} v_i x_i'(0) > 1$, Theorem 2 implies that POA > 1.

As before, we illustrate for the case of two CPs and $x_i(\mu) = a_i \phi(\mu)$. (Recall that $\theta_i := a_i v_i$.) Further, for an explicit characterization of the value of the POA, we consider two example forms for $\phi(\cdot)$. For the logarithmic usage function $x_i(\mu) = a_i \log(1 + b\mu)$ (where $b > 0$), it can be shown that (assuming $\theta_1 + \theta_2 > \frac{1}{\phi'(0)} = \frac{1}{b}$)

$$\mathrm{POA} = \frac{b(\theta_1 + \theta_2) - 1}{b \max(\theta_1, \theta_2) - 1}.$$

For the bounded exponential usage function $x_i = a_i(1 - e^{-b\mu})$ (where $b > 0$), it can be similarly shown that (assuming again that $\theta_1 + \theta_2 > \frac{1}{\phi'(0)} = \frac{1}{b}$)

$$\mathrm{POA} = \frac{\log\left(b(\theta_1 + \theta_2)\right)}{\log(b \max(\theta_1, \theta_2))}. \tag{2}$$

In both of the above cases, note that even within the region $\max(\theta_1, \theta_2) > \frac{1}{b}$, the POA is unbounded. Interestingly, the POA approaches 1 when $\max(\theta_1, \theta_2) >> \min(\theta_1, \theta_2)$, i.e., when the CP sizes are highly asymmetric.

We explore the POA numerically in Section 4.

3.3. Bargaining framework

Finally, we consider the setting where the CPs 'bargain' to arrive at an agreement on their capacity contributions. Note that this requires that the CPs are able to communicate with one another. We invoke the classical Nash bargaining solution from the bargaining literature to capture the agreement between the CPs. Our main result is that the bargaining solution is even more efficient than the cooperative regime with respect to capacity expansion. As discussed before, this has significant implications from a regulatory standpoint.

To define the Nash bargaining solution (NBS), we first define the set of feasible, non-negative surplus vectors

$$\mathcal{F} := \{(f_1(Q), \ldots, f_K(Q)) \mid Q \geq 0\} \cap \mathbb{R}_+^K.$$

A Nash bargaining solution (NBS) $f^B = (f_1^B, f_2^B, \cdots, f_K^B)$ is defined to be a solution of the following maximization.

$$\begin{aligned} \max \quad & \prod_{i=1}^{K} \hat{f}_i \\ \text{such that} \quad & (\hat{f}_1, \ldots, \hat{f}_K) \in \mathcal{F} \end{aligned} \tag{3}$$

Note that we are maximizing the product of the CP surpluses, subject to the constraint that each surplus is non-negative.[5] It is important to note that the axiomatic development of the Nash bargaining formulation assumes that the payoff space, which is the set of all payoff vectors of both players, is convex. This in turn implies that the NBS is unique. In the present setting, we are unable to prove the convexity of the set $\mathcal{F}$ (although numerical experiments suggest that the set is indeed convex). However, we prove via direct arguments that the optimization (3) has a unique maximizer (see Lemma 3).

The Nash bargaining framework involves a *disagreement outcome*, which is the vector of payoff pairs if the players (CPs) fail to arrive at an agreement. In the present setting, it is natural to take the *disagreement outcome* to be $(0, 0, \cdots, 0)$, which corresponds to no capacity investment from the CPs.[6] Note that under the Nash bargaining framework, the CPs arrive at an agreement if and only if the optimization (3) has an optimal value that is strictly positive.

Our first result characterizes the condition for a non-trivial NBS solution.

Lemma 2. *The NBS solution is non zero if and only if* $\sum_{i=1}^{K} v_i x_i'(0) > 1$.

Proof. Suppose that $\sum_{i=1}^{K} v_i x_i'(0) > 1$. Then there exists $\bar{\mu} > 0$ such that

$$\sum_{i=1}^{K} v_i x_i(\bar{\mu}) - \bar{\mu} > 0.$$

[5]While Nash originally framed the bargaining solution for two players, the above generalization to K players is well known [24]. However, this generalization to more than two players is not widely used because it ignores the possibility that subsets of the K players can form coalitions (see [24]). However, we believe that that is not an issue in our problem because each of the CPs is making voluntary contributions towards the common goal of capacity expansion.

