

Live Streaming of Sport Events: Pricing, Quality, and Side Payments

M. Sereno and P. Castagno
Università di Torino
Torino, Italy
{matteo.sereno,paolo.castagno}@unito.it

V. Mancuso
Università di Palermo
Palermo, Italy
and
IMDEA Networks Institute
Leganés, Madrid, Spain
vincenzo.mancuso@ieee.org

M. Ajmone Marsan
IMDEA Networks Institute
Leganés, Madrid, Spain
marco.ajmone@imdea.org

Abstract—The commercial success of live streaming—especially for major sports events—is putting increasing pressure on Internet Service Providers (ISPs) and raises questions about pricing, service quality, and network neutrality. A key open problem is understanding how congestion, pricing, and side payments affect user choices and market outcomes.

This paper introduces a modeling framework to study interactions among users, ISPs, and content providers (CPs) in congestion-sensitive environments. The framework combines quality-of-experience modeling, discrete choice theory, and evolutionary population dynamics, with particular attention to side payments from CPs to ISPs and their impact on user allocation and revenues.

Using this framework, we analyze how side payments may affect ISP market shares and give rise to potential neutrality and fairness concerns. Results show that side payments can benefit both CPs and ISPs, but that their allocation among ISPs plays a critical role in limiting induced market distortions.

Index Terms—Live Streaming, QoE-aware Congestion Modeling, Network Neutrality, Game-Theoretic Modeling.

I. INTRODUCTION

The continuous growth of Internet traffic, largely driven by video streaming—which already accounts for approximately 65% of total traffic, with platforms such as Netflix (15%) and YouTube (11.6%) among the top contributors [1]—poses significant challenges for network infrastructures and their operators. In particular, *live* video streaming is rapidly expanding, especially for sports and other real-time events, and poses a major challenge due to the simultaneous demand from large user populations.

As noted in [1]: “Live streaming is the next big capacity driver for the foreseeable future, and network operators will need to plan for these new peaks.”

The significant impact of live streaming on network infrastructures, combined with economic pressure on operators, has led to proposals for financial contributions—commonly

referred to as *side payments*—from content providers (CPs) to Internet Service Providers (ISPs). These payments aim to partially offset the cost of network upgrades. Proposals to introduce side payments have triggered a tussle between operators and content providers, with regulators struggling to identify a balanced and sustainable response. At the same time, consumer advocacy groups have raised concerns that side payments may violate network neutrality (NN), similarly to other contested practices such as traffic prioritization and zero-rating.

Side payments in the context of NN have been widely studied from multiple perspectives, including economic modeling [2]–[4], network engineering [5]–[7], regulatory concerns [8], [9], and legal theory [10], [11].

Among regulatory approaches, we focus on the Italian case, which emphasizes *fair cooperation* between CPs and ISPs—a central issue in side payment design and explicitly addressed by the national regulatory authority AGCOM¹. Official AGCOM resolutions require CPs to collaborate with *all* ISPs to ensure efficient and non-discriminatory content delivery [12], [13], yet their practical implementation remains non-trivial. In particular, it remains unclear how side payments interact with congestion, pricing, and user behavior, how they affect ISP market shares and neutrality outcomes, and under which conditions they may distort relative ISP positions or be designed to support fair cooperation without undermining network neutrality.

To evaluate these aspects, we develop a mathematical framework capturing CP–ISP interactions and the relations among pricing, user-perceived quality of experience (QoE), side payments, and network neutrality, with a focus on how side payments may alter ISP market shares.

The main contributions of this work are threefold. First, we propose a model that combines congestion-sensitive quality of experience, user choice, and evolutionary population dynamics to study CP–ISP interactions. Second, we use this framework to analyze how side payments can affect user allocation, stakeholder revenues, and ISP market shares under neutrality

This work has been supported by the the RESTART Program, financed by the Italian government with the resources of the Italian Recovery, Transformation and Resilience Plan – Mission 4, Component 2, Investment 1.3, theme 14 “Telecommunications of the future” (PE00000001 - program “RESTART”, CUP D93C22000910001), and by project TUCAN6-CM (TEC-2024/COM-460), funded by CM (ORDEN 5696/2024).

¹<https://www.agcom.it>

constraints. Third, we show that side payments may induce non-trivial shifts in market outcomes, highlighting potential neutrality and fairness concerns.

After reviewing related work in Section II, Section III introduces the system model, including CPs, ISPs, user fees, side payments, and user subscription choices. Section IV defines a QoE model distinguishing general Internet access from live streaming, while user dynamics are modeled as a population game in Section V. Section VI defines ISP and CP utilities and analyzes the impact of side payments. Section VII presents numerical results illustrating the proposed framework and highlighting non-trivial behaviors. Conclusions follow in Section VIII.

II. RELATED WORK

The interaction between content providers (CPs), Internet service providers (ISPs), and end-users under network neutrality (NN) constraints has been widely studied in the economic and game-theoretic literature. Seminal works such as [2], [3] laid the foundations for analyzing the implications of NN and non-NN policies using classical economic models and Stackelberg game formulations. In this context, a side payment refers to a monetary transfer from CPs to ISPs in exchange for prioritized or guaranteed content delivery. This mechanism was first explicitly formalized by Musacchio and Walrand [3] in a two-sided market framework, where ISPs serve both CPs and end-users, and pricing strategies can influence investment incentives and overall welfare. Building upon these foundations, a line of research has emerged that incorporates discrete choice theory and congestion-insensitive utility modeling to evaluate the social and economic impact of pricing strategies and side payments [14], [15]. Recent contributions, including [16]–[18], have developed models to compare the outcomes of NN-compliant and non-compliant policies in terms of metrics such as user welfare, global welfare, and provider revenues. These works primarily aim at identifying under which conditions side payments can improve overall system efficiency or harm user satisfaction. In these models, utility functions are typically exogenous and do not depend on network congestion or user distribution.

Beyond these foundational and mid-stage contributions, a more recent strand of literature has begun to reframe the role of side payments in networked systems. For instance, [19] provides empirical evidence that rigid NN regulation discourages infrastructure investment, particularly in contexts with high demand for capacity, such as mobile or broadband networks. In this view, side payments and differentiated service levels can act as incentives for ISPs to invest in network capacity, ultimately improving aggregate welfare.

