Quality Competition Among Internet Service Providers

Simon Scherrer[a], Seyedali Tabaeiaghdaei[a], Adrian Perrig[a]

[a]*ETH Zurich*

Abstract

Internet service providers (ISPs) have a variety of quality attributes that determine their attractiveness for data transmission, ranging from quality-of-service metrics such as jitter to security properties such as the presence of DDoS defense systems. ISPs should optimize these attributes in line with their profit objective, i.e., maximize revenue from attracted traffic while minimizing attribute-related cost, all in the context of alternative offers by competing ISPs. However, this attribute optimization is difficult not least because many aspects of ISP competition are barely understood on a systematic level, e.g., the multi-dimensional and cost-driving nature of path quality, and the distributed decision making of ISPs on the same path.

In this paper, we improve this understanding by analyzing how ISP competition affects path quality and ISP profits. To that end, we develop a game-theoretic model in which ISPs (i) affect path quality via multiple attributes that entail costs, (ii) are on paths together with other selfish ISPs, and (iii) are in competition with alternative paths when attracting traffic. The model enables an extensive theoretical analysis, surprisingly showing that competition can have both positive and negative effects on path quality and ISP profits, depending on the network topology and the cost structure of ISPs. However, a large-scale simulation, which draws on real-world data to instantiate the model, shows that the positive effects will likely prevail in practice: If the number of selectable paths towards any destination increases from 1 to 5, the prevalence of quality attributes increases by at least 50%, while 75% of ISPs improve their profit.

1. Introduction

In today's Internet, the Border Gateway Protocol (BGP) supplies ISPs with potentially multiple paths towards an IP prefix. When selecting among these paths, ISPs decide on the basis of price and quality of the available paths. This path quality is determined by multiple *quality attributes* of potentially multiple on-path ISPs. Such quality attributes may include conventional performance metrics (e.g., bandwidth, latency, loss rate, jitter) or security features (e.g., presence of security middleboxes), but also properties that traditionally receive less attention, e.g., environmental, social, and governmental (ESG) properties such as carbon emission from data transmission [49] or geopolitical concerns regarding on-path ISPs [13]. Transit ISPs invest in their attributes, communicate them in path announcements, and thereby attract traffic from selecting ISPs. However, improving these attributes comes at a cost, which may exceed the additional revenue from attracted traffic, especially if ISPs on competing paths also raise their quality level.

Given this competitive setting, ISPs today face complex strategic questions when optimizing profit: What quality attributes should be invested in, and to what extent? How should prices be determined? And how are these decisions affected by ISPs on competing paths *and* ISPs elsewhere on the provided paths?

Well-informed strategic decisions thus require a fundamental understanding of ISP competition under path selection, not only on an intuitive, but also on a rigorous analytical level. While such an understanding has been furthered by previous academic research [25, 29, 36, 41, 47], many open questions of practical relevance remain, e.g., regarding the multi-attribute nature of path quality, the dependence of fixed and variable ISP cost on provided quality, the feasibility of cooperation among ISPs on the same path, and the impact of differing degrees of competition intensity (cf. §6).

To address these questions, we present a new game-theoretic model, enabling a rigorous investigation of quality competition among ISPs. We perform this investigation through theoretical analysis and simulation:

Theoretical analysis. We conduct an extensive theoretical analysis to systematically understand the effect of ISP competition on path quality, and ISP profits (Path price constitutes a quality attribute in our model). In particular, we identify closed-form solutions for the Nash equilibria of the competition dynamics, prove the stability of these equilibria, and contrast them for varying degrees of path diversity and ISP heterogeneity. On the one hand, this theoretical analysis confirms intuitive insights, namely that competition tends to raise the prevalence of valuable attributes. On the other hand, our model reveals counter-intuitive insights, namely that the cooperation between ISPs on the same path suffers from a prisoner's dilemma, that ISP profits can increase under intensified competition, and that additional paths may decrease the prevalence of quality attributes if unchangeable path attributes are starkly different.

Simulation-based case study. To determine which competition effects are significant in practice, we leverage our model for a simulation-based case study. In this case study, we investigate the competition dynamics in the Internet core with respect to two attributes (internal bandwidth and the share of clean energy used by an ISP). This simulation requires a numerical instantiation of the model, based on real-world data. For this model instance, our simulations yield robust evidence that competition raises the prevalence of valuable attributes, the quality of available paths, and the profits of most ISPs.

In summary, our work includes the following contributions:

- **Game-theoretic competition model:** Our new ISP-competition model (§2) departs from previous competition models by representing both inter-path competition and intra-path cooperation, accommodating a multi-faceted notion of path quality, revealing the effect of path diversity, and reflecting realistic ISP cost structures (§6).
- **Theoretical analysis:** We conduct a rigorous theoretical analysis by reasoning from basic competition scenarios that showcase the fundamental effects in ISP competition (§3). In particular, we contrast monopolistic and competitive scenarios in ISP path selection, investigate networks with varying similarity in ISP profit functions, and identify asymptotically stable equilibria and social optima of the competition. Our analysis suggests that ISP competition has nuanced effects on ISP profits and path quality, going beyond the predictions of basic economic theory.
- **Large-scale simulation:** We demonstrate how to instantiate our model based on real-world data, with the goal of predicting competition effects in the Internet core (§4). These predictions are generated with simulations, which rely on randomization to achieve robust results, represent the competition behavior with better-response dynamics, and are executed for varying path diversity. The simulation results suggest that competition, induced by path diversity, has positive effects for a majority of ISPs on multiple tiers of the Internet, i.e., raises ISP profits and path quality (§5).

2. Model and First Insights

In the following, we present a game-theoretic model, which we employ to investigate the competition dynamics under attribute-oriented path selection. While our model reflects common characterizations of inter-domain network economics, it is more general than previous models (cf. §6).

Network and paths. We abstract the network as a set N of ISPs, which represent the players in the competition game. Each ISP $n \in N$ is assumed to be fully rational. The ISPs form paths, where each path $r \subseteq N$ is a set of ISPs. All usable paths in a network are collected in the path set R, and all usable paths between selecting ISP n_1 and destination ISP n_2 constitute the set $R(n_1, n_2)$. Throughout this work, we study how ISPs affect the quality of paths as given by path set R, not how ISPs strategically adapt the set R of usable paths via interconnection agreements and announcements, which is a related but distinct problem [26, 35].

Attributes. We consider a network with a set K of ISP attributes, $|K| \geq 1$, that are relevant in path selection. Hence, each ISP n is associated with an attribute vector $\mathbf{a}_n \in \mathbb{R}^K_{\geq 0}$, where $a_{nk} \in \mathbb{R}_{\geq 0}$ denotes the prevalence of attribute $k \in K$ in ISP n. As a player in the competition game, each ISP n strives to choose its attributes $\mathbf{a}_n$ in order to optimize its profit (see below). Since the lowest possible degree of attribute prevalence is attained if an ISP does not possess an attribute at all, we restrict the attribute values to non-negative real numbers: $\forall n \in N, \forall k \in K.\ a_{nk} \geq 0$. For convenience of notation, we also define an attribute matrix $\mathbf{A} \in \mathbb{R}_{\geq 0}^{|N|\times|K|}$, with the n-th row being $\mathbf{a}_n$.

Path valuations. The attributes of an ISP n determine the attractiveness of using paths including that ISP. Hence, we define the attractiveness of available options on the level of paths, specifically by *valuation functions* $\{v_r\}_{r\in R}$. The valuation v_r for path r then depends on all attributes $\mathbf{a}_n$ of all on-path ISPs $n \in r$. Since we consider desirable attributes in our model, every function that is monotonically increasing in all attribute values is a suitable valuation function. Throughout this paper, we use affine functions:

$$v_r(\mathbf{A}) = \sum_{n\in r}\sum_{k\in K} \alpha_{rnk} a_{nk} + \alpha_{r0}, \tag{1}$$

where each $\alpha_{rnk} > 0$ determines how strongly attribute k of ISP n affects the valuation of path r, and $\alpha_{r0} \geq 0$ is the *base valuation* of path r. This formulation captures several real-world aspects of path valuations, as the variation in α_{rnk} captures that the attributes and ISPs associated with a path may have varying importance for path valuation, e.g., ISPs providing a large segment of the path might be more relevant for the valuation. The linear formulation might be counter-intuitive given that the marginal utility of attribute prevalence is likely decreasing; we rely on the formulation for path-selection probability below to capture that the volume of attracted demand on a path is sub-linear in path attributes. Moreover, we show by simulation that the model predictions do not strongly rely on the affine formulation (cf. §5.2.3).

Path-selection probability. Path valuations inform the path selection at the selecting ISP, and thus determine the probability of each path being selected. More precisely, when a selecting ISP n_1 selects a path towards a prefix hosted by ISP n_2, each path r among the available paths $R(n_1, n_2)$ is selected for transit with probability $p_r(\mathbf{A})$. Inspired by the popular logit-demand model [2], we consider the selection probability p_r to be proportional to the *relative attractiveness* of path r compared to alternative paths:

$$\forall (n_1, n_2) \in N \times N,\ r \in R(n_1, n_2). \quad p_r(\mathbf{A}) = \frac{v_r(\mathbf{A})}{1 + \sum_{r'\in R(n_1,n_2)} v_{r'}(\mathbf{A})}. \tag{2}$$

Crucially, the addition term 1 in the fraction denominator captures *demand elasticity*, i.e., selecting ISP n_1 might not select any path in $R(n_1, n_2)$ at all if the available paths are generally unattractive. Instead, selecting ISP n_1 might not offer its customers any path to n_2, create a new path to n_2 by concluding a peering agreement, or obtain the desired content from another destination ISP than n_2. Not least, this demand elasticity also avoids a singularity in the model when all paths are worthless, i.e., $v_r = 0\ \forall r$.

Demand. The path-selection probabilities above determine the demand volume D_n that is obtained by any ISP n, which we formalize by:

$$D_n(\mathbf{A}) = \sum_{r\in R.\ n\in r} \delta_r(\mathbf{A}) = \sum_{\substack{r\in R.\ n\in r \\ r\in R(n_1,n_2)}} p_r(\mathbf{A}) \cdot d_{(n_1,n_2)}. \tag{3}$$

Due to the elasticity of demand, actual total demand (i.e., $\delta_r(\mathbf{A})$ summed over all $r \in R\,(n_1, n_2)$) is strictly below the *demand limit* $d_{(n_1,n_2)}$.

The practical interpretation of path demand δ_r in Eq. (3) depends on the transit behavior of the selecting ISP n_1. If n_1 is a stub AS, then traffic originates within n_1 and can be split across multiple paths towards a given prefix. If n_1 thus selects multiple paths, δ_r denotes the actual demand allocated to path r. In contrast, if ISP n_1 is a transit AS, the traffic transited by n_1 must follow the single path announced by ISP n_1 to neighboring ISPs, as BGP transit loops might arise otherwise. If n_1 thus selects only a single path for transit, δ_r corresponds to the *expected* demand allocated on path r.