[6]Another natural candidate for the disagreement outcome is the Nash equilibrium between CPs. In the appendix, we show that the main conclusions of this section remain valid under this alternative definition of disagreement outcome.

It is easy to see that one can find $\bar{Q}_1, \ldots, \bar{Q}_K \geq 0$, such that

$$\sum_{i=1}^{K} \bar{Q}_i = \bar{\mu},$$
$$\bar{f}_i := v_i x_i(\bar{\mu}) - \bar{Q}_i > 0.$$

Since we have demonstrated a point $(\bar{f}_1, \bar{f}_2, ...\bar{f}_K) \in \mathcal{F}$ that yields a strictly positive objective value for the optimization (3), it follows that we have a non-trivial NBS solution.

Next, suppose that there exists a solution to (3). This implies that there exists $\hat{Q}_1, \hat{Q}_2, ...\hat{Q}_K \geq 0$ such that $\hat{f}_i := f_i(\sum_{i=1}^{K} \hat{Q}_i) > 0$ for $i = 1, 2, ...K$. This in turn implies that

$$\sum_{i=1}^{K} v_i x_i(\sum_{j=1}^{K} \hat{Q}_j) - (\sum_{i=1}^{K} \hat{Q}_i) > 0,$$

which proves that the capacity optimization under the cooperative regime has a positive solution. It now follows that $\sum_{i=1}^{K} v_i x_i'(0) > 1$ (see Lemma 1). □

Note that existence of a non-trivial NBS solution is equivalent to the capacity expansion under the cooperative regime being strictly positive (see Lemma 1). Thus, the set of system parameters over which we have a non-trivial NBS solution is a strict superset of the set of system parameters over which a non-trivial Nash equilibrium exists.

Our next result establishes the uniqueness of the NBS solution.

Lemma 3. *If* $\sum_{i=1}^{K} v_i x_i'(0) > 1$, *then the optimizer* $f^B = (f_1^B, f_2^B, ...f_K^B)$ *of* (3) *is unique.*

Proof. Suppose, for the purpose of obtaining a contradiction, that there exist two optimizers $\hat{f} = (\hat{f}_1, \hat{f}_2, ...\hat{f}_K)$ and $\bar{f} = (\bar{f}_1, \bar{f}_2, ...\bar{f}_K)$ of (3). Clearly, there exists $(\hat{Q}_1, \hat{Q}_2, ...\hat{Q}_K)$ and $(\bar{Q}_1, \bar{Q}_2, ...\bar{Q}_K)$ such that

$$f_i(\hat{Q}) = \hat{f}_i, \; f_i(\bar{Q}) = \bar{f}_i, \; i \in \{1, 2, 3...K\}$$

It is easy to see that

$$\left(\prod_{i=1}^{K} \frac{\hat{f}_i + \bar{f}_i}{2} \right) > \prod_{i=1}^{K} \hat{f}_i = \prod_{i=1}^{K} \bar{f}_i > 0. \tag{4}$$

Consider now, for $i = 1, 2, \dots, K$

$$\begin{aligned}\tilde{f}_i &:= f_i \left(\frac{\hat{Q}_1 + \bar{Q}_1}{2}, \frac{\hat{Q}_2 + \bar{Q}_2}{2}, \cdots, \frac{\hat{Q}_K + \bar{Q}_K}{2} \right) \\ &= v_i x_i \left(\frac{\hat{Q}_1 + \bar{Q}_1}{2}, \frac{\hat{Q}_2 + \bar{Q}_2}{2}, \cdots, \frac{\hat{Q}_K + \bar{Q}_K}{2} \right) - \frac{\hat{Q}_i + \bar{Q}_i}{2} \\ &\geq \frac{1}{2} v_i x_i (\sum_{j=1}^{K} \hat{Q}_j) + \frac{1}{2} v_i x_i (\sum_{j=1}^{K} \bar{Q}_j) - \frac{\hat{Q}_i + \bar{Q}_i}{2} \\ &= \frac{\hat{f}_i + \bar{f}_i}{2}.\end{aligned}$$

The bounding above invokes Jensen's inequality. It now follows, from (4), that

$$\prod_{i=1}^{K} \tilde{f}_i > \prod_{i=1}^{K} \hat{f}_i = \prod_{i=1}^{K} \bar{f}_i > 0.$$

Since $(\tilde{f}_1, \tilde{f}_2, \cdots, \tilde{f}_K) \in \mathcal{F}$, we have a contradiction. □