Our work builds upon this literature by adopting the same foundational framework of discrete choice theory as in [17], [18], while introducing a fundamental shift: we propose a model where utility functions are *congestion-dependent*, directly reflecting the relationship between perceived service quality and the actual load in the system. Starting from a high-level description of the live streaming service architecture, our




Fig. 1: Live streaming distribution scenario

modeling approach captures the interdependencies between users' choices, pricing, and congestion effects, which are especially relevant in the context of massive-scale live sports streaming.

In this scenario, we model side payments not as optional pricing tools but as *necessary mechanisms* that enable ISPs to sustain a satisfactory QoE for a user base that would otherwise exceed the system's serviceable capacity. This functional interpretation aligns with the evolving literature that views side payments as enablers of scalability and investment, rather than sources of inefficiency.

From a modeling perspective, this view emphasizes the need to represent how users adapt their choices over time in response to congestion-sensitive QoE and pricing conditions.

We adopt a *population game* framework, which better accommodates the dynamic and interdependent nature of congestion-sensitive services. In this setting, user behaviors evolve over time through payoff-based revision protocols, modeled via logit dynamics [20].

Furthermore, our framework explores a preliminary direction for quantifying the degree of network neutrality compliance of different pricing and side payment policies. Unlike prior work, which tends to classify policies as either NN-compliant or not, we introduce a numerical perspective that may contribute to a more nuanced evaluation. Although not conclusive, this effort represents a first step toward a more refined understanding of policy implications in large-scale content delivery systems.

Our model builds on a large body of work investigating CP–ISP interactions in the presence of congestion and regulatory constraints. Recent studies address these interactions from different perspectives, including pricing, platform competition, and network management, while explicitly accounting for congestion effects [21]–[23]. Differently from these approaches, in our setting congestion directly drives users' choice behavior, coupling quality of experience with population dynamics.

III. HIGH-LEVEL SYSTEM DESCRIPTION

We consider a live streaming distribution system comprising an origin server that generates live streaming content and distributes it over the Internet to a set of Internet Service Providers (ISPs). These ISPs, in turn, deliver the streamed content to their respective end users. A schematic representation of this system is presented in Fig. 1.




Fig. 2: Economic relations among stakeholders for the case of two ISPs and one CP (indicated flows are meant per subscriber)

It is important to emphasize that, for the purposes of this paper, the term origin server encompasses any type of distribution infrastructure responsible for forwarding streaming content from the locations where events take place to the entry points of various ISPs. We assume that the origin server is appropriately dimensioned and that no bottleneck is present between the stream origin and the ISP network ingress.

Bottlenecks may arise due to the limited bandwidth available at the interconnections between ISPs and core network, as well as within ISP networks. The latter can be mitigated by deploying proxy caches within ISPs or implementing multicast solutions (always within the ISP). Alternatively, if the ISP's networks have scarce resources, additional capacity can be deployed. In all cases, an upgrade of the ISP infrastructure is needed, which, of course, entails substantial investments.

A. System model

We consider a single Content Provider (CP) offering a live streaming service and one or more ISPs delivering it to end users. Users subscribing to the CP service pay a flat fee p_i to ISP i , $i \in \{A, B, \dots\}$, and a subscription fee p_c to the CP. Following standard formulations on non-neutral networks (e.g., [17], [18]), we introduce a *side payment* q_i , representing a per-user transfer from the CP to ISP i .

In the non-neutrality literature, side payments typically model financial transfers from CPs to ISPs, often in exchange for prioritized delivery or access to user bases [17], [18]. In our model, side payments are reinterpreted as a mechanism to support additional delivery capacity within the ISP network. Specifically, we view side payments as a means for ISPs to improve service quality by acquiring additional infrastructure resources (e.g., renting capacity or radio resources as demand increases), thereby reducing congestion or serving more users. We assume that side payments are fully reinvested by ISPs to improve live streaming capacity. These resources are fully reserved for live streaming traffic generated by the CP.

We model users as strategic agents, focusing on how the predefined CP–ISP setup—including service fees, subscription fees, and side payments—affects user populations.

All prices p_c , p_i , and q_i , $i \in \{A, B, \dots\}$, are positive and expressed per user. Payment flows are summarized in Fig. 2 for the case of two ISPs labeled A and B . While p_c , p_A , and p_B are paid by subscribers, side payments q_A and q_B are inter-provider transfers supporting enhanced service quality (see Section VI).

B. User populations

The simplified scenario described above allows us to study how side payments affect user choices, network usage, rev-

enues of different actors in the system, and, eventually, the network neutrality effects of side payments. To this extent, the key aspect to model is how users make a decision on which service to subscribe. Specifically, users can be partitioned into mutually disjoint subsets, according to the option s selected from available alternatives:

- Option (c, i) with $i \in \{A, B, \dots\}$: a user subscribes to ISP i and to the CP for access to streaming services.
- Option $(-, i)$, with $i \in \{A, B, \dots\}$: a user subscribes exclusively to ISP i , not to the CP.
- Option 0 : a user does not subscribe to any service.

If $\mathcal{S}$ indicates the set of possible user choices, i.e.,

$$\mathcal{S} = \{0, (-, A), (-, B), \dots, (c, A), (c, B), \dots\}, \quad (1)$$

we can denote by θ_s , with $s \in \mathcal{S}$, the fraction of users associated with each choice, which, equivalently, represents the probabilities to find a user in each of the possible disjoint subsets of the overall population, obtaining a population vector:

$$\theta = [\theta_0, \theta_{(-, A)}, \theta_{(-, B)}, \dots, \theta_{(c, A)}, \theta_{(c, B)}, \dots]. \quad (2)$$

While the model can be extended to multiple competing CPs, we focus on a single CP to isolate the impact of side payments.

Next, we derive a model for the QoE of users (cf. Section IV), which will allow us to model the rational behavior of users, hence their choice to subscribe to a service or not (cf. Section V), thus allowing the generation of vector θ . Eventually, with the derived models we will build expressions for the utility of ISPs and CP, and a metric to evaluate how side payments perturb the system (cf. Section VI).