Profit. Given the demand model, an ISP n can affect attracted demand D_n with an appropriate choice of $\mathbf{a}_n$. However, the profit π_n depends not only on the volume of attracted demand, but also on the cost for provision of the attributes. We thus consider the profit function π_n of ISP n to have three components. First, the ISP profit is increased by a revenue function $\mathcal{R}_n$, which is a function of D_n. Second, the profit is reduced by a demand-dependent cost function Φ_n, which as well depends on D_n, but also directly on $\mathbf{a}_n$, as the cost of transmitting a unit of demand depends on the chosen attributes. Third, the profit is reduced by a demand-independent cost function Γ_n, which depends only directly on the chosen attributes $\mathbf{a}_n$ and thus corresponds to the 'fixed cost' of possessing certain attributes.

While these component functions could in principle be any monotonically increasing function, we use the following formulations for the component functions throughout this paper:

$$\mathcal{R}_n(\mathbf{A}) = \rho_n \cdot D_n(\mathbf{A}) \qquad \Phi_n(\mathbf{A}) = \left(\sum_{k\in K} \phi_{nk} a_{nk} + \phi_{n0}\right) \cdot D_n(\mathbf{A}) \qquad \Gamma_n(\mathbf{a}_n) = \sum_{k\in K} \gamma_{nk} a_{nk} + \gamma_{n0} \tag{4}$$

where the coefficient $\rho_n > 0$ is the maximum per-unit transit price of ISP n (actual transit prices are additionally subsumed within quality attributes, as described below). Furthermore, the coefficients $\phi_{nk} \geq 0$ and $\gamma_{nk} \geq 0$ determine the attribute-specific increase in demand-dependent and demand-independent cost, respectively, and the intercepts $\phi_{n0} \geq 0$ and $\gamma_{n0} \geq 0$ express the attribute-independent basic values for the respective cost terms. Throughout the paper, we assume $\rho_n \geq \phi_{n0}$ for all ISPs $n \in N$, as a rational ISP that loses money by attracting demand even with the most cost-saving strategy (i.e., $\mathbf{a}_n = \mathbf{0}$) would in fact go out of business. The affine formulations of Φ_n and Γ_n predict qualitatively similar competition effects as quadratic functions, as we demonstrate by simulation in §5.2.3.

In summary, the profit function π_n is of the following form in our investigations:

$$\pi_n(\mathbf{A}) = \mathcal{R}_n(\mathbf{A}) - \Phi_n(\mathbf{A}) - \Gamma_n(\mathbf{a}_n) = D_n(\mathbf{A}) \cdot \left(\rho_n - \sum_{k\in K} \phi_{nk} a_{nk} - \phi_{n0}\right) - \sum_{k\in K} \gamma_{nk} a_{nk} - \gamma_{n0}. \tag{5}$$

Undesirable attributes. So far, our model formulation assumes *desirable* attributes, which increase path attractiveness in high quantities and are costly to increase, e.g., bandwidth. However, many relevant ISP attributes are *undesirable* in high quantities and are challenging to decrease, e.g., transit price or latency. To accommodate undesirable attributes in the model, a naive approach would consist of allowing negative coefficients α_{rnk}, ϕ_{nk}, and γ_{nk} for any undesirable attribute k. However, such negativity would entail nonsensical model predictions, such as potentially negative path-selection probabilities from Eq. (2), or infinite profit given an undesirable attribute (n, k) with $a_{nk} = \infty$ and $\gamma_{nk} < 0$ (cf. Eq. (5)).

To avoid such nonsensical predictions and preserve model tractability, we suggest to convert undesirable attributes into their desirable counterparts. For example, the actual transit-price attribute $a'_{nk} \in [0, \rho_n]$ could be translated to a non-negative *cheapness* attribute $a_{nk} = \rho_n - a'_{nk}$. The cheapness attribute a_{nk} then formally contributes to the costs in Φ_n, while actually determining the traffic-unit revenue in the profit function π_n. For illustration, consider a monopolistic ISP n that charges a price a'_{nk} for forwarding one unit of traffic and provides a single path r connecting the ISPs n_1 and n_2. If the price a'_{nk} is the only relevant attribute for path valuation, the profit π_n of ISP n is affected as follows by the price a'_{nk}:

$$\begin{aligned}\pi_n(a'_{nk}) &= \mathcal{R}_n(a'_{nk}) - \Phi_n(a'_{nk}) - \Gamma_n(a'_{nk}) \\ &= \underbrace{\left(\frac{\alpha_{rnk}(\rho_n - a'_{nk})}{1 + \alpha_{rnk}(\rho_n - a'_{nk})} \cdot d_{(n_1,n_2)}\right)}_{\text{decreases in price}} \cdot \underbrace{(\rho_n - (\rho_n - a'_{nk}))}_{\substack{=a'_{nk} \text{ (price)} \\ \text{increases in price}}} - \gamma_{n0},\end{aligned} \tag{6}$$

which suggests some optimal price $a'_{nk} \neq \infty$. Note that the price a_{nk} affects the valuation of path r with a certain weight α_{rnk}, directly affects the profit per traffic unit (i.e., $\phi_{nk} = 1$), and does not affect the demand-independent cost (i.e., $\gamma_{nk} = 0$).

Nash Equilibrium. The competition dynamics in attribute-oriented path selection can be characterized by their Nash equilibria. In our setting, such a Nash equilibrium is a choice of attributes in which each ISP has optimal attributes (w.r.t. its profit) given the attributes of all other ISPs:

Definition 1. *A choice of attributes* $\mathbf{A}^{+}$ *form a* ***Nash equilibrium*** *if and only if*

$$\forall n \in N.\ \mathbf{a}_n^{+} = \mathbf{a}_n^{*}(\mathbf{A}_{-n}^{+}) = \arg\max_{\mathbf{a}_n \in \mathbb{R}_{\geq 0}^{|K|}} \pi_n\big(\mathbf{a}_n \oplus_n \mathbf{A}_{-n}^{+}\big) \tag{7}$$

where $\oplus_n$ *combines the attribute choice* $\mathbf{a}_n$ *of ISP* n *with the equilibrium attribute choices* $\mathbf{A}_{-n}^{+}$ *of the remaining ISPs.*

In this abstract form, the Nash equilibria offer little opportunity for analytical characterization. However, if the attributes a_{nk} are restricted to $[0, a_{\max}]$ rather than to $\mathbb{R}_{\geq 0}$ (e.g., if there is an upper bound $a_{\max}$ on all attribute values), the existence of Nash equilibria is guaranteed by Brouwer's fixed-point theorem [5]. This guarantee holds because the restriction of the attribute matrix $\mathbf{A}$ to a compact convex set necessitates that at least one fixed point $\mathbf{A}^{+}$ is preserved by the iteration function $\mathbf{A}^{*}(\mathbf{A}) = [\mathbf{a}_{n_1}^{*}(\mathbf{A}_{-n_1}), ..., \mathbf{a}_{n_N}^{*}(\mathbf{A}_{-n_N})]^{\top}$. To gain a deeper understanding of Nash equilibria beyond that special case, we concretize equilibria in this work, and investigate these equilibria with respect to existence, uniqueness, stability, and efficiency.

Social Optimum. To assess the efficiency of Nash equilibria, we compare these equilibria to *social optima*. In our setting, such a social optimum optimizes a metric that aggregates the perspectives of all agents involved in the competitive dynamics. Our model contains two types of agents, namely *selecting ISPs* and *transit ISPs*, with non-aligned interests, which warrants two different formalizations of the social optimum.

First, selecting ISPs are interested in path quality. Hence, the social efficiency for selecting ISPs is simply measured by the aggregate valuation V of all paths in the network, given a choice of attributes $\mathbf{A}$:

$$V(\mathbf{A}) = \sum_{r \in R} v_r(\mathbf{A}). \tag{8}$$

Since the valuation functions v_r are assumed to be linear and therefore unbounded in this paper, a finite social optimum for selecting ISPs only exists if all attributes are restricted to a finite domain.

Second, transit ISPs are interested in profit. To characterize the social optimum from the perspective of transit ISPs, we rely on the conditions of the *Nash bargaining solution (NBS)*, i.e., the conditions that a global attribute choice $\mathbf{A}$ would have to fulfill if ISPs had to agree on it in cooperative bargaining [30]. The two most important NBS conditions are *Pareto-optimality*, i.e., no ISP can increase its profit without any other ISP experiencing a decrease in its profit, and *symmetry*, i.e., among Pareto-optimal profit distributions, the fairest distribution is preferred. These conditions are achieved if the attribute choice $\mathbf{A}$ optimizes the *Nash bargaining product*:

Definition 2. *A choice of attributes* $\mathbf{A}^{\circ}$ *forms a* ***social optimum*** *from the perspective of transit ISPs if it corresponds to the* ***Nash bargaining solution (NBS)****, i.e.,*

$$\mathbf{A}^{\circ} = \arg\max_{\mathbf{A} \in \mathbb{R}_{\geq 0}^{|N| \times |K|}} \Pi_{n \in N}\ \pi_n(\mathbf{A}). \tag{9}$$

3. Theoretical Analysis

In this section, we theoretically analyze the competition dynamics in path selection. For that purpose, we focus on an individual *market* in isolation, i.e., the competition between transit ISPs for traffic between a single source-destination pair (n_1, n_2). As a result, we write $R = R(n_1, n_2)$ and $d = d_{(n_1,n_2)}$ throughout this section.

3.1. Optimal Attribute

To analyze the competition dynamics, we first investigate how any single ISP n should choose its attribute k in response to the publicly observable attribute choices of all other ISPs. This optimal attribute is given by the following closed-form solution:

Theorem 1. *Best-Response Attribute.* *In an individual market with arbitrarily overlapping paths, the optimal admissible attribute* a^*_{nk} *of ISP* n *given the remaining attributes* $\mathbf{A}_{-nk}$ *is*

$$a^*_{nk}(\mathbf{A}_{-nk}) = \begin{cases} \hat{a}^*_n(\mathbf{A}_{-nk}) & \text{if } \hat{a}^*_n(\mathbf{A}_{-nk}) \in \mathbb{R} \text{ and } \hat{a}^*_n(\mathbf{A}_{-nk}) \geq 0, \\ 0 & \text{otherwise,} \end{cases} \tag{10}$$

where $\hat{a}^*_{nk}(\mathbf{A}_{-nk})$ *is the optimal unrestricted (i.e., potentially complex or negative) attribute:*

$$\hat{a}^*_{nk}(\mathbf{A}_{-nk}) = \frac{1}{\alpha_{nk}} \left(\sqrt{\frac{d\big(1 + v_{-r(n)}(\mathbf{A}_{-nk})\big)}{d\phi_{nk} + \gamma_{nk}} \big(\phi_{nk}(1 + v_{-nk}(\mathbf{A}_{-nk})) + \alpha_{nk}(\rho_n - \Phi_{-nk}(\mathbf{A}_{-nk}))\big)} - \big(1 + v_{-nk}(\mathbf{A}_{-n})\big) \right) \tag{11}$$

Eq. (11) *uses the following abbreviations:*

$$\alpha_{nk} = \sum_{\substack{r \in R. \\ n \in r}} \alpha_{rnk} \qquad v_{-r(n)}(\mathbf{A}) = \sum_{\substack{r' \in R. \\ n \notin r'}} v_{r'}(\mathbf{A}) \tag{12}$$

$$v_{-nk}(\mathbf{A}) = \sum_{r \in R} \alpha_{r0} + \sum_{\substack{(n',k') \in N \times K. \\ (n',k') \neq (n,k)}} \alpha_{n'k'} a_{n'k'} \qquad \Phi_{-nk}(\mathbf{A}) = \sum_{k' \in K \setminus k} \phi_{nk'} a_{nk'} + \phi_{n0} \tag{13}$$