Having established the uniqueness of the NBS solution, the next step is to characterize the capacity expansion under the Nash bargaining scheme. However, it turns out that the q-charge under NBS, denoted by $f^B = (f_1^B, \dots, f_K^B)$, is not necessarily unique. Indeed, we show that the solution is associated with at least one and at most two values of the capacity expansion parameter μ. When the solution is unique, we show that the resulting μ, denoted by μ^B, is equal to the cooperative capacity expansion parameter μ^C. When there are two possible values, the smaller value is less than μ^C, whereas the greater value exceeds μ^C. If we thus follow the convention that for the same surplus vector, the CPs choose the greater capacity expansion (resulting in a greater benefit to the users), we conclude that the capacity expansion under NBS actually exceeds that under the cooperative regime. This is formalized in the following theorem.

Theorem 3. *If $\sum_{i=1}^{K} v_i x_i'(0) > 1$, then there exists a unique $\mu^B \geq \mu^C$, $Q_1^B, Q_2^B, ...Q_K^B \geq 0$, such that*

$$\mu^B = \sum_{i=1}^{K} Q_i^B,$$

$$f_i(Q_1^B, Q_2^B, ...Q_K^B) = f_i^B, \quad i \in \{1, 2, ...K\}$$

Proof. Let $g(\mu) := \sum_{i=1}^{K} v_i x_i(\mu) - \mu$ denote total surplus of all CPs. Clearly, $\sum_{i=1}^{K} f_i^B \leq g(\mu^C)$.

Now, suppose that $\sum_{i=1}^{K} f_i^B < g(\mu^C)$. Given the convexity properties of g, it follows that there exist unique values $\underline{\mu}, \bar{\mu}$ such that $0 < \underline{\mu} < \mu^C < \bar{\mu}$, and $g(\underline{\mu}) = g(\bar{\mu}) = \sum_{i=1}^{K} f_i^B$. Note that $\underline{\mu}$ and $\bar{\mu}$ are the only possible μ under f^B.

Now, let $(\underline{Q}_1, \underline{Q}_2, ...\underline{Q}_K)$ and $(\bar{Q}_1, \bar{Q}_2, ...\bar{Q}_K)$ be defined as follows. For $i \in \{1, 2, ...K\}$

$$f_i^B = v_i x_i(\underline{\mu}) - \underline{Q}_i, \quad f_i^B = v_i x_i(\bar{\mu}) - \bar{Q}_i.$$

Note that $(\underline{Q}_1, \underline{Q}_2, ...\underline{Q}_K)$ and $(\bar{Q}_1, \bar{Q}_2, ...\bar{Q}_K)$ are the q-charges corresponding to $\underline{\mu}$ and $\bar{\mu}$ respectively, if feasible. Clearly, $(\bar{Q}_1, \bar{Q}_2, ...\bar{Q}_K) > (\underline{Q}_1, \underline{Q}_2, ...\underline{Q}_K)$. Since at least one of $\underline{\mu}$ and $\bar{\mu}$ is feasible, it follows that $(\bar{Q}_1, \bar{Q}_2, ...\bar{Q}_K) \geq (0, 0, \cdots, 0)$. The proof is completed by setting $\mu^B = \bar{\mu}$, and $(Q_1^B, Q_2^B, ...Q_K^B) = (\bar{Q}_1, \bar{Q}_2, ...\bar{Q})_K$.

The case of $\sum_{i=1}^{K} f_i^B = g(\mu^C)$ is trivial; in this case, we have $\mu^B = \mu^C$. This completes the proof. □

Our next result sheds further light on the relationship between μ^B (defined by Theorem 3) and μ^C. Specifically, it states that whenever all providers make a positive contribution under the NBS, the corresponding capacity expansion equals the cooperative capacity expansion.

Theorem 4. *If $Q_i^B > 0 \ \forall i \in \{1, 2, ...K\}$,*

$$\mu^B = \mu^C.$$

Proof. Recall that the objective function is

$$R = \prod_{i=1}^{K} f_i = \prod_{i=1}^{K} (v_i x_i(\mu) - Q_i).$$

We use first order conditions; taking the partial derivative w.r.t. Q_i, we get,

$$\frac{\partial R}{\partial Q_i} = \frac{v_i x_i'(\mu) - 1}{v_i x_i(\mu) - Q_i} R + R \sum_{j=1, j \neq i}^{K} \frac{v_j x_j'(\mu)}{v_j x_j(\mu) - Q_j}.$$