IV. QOE MODELING

In this section, we define how to model the QoE associated with the two categories of services offered by the ISP and the CP: basic connectivity (non-streaming) offered to all subscribers of an ISP and live streaming offered to subscribers of an ISP that also subscribe to the CP. Our modeling approach assumes that each service (either basic connectivity or live streaming) is dimensioned for a maximum number of users, denoted generically by $N_{\max}$. When needed, we will distinguish between the two types of services by using superscripts. We reasonably assume that, when the number of users utilizing a service does not exceed $N_{\max}$, QoE is optimal. Instead, if the number of users exceeds $N_{\max}$, the service quality degrades according to a logarithmic law.

Logarithmic degradation for ISP services is justified by network management policies that allow the system to tolerate moderate congestion while still providing acceptable QoE. On the contrary, increasing resources for networking services, yields progressively diminishing improvements in user satisfaction beyond appropriate initial thresholds [24]. Furthermore, it is worth emphasizing that, in the context of streaming systems, the logarithmic relationship between the resources allocated to streaming and the perceived quality is supported in the literature, e.g., in [25]–[27], where the Weber–Fechner law is invoked to justify the logarithmic sensitivity of human perception.

As a further hypothesis, we assume that each ISP dedicates a portion of its resources specifically to live streaming services, and that basic connectivity and live streaming services use orthogonal resource sets. Still, a user can access both types of services through the same ISP. This assumption enables a modular modeling of QoE, while ensuring consistency; it also avoids unrealistic cases such as an ISP with poor QoE for basic connectivity, but excellent streaming performance.

A. QoE for basic connectivity

Each ISP i is dimensioned to deliver basic connectivity to a maximum number $N_{\max}^{(i,b)}$ of users out of a population of cardinality N_{tot} . When the current number of subscribers $(\theta_{(-,i)} + \theta_{(c,i)})N_{\text{tot}}$ is below $N_{\max}^{(i,b)}$, service quality is optimal, which is expressed with the relation $\text{QoE} = 1$, while values lower than 1 indicate sub-optimality. Indeed, the QoE perceived by users depends on network load, since a growing number of subscribers may lead to congestion and quality degradation. We therefore define the per-user share of available network resources at ISP i as:

$$x_i = \frac{N_{\max}^{(i,b)}}{(\theta_{(-,i)} + \theta_{(c,i)}) N_{\text{tot}}}.$$

With the above notation, user satisfaction is a logarithmic function of x_i . Specifically, and without loss of generality, we use a simple mapping between x_i and the associated QoE:

$$\text{QoE}^{(-,i)}(x_i) = \ell \log(1 + k \cdot x_i), \quad (3)$$

where $\ell = 1/\log(1 + k)$ so that $\text{QoE}^{(-,i)}(1) = 1$, and k is a shaping parameter that can be chosen to satisfy specific behavioral constraints. Note that $\text{QoE}^{(-,i)}(x_i) > 1$ for $x > 1$, which means that QoE values are not bounded to 1. In practice, we shape the value of k such that the QoE function be practically flat for x_i that exceeds 1 by reasonable amounts. In any case, QoE values larger than 1 can be interpreted as levels of experience that exceed expectations, which we can take into account in the users subscription choices.

B. QoE for live streaming

The QoE for streaming services depends on the streaming bitrate available to each user. Assume that ISP i is provisioned to support up to $N_{\max}^{(i,s)}$ concurrent streaming users at a target optimal bitrate. Let the actual number of subscribed users be $\theta_{(c,i)}N_{\text{tot}}$. Then, a subscriber will receive at least the optimal bandwidth if $\theta_{(c,i)}N_{\text{tot}} \leq N_{\max}^{(i,s)}$. We can therefore use the following variable to identify which fraction of the optimal bitrate subscribers will receive:

$$x_{(c,i)} = \frac{N_{\max}^{(i,s)}}{\theta_{(c,i)} N_{\text{tot}}}. \quad (4)$$

As in the case of basic connectivity service, we model the QoE for streaming service through a logarithmic function, which takes as argument the normalized bitrate $x_{(c,i)}$:

$$\text{QoE}^{(c,i)}(x_{(c,i)}) = \ell \log(1 + k \cdot x_{(c,i)}), \quad (5)$$

where k and ℓ are the same parameters introduced in (3), and ℓ is a function of k only, so that k is the only free parameter.

V. MODELING CHOICES AND POPULATION DYNAMICS

End users select an ISP depending on a number of factors, the most important of which are the price charged by the ISP and the obtained QoE. Cost considerations might also prevent users from subscribing any ISP service. Possibly, some users might want to also subscribe to a CP after they subscribe to an ISP. Hence, all possible alternatives $s \in \mathcal{S}$ must be evaluated jointly in terms of total price and QoE for the desired service.

While price is an input of the system that we consider, QoE depends on ISP populations, as discussed in the previous section. To formalize this, we define utility functions for users accessing a service.

A. Rational user behavior

We consider users as rational agents with some slackness. If we were to consider strictly rational agents, a slight variation in either quality or price might induce threshold effects, such as a sudden migration of all users from one ISP to another, following a marginal change in service conditions. However, such threshold phenomena are not common in practice. Instead, the dynamics we observe in reality (e.g., users switching ISP, as reported by independent agencies, consumer associations and authorities) exhibit gradual transitions rather than abrupt shifts. These gradual behaviors can be modeled using **discrete choice models** [14], [15], which assume that users exhibit heterogeneous behaviors when making choices. Some users may be highly sensitive to variations in price and quality, while others may stick with their initial choice out of inertia, even if it is not optimal. At times, changes may only occur when variations in price or quality become significant enough.