The proof of Theorem 1 is provided in Appendix A. To provide intuition about the formula in Eq. (11), we transform it to a simplified version:

$$\hat{v}^*_{r(n)} = \sum_{\substack{r \in R. \\ n \in r}} \hat{v}^*_r = \sqrt{\Big(\underbrace{\frac{d\phi_{nk}(1 + v_{-nk})}{d\phi_{nk} + \gamma_{nk}}}_{①} + \underbrace{\frac{d\alpha_{nk}(\rho_n - \Phi_{-nk})}{d\phi_{nk} + \gamma_{nk}}}_{②} \Big) \underbrace{\big(1 + v_{-r(n)}\big)}_{③}} - \underbrace{\big(1 + v_{-r(n)}\big)}_{③} \tag{14}$$

where $\hat{v}^*_{r(n)}$ contains the sum of unrestricted valuations of all paths containing ISP n that would be optimal for n given $\mathbf{A}_{-nk}$. This $\hat{v}^*_{r(n)}$ (and thus also $\hat{a}_{nk}$) correlates positively with term ①, which relates to the share of demand-dependent cost ($\propto d\phi_{nk}$) among total cost ($\propto d\phi_{nk} + \gamma_{nk}$) with respect to attribute (n, k). This correlation suggests that ISPs should champion attributes with low demand-independent cost compared to demand-dependent cost. Moreover, $\hat{v}^*_{r(n)}$ correlates with term ②, which relates to the revenue from attribute (n, k) per unit of cost from the attribute, i.e., the 'return' on attribute (n, k). Term ③, which describes the attractiveness of paths avoiding n, can have a positive effect on $\hat{v}^*_{r(n)}$ up to a point, as competition incentivizes ISP n to raise its attribute values. However, from a certain point onwards, term ③ has a negative effect on $\hat{v}^*_{r(n)}$, as detracting traffic from highly attractive alternatives becomes too costly compared to the achievable revenue.

For an individual market, the equilibrium condition from Definition 2 can thus be concretized based on Theorem 1: A choice of attributes $\mathbf{A}^+$ is a Nash equilibrium if and only if

$$\forall n \in N, \forall k \in K. \quad a^+_{nk} = a^*_{nk}(\mathbf{A}^+_{-nk}). \tag{15}$$

For the general case, we find that deriving equilibria based on this condition is intractable. For example, when considering a market with two disjoint paths, a single attribute, and a single ISP with arbitrary parameters on each path, the equilibrium must be found by solving a quartic equation, which impedes an analysis even for that simple network. Fortunately, we identify two types of markets that allow deriving closed-form equilibria and therefore analytic insights, while still capturing the fundamental characteristics of ISP competition, i.e., inter-path competition, intra-path cooperation, and ISP heterogeneity. Concretely, we separately analyze homogeneous markets (cf. §3.2) and heterogeneous markets with attribute-independent traffic-unit cost (cf. §3.3), both with disjoint paths.

3.2. Homogeneous Markets

By homogeneous markets, we refer to topologies of $Q := |R| > 0$ disjoint paths in competition, each of which accommodates the same number $I = |N|/Q$ of ISPs. All ISPs are identical and all attributes are identically valuable and costly, i.e., for all ISPs $n \in N$, it holds that $\rho_n = \rho$ and $\phi_{n0} = \phi_0$, and $\forall k \in K$, it holds that $\alpha_{nk} = \alpha_1$, $\phi_{nk} = \phi_1$, and $\gamma_{nk} = \gamma_1$. Moreover, the path-valuation functions for all paths are identical as well, i.e., $\forall r \in R.\ \alpha_{r0} = \alpha_0$. While artificial, such competition among completely equal goods (here: paths) and firms (here: ISPs) is a virtually universal assumption in competition models, as homogeneity allows isolating pure competition effects, i.e., effects that are not due to pre-existing differences between competitors [12, 14].

In our case, the homogeneity also permits to identify the Nash equilibria of the competition dynamics (§3.2.1), to investigate the convergence to these equilibria (§3.2.2), to compare these equilibria to social optima (§3.2.3), and to evaluate the effect of competition intensity (§3.2.4).

3.2.1. Equilibria

The symmetry of the homogeneous markets allows finding a competition equilibrium, which is guaranteed to exist, and can be efficiently computed for arbitrary topologies thanks to its closed-form representation:

Theorem 2. ***Nash Equilibrium in Homogeneous Markets.*** *The Nash equilibrium of a homogeneous market is given by an attribute sum a^+ such that $\sum_k a^+_{nk} = a^+\ \forall n \in N$, where $a^+ = \max(0,\ \hat{a}^+)$ with*

$$\hat{a}^+ = \frac{\sqrt{T_2^2 - 4T_1T_3} - T_2}{2T_1}. \tag{16}$$

Eq. (16) *uses the following abbreviations:*

$$T_1 = Q^2I^2\alpha_1^2 - \frac{d}{d\phi_1 + \gamma_1}(QI - 1)(Q - 1)I\alpha_1^2\phi_1, \tag{17}$$

$$T_2 = 2QI\alpha_1\left(1 + Q\alpha_0\right) - \frac{d}{d\phi_1 + \gamma_1}\cdot \tag{18}$$

$$\Big(\alpha_1\phi_1\left(QI - 1\right)\left(1 + (Q - 1)\alpha_0\right) + I\alpha_1\left(Q - 1\right)\left(\phi_1\left(1 + Q\alpha_0\right) + \alpha_1\left(\rho - \phi_0\right)\right)\Big), \text{ and}$$

$$T_3 = (1 + Q\alpha_0)^2 - \frac{d}{d\phi_1 + \gamma_1}(1 + (Q - 1)\alpha_0)\big(\phi_1(1 + Q\alpha_0) + \alpha_1(\rho - \phi_0)\big). \tag{19}$$

The proof of Theorem 2 is provided in Appendix B. Note that the equilibrium in Theorem 2 is only unique with respect to the attribute sum a^+ of any ISP and hence also with respect to path valuations, but not necessarily with respect to individual attribute values a_{nk}.

3.2.2. Stability

The Nash equilibrium from Theorem 2 is an interesting fixed point of the competitive dynamics in homogeneous markets. However, the equilibrium is only relevant if the distributed profit optimization by the ISPs converges to it. Hence, the equilibrium must be additionally investigated with respect to its *stability*, i.e., its attractive effect on the competition dynamics. To investigate this stability, we formally describe the competition by the following system of ordinary differential equations (ODEs):

$$\forall n \in N. \quad \dot{a}_n(t) = a_n^*\big(\mathbf{A}_{-n}(t)\big) - a_n(t) \tag{20}$$

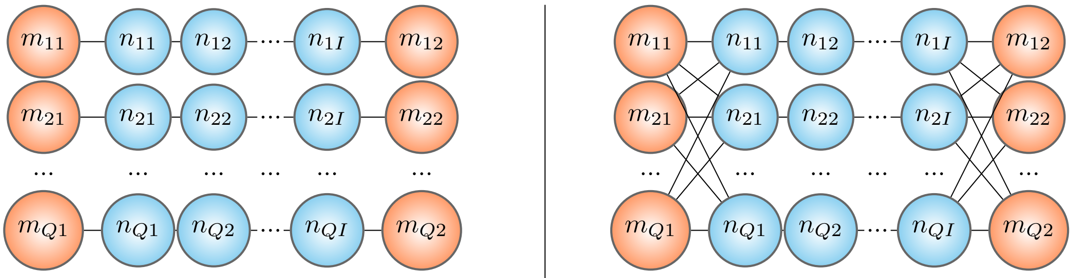

(a) Competition-free network $\mathcal{N}_1$.

(b) Competitive network $\mathcal{N}_2$.

Figure 1: Homogeneous markets with and without inter-path competition.

Intuitively, one ODE in this system describes the behavior of an ISP n which continuously adjusts its attribute-sum value a_n towards the optimal choice a_n^* given the contemporary attribute values of all other ISPs. Given this dynamic process, we can show the following property:

Theorem 3. ***Stability of Homogeneous Equilibrium.*** *The Nash equilibrium from Theorem 2 is an asymptotically stable equilibrium of the competition dynamics in Eq.* (20)*.*

The proof of Theorem 3 is provided in Appendix C.

3.2.3. Intra-Path Dynamics

The equilibrium formalization from Theorem 2 also applies to the case where an ISP pair is only connected by a single usable path. Such a single-path scenario represents a monopoly in economic terms. Crucially, the ISPs on the same single path are supposed to *cooperate* rather than compete, as the decisions by each ISP contribute to path attractiveness, which in turn benefits all ISPs. For that single-path case, we can make the following interesting observation about the cooperation among on-path ISPs:

Theorem 4. ***Suboptimality of Homogeneous Equilibrium.*** *On a single path with I identical ISPs, the equilibrium attribute sum a^+ is generally lower than the NBS attribute sum a°, i.e., $\forall I \in \mathbb{N}, I \geq 1.\ a^+ \leq a^\circ$.*

The proof of Theorem 4 is provided in Appendix D.

Intuitively, Theorem 4 states that the cooperation by on-path ISPs suffers from inefficiency caused by individual selfishness, similar to a prisoner's dilemma [33]. More precisely, the NBS attribute sum a°, which would optimize every ISP's profit if chosen universally, is not a rational choice for an individual ISP. In particular, if an ISP n chooses $a_n = a^\circ$, ISP n enables another ISP m to optimize its profit π_m by choosing a lower attribute sum $a_m < a^\circ$, and thus to free-ride on the path attractiveness created by ISP n. Because of this selfish deviation from the global optimum, the on-path ISPs converge to the equilibrium attribute prevalence a^+, which is generally lower than the NBS attribute sum a°, prevents the transit ISPs from reaping optimal profit, and also reduces the path attractiveness for the path-selecting ISP.

3.2.4. Competition Effects

After investigating intra-path cooperation in the preceding section, we now investigate the effect of *inter-path competition* on attribute prevalence. In particular, we are interested in the dependence of the equilibrium attribute a^+ on the number of available paths between an origin-destination pair.

To characterize this dependence, we compare the Nash equilibria in two homogeneous markets. First, we consider the competition-free network $\mathcal{N}_1$ in Fig. 1a, which is partitioned between Q origin-destination pairs $\{(m_{q1}, m_{q2})\}_{q=1,...,Q}$, each connected by a single path with I ISPs and obtaining the same demand limit d'. Effectively, this network is a set of homogeneous markets, each with only one available path and no competition. Second, we consider the competitive network $\mathcal{N}_2$ in Fig. 1b, where each of the Q origin-destination pairs can use all Q available paths.

By identifying the equilibrium attribute value $a^+(\mathcal{N})$ for each network $\mathcal{N}$, we find that competition has a consistently positive effect on attribute prevalence:

Theorem 5. ***Attribute Improvement under Homogeneous Competition.*** *The equilibrium attribute prevalence is never lower in the competitive network $\mathcal{N}_2$ than in the competition-free network $\mathcal{N}_1$, i.e., $a^+(\mathcal{N}_2) \geq a^+(\mathcal{N}_1)$ for all $Q \in \mathbb{N} \geq 1$.*

The proof of Theorem 5 is provided in Appendix E.