Equating the partial derivatives to 0,

$$\sum_{j=1}^{K} \frac{v_j x_j'(\mu^B)}{v_j x_j(\mu^B) - Q_j} - \frac{1}{v_i x_i(\mu^B) - Q_i} \;=\; 0$$

which implies

$$\sum_{j=1}^{K} \frac{v_j x_j'(\mu^B)}{v_j x_j(\mu^B) - Q_j} = \frac{1}{v_i x_i(\mu^B) - Q_i} \qquad \forall i \in \{1, 2, ...K\}$$

This in turn means that

$$\begin{aligned} v_i x_i(\mu^B) - Q_i \;&=\; \text{constant} \qquad \forall i \in \{1, 2, ...K\}, \qquad \text{and hence} \\ \sum_{j=1}^{K} v_j x_j'(\mu^B) \;&=\; 1, \qquad \text{implying} \\ \mu^B \;&=(\sum_{j=1}^{K} v_j x_j')^{-1}(1) \;=\; \mu^C. \end{aligned}$$

This completes the proof. □

Note that while the cooperative regime does generate a higher aggregate surplus for the CPs, the bargaining framework above results in a higher capacity expansion. This means that even though strategic interaction between the CPs can lower CP surplus compared to the cooperative case, it is beneficial to the user base.

At this point, we note that the bargaining solution as characterized via the preceding results should not be interpreted literally as the precise outcome of a negotiation between the CPs. In practice, bargaining between strategic firms is a complex process that is influenced by several factors, including the barganing power of each agent, the information assymetry between agents, their risk-taking abilities, and so on. Thus, one must view the bargaining solution characterized here as being *qualitatively indicative* of the arrangement that would occur in practice [25].

We conclude by defining the benefit of bargaining (BOB), which captures the relative benefit of the bargaining solution over the cooperative model with respect to capacity expansion:

$$\text{BOB} = \frac{\mu^B}{\mu^C}.$$

Note that the BOB is well defined for $\sum_{i=1}^{K} v_i x_i'(0) > 1$, and is lower bounded by 1. However, a closed form characterization of the BOB is infeasible for even the simplest usage models. In the next section, we thus resort to numerical experiments to gain additional insights on the efficiency of the bargaining based solution.

4. q-charges: Numerical illustrations

Given the analytical results in Section 3, the goal of this section is to gain additional insights on the impact of non-cooperation and bargaining between the CPs via numerical experiments. Toward this, we consider the case of two CPs, and the bounded exponential usage functions $x_i(\mu) = a_i(1 - e^{-b\mu})$, where $b > 0$. Recall that $\theta_i := a_i v_i$.

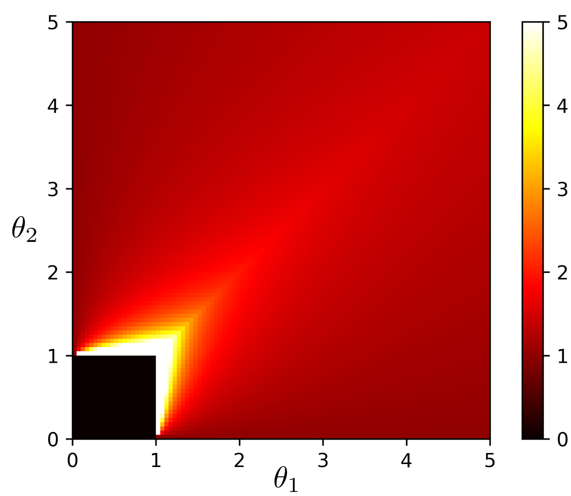

Figure 2: Price of anarchy

We first consider the price of anarchy (POA). Figure 2 illustrates how the POA varies across the $\theta_1 \times \theta_2$ space. In this heat map, a lighter shade indicates a higher value of POA. Note that the POA grows unboundedly as one approaches the unit square $[0, 1] \times [0, 1]$ from outside (see (2)). However, to avoid saturating the colour scheme, we have capped the dynamic range of the dependent variable by plotting $\min(POA, 5)$. Note that the POA is maximum when $\theta_1 \approx \theta_2$, i.e., when the CPs are of comparable size (notice the