The approach we employ in this paper is based on the logit model, as described in [14], [15], and applied in [18] and other works by the same authors. In this framework, we assume that users are not strictly deterministic in their choice but are sensitive to both price and QoE. A utility function expressing the users' satisfaction level that guides their choices incorporates these two aspects, although some unobserved factors, such as the reputation of the service provider, loyalty effects, or simple inertia, also influence the decision. These *hidden* effects are captured by a random variable κ_s , such that user's utility (or satisfaction level), of choosing alternative $s \in \mathcal{S}$, is modeled as $V_s = v_s + \kappa_s$, where v_s denotes a deterministic component of the utility, and κ_s is assumed to follow a Gumbel distribution independently across options [14], [15].

This assumption leads to the well-known multinomial logit model [28], where the probability of choosing option s is:

$$\theta_s = e^{v_s} / \sum_{j \in \mathcal{S}} e^{v_j}. \quad (6)$$

The above formula depends only on the deterministic part of the user's utility, whose derivation is presented next.

B. Value of choices

By combining the considerations made in the previous section with those introduced here, we can define the deterministic component of the utility for each alternative $s \in \mathcal{S}$, on a case-by-case basis as follows:

$$\begin{aligned}
v_0 &\stackrel{\text{def}}{=} 0, \\
v_i(x_i) &\stackrel{\text{def}}{=} \alpha \ell \log(1 + k x_i) - \beta \log(p_i), \\
v_{(c,i)}(x_{(c,i)}) &\stackrel{\text{def}}{=} \alpha \ell \log(1 + k x_{(c,i)}) - \beta \log(p_c),
\end{aligned} \tag{7}$$

where we have assumed that price dependency is also logarithmic, that the importance of QoE and price is different and is expressed through parameters α and β , and that the utility of CP subscribers only depends on live streaming services (because of the stringent requirements of live streaming).

By rearranging the terms in the above expressions, considering that quantities x_i and $x_{(c,i)}$ are functions of vector θ defined in (2), we introduce auxiliary functions that allow for a compact definition of utility components and facilitate derivations based on the application of logit dynamics:

$$\begin{aligned}
f_0(\theta) &\stackrel{\text{def}}{=} 1, \\
f_i(\theta) &\stackrel{\text{def}}{=} (1 + k x_i)^\ell, \\
f_{(c,i)}(\theta) &\stackrel{\text{def}}{=} (1 + k \cdot x_{(c,i)})^\ell.
\end{aligned}$$

If we also introduce the cost-related quantities $p_{(c,i)} = p_c + p_i$, which is the total price paid to access live streaming through an ISP, we can express the deterministic utility component in a compact form as follows:

$$v_s = \alpha \log(f_s(\theta)) - \beta \log(p_s), \quad \forall s \in \mathcal{S}. \tag{8}$$

C. Population game dynamics

With the analysis derived in the previous subsection, (6) can be rewritten as follows:

$$\theta_s = \frac{(f_s(\theta))^\alpha}{(p_s)^\beta} \bigg/ \left(\sum_{i \in \mathcal{S}} \frac{(f_i(\theta))^\alpha}{(p_i)^\beta} \right), \tag{9}$$

which represents the probabilities to choose an option through an implicit function and hints to the fact that we will need numerical iterations to find the distribution of users' choices.

To find the distribution of users over the possible choices, i.e., to compute θ , we resort to a game theoretic framework in the context of population games proposed in [20]. By so doing, the properties of population games guarantee convergence and uniqueness of the equilibrium solution, at the same time allowing the evaluation of population dynamics in the case of sudden changes in operational conditions (e.g., when network parameters or prices change, or when users join or leave the network).

These dynamics capture the feedback loops between individual decisions and the overall system state, making them a fundamental component in the analysis of population games. Various types of dynamics have been proposed in the literature, each characterized by different assumptions about agent behavior (see [20], Chapter 5 for details).

As we already said, in this paper we focus on **Logit Dynamics**, which introduce a stochastic component to the adjustment process. With this approach, population dynamics can be described by the following iterative update rule:

$$\theta^{[k+1]} = \theta^{[k]} + \gamma \left(\phi(\theta^{[k]}) - \theta^{[k]} \right), \tag{10}$$

where $\phi(\cdot)$ is computed according to (9) and represents the drift of vector θ , which is smoothed through a parameter γ .

Specifically, parameter γ in the update rule plays a critical role in shaping the speed and stability of the population dynamics. Small values of γ lead to slow, gradual adjustments in the population distribution, capturing the idea that users are cautious and tend to maintain their current strategy unless there is a strong incentive to switch. In contrast, large values of γ result in rapid adjustments, making the population more responsive to changes in payoffs, but potentially leading to overshooting and oscillations around the equilibrium. This balance is essential for achieving both stability and flexibility in the population dynamics.

We briefly outline the main arguments underlying existence, stability, uniqueness, and convergence of the population dynamics in (10); full proofs are omitted due to space constraints.

A first key point concerns the structure of the functions $f_s(\theta)$, which depend on $1/\theta_s$ as a consequence of congestion effects. These expressions become undefined if any component θ_s approaches zero. To ensure mathematical consistency, we therefore restrict the domain of the update function to the interior of the simplex, i.e., to population states in which all choices in $\mathcal{S}$ retain a strictly positive share of users. This assumption is not restrictive in practice, as in realistic deployments, all service options typically maintain a non-negligible user base.

Under this restriction, each function $f_s(\theta)$ is well-defined, strictly positive, and continuously differentiable. These properties ensure continuity of the induced payoff functions and, together with compactness of the domain, guarantee the existence of at least one equilibrium point for the induced population game, following standard results for population games [20].

Moreover, the payoff associated with each strategy is strictly decreasing in the corresponding population share θ_s , reflecting congestion effects. This monotonicity property implies that the induced population game belongs to the class of stable games in the sense of population game theory [20], [29]. Stability is a structural property of the game and does not depend on the specific adjustment dynamics. For stable population games, the equilibrium is unique [29].

Finally, the update rule in (10) corresponds to a multinomial logit revision protocol in the sense of [20], i.e., a stochastic, payoff-based adjustment process in which users revise their choices with probabilities that increase smoothly with perceived payoffs. For stable population games, established results show that logit dynamics converge globally to the unique equilibrium from any interior initial condition [20], [29].