Surprisingly, the higher attribute values in the competitive equilibrium do not necessarily come at the cost of lower ISP profits. Instead, the profit of the ISPs may even increase under competition, which is important because the ISPs partially control whether they engage in competition at all, namely through path announcements. For example, an increase in equilibrium profits through competition happens if competition causes only a modest increase in attribute values:

Theorem 6. ***Profit Improvement under Homogeneous Competition.*** *The equilibrium profit $\pi^+(\mathcal{N}_2)$ in the competitive network preserves or exceeds the equilibrium profit $\pi^+(\mathcal{N}_1)$ of the uncompetitive network if $a^+(\mathcal{N}_2) \in [a^+(\mathcal{N}_1), a^\circ(\mathcal{N}_1)]$, i.e., the equilibrium attribute sum $a^+(\mathcal{N}_2)$ from the competitive network is between the equilibrium attribute sum $a^+(\mathcal{N}_1)$ of the uncompetitive network and the corresponding NBS attribute sum $a^\circ(\mathcal{N}_1)$.*

The proof of Theorem 6 is provided in Appendix F. In summary, we conclude that inter-path competition in homogeneous markets is always desirable from the perspective of path-selecting ISPs, and potentially desirable from the perspective of transit ISPs.

3.3. Heterogeneous Markets

The homogeneous markets discussed in the previous section can reflect competition dynamics among arbitrarily many paths. However, these market models cannot represent differences between paths that go beyond attribute values. In reality, on-path ISPs may differ in their importance for the path valuation, in their revenue per traffic unit, and in their attribute-specific costs. For example, consider a quality attribute k corresponding to the internal bandwidth of an ISP, and consider an ISP n with a relatively wide-spread internal network infrastructure. Thanks to its size, this ISP n likely provides relatively large segments of paths, making its ability to avoid congestion and thus its internal bandwidth more valuable than the internal bandwidth of smaller ISPs. However, ISP n also likely incurs a higher cost for bandwidth upgrades due to its more widespread infrastructure. The profit optimization of ISP n can thus be better represented by the model if ISP n is assigned a larger valuation weight α_{nk} and a larger fixed-cost weight γ_{nk} than other ISPs.

In this section, we therefore study competition among heterogeneous ISPs, i.e., every ISP n has arbitrary parameters α_{nk} $\forall k \in K$, ρ_n, ϕ_{n0}, γ_{nk} $\forall k \in K$, and γ_{n0}. To achieve tractability despite this additional complexity, we restrict our analysis to markets with at most two paths. Moreover, since traffic-unit cost is commonly considered negligible for ISPs [43], we consider especially the attribute-dependent part of this traffic-unit cost to be negligible, i.e., $\phi_{nk} = 0$ $\forall n \in N, k \in K$.

3.3.1. Intra-Path Dynamics

To characterize the attribute-choice dynamics among ISPs in a heterogeneous market, we first consider a single path in isolation, i.e., a monopoly scenario. As before, ISPs on a path collectively determine the attractiveness of the path, but optimize only their individual profit. This selfishness may lead to a sub-optimal global outcome, both regarding ISP profits and path valuations. To quantify this shortfall, we first identify the Nash bargaining solution (NBS) for the attribute choices, i.e., the attribute values that all on-path ISPs would agree on if they collectively negotiated and if they were bound by the result of the negotiation. This Nash bargaining solution represents the global optimum with respect to the ISP profits.

Theorem 7. ***Profit Optimum on Heterogeneous Paths.*** *On a path r with heterogeneous ISPs, the attributes $\mathbf{A}^\circ$ form a Nash bargaining solution if and only if these attributes optimize the product of all ISP profits while*

1. *leading to the NBS path valuation v_r°,*

2. *containing non-zero attribute values only for attributes (n,k) with maximal ratio $\frac{\alpha_{nk}}{\gamma_{nk}}$,*
3. *and containing zero attribute values for all other attributes.*

More formally, the conditions on $\mathbf{A}^\circ$ can be stated as follows:

$$\mathbf{A}^\circ = \arg\max_{\mathbf{A}\in\mathbb{R}_{\geq 0}} \Pi_{n\in r}\pi_n(\mathbf{A}) \tag{21}$$

$$\text{subject to} \qquad v_r(\mathbf{A}) = v_r^\circ \qquad \forall(n,k)\in K_r^\circ.\ a_{nk}\geq 0 \qquad \forall(n,k)\notin K_r^\circ.\ a_{nk}=0$$

$$\text{where} \qquad v_r^\circ = \max\left(\alpha_{r0},\ \max_{(n',k')\in r\times K}\sqrt{\frac{\alpha_{n'k'}}{\gamma_{n'k'}}}\sqrt{d\sum_{n\in r}(\rho_n-\phi_{n0})}-1\right),\ \text{and} \tag{22}$$

$$K_r^\circ = \left\{(n,k)\mid (n,k)=\arg\max_{(n',k')\in r\times K}\frac{\alpha_{n'k'}}{\gamma_{n'k'}}\right\}. \tag{23}$$

The proof of Theorem 7 is provided in Appendix G. Notably, the social optimum $\mathbf{A}^\circ$ can be computed with low computational cost: Even if path r includes a high number of ISPs with numerous attributes, the number $|K_r^\circ|$ of non-zero attributes is generally much lower, and typically even only 1. This low number of relevant variables lowers the cost of the convex optimization problem in Eq. (21).

To optimize aggregate profit, the on-path ISPs should thus only upgrade the attribute(s) with maximal 'return' α_{nk}/γ_{nk} while minimizing the prevalence of all other attributes. This return ratio α_{nk}/γ_{nk} yields the valuation for attribute k of ISP n, compared to the cost that ISP n incurs for adopting that attribute. The return ratio also correlates with the optimal path valuation (cf. Eq. (22)).

However, aggregate profit is not the objective of selfish ISPs when determining attribute values. Instead, selfish ISPs optimize their individual profit, and eventually arrive at the following equilibrium by their non-aligned optimization behavior:

Theorem 8. ***Nash Equilibrium on Heterogeneous Paths.*** *On a path r with heterogeneous ISPs, the attributes $\mathbf{A}^+$ form a Nash equilibrium if and only if these attributes*

1. *lead to the equilibrium path valuation v_r^+,*
2. *contain non-zero attribute values only for attributes (n,k) with maximal ratio $\frac{\alpha_{nk}(\rho_n-\phi_{n0})}{\gamma_{nk}}$,*
3. *and contain zero attribute values for all other attributes.*

More formally, the conditions on $\mathbf{A}^+$ can be stated as follows:

$$v_r(\mathbf{A}^+) = v_r^+ \qquad \forall(n,k)\in K_r^+.\ a_{nk}^+\geq 0 \qquad \forall(n,k)\notin K_r^+.\ a_{nk}^+=0 \tag{24}$$

$$\text{where} \qquad v_r^+ = \max\left(\alpha_{r0},\ \max_{(n',k')\in r\times K}\sqrt{\frac{\alpha_{n'k'}}{\gamma_{n'k'}}}\sqrt{d(\rho_{n'}-\phi_{n'0})}-1\right),\ \text{and} \tag{25}$$

$$K_r^+ = \left\{(n,k)\mid (n,k)=\arg\max_{(n',k')\in r\times K}\frac{\alpha_{n'k'}(\rho_{n'}-\phi_{n'0})}{\gamma_{n'k'}}\right\}. \tag{26}$$

The proof of Theorem 8 is provided in Appendix H.

Interestingly, the equilibrium in Theorem 8 is similar to the Nash-bargaining solution in Theorem 7, but contains one crucial difference: The return ratio associated with cultivated attributes includes the net revenue per unit of traffic $\rho_n-\phi_{n0}$ of ISP n (Eq. (23) vs. Eq. (26)). This inclusion reflects that each ISP n optimizes its individual profit rather than the aggregate profit: When optimizing an attribute (n,k) for individual profit, an ISP n only considers its individual net revenue per traffic unit, not the *aggregate* net revenue per traffic unit, which would be relevant for aggregate profit.

This difference, albeit subtle, generally leads to different attribute choices in equilibrium than postulated by the Nash-bargaining solution, meaning that the transit ISPs generate sub-optimal profits. Unfortunately, also the path-selecting ISP suffers from this selfishness, as the individual-profit optimization leads to less valuable paths:

Theorem 9. ***Suboptimality of Heterogeneous Equilibrium.*** *On a path r with heterogeneous ISPs, the equilibrium path valuation $v_r(\mathbf{A}^+)$ never exceeds the NBS path valuation $v_r(\mathbf{A}^\circ)$, i.e., $v_r(\mathbf{A}^+) \leq v_r(\mathbf{A}^\circ)$.*

The proof of Theorem 9 is provided in Appendix I.

3.3.2. Two-Path Equilibria

In the preceding section, the social optimum and the Nash equilibrium are characterized for a single-path scenario, which is exclusively informed by (failing) intra-path cooperation among selfish ISPs. Since we are also interested in the effect of inter-path competition, we now consider heterogeneous markets in which the path-selecting ISP can select between two disjoint paths. For these networks, the single-path equilibrium in Theorem 8 can be adjusted to the following equilibrium, which is guaranteed to exist and efficiently computable for the reasons noted in §3.3.1:

Theorem 10. ***Nash Equilibrium in Heterogeneous Markets.*** *In a two-path heterogeneous market, the attribute values $\mathbf{A}^+$ form a Nash equilibrium if and only if the attribute values $\mathbf{A}^+$ satisfy the conditions from Theorem 8, but with modified equilibrium path valuation v_r^+:*

$$v_r^+ = \max\left(\alpha_{r0},\ \hat{v}_r^*\left(\max\left(\alpha_{\bar{r}0}, \hat{v}_{\bar{r}}^+\right)\right)\right) \tag{27}$$

$$\text{where}\quad \bar{r} \text{ is the alternative path to } r, \quad \hat{v}_r^*(v_{\bar{r}}) = \psi_r\sqrt{d}\sqrt{1+v_{\bar{r}}} - (1+v_{\bar{r}}),$$

$$\hat{v}_r^+ = \frac{\psi_r^3\psi_{\bar{r}}}{(\psi_r^2+\psi_{\bar{r}}^2)^2}\left(\sqrt{d\left(\psi_r^2+\psi_{\bar{r}}^2\right)+\frac{1}{4}\psi_r^2\psi_{\bar{r}}^2d^2}+\frac{d}{2}\psi_r\psi_{\bar{r}}\right) - \frac{\psi_{\bar{r}}^2}{\psi_r^2+\psi_{\bar{r}}^2}, \text{ and} \tag{28}$$

$$\psi_r = \max_{\substack{n\in r\\ k\in K}} \sqrt{\frac{\alpha_{nk}(\rho_n - \phi_{n0})}{\gamma_{nk}}}. \tag{29}$$

The proof of Theorem 10 is provided in Appendix J. We note that ψ_r from Eq. (29) is the square root of the maximum individual return ratio discussed in the previous section, albeit only among the attributes of path r. In the following, we refer to ψ_r as the *characteristic ratio* of path r.