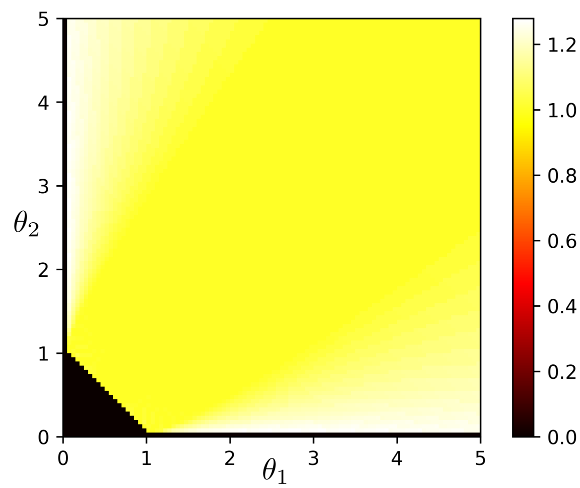

Figure 3: Benefit of bargaining

lighter shade along the diagonal $\theta_1 = \theta_2$). This is because in this case, only one of the CPs ends up contributing towards capacity expansion under any Nash equilibrium, whereas both CPs make comparable contributions under the cooperative regime. It is also worth noting that the POA decreases as the CP sizes become more asymmetric ($\theta_1 \gg \theta_2$ or $\theta_2 \gg \theta_1$). This is to be expected, since even under the cooperative regime, the larger CP would make the dominant contribution in this case.

Next, we consider the benefit of bargaining (BOB), which captures the relative improvement in capacity expansion under bargaining-based peering compared with the cooperative setting. Figure 3 shows how the BOB varies across the $\theta_1 \times \theta_2$ space. As before, a lighter shade indicates a higher value of BOB. We observe that the BOB is close to one when the CPs are of comparable size, i.e., $\theta_1 \approx \theta_2$. On the other hand, the benefit of bargaining grows as the CP sizes are highly asymmetric ($\theta_1 \gg \theta_2$ or $\theta_2 \gg \theta_1$). This is further illustrated in Figure 4, which shows the ratio of the minimum to the maximim q-charge under the bargaining framework (as defined by Theorem 3); recall from Theorem 4 that when this ratio is positive, BOB $= 1$. As we see in Figure 4, BOB exceeds one when the CPs are of asymmetric size (the region in black).

Interestingly, the above observations imply that the capacity expansion

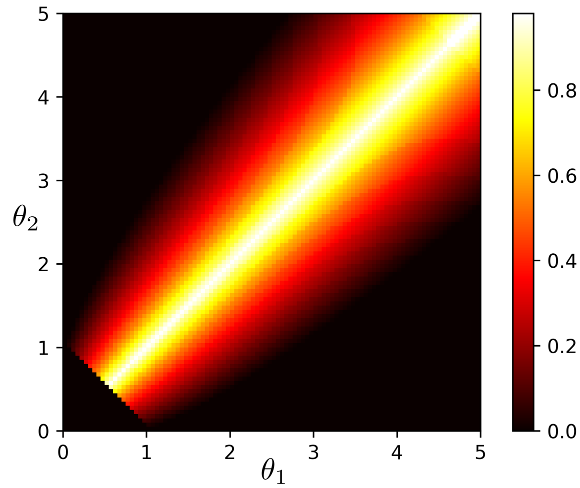

Figure 4: Ratio of minimum to maximum contribution under NBS

under bargaining-based peering is always substantial relative to that under the non-cooperative peering: When $\theta_1 \approx \theta_2$, we have a large POA but BOB ≈ 1, whereas when $\max(\theta_1, \theta_2) \gg \min(\theta_1, \theta_2)$, we have POA ≈ 1 but a large BOB. Figure 5 depicts the product of BOB and POA, which is the ratio of μ^B and μ^N. It is instructive that this product assumes a higher value when the CPs are symmetric. Note that close to the unit square $[0, 1] \times [0, 1]$ the product tends to ∞. As before, we have capped the value of POA $\times$ BOB to 5 for clarity.

In summary, the results of this section, in conjunction with the results of the previous section, confirm the efficiency of bargaining based capacity investments by CPs.