D. A note on the price of disconnection

In our model, p_0 denotes the cost of the null option, i.e., users do not subscribe to any ISP and cannot access online content. Although this option involves no payment, it represents a reference level of non-utility or missed opportunity cost. Following prior work (e.g., [18]), we set $p_0 = 1$ as

a normalized benchmark. This normalization does not affect generality, allows all other prices—such as p_i and p_c —to be interpreted relative to no connectivity, and is consistent with the logit framework, where utility is defined up to an additive constant.

VI. SIDE PAYMENTS AND UTILITIES

Having analyzed user population responses to prices and QoE, we now focus on how side payments affect population games and the resulting ISP and CP utilities under congestion-dependent user choices.

A. Side payments

In our model, side payments are interpreted as a mechanism through which ISPs can improve the quality of live streaming services offered to CP users. Specifically, we assume that resources acquired via side payments are fully reinvested to support CP-generated live streaming traffic, for example, by acquiring or renting additional capacity or by enabling service delivery to a larger number of users.

Building on this interpretation, a straightforward mechanism assumes that side payments are proportional to the CP subscription price and scale with the number of CP subscribers served by each ISP. We denote the per-subscriber side payment from the CP to ISP i as $q_i = \delta_i p_c$, with $0 \leq \delta_i < 1$.

This class of *proportional side payment* mechanisms is appealing due to its simplicity and transparency, as it links transfers to actual CP demand without requiring complex negotiations, dynamic feedback mechanisms, or central coordination. In its most general form, it only requires the CP and ISP i to agree on the proportionality factor δ_i .

Within this class, a natural and commonly adopted choice is to let δ_i depend on the market share of each ISP. In this work, we consider a proportional allocation based on the share of CP subscribers served by each ISP, which yields the following expression:

$$\delta_i = \frac{\theta_{(c,i)}}{\sum_{j \in \{A,B,\dots\}} \theta_{(c,j)}} \frac{q_{\text{tot}}}{p_c}, \quad \text{with } 0 \leq q_{\text{tot}} < p_c. \quad (11)$$

Summing over all ISPs, the above rule yields a total side payment equal to q_{tot} per CP subscriber. Hence, q_{tot} represents the fraction of the CP's per-user revenue allocated to side payments.

We emphasize that this allocation is not intended to be the unique or optimal side-payment mechanism. Rather, it serves as a simple and transparent baseline that allows us to study the impact of side payments on QoE, congestion, and user dynamics under a demand-aligned rule. Alternative mechanisms—such as negotiation-based bargaining solutions, regulator-imposed allocations, or cost-sharing schemes based on technical or economic metrics—are outside the scope of this work. Accordingly, in what follows we focus on the proportional side payment scheme defined in (11).

B. Utility of ISPs and CP

Since we assume that side payments are immediately reinvested to enhance the ISP delivery capacity, and all CP subscribers need to subscribe to an ISP, the utility of ISP i at equilibrium is given by the product of the ISP subscription price p_i times the total number of ISP subscribers $(\theta_{(-,i)} + \theta_{(c,i)}) N_{\text{tot}}$, which includes users with basic connectivity and with access to the live streaming service of CP. Hence, considering the utility normalized to the cardinality of the population N_{tot} , we obtain the following relation:

$$U_i = p_i (\theta_{(-,i)} + \theta_{(c,i)}). \quad (12)$$

This assumption marks a conceptual difference from prior models such as those in [17], [18], where side payments directly increase the ISP's revenue. In contrast, we interpret them as cost-offsetting instruments used to maintain acceptable QoE levels under high load, which is particularly relevant in the context of massive-scale live streaming. However, side payments indirectly cause ISP's utility alterations due to the fact that they affect QoE, hence the user distribution of θ .

In turn, the utility of the CP, normalized to the cardinality of the population N_{tot} , can be written as a function of the difference between subscription fees and side payments per CP subscriber:

$$U_c = \sum_{i \in \{A,B,\dots\}} (p_c - q_i) \theta_{(c,i)}. \quad (13)$$

In the above utility formulas, we have intentionally omitted production and network costs, which do not affect utility variations since they are practically constant, as reported by the BEREC in a recent study [30].

C. Impact of side payments on QoE

A key modeling choice concerns the conversion of monetary side payments into network capacity, which in turn can be seen as an increase in the value of $N_{\text{max}}^{(c,i)}$ for ISPs that receive side payments. This has a positive effect on the argument of the log function in (5), hence it causes a convex increase in the QoE for live streaming, with progressively diminishing returns. Hence, excessively high side payments might be counter productive for the CP, as we will analyze through numerical evaluation of system performance.

While in reality the relationship between side payments and network capacity may depend on the specific market, geography, or technology adopted by the ISP, unless otherwise specified, we assume a super-linear transformation, which models progressive discounts typically offered by infrastructure providers for bundle contracts. In practice, we will use the following relation between side payments q_i (received by ISP i per each CP subscriber it hosts) and increase of live streaming capacity $\Delta N_{\text{max}}^{(c,i)}$:

$$\Delta N_{\text{max}}^{(c,i)} = (\zeta q_i \theta_{(c,i)} N_{\text{tot}})^\eta, \quad (14)$$

where ζ is the money-to-capacity conversion factor and $\eta > 1$ models the superlinear discount effect.




Fig. 3: Evolution of the metric U_c as a function of the price p_c (results obtained by varying α with fixed $\beta = 3$); black markers are used to report peak values




Fig. 4: Distribution of the three user groups $\{0, (-, A), (c, A)\}$

D. Market perturbations caused by side payments

Side payments translate into increased ISP streaming capacity, hence they affect QoE. In turn, we have seen how QoE affects user's choices, which is reflected in θ . Hence, side payments ultimately perturb the market share of ISPs.

To measure the perturbation due to side payments, we introduce a metric aimed at assessing whether, and to what extent, the allocation of incentives to ISPs impacts the *status quo*—that is, the equilibrium distribution of users in the absence of such incentives.