Similar to §3.2, we are again interested in the stability of the equilibrium w.r.t. the process:

$$\forall n \in N, k \in K. \quad \dot{a}_{nk}(t) = a_{nk}^*(\mathbf{A}_{-nk}(t)) - a_{nk}(t). \tag{30}$$

However, stability analysis in the case of heterogeneous two-path networks is complicated by the fact that the equilibrium from Theorem 10 is only unique in the path valuations $\{v_r\}_{r\in R}$, but not necessarily unique in the attribute choices $\mathbf{A}$ by the ISPs. Therefore, if the equilibrium is not unique in $\mathbf{A}$, no single equilibrium $\mathbf{A}^+$ is asymptotically stable in a narrow sense, as the process in Eq. (30) does not converge to $\mathbf{A}^+$ from $\mathbf{A}(t)$ in case $\mathbf{A}(t)$ already represents a different equilibrium.

Therefore, we focus on the stability of unique equilibria:

Theorem 11. ***Stability of Heterogeneous Equilibrium.*** *The Nash equilibrium $\mathbf{A}^+$ from Theorem 10 is an asymptotically stable equilibrium of the competition dynamics in Eq.* (30) *if the equilibrium $\mathbf{A}^+$ is unique, i.e., if there is only one attribute on every path which has potentially non-zero prevalence ($|K_r^+| = |K_{\bar{r}}^+| = 1$).*

The proof of Theorem 11 is provided in Appendix K.

3.3.3. Competition Effects

Based on the equilibria for single-path and two-path markets, we now investigate the effect of inter-path competition in heterogeneous markets. For this investigation, we use a similar approach as in §3.2.4: We contrast a competition-free network $\mathcal{N}_3$, which consists of two paths r and $\overline{r}$, each connecting one origin-destination pair, with a competitive network $\mathcal{N}_4$, where both origin-destination pairs are connected by both paths. The origin-destination pair connected by path r in the competition-free network $\mathcal{N}_3$ has demand limit d_r; hence, the total demand limit $d = d_r + d_{\overline{r}}$ is distributed over both paths in the competitive network $\mathcal{N}_4$. The networks $\mathcal{N}_3$ and $\mathcal{N}_4$ thus differ in the same manner as the networks $\mathcal{N}_1$ and $\mathcal{N}_2$ from Fig. 1, except that the different paths may have different length in ISPs, each ISP may have different parameters, and each origin-destination pair may have a different demand limit. When contrasting these two networks, we gain the following insight:

Theorem 12. ***Attribute Improvement under Heterogeneous Competition.*** *For any competition-free network $\mathcal{N}_3$ and the corresponding competitive network $\mathcal{N}_4$, a demand limit d exists such that the competitive network $\mathcal{N}_4$ has a higher equilibrium valuation than the competition-free network $\mathcal{N}_3$ independent of the demand distributions $(d_r, d_{\overline{r}})$, i.e.,*

$$\exists d \ s.t. \ \forall d_r, d_{\overline{r}} \ with \ d_r + d_{\overline{r}} = d. \ V^+(\mathcal{N}_4) \geq V^+(\mathcal{N}_3) \tag{31}$$

The proof of Theorem 12 is provided in Appendix L.

In simplified terms, inter-path competition thus affects the attribute values and the path valuations positively for high-enough demand, given the remaining network parameters. This condition on demand, however, raises the question whether competition reduces the network valuation in some circumstances. Indeed, we find that such a counter-intuitive effect can arise at every demand level if the remaining network parameters are unfavorable:

Theorem 13. ***Attribute Decline under Heterogeneous Competition.*** *For every demand distribution $(d_r, d_{\overline{r}})$, there exist characteristic ratios $(\psi_r, \psi_{\overline{r}})$ and path base valuations $(\alpha_{r0}, \alpha_{\overline{r}0})$ such that the competitive network $\mathcal{N}_4$ has a lower equilibrium valuation than the competition-free network $\mathcal{N}_3$, i.e.,*

$$\forall d_r, d_{\overline{r}}. \ \exists \psi_r, \psi_{\overline{r}}, \alpha_{r0}, \alpha_{\overline{r}0} \ s.t. \ V^+(\mathcal{N}_4) < V^+(\mathcal{N}_3). \tag{32}$$

The proof of Theorem 13 is provided in Appendix M.

To understand this effect intuitively, we note that an ISP n optimizes its profit by balancing the *marginal revenue* and the *marginal cost* with respect to attribute prevalence, i.e., adjusts attribute prevalence as long as the adjustment generates more revenue than cost. In the competition-free scenario of $\mathcal{N}_3$, the marginal revenue and cost of an ISP n with respect to attribute (n, k) are:

$$\frac{\partial \mathcal{R}_n}{\partial a_{nk}} = \frac{d_r \alpha_{nk}}{(1+v_r)^2} \cdot \rho_n \qquad \frac{\partial}{\partial a_{nk}} (\Phi_n + \Gamma_n) = \frac{d_r \alpha_{nk}}{(1+v_r)^2} \cdot \phi_{n0} + \gamma_{nk} \tag{33}$$

In contrast, the corresponding terms for the competitive scenario in network $\mathcal{N}_4$ are as follows:

$$\frac{\partial \mathcal{R}_n}{\partial a_{nk}} = \frac{d\alpha_{nk} \cdot (1+v_{\overline{r}})}{(1+v_r+v_{\overline{r}})^2} \cdot \rho_n \qquad \frac{\partial}{\partial a_{nk}} (\Phi_n + \Gamma_n) = \frac{d\alpha_{nk} \cdot (1+v_{\overline{r}})}{(1+v_r+v_{\overline{r}})^2} \cdot \phi_{n0} + \gamma_{nk} \tag{34}$$

On the one hand, competition has a positive effect on marginal revenue $\partial\mathcal{R}_n/\partial a_{nk}$ by increasing the total amount of attractable demand from d_r to $d > d_r$. On the other hand, the new competition embodied by the alternative-path valuation $v_{\overline{r}}$ has a negative effect on marginal revenue. The negative effect predominates if the alternative-path valuation $v_{\overline{r}}$ is relatively large and unresponsive to competition, as the proof of Theorem 13 demonstrates. If marginal revenue in fact decreases, marginal cost decreases less strongly as $\rho_n \geq \phi_{n0}$. Given negative marginal profit, the profit of ISP n is thus optimized by a lower attribute prevalence a_{nk}, which translates into decreasing path value.

4. A Model Instance Based on Real-World Data

In this section, we demonstrate how to instantiate our competition model from §2 to investigate a large-scale network containing multiple intertwined markets. To that end, we construct a topology approximating the Internet core and a corresponding traffic matrix in §4.1. Furthermore, we consider two ISP attributes in the competitive dynamics, namely *internal bandwidth* and *clean-energy share*, and estimate appropriate model parameters in §4.2 and §4.3, respectively. Attribute-independent parameters are estimated in §4.4.

Importantly, we note that estimating highly realistic parameters for the model goes beyond the scope of this paper, as the scarcity of publicly available data and the complexity of real-world business practices considerably complicates this estimation. Therefore, the goal of the following parameter estimation is to place the parameters in the right order of magnitude, especially in relation to each other, rather than to determine each parameter highly realistically. Interestingly, our sensitivity analysis in §5 suggests that such an approximate estimation might be sufficient to yield useful predictions.

4.1. Network Topology and Demand

To investigate the effects of competition in practically interesting, large-scale settings while keeping the complexity of the simulation manageable, we extract a network topology that roughly approximates the Internet core from a public dataset. In particular, we rely on a CAIDA dataset containing 12 300 autonomous systems (ASes), their economic relationships, and the geolocation of their interconnections (i.e., inter-domain interfaces) [7]. From this dataset, we extract the topology of the 2000 most interconnected ASes by iteratively removing the lowest-degree ASes.

In this reduced topology, we aim at finding the 5 shortest paths between every origin-destination pair of ASes. For scalability, we can only consider AS paths with at most 4 AS hops, which is not a strong limitation: The paths in our topology only represent the core-traversing segments of whole Internet paths, which have an average length of around 5 hops (and decreasing) [18]. Moreover, for both scalability and practical relevance, we only consider paths that are Gao-Rexford-compliant [16], i.e., are compatible with the economic self-interest of ASes regarding monetization of traffic. The restriction to Gao-Rexford-compliant paths also ensures that the intermediate ASes in a discovered path are in fact transit providers rather than content providers or eyeball ISPs. With these constraints, we can identify 5 paths for $\sim 52.4\%$ of AS pairs in the topology.

While only a subset of all AS pairs, these pairs of closely located ASes are disproportionately relevant for the competition dynamics, as they account for a substantial share of traffic given the gravity-like nature of Internet traffic [34]. Gravity models make the simplifying assumption that the traffic demand $d_{(n_1,n_2)}$ between two ISPs n_1 and n_2 is proportional to the product of the 'masses' $m_1 \cdot m_2$ of the two ISPs divided by the squared distance r_{12}^2 between the ISPs:

$$d_{(n_1,n_2)} \propto G_{12} = \frac{m_1 \cdot m_2}{r_{12}^2}. \tag{35}$$

In order to synthesize a traffic matrix for our purpose, we concretize this gravity model as follows. First, we calculate the mass m_n of an AS n as the number of distinct IPs in all prefixes owned by AS n and by the ASes in the customer cone of AS n. This information is available via the datasets 'Routeviews Prefix-to-AS Mapping' [8] and 'AS Relationships' [9], both from CAIDA. Second, we determine the distance r_{12} for each AS pair (n_1, n_2) as the average number of hops in the 5 paths connecting the AS pair. Third, we calculate the gravity G_{12} according to Eq. (35) for every AS pair (n_1, n_2). Finally, we allocate the total Internet traffic volume of 170 Tbps [27] to the AS pairs (n_1, n_2) according to the relative size of G_{12}.

4.2. Attribute 1: Internal Bandwidth

To instantiate the model, we define the ISP attributes K that are affected by the competitive dynamics, and the corresponding model parameters. As an intuitive example of desirable ISP attributes, we consider the *internal bandwidth* of an ISP (in Gbps) the first such attribute ($k = 1$). If the ISPs along a path have a large bandwidth capacity, these ISPs are likely able to absorb sudden traffic surges, tolerate equipment failures, handle large traffic flows, and in general deliver a high quality of service; hence, the internal bandwidth of on-path ISPs correlates with the attractiveness of the given path.

4.2.1. Valuation

This valuation by path-selecting ISPs is quantified by the valuation function v_{r1}, quantifying the valuation of a bit traversing path r given the internal bandwidth of on-path ISPs. This valuation function v_{r1} is characterized by the parameters α_{rn1}, giving the valuation of a bit traversing ISP n on path r for a unit of the internal bandwidth of ISP n. For this quantification, we rely on two empirical findings. First, the average US consumer transmits 536.3 GB of data per month [31], and is willing to pay 94 USD per month for a 1Gbps connection [22]. Hence, we arrive at a monthly willingness-to-pay of around $w = 0.17$ USD per GB at the quality of a 1Gbps connection. With this willingness-to-pay w, we determine the bandwidth valuation parameters α_{rn1}, namely by defining $\alpha_{rn1} = w/(|r| \cdot m_n)$, where $|r|$ is the number of ASes on path r (averaging the internal bandwidth across on-path ISPs) and m_n is the number of IPs in the customer cone of AS n (correcting for the number of end-points sharing the bandwidth). Multiplied with the internal bandwidth a_{n1}, these valuation parameters thus approximate the valuation per bit traversing ISP n given the internal bandwidth of ISP n. Notably, this choice of α_{rn1} assumes that the complete willingness-to-pay w of end-hosts can be exploited by transit ISPs; in reality, a substantial share of this willingness-to-pay is already captured by eyeball ISPs. However, our sensitivity analysis (cf. §5.1) indicates that our results hold even under lower α_{rn1}, i.e., for a lower willingness-to-pay relevant to transit providers.