5. Concluding remarks

In this paper, we propose network neutral capacity expansion funded by voluntary contributions by CPs. Our results show that this is indeed feasible. Moreover, the observation that bargaining based peering results in the highest capacity expansion suggests that policy makers could set up a transparent platform for CPs to make commitments for Internet infrastructure expansion. In contrast, the present practice of confidential and bilateral peering arragements between CPs and ISPs leads not only to potentially

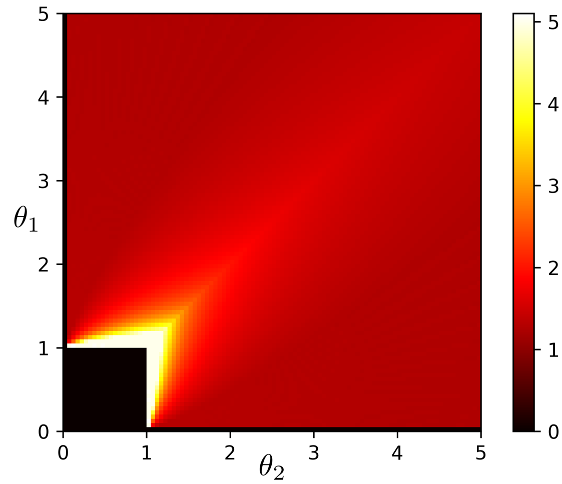

Figure 5: Product of POA and BOB

lower infrastructure investments, but also to non-neutral internet access for users. Of course, for the proposed mechanism to work in practice, one would require regulatory intervention that (i) enforces strong network neutrality [6], disallowing (or at least discouraging) confidential bilateral peering arrangements between CPs and ISPs, and (ii) encourages transparent platforms for ISPs to contribute directly to internet capacity expansion.

An immediate extension of our model is to consider the case when the users pay an internet access fee that is proportional to usage. Mathematically, it turns out that this case is identical to the case of flat user pricing, and all our previous conclusions apply. To see this, note that the CP surplus functions remain unchanged from the flat access price case. Thus, if $\mu = \sum_{i=1}^{K} Q_i$, then the analysis and conclusions of Section 3 apply.

Another extension that is possible, is to assume that the CPs pay a volume based $q-$charge, i.e., they pay per byte. Specifically, CP i commits to pay the ISP q_i per unit consumption (over the baseline), where $q_i \in [0, v_i)$ towards capacity expansion. Under this model, the realised capacity expansion is determined by the fixed point equation $\sum_{i=1}^{K} q_i x_i(\mu) = \mu$, and the surplus of CP i is given by $(v_i - q_i)x_i(\mu)$. While this model seems interesting, it is substantially less tractable than the 'flat' q-charge model considered in this paper. This is because the CP payoff functions are no longer concave with respect to the actions, making it problematic to define the non-cooperative/bargaining-based capacity expansion uniquely. Interestingly, we are still able to prove a result analogous to Theorem 4 for this model.

Appendix A. Bargaining framework with disagreement outcome being a Nash equilibrium

In our development of the bargaining framework for network neutral capacity expansion in Section 3.3, it was assumed that the disagreement outcome is the *zero* vector. Another natural choice for the disagreement outcome is the Nash equilibrium (see, for example, [26]). In this section, we note that the main conclusions of Section 3.3 hold even when the disagreement outcome is taken to be a Nash equilibrium.

Note that even though the capacity expansion under the Nash equilibrium is uniquely defined, the CP payoffs are not necessarily so, specifically when there is more than one maximizer of $\max_i \left[(x_i')^{-1}(\frac{1}{v_i})\right]$ (see Theorem 1). Accordingly, *let us assume that the disagreement outcome is a certain uniquely defined vector of CP payoffs* $f^N = (f_i^N, 1 \leq i \leq K)$ *consistent with a Nash equilibrium.*

For example, if $\max_i[v_i x_i'(0)] > 1$ and $\{i_1, i_2, ...i_k\} = \arg\max_i \left[(x_i')^{-1}(\frac{1}{v_i})\right]$, we may define

$$f_i^N = \begin{cases} v_i x_i(\mu^N) - \frac{\mu^N}{k} & \text{for } i \in \{i_1, i_2, ...i_k\} \\ v_i x_i(\mu^N) & \text{for } i \notin \{i_1, i_2, ...i_k\} \end{cases}.$$

With this disagreement outcome, we define the Nash bargaining solution f^B as a solution of the following maximization.

$$\begin{aligned} \max \qquad & \prod_{i=1}^{K} (\hat{f}_i - f_i^N) \\ \text{such that} \qquad & (\hat{f}_1, \ldots, \hat{f}_K) \in \mathcal{F}, \end{aligned}$$

where

$$\mathcal{F} := \{(f_1(Q), \ldots, f_K(Q)) \;:\; Q \geq 0,\; f_i(Q) \geq f_i^N \;\forall\, i\}.$$

We note that the above modification would in general result in a different bargaining solution as compared to solution where the disagreement outcome is taken to be zero. However, it is not hard to show that the statements of

Lemma 2, Lemma 3, and Theorem 3 continue to hold with this modification.[7] The proofs follow along similar lines to the ones presented in Section 3.3.