Specifically, we define the *market perturbation* due to side payments $\mathbf{q} = \{q_A, q_B, \dots\}$, with $\sum_{j \in \{A, B, \dots\}} q_j = q_{\text{tot}}$, hereafter denoted as $M_i(\mathbf{q})$ for each ISP i , as the relative variation caused to the number of ISP subscribers, i.e.:

$$M_i(\mathbf{q}) = \frac{\theta_{(-,i)}^{(\mathbf{q})} + \theta_{(c,i)}^{(\mathbf{q})}}{1 - \theta_0^{(\mathbf{q})}} \cdot \frac{1 - \theta_0^{(0)}}{\theta_{(-,i)}^{(0)} + \theta_{(c,i)}^{(0)}} - 1, \quad (15)$$

where superscript $(\mathbf{q})$ denotes choice probabilities affected by side payments $\mathbf{q}$, whereas (0) indicates probabilities in the absence of side payments.

Another interesting metric is represented by the amount of side payment that should be added (or subtracted) to an ISP to neutralize the market perturbation. We call such quantity *equity compensation*, which can be formally defined as

$$E_i(\mathbf{q}) = q'_i - q_i \mid M_i(\mathbf{q}') = 0 \quad \forall i \in \{A, B, \dots\}. \quad (16)$$

VII. NUMERICAL EVALUATION

Our model is characterized by several parameters, summarized in Table I together with their meanings and the

TABLE I: Parameters used in numerical experiments.

Symbol	Meaning	Value
p_0	Reference cost associated with option 0	1
p_A, p_B	Connectivity service user cost (paid to the ISP)	2
p_c	Live streaming service user cost (paid to the CP)	3 ^[a]
α	User's sensitivity to QoE.	4 ^[b]
β	User's sensitivity to prices.	3
γ	Population game update speed	0.01
δ_i	Fraction of p_c given to ISP i as side payment	0 to 1
ζ	Money-to-capacity conversion factor	1
η	Superlinear discount factor	1.05
θ_s, θ	Option selection probabilities	0 to 1
k	QoE log-shape factor	100
N_{tot}	Population size	20M ^[c]
$\sum_i N_{\text{max}}^{(i,b)}$	Total service delivery capacity for connectivity	14M ^[c]
$\sum_i N_{\text{max}}^{(i,s)}$	Total service delivery capacity for live streaming	0.6M ^[c]

^[a] In some experiments, p_c varies in the range $[0, \dots, 10]$, while in others it is fixed at 3.

^[b] In some experiments, values vary in the range $[1, \dots, 4]$.

^[c] Expressed in terms of the number of users.

values used in the numerical experiments. We first analyze a simplified scenario with a single CP and one ISP, denoted A , which allows us to examine the model's core parameters and the effects of side payments without ISP competition. We then extend the analysis to a scenario with two ISPs, A and B , and a single CP.

In the single ISP case, before receiving side payments, ISP_A is assumed to be capable of offering connectivity to around 0.6 million streaming users and basic connectivity to 14 million users, out of a population of 20 million potential users.

The experiments involving two ISPs aim to explore their interactions as they compete to acquire users. The CP, through side payments, contributes to increasing the ISPs' capacity for live streaming services. In most of the numerical experiments, we considered ISPs of different sizes to examine how the CP's distribution of side payments affects the two ISPs. Specifically, the allocation of resources for connectivity and streaming was structured so that the ratio between the two ISPs is about 15% vs. 85% in terms of resources and about 30% vs. 70% in terms of market share, respectively for ISP A and B, while maintaining the aggregate values of the single ISP case.

A. Single ISP scenario

Fig. 3 illustrates the utility of the CP as a function of its service subscription fee p_c . Each curve in the figure corresponds to a different value of parameter α , with fixed value $\beta = 3$, capturing how the interplay between the subscription price p_c and users' willingness to pay for the streaming service influences U_c . When $\alpha < \beta$, users' choices prioritize cost over QoE. When users are largely insensitive to QoE ($\alpha = 1$, corresponding to the black solid line), their valuation of the service option (c, A) is primarily determined by its price p_c , as described in (7). While lower values of α provide the CP with an utility skewed towards zero, as α increases (with β




Fig. 5: Utility of CP (U_c) versus CP subscription cost for different side payment volumes ($\alpha = 4$, $\beta = 3$)

held constant), the relative importance of price in the user's choice diminishes, while the perceived QoE gains greater influence. This shift enhances the users' overall valuation of the streaming service, leading to a higher utility for the CP. Consequently, the peak of each utility curve—indicated by a black marker—moves both upward and to the right, meaning that users are increasingly willing to accept higher subscription fees as QoE becomes a more significant factor in their choices.

In the experiments that follow, we focus our discussion on the scenario with $\alpha = 4$ and $\beta = 3$. For this case, Fig. 4 illustrates how the population distribution vector θ evolves as the streaming subscription price p_c increases. With low values of p_c , the distribution shows a majority of users choosing to subscribe to ISP A (blue curve), and the remaining population dividing between clients opting out of both services (black curve) and clients subscribing to both the ISP and CP (red curve). When p_c grows, as expected, we see that the blue curve gets higher, while the red curve goes down and the black curve exhibits a modest increase.

Fig. 5 shows the impact of the amount of side payments on the CP's utility, U_c . The solid black line is the same as the yellow dashed line in Fig. 3 with $\alpha = 4$ and $\beta = 3$, that might lead the CP to set $p_c = 3$, roughly at the peak of the curve. It is very interesting to observe that the CP gains in utility by deciding to share a part of its income with the ISP through side payments, thus enabling the ISP to enhance its infrastructure in such a way that the perceived QoE of the streaming service increases (and the number of subscribers consequently grows). In particular, if the content provider agrees to a side payment of $q_{tot} = q_A = 0.4 p_c$, the resulting utility increases by approximately 70%. Indeed, at $p_c = 3$, the CP's utility U_c rises from 0.14—in the absence of side payments—to 0.24 when side payments are introduced. Furthermore, observing the curves of U_c as a function of the price p_c for different amounts of side payment q_A , a clear trend emerges: side payments are beneficial for the CP utility, but the utility gain diminishes as side payments increase. This suggests that diverting increasingly higher fractions of p_c to ISPs through q_A may lead to diminishing or even negative returns in terms of CP utility—which indeed happens in the case of $q_A = 0.5 p_c$.




Fig. 6: Perturbation of overall market share (M_i) and of streaming market share due to side payments. The vertical dotted line marks the overall market share before side payments.




Fig. 7: Effect of side payments and baseline network capacity share on market perturbations (left) and corresponding equity compensation (right)

B. Two-ISP scenario

Considering multiple ISPs enables us to evaluate the effect of side payments on ISP market shares, with a focus on the possible market perturbations induced by side payments. In the following, for the sake of simplicity, we restrict our analysis to the case of two ISPs.

The smaller ISP (i.e., A), with no side payments, has an infrastructure with about 15% of the total installed capacity and a market share of about 30% of the population of streaming customers. Side payments translate into network infrastructure upgrades, with corresponding improvements of the capability of both ISPs to provide live streaming with high QoE. Fig. 6 shows the market share perturbation of the two ISPs (curves for ISP A are black, while blue is for ISP B) for increasing values of the normalized side payment to ISP A, q_A/p_c (the rest going to ISP B). With market perturbation we mean the percentage increase (or decrease when negative) of client population gained (or lost) due to the network upgrades to the ISPs infrastructure enabled by side payments. This corresponds to the expression in (15), which refers to deviations from the ISP initial market share. In the figure, solid lines refer to the whole market (basic access + streaming), while the broken lines refer to the streaming market. The vertical dotted red line marks the point corresponding to side payments proportional to the initial market share of 30%–70% of ISP A and B respectively. The values of the curves at the intersections with the vertical




Fig. 8: Effect of side payments on the sizes of the option sets, and CP and ISP utilities, vs. the total side payment.

red line indicate the market distortion that would be attained in case of an assignment of side payment proportional to the initial market share. Such allocation criterion, although straightforward, clearly violates the NN principle by greatly increasing the smaller ISP market share, and causing a market distortion of about 25%. On the contrary, the points where the two curves cross each other and intersect the horizontal axis—either in the case of overall market deviation or streaming market deviation—represent situations where side payments introduce no distortion in the market. Being able to quantify the amount of market distortion induced by side payments, and even to identify allocations of side payments that avoid distortion is of great value for national regulators. When comparing the two market deviations—namely overall and streaming—it is clear that system-wide deviations are smaller, with values ranging from -4 to 13% , while streaming market deviations get up to 150% . Such difference depends on the fact that in this setting we have small $\theta_{(c,i)}$ and big $\theta_{(-,i)}$, and movements of a same number of users impact differently the two percentage market deviations.

In Fig. 7 we show in the left plot the market perturbation of (15) for ISP A versus the ISP capacity share for two values of q_{tot} . In the right plot we show the corresponding (negative) equity compensation of (16). We can see that the imbalance in the ISP resources produces larger deviations in favor of the smaller ISP. Deviations could be compensated by reducing the amount of side payments going to the smaller of the two ISPs. However, the equity compensation is not a monotonic function due to two contrasting effects: infrastructure imbalance is further exacerbated by side payments, but user's QoE grows faster with infrastructure upgrades when the network is small. In general, our results reveal that side payments play a constructive role in facilitating ISP market re-equilibration while at the same time increasing the QoE of users.

Finally, Fig. 8 illustrates the impact of side payments on both user population shares and utilities. The left plot reports the distribution of users across the available options—no service (θ_0 , black lines), streaming over ISP A ($\theta_{(c,A)}$, blue lines), and streaming over ISP B ($\theta_{(c,B)}$, red lines)—as a function of the fraction of the CP fee allocated to side payments, which are distributed according to the proportional rule in (11). Beyond the population shifts already discussed, this plot also highlights the positive effect of side payments on users' experienced QoE. This is evident from the comparison between the trajectories of $\theta_{(c,A)}$ and $\theta_{(c,B)}$ with and without

side payments, and stems from the explicit dependence of QoE on side payments in Eq. (4), Eq. (5), and Eq. (7).

Horizontal lines indicate the population shares in the absence of side payments. Both $\theta_{(c,A)}$ and $\theta_{(c,B)}$ increase almost linearly as side payments grow. However, ISP A exhibits a larger relative increase—despite a lower initial share—consistent with the trends already observed in Fig. 6. Moreover, the slight decrease in the θ_0 group suggests that improved QoE attracts some previously non-subscribing users back into the system.

The middle plot shows the percentage change in the utilities of the CP and the two ISPs as a function of the fraction of the CP fee allocated to side payments. While the utility of ISP B slightly decreases and that of ISP A increases by a few percentage points, the CP utility can grow by up to 110% for large side payments, driven by the increase in CP subscriptions observed in the left plot.

Since side payments reduce ISP B's utility, ISP B may choose to refuse them, leaving the payment with the CP. This case is illustrated in the right plot of Fig. 8, where ISP B performs worse than in the baseline scenario and the CP benefits much less than before, while ISP A emerges as the only winner, with utility gains almost doubling.

VIII. CONCLUSIONS AND FURTHER DEVELOPMENTS

We developed a modeling framework capturing the interactions between ISPs, users, and a live streaming CP, focusing on pricing policies and side payment allocations. Based on population games, it enables the analysis of how these policies affect resource distribution and user QoE under congestion-dependent conditions, and introduces quantitative network neutrality metrics to assess how side payments may disrupt neutrality and induce market distortions, which in our setting contribute to ISP market re-equilibration.

The framework has practical relevance for live streaming of sports events, where subscription fees are typically fixed over long periods, supporting CPs in price setting and side payment evaluation, and assisting regulators in assessing compliance with network neutrality principles. To the best of our knowledge, no prior modeling framework quantitatively supports these tasks.

Future work will address NN perturbations induced by practices such as zero-rating, bundled offers, and vertical integration, as well as alternative side payment allocation mechanisms accounting for ISP-CP power dynamics or regulatory intervention.

REFERENCES

- [1] AppLogic Networks, “Global Internet Phenomena Reports 2025,” AppLogic Networks, Tech. Rep., 2025. [Online]. Available: https://www.sandvine.com/hubfs/AppLogic_Networks/Collateral/Global%20Internet%20Phenomena%20Reports/GIPR%202025.pdf
- [2] N. Economides and J. Tåg, “Network neutrality on the Internet: A two-sided market analysis,” *Information Economics and Policy*, vol. 19, no. 2, pp. 210–226, 2007.
- [3] 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, 2009.
- [4] J. Krämer, L. Wiewiorra, and C. Weinhardt, “Net neutrality: A progress report,” *Telecommunications Policy*, vol. 37, no. 9, pp. 794–813, 2013.
- [5] R. T. B. Ma, D. M. Chiu, J. C. S. Lui, V. Misra, and D. Rubenstein, “On Cooperative Settlement Between Content, Transit, and Eyeball Internet Service Providers,” *IEEE/ACM Transactions on Networking*, vol. 19, no. 3, pp. 802–815, 2011.
- [6] A. Nikkiah and S. Jordan, “A two-sided model of paid peering,” *Telecommunications Policy*, vol. 46, no. 6, p. 102352, 2022.
- [7] B. Frank, I. Poese, Y. Lin, G. Smaragdakis, A. Feldmann, B. Maggs, J. Rake, S. Uhlig, and R. Weber, “Pushing CDN-ISP collaboration to the limit,” *SIGCOMM Comput. Commun. Rev.*, vol. 43, no. 3, 2013.
- [8] T. Wu, “Network Neutrality, Broadband Discrimination,” *Journal of Telecommunications and High Technology Law*, vol. 2, pp. 141–179, 2003. [Online]. Available: <https://ssrn.com/abstract=388863>
- [9] B. van Schewick, *Network Neutrality: From Policy to Law to Regulation*. Oxford University Press, 2021.
- [10] M. A. Lemley and L. Lessig, “The end of end-to-end: Preserving the architecture of the internet in the broadband era,” *UCLA Law Review*, vol. 48, no. 4, pp. 925–972, 2001, reprinted and cited extensively in policy debates on net neutrality. [Online]. Available: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=247737
- [11] C. T. Marsden, *Net Neutrality: Towards a Co-Regulatory Solution*. London, UK: Bloomsbury Publishing, 2010.
- [12] AGCOM, “Delibera n. 206/21/CONS: Atto di indirizzo per il corretto dimensionamento e la dislocazione geografica della rete di distribuzione (CDN) delle partite di calcio di serie A per le stagioni 2021-2024 in live streaming,” 2021, <https://www.agcom.it/sites/default/files/migration/delibera/Delibera%20206-21-CONS.pdf>.
- [13] —, “Approvazione della proposta definitiva di impegni relativa al procedimento sanzionatorio n. 1/22/DTC presentata dalla società Dazn Limited Ltd ai sensi della legge 4 agosto 2006, n. 248 e del regolamento allegato alla delibera n. 697/20/CONS,” 2022, <https://www.agcom.it/provvedimenti/delibera-250-22-cons>.
- [14] M. E. Ben-Akiva and S. R. Lerman, *Discrete Choice Analysis: Theory and Application to Travel Demand*. The MIT Press, 1985.
- [15] K. E. Train, *Discrete Choice Methods with Simulation*. Cambridge University Press, 2003.
- [16] E. Altman, A. Legout, and Y. Xu, “Network Non-neutrality Debate: An Economic Analysis,” in *NETWORKING 2011*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2011, pp. 68–81.
- [17] P. Coucheney, P. Maillé, and B. Tuffin, “Network neutrality debate and ISP inter-relations: traffic exchange, revenue sharing, and disconnection threat,” *NETNOMICS: Economic Research and Electronic Networking*, vol. 15, p. 155–182, 2014.
- [18] P. Maillé and B. Tuffin, *From Net Neutrality to ICT Neutrality*. Springer Nature, 2022.
- [19] W. Briglauer, C. Cambini, K. Gugler, and V. Stocker, “Net neutrality and high-speed broadband networks: evidence from OECD countries,” *European Journal of Law and Economics*, vol. 54, no. 3, pp. 327–357, 2022.
- [20] W. H. Sandholm, *Population Games and Evolutionary Dynamics*. Cambridge, MA: MIT Press, 2010.
- [21] L. Guijarro, V. Pla, and J. R. Vidal, “Isp pricing and platform pricing interaction under net neutrality,” *arXiv preprint*, vol. arXiv:2401.14791, 2024.
- [22] C. Jiang, L. Gao, F. Hou, and J. Li, “Economic analysis of edge caching enabled mobile internet ecosystem,” *IEEE Transactions on Mobile Computing*, 2024.
- [23] A. Gautier, J.-C. Poudou, and M. Roland, “Net neutrality and universal service obligations: It’s all about bandwidth,” *Review of Industrial Organization*, 2024.
- [24] P. Reichl, B. Tuffin, and R. Schatz, “Logarithmic laws in service quality perception: where microeconomics meets psychophysics and quality of experience,” *Telecommunication Systems*, vol. 52, no. 2, pp. 587–600, 2013.
- [25] P. Reichl, B. Tuffin, and P. Maillé, “Economics of Quality of Experience,” in *Telecommunication Economics*, 2012.
- [26] P. Lebreton and K. Yamagishi, “Transferring Adaptive Bit Rate Streaming Quality Models from H.264/HD to H.265/4K-UHD,” *IEICE Transactions on Communications*, vol. E102-B, no. 12, pp. 2226–2241, 2019.
- [27] T. Kimura, T. Kimura, A. Matsumoto, and K. Yamagishi, “Balancing Quality of Experience and Traffic Volume in Adaptive Bitrate Streaming,” *IEEE Access*, vol. 9, pp. 15 530–15 545, 2021.
- [28] D. McFadden, “Conditional Logit Analysis of Qualitative Choice Behavior,” in *Frontiers in Econometrics*, P. Zarembka, Ed. New York: Academic Press, 1974, pp. 105–142.
- [29] J. Hofbauer and W. H. Sandholm, “Stable games and their dynamics,” *Journal of Economic Theory*, vol. 144, no. 4, pp. 1665–1693, 2009.
- [30] BERECA via Communications Chambers, “Network traffic growth does not cause significant incremental costs,” Tech. Rep., 2023, traffic-insensitive network cost analysis.