Furthermore, the bandwidth valuation function v_{r1} is also characterized by the base valuation α'_{r0} of path r. However, since a path only has value in terms of bandwidth if the on-path ASes have non-zero internal bandwidth, we choose $\alpha'_{r0} = 0$.

4.2.2. Cost

Apart from increasing valuation by path-selecting ISPs, providing bandwidth also has a cost. However, it is difficult to quantify the cost of providing a Gbps of internal bandwidth, as this cost heavily depends on the way of provision (leasing or physically installing new capacity), on the necessary installation procedures (e.g., length of cables to be newly laid), on the location where capacity should be added, and on numerous other aspects. Hence, we rely on the simple insight that the cost of providing a Gbps of connectivity is likely lower than the corresponding willingness-to-pay by consumers (94 USD per Gbps per month [22]), as ISPs would go out of business otherwise. Hence, we randomly vary the cost parameter γ_{n1} between 0 and 94 USD per Gbps per month in our simulations, for all $n \in N$. Importantly, the provision of bandwidth only affects the demand-independent cost Γ_n of an ISP n, as providing a certain bandwidth capacity causes the same cost independent of the actually experienced demand. Hence, we can also define the demand-dependent cost parameter for the bandwidth attribute $k = 1$: $\phi_{n1} = 0$ for all $n \in N$.

4.2.3. Attribute Bounds.

Using internal bandwidth as one of multiple attributes leads to an implausible model prediction in the case where all ASes on a path r have zero internal bandwidth ($a_{n1} = 0\ \forall n \in r$), but some non-zero values for other attributes. In that case, the valuation function v_r might still assign some non-zero valuation and some demand to path r, although the zero-bandwidth path r is clearly worthless. To avoid this implausible case of the model, we place a lower bound on the bandwidth attribute $a_{n1}\ \forall n \in N$. This lower bound is given by 10% of the demand experienced by AS n if the demand of every origin-destination pair was equally distributed among the available 5 paths:

$$\forall n \in N.\ a_{n1} \geq \frac{0.1}{5} \cdot \sum_{\substack{r \in R.\ n \in r \\ r \in R(n_1,n_2)}} d_{(n_1,n_2)} \tag{36}$$

4.3. Attribute 2: Clean-Energy Share

Path-selection preferences are not exclusively related to transmission performance (such as internal bandwidth of on-path ISPs), but may also reflect ESG considerations [11, 23]. For example, in carbon-intelligent routing [39, 49], path selection takes into account the carbon emission that results from data transmission. More precisely, the path-specific *transmission carbon intensity*, i.e., the volume of carbon emission per bit of transmitted data on a given path, affects path selection. To investigate the effect of competition on this carbon intensity, we choose the *share of clean energy* used by an ISP (in percent) as the second attribute ($k = 2$) for our simulations, i.e., $a_{n2} \in [0, 1]\ \forall n \in N$.

4.3.1. Transmission Carbon Intensity

The clean-energy share attributes of on-path ASes determine the carbon intensity of a path as follows. First, any AS-level path r must be transformed into a router-level path s_r, which is possible by means of the CAIDA ITDK dataset [10]. For simplicity, we assume that the intra-AS router-level path s_{rn} in AS n is the shortest router-level path between the two AS interconnections derived from the AS-level path r. For any intra-AS path s_{rn}, we determine the energy intensity e_{rn}, i.e., the amount of consumed electricity per bit transmitted on path s_{rn}. This energy intensity e_{rn} can be calculated from the number of routers and the covered distance of path s_{rn}, given by the CAIDA ITDK dataset, and the energy-intensity values for various devices, as reported by Heddeghem at al. [17]. Then, we calculate the maximum transmission carbon intensity $c_{rn,\max}$ of any intra-AS path s_{rn} by multiplying the corresponding energy intensity e_{rn} with the the energy carbon intensity $c_{\max}$ of the most carbon-intensive electricity, namely 875 gCO_2/kWh for coal-generated electricity [19]. This maximum transmission carbon intensity $c_{rn,\max}$ thus quantifies the carbon emission associated with the transmission of a bit across path s_{rn} if AS n used maximally carbon-intensive electricity. Finally, we derive the *actual* transmission carbon intensity c_{rn} of any intra-AS path s_{rn} as the product of the maximum transmission carbon intensity $c_{rn,\max}$ and the dirty-energy share of ISP n, i.e., $1 - a_{n2}$. The carbon intensity c_r of a path r is the sum of carbon-intensity values c_{rn} of the constituting intra-AS paths $s_{rn}\ \forall n \in r$:

$$c_r(\mathbf{A}) = \sum_{n\in r} c_{rn}(\mathbf{A}) = \sum_{n\in r} c_{rn,\max} \cdot (1 - a_{n2}) = \sum_{n\in r} e_{rn} \cdot c_{\max} \cdot (1 - a_{n2}). \tag{37}$$

4.3.2. Valuation

This carbon-intensity calculation also informs the valuation v_{r2}, which quantifies the valuation of path r exclusively with respect to carbon emissions. In fact, we understand v_{r2} as an affine function of the *negative* carbon intensity of path r:

$$v_{r2}(\mathbf{A}) = \sum_{n\in r} \alpha_{rn2} a_{n2} + \alpha''_{r0} = -\sum_{n\in r} p_{CO_2} c_{rn}(\mathbf{A}) + q_r = \sum_{n\in r} (p_{CO_2} \cdot c_{rn,\max} \cdot a_{n2} - p_{CO_2} \cdot c_{rn,\max}) + q_r. \tag{38}$$

where p_{CO_2} is the cost of emitted CO_2, chosen as 90 USD per ton according to the EU emission-trading scheme [15], and q_r is a constant that ensures the non-negativity and comparability of the valuation (see below). From Eq. (38), we can determine the valuation parameters α_{rn2}, describing the valuation of ISP n's clean-energy share on path r, as $p_{CO_2} \cdot c_{rn,\max}$. The base valuation α''_{r0} is determined based on two considerations. First, the valuation function v_{r2} must be consistently non-negative. Second, the valuation function must allow a meaningful comparison between paths $R(n_1, n_2)$ connecting the same AS pair (n_1, n_2): For example, if all ISPs use zero clean energy, a path with higher energy intensity should still be valued less than a path with lower energy intensity. Conversely, if all ISPs use perfectly clean energy, all paths should be valued identically. To achieve these properties, we determine α''_{r0} as follows:

$$\alpha''_{r0} = -\sum_{n\in r} p_{CO_2} \cdot c_{rn,\max} + q_r = -\sum_{n\in r} p_{CO_2} \cdot c_{rn,\max} + \max_{\substack{r'\in R(n_1,n_2)\\ r\in R(n_1,n_2)}} \sum_{n'\in r'} p_{CO_2} \cdot c_{r'n',\max}. \tag{39}$$

With such determined v_{r2}, we can formalize the complete path-valuation v_r as the sum of the attribute-specific valuation functions v_{r1} and v_{r2}. Since $\alpha_{rnk} a_{nk}$ yields a valuation per bit for both attributes $k \in \{1, 2\}$, the attribute-specific valuation functions are compatible.

4.3.3. Cost

To estimate the costs associated with the clean-energy share of an ISP n, we rely on the analysis of the levelized cost of energy (LCOE) of different electricity-generation technologies, performed by Lazard [21].

According to the Lazard analysis, electricity from low-carbon sources (solar, wind, nuclear) is on average $g = 3.375$ USD per MWh more expensive than electricity from high-carbon sources (coal, gas). This cost penalty, together with the average energy intensity of all intra-AS paths in AS n, yields the parameter ϕ_{n2} relevant for demand-dependent cost:

$$\phi_{n2} = g \cdot \frac{1}{|R(n)|} \cdot \sum_{r \in R(n)} e_{rn}, \tag{40}$$

where $R(n) = \{r \in R \mid n \in R\}$. Multiplied with the clean-energy share attribute a_{n2}, the parameter ϕ_{n2} yields the extra cost per transported bit that AS n incurs by using clean energy.

Regarding demand-independent cost, we note that the idle-power requirement of network devices plays an important role, as this requirement generates electricity bills even in absence of demand. The idle-power consumption u_n of a complete AS n can be estimated from the number of devices in AS n [10], the power consumption of network devices [17], and an average idle-power requirement of 85% [20]. This idle-power consumption u_n, together with the extra cost g for clean energy, determines the parameter γ_{n2} relevant for demand-independent cost: $\gamma_{n2} = g \cdot u_n$.

4.4. Attribute-Independent Parameters

In addition to the attribute-specific parameters in Sections 4.2 and 4.3, the attribute-independent parameters ρ_n, ϕ_{n0} and γ_{n0} also appear in our model.

The parameter ρ_n quantifies the revenue per transported bit of AS n. To estimate this parameter, we use a top-down approach: We divide the global annual revenue of wholesale Internet backbone providers (45.2 billion USD in 2019 [45]) by the amount of global annual Internet traffic (433 exabyte in 2019 [27]), and arrive at an average revenue of $\rho = 0.104$ USD per GB. For simplicity, we use this ρ as revenue parameter ρ_n for every ISP $n \in N$.

The parameter ϕ_{n0} describes the marginal cost of AS n per transported bit, excluding extra marginal cost due to clean-energy usage (cf. §4.3). As this marginal cost is commonly understood to be 'essentially zero' [43], we determine $\phi_{n0} = 0 \; \forall n \in N$.

Conversely, the attribute-independent fixed cost γ_{n0} of AS n can be quite substantial. However, since we are mainly interested in the attribute-optimization behavior of ASes under competition, and γ_{n0} does not affect this optimization behavior, we abstain from estimating γ_{n0}, i.e., use $\gamma_{n0} = 0$ in our simulations. As a result, the absolute value of the profit function π_n is not meaningful, which we take into account for the result discussion in §5.2.

5. Simulation

Section 3 theoretically illustrates the diverse results of quality competition among ISPs. These results include both positive and negative effects of competition on attribute prevalence and profits, depending on the concrete topologies of competing paths in the considered markets. In this section, we investigate which types of effect are observable if competition is introduced in a large-scale topology where transit ASes (ASes, corresponding to ISPs) simultaneously compete in multiple markets. To that end, we run simulation experiments described in §5.1 for the instance of the competition model constructed in §4, and discuss the results in §5.2.

5.1. Experiments

Since the parameters estimated in §4 are afflicted with considerable uncertainty, we conduct our experiments without being overly reliant on the exact estimated parameter values. More precisely, we generate 20 different sets of model parameters by randomly modifying each model parameter y such that it lies between 0 and $2y$ in virtually all cases.

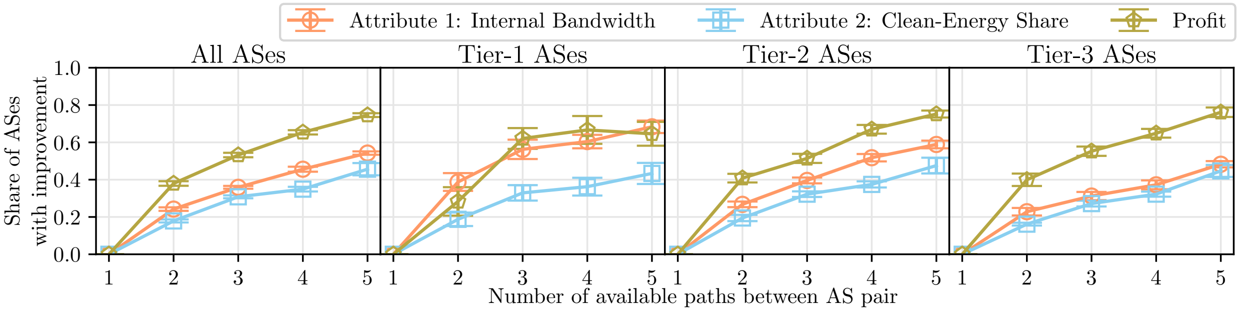

Figure 2: their profit compared to a single-path monopoly, independent of their topology level.

We achieve this modification by randomly sampling each model parameter y' from the following restricted normal distribution, based on the corresponding estimated parameter y:

$$y' \sim \max\left(\mathcal{N}\left(y, \frac{y^2}{9}\right), 0\right). \tag{41}$$

For each random sample of parameters, we investigate the effect of increasing *intensity of competition* on the attribute-value choice of transit ASes. In our experiments, the intensity of competition corresponds to the number of usable paths between any AS pair. This number of usable paths is varied between 1 and 5 across experiments, where the case of 1 path corresponds to a monopoly scenario. In each simulation experiment, we thus simulate the competitive dynamics given a set of randomly sampled model parameters and given a certain intensity of competition.

Each simulation experiment amounts to computing round-robin better-response dynamics [6], where all ASes consecutively adjust their attribute values in the direction which increases their profit. This discrete process is an approximation of the continuous ODE process in Eq. (20). Moreover, the process can be understood as reflecting bounded rationality [37], as we assume that ASes can only identify profit-improving rather than profit-optimal attribute values. The simulation is terminated once the competitive dynamics have converged upon an equilibrium, i.e., each round only causes negligible change in the attribute values $\mathbf{A}$. The attribute values $\mathbf{A}^+$ in the equilibrium then represent the results of the experiment.

5.2. Results

The results of the experiments described in §5.1 are visualized in Figs. 2–4. The error bar of any data point in these figures illustrates the variance of the respective aggregate result across the 10 random parameter samples. Interestingly, the variance of the aggregate results is quite limited, although the variance in individual parameters is considerable. This observation indicates that our results are not highly sensitive to the model parameters from §4, and suggests that an approximate estimation of model parameters might be sufficient to yield useful predictions.

The central question in our analysis is: How does the intensity of competition affect the attribute values and the ISP profits? Our theoretical investigation in §3 indicates that the competitive dynamics can both increase and decrease these indicators compared to a monopoly scenario, depending on network properties. Hence, we investigate which type of effect is predominant for the large-scale network from §4.

In this analysis, we distinguish three groups of ASes that differ in their topology rank, namely tier-1 ASes (ASes that have no provider), tier-2 ASes (ASes that have only tier-1 providers), and tier-3 ASes (ASes that have only tier-1 and tier-2 providers). Since these AS groups differ in their market power, the effect on competition on attributes and profits for these groups may be different.

5.2.1. Attribute Prevalence

Regarding competition effects on attribute prevalence, we gain the following insights:

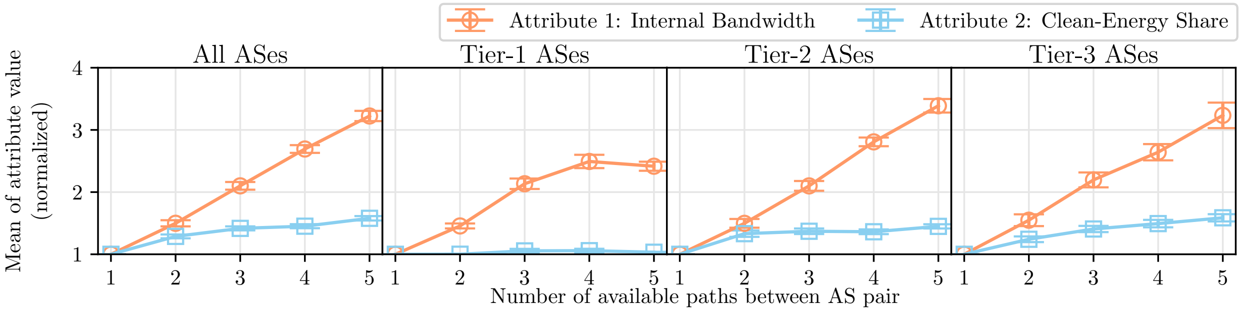

Figure 3: compared to a monopoly situation. However, attributes may be affected differently by competition due to differences in valuation and cost.

Effects on transit ASes. Fig. 2 confirms that an increasing number of options for the path-selecting ASes intensifies competition between transit ASes, which then improve their attribute values in response. In particular, about half of ASes improve both their attribute values given 5 available paths, whereas only 20% improve their attributes in a duopoly scenario (compared to a monopoly scenario). Note that some ASes may decrease their attribute values under competition for the counter-intuitive reasons explained in §3.3.3.

Moreover, the general level of attribute prevalence is raised by competition, which is implied by the increasing global average of attribute values in Fig. 3. The internal-bandwidth attribute is more strongly affected by this average gain than the clean-energy–share attribute, because (i) the bandwidth attribute is not upper-bounded, and (ii) the bandwidth attribute has zero demand-dependent cost.

Effects on selecting ASes. The average improvement in attributes translates into a more attractive offer for path-selecting ASes. More precisely, the most attractive path between each AS pair tends to become more attractive as competition increases: Fig. 4 shows that 75% of AS pairs obtain access to a more valuable path if two paths instead of a single path are available (increasing maximum valuation), irrespective of the tier of the path-selecting AS. Notably, we would expect that around 50% of AS pairs obtain a second path of higher quality in the absence of dynamic competition effects. Hence, the increasing maximum valuation is a combined effect of multi-path availability and competition. Moreover, in absence of competition effects, a second path can only decrease, but not raise the minimum valuation across available paths. However, we observe that for 40% of AS pairs, both paths in a two-path scenario are more highly valued than the single path in a monopoly scenario, which suggests that competition raises the value of the previously monopolistic path. However, as the number of available paths increases, the tendency of additional paths to decrease the minimum quality becomes more visible. Finally, these offer improvements materialize for all tiers of path-selecting ASes.

Differences in market power. Intriguingly, the higher market power of tier-1 ASes is not visible in Fig. 2, as tier-1 ASes are equally likely as lower-tier ASes to improve their attributes. However, the market power of tier-1 ASes becomes apparent in Fig. 3, which indicates that tier-1 ASes in competition improve their attributes to a lower extent than ASes on lower tiers.

5.2.2. Profits

Regarding competition effects on profits, we make the following observations:

Effects on transit ASes. Increasing path diversity and competition lead to increasing profits for a substantial share of ASes (cf. Fig. 2). At 5 available paths per AS pair, 75% of ASes increase their profits. This insight is surprising, given that competition is traditionally expected to increase consumer welfare and to reduce producer profits. In an ISP market, however, profits may increase because competition is modulated by path diversity. Such path diversity not only allows selecting ASes to select more different paths, but also allows transit ASes to attract and monetize traffic from more selecting ASes, increasing profit. Importantly, such an increase in attracted demand for an ISP does not necessarily reduce the attracted demand of another ISP, as the volume of total demand is elastic in our model.

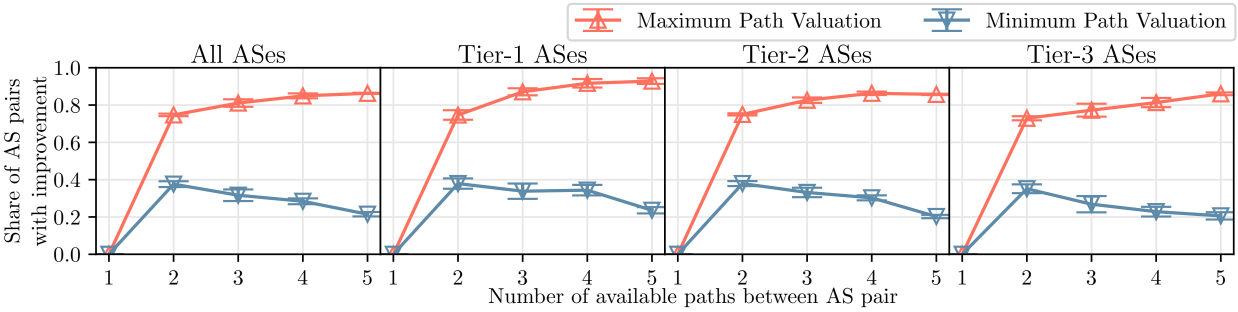

Figure 4: On average, competition raises the the attractiveness of the most *and* the least attractive path that connect two ASes, independent of the source-AS tier.

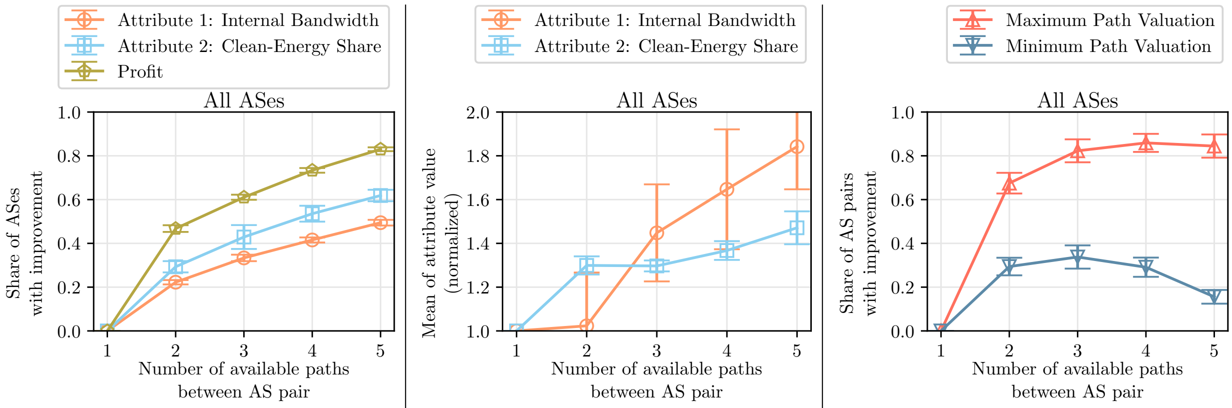

(a) Improvements in transit ASes.
(b) Mean of attribute values.
(c) Improvements for selecting ASes.

Figure 5: Equilibrium results for competitive dynamics with non-affine functions for valuation and cost.

Differences in market power. Interestingly, the profit increase under competition among 5 paths is more pronounced for tier-2 and tier-3 ASes than for tier-1 ASes. The reason is that the tier-1 ASes become more likely to be circumvented as path-selecting ASes obtain additional path options, and lower-tier ASes can thus attract more demand.

5.2.3. Sensitivity to Model Functions

In our model, we assume affine functions for both the path-valuation functions v_r and the cost functions Φ_n and Γ_n. To assess the impact of this assumption, we rerun the simulations by replacing linear dependencies within these functions. In particular, we replace a_{nk} in v_r by sub-linear $\sqrt{a_{nk}}$ (cf. Eq. (1)), and replace a_{nk} in Φ_n and Γ_n by super-linear a_{nk}^2 (cf. Eq. (4)). The results are presented in Fig. 5.

Intriguingly, the results for the non-affine functions closely match the results for affine functions in a qualitative sence, i.e., competition improves the attribute values, profits, and path options for a large share of ASes. Quantitatively, the largest difference concerns the mean increase in attribute values (cf. Fig. 5b vs. Fig. 3), which is considerably lower for the bandwidth attribute for the non-affine functions. However, this effect is to be expected because the non-affine functions make attributes both less valuable and more costly (if $a_{nk} > 1$, as for the bandwidth attribute), and thus less attractive to invest in.

6. Related Work

General competition models. Internet competition is frequently studied by means of the three fundamental competition models from the economic literature. First, Cournot competition [12] describes multiple firms that produce the same homogeneous good, individually determine the quantity to be produced, and thus indirectly determine the market price. Cournot competition suggests that in comparison with a monopoly, a duopoly increases the quantity of the good and reduces its price, indicating that competition benefits consumers. In the second competition model, Bertrand competition [14], firms also produce a homogeneous good, but directly set a price instead of a quantity. Moreover, the firm with the lowest price acquires the whole market. Hence, Bertrand competition is considered more suitable to analyze highly competitive markets. Lastly, Stackelberg competition [44] is similar to Cournot competition, but is suitable for hierarchical markets in which follower firms determine their production quantity after observing the quantity produced by a leader firm. All competition models have been adapted to *networked markets*, i.e., markets in which each segment of consumers can only be served by a corresponding subset of firms [1, 3, 28].

Our competition model is more strongly inspired by the logit-demand model [2], which originates from econometrics, can more directly represent competition between goods with different characteristics, and has been used in research on Internet transit pricing [42]. Still, the market in our model is networked, as determined by the network topology.

Internet competition models. To study Internet competition in particular, Shakkottai and Srikant [36] leverage Bertrand and Stackelberg competition to theoretically analyze the effect of competition in different levels of the Internet, i.e., for tier-1, tier-2, and local ISPs. Their model shows that competition may exert downwards pressure on prices, and an assimilating pressure on the quality-of-service (QoS) levels offered by different ISPs. These insights are confirmed by a subsequent line of simulation-based research [25, 41]. These studies analyze the ISP competition induced by virtualized access networks and by the ChoiceNet proposal [46], which describes a marketplace for transit services. Using both theoretical analysis and simulation, Nagurney and Wolf [29] expand on a study by Zhang et al. [47] to investigate the intertwined competition dynamics among service providers (in Bertrand competition) and among network providers (in Cournot competition). In this analysis, the offers of service providers and network providers differ in both price and quality, and converge to an equilibrium describable by variational-inequality conditions.

In this paper, we extend the previous work in a number of aspects. First, our model acknowledges that path quality may be determined collectively by multiple selfish on-path ISPs, and reveals the inefficient cooperation within a path due to selfishness. In contrast, previous work assumed that network service is bought from a single access/transit provider, and is thus effectively limited to one-hop paths. Second, path quality in our model depends on multiple underlying attributes, whereas previous work abstracts path quality in a single attribute. This fine-grained view of quality attributes is valuable, as it reveals how different attributes are affected differently by competition (cf. §3.3). Third, our model represents the internal cost structure of ISPs in a detailed manner, as it (i) distinguishes demand-dependent and demand-independent cost, and (ii) formalize the cost dependence on quality attributes, unlike previous work. Fourth, we make an effort to find realistic parameters for our large-scale simulations, whereas the parameters in previous simulation-based works are arbitrary. Finally, we explicitly investigate the differences between differing degrees of competition, and find network examples in which more intense competition lead to previously undocumented effects, namely increasing profits and decreasing path quality.

ISP cooperation. The economic dynamics between between network entities that collectively provide connectivity has been studied with the lens of cooperative game theory [4], i.e., assuming that agents within a group choose rules which are enforced thereafter. Such considerations can inform financial settlement among ISPs in a coalition, where settlement mechanisms based on the Shapley value [24] or ISP characteristics [48] have been proposed. In our work, we discuss intra-path dynamics using non-cooperative game theory, as setting up a binding contract among the ISPs on every path is difficult in practice. Moreover, our model also reflects that multiple coalitions (paths) are in competition, which is missing from previous work.

Path-selection models. Our model does not only predict how ISPs adapt their quality attributes under competition, but also how ISPs select among available paths. In that latter aspect, our work has a connection to path-selection models such as *hot-potato routing* [38, 40]. The model of hot-potato routing is based on the fundamental assumption that ISPs generally select the path with the geographically closest ISP egress, and predicts actual path selection in the Internet with substantial accuracy [38]. While the hot-potato–routing model is considerably simpler than our model, the higher complexity of our model allows to complement the hot-potato model in two valuable respects. First, our model can accommodate the geographic proximity of ISP interfaces as a quality attribute with high valuation weight; additional attributes may then help explain why actual path selection sometimes deviates from hot-potato routing. Second, our model may additionally explain how ISPs *react* to hot-potato routing, i.e., by strategically placing their ingresses to attract traffic.

7. Conclusion

ISPs determine the quality attributes of their connectivity offer (e.g., performance metrics, security features, sustainability properties) in line with their profit objective and alternative offers by other ISPs. The presence of such alternative offers (i.e., competition) tends to improve path quality, as we demonstrate in this paper. We provide evidence for this positive effect of ISP competition with an extensive theoretical analysis, based on a new game-theoretic model, and a large-scale simulation, based on empirical data. Our theoretical analysis suggests that an augmented path choice incentivizes transit ISPs to improve path quality, especially if ISPs have similar cost structures (Theorems 5 and 12). Interestingly, this higher investment in quality attributes does not necessarily reduce transit ISP profits, as entering competition (by connecting to new customers) also allows transit ISPs to attract revenue-generating traffic from new customer segments (Theorem 6). While these positive effects do not materialize in unfavorable circumstances (Theorem 13), our simulation-based case study indicates largely positive effects of competition in practice.

Our analysis does not only reveal the macroscopic effects of competition, but also formalizes the rational attribute choice for ISPs, which can inform ISP business strategies. In particular, we obtain three main recommendations for quality investment from our analysis:

- *Invest in attributes with low fixed cost:* Theorem 1 suggests that the optimal extent of a quality attribute correlates with the ratio of demand-dependent attribute-specific cost to total attribute-specific cost. Hence, the lower the demand-independent (fixed) cost of an attribute, the higher the optimal investment in the attribute. For example, renting internal bandwidth on-demand tends to improve profit more than a fixed bandwidth installation.

- *Invest exclusively in attributes with high return:* Theorem 1 also shows a correlation between the optimal attribute extent and the attribute return, i.e., the attribute-specific net revenue per traffic unit, divided by the attribute-specific cost. Theorem 10 even suggests *specialization* in heterogeneous markets with negligible demand-dependent cost, i.e., only the path attribute with the highest return should be invested in, while all less attractive attributes should be abandoned.

- *Engage in competition and on-path cooperation:* Both our theoretical analysis (Theorem 6) and our simulations (§5) show that engaging in competition by connecting to new customers tends to increase transit ISP profit, as the revenue from newly attracted traffic generally outweighs the costs of competing in the new markets. Furthermore, transit ISPs should also engage in cooperation with other on-path ISPs by coordinating attribute investment. Such coordination leads to achievement of the Nash bargaining solution, and therefore to higher profits and higher path quality, which also benefits path users (Theorems 4 and 9). However, to achieve stable cooperation among selfish ISPs, additional work based on mechanism design is needed.

Finally, we emphasize that our new competition model is not only applicable to ISP competition, but in general to settings in which coalitions of selfish entities stand in competition. While ISPs and paths represent the entities and coalitions in the ISP market, competition in other markets arises between firms that form value chains. Our model allows investigating complex economic phenomena such as the interaction between firms along a value chain, or the effect of overlapping value chains. Hence, we are intrigued to investigate whether our findings for the ISP market translate to other economic sectors.

References

[1] Elliot Anshelevich and Shreyas Sekar. 2015. Price competition in networked markets: How do monopolies impact social welfare?. In *International Conference on Web and Internet Economics*. Springer, 16–30.

[2] David Besanko, Sachin Gupta, and Dipak Jain. 1998. Logit Demand Estimation under Competitive Pricing Behavior: An Equilibrium Framework. *Management Science* 44, 11-part-1 (1998), 1533–1547.

[3] Kostas Bimpikis, Shayan Ehsani, and Rahmi Ilkılıç. 2019. Cournot competition in networked markets. *Management Science* 65, 6 (2019), 2467–2481.

[4] Rodica Branzei, Dinko Dimitrov, and Stef Tijs. 2008. *Models in cooperative game theory*. Vol. 556. Springer Science & Business Media.

[5] LEJ Brouwer. 1910. Über Abbildungen von Mannigfaltigkeiten. *written at the Univ of Amsterdam* (1910).

[6] Antonio Cabrales and Roberto Serrano. 2011. Implementation in adaptive better-response dynamics: Towards a general theory of bounded rationality in mechanisms. *Games and Economic Behavior* 73, 2 (2011), 360–374.

[7] Center for Applied Internet Data Analysis (CAIDA). 2019. AS Relationships – with Geographic Annotations. `https://www.caida.org/data/as-relationships-geo/` (last visited Jan 10, 2022), archived at `https://perma.cc/S4XX-Y2EK`.

[8] Center for Applied Internet Data Analysis (CAIDA). 2019. Routeviews Prefix to AS Mappings Data Set (pfx2as) for IPv4 and IPv6. `https://www.caida.org/data/routing/routeviews-prefix2as.xml` (last visited Jan 10, 2023), archived at `https://perma.cc/RB6L-PEVA`.

[9] Center for Applied Internet Data Analysis (CAIDA). 2020. AS-Relationships. `https://www.caida.org/data/as-relationships/` (last visited Jan 10, 2023), archived at `https://perma.cc/RYN2-BSNB`.

[10] Center for Applied Internet Data Analysis (CAIDA). 2022. Macroscopic Internet Topology Data Kit (ITDK). `https://www.caida.org/data/internet-topology-data-kit/` (last visited Jan 10, 2023), archived at `https://perma.cc/CU8X-7GRU`.

[11] Niangjun Chen, Xiaoqi Ren, Shaolei Ren, and Adam Wierman. 2015. Greening multi-tenant data center demand response. *ACM SIGMETRICS Performance Evaluation Review* 43, 2 (2015), 36–38.

[12] Antoine-Augustin Cournot. 1838. *Recherches sur les principes mathématiques de la théorie des richesses*. L. Hachette.

[13] Alex Davidson, Matthias Frei, Marten Gartner, Hamed Haddadi, Adrian Perrig, Jordi Subirà Nieto, Philipp Winter, and François Wirz. 2022. Tango or square dance? how tightly should we integrate network functionality in browsers?. In *Proceedings of the 21st ACM Workshop on Hot Topics in Networks*. 205–212.

[14] Francis Y Edgeworth. 1925. The pure theory of monopoly. *Papers relating to political economy* 1, 111 (1925), 42.

[15] Ember Climate. 2022. EU Carbon