References

[1] A. Kalvit, G. Kasbekar, D. Manjunath, J. Nair, Capacity expansion of neutral ISPs via content peering charges: The bargaining edge, in: Proceedings of the 11th EAI International Conference on Performance Evaluation Methodologies and Tools, 2017.

[2] G. Goth, New internet economics might not make it to the edge, IEEE Internet Computing (2010) 7–9.

[3] C. Dovrolis, The evolution and economics of internet interconnections, Submitted to Federal Communications Commission (2015).
URL `www.cc.gatech.edu/~dovrolis/Papers/FCC-declaration.pdf`

[4] S. Sen, C. Joe-Wong, S. Ha, M. Chiang, A survey of smart data pricing: Past proposals, current plans, and future trends, ACM Computing Surveys 46 (2).

[5] R. T. B. Ma, D. M. Chiu, J. C. S. Lui, V. Misra, On cooperative settlement between content, transit and eyeball internet service providers, IEEE/ACM Transactions on Networking 19 (3) (2011) 802–815.

[6] R. Hastings, Internet tolls and the case for strong net neutrality, `https://media.netflix.com/en/company-blog/internet-tolls-and-the-case-for-strong-net-neutrality`, [Online; accessed 31-July-2017] (2014).

[7] D. Clark, W. Lehr, S. Bauer, Interconnection in the internet: The policy challenge, in: 39th Research Conference on Communication, Information and Internet Policy, 2011.

[8] J. Musacchio, G. Schwartz, J. Walrand, A two-sided market analysis of provider investment incentives with an application to the net-neutrality issue, Review of Network Economics 8 (1).

[7] However, Theorem 4 does not hold under the above modification of the disagreement outcome, since the exact value of the bargaining-based capacity expansion is sensitive to the disagreement outcome.

[9] N. Economides, J. Tag, Network neutrality on the internet: A two-sided market analysis, Information Economics and Policy 24 (2) (2012) 91–104.

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

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

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

[13] L. Zhang, W. Wu, D. Wang, Sponsored data plan: A two-class service model in wireless data networks, in: Proceedings of ACM SIGMETRICS, 2015.

[14] R. T. B. Ma, D. M. Chiu, J. C. S. Lui, V. Misra, D. Rubenstein, Internet economics: the use of shapley value for isp settlement, IEEE/ACM Transactions on Networking 18 (3) (2010) 775–787.

[15] J. J.Laffont, S. Marcus, P. Rey, J. Tirole, Internet peering, American Economic Review 91 (2) (2001) 287–291.

[16] G. Shrimali, S. Kumar, Paid peering among internet service providers, in: Proceedings of GameNets, 2006.

[17] G. Shrimali, S. Kumar, Bill-and-keep peerng, Telecommunications Policy 32 (1) (2007) 19–32.

[18] D. Zarchy, A. Dhamdhere, C. Dovrolis, M. Schapira, arXiv preprint (October 2016). [link].
URL `https://arxiv.org/abs/1610.01314`

[19] C. Courcoubetis, L. Gyarmati, N. Laoutaris, P. Rodriguez, K. Sdrolias, Negotiating premium peering prices: A quantitative model with applications, ACM Transactions on Internet Technology 16 (2) (2016) 1–22.

[20] C. Courcoubetis, K. Sdrolias, R. Weber, Paid peering: pricing and adoption incentives, Journal of Communications and Networks 18 (6).

[21] R. T. B. Ma, Pay or perish: The economics of premium peering, IEEE Journal on Selected Areas in Communications 35 (2) (2017) 353–366.

[22] D. Fudenberg, J. Tirole, Game Theory, MIT press, 1991.

[23] Y. Narahari, Game Theory and Mechanism Design, World Scientific and IISc Press, 2014.

[24] R. B. Myerson, Game theory, Harvard university press, 2013.

[25] K. Chatterjee, Game theory and the practice of bargaining, Group Decision and Negotiation 5 (4-6) (1996) 355–369.

[26] A. Okada, A noncooperative approach to the nash bargaining problem, in: Game Equilibrium Models III, Springer, 1991, pp. 7–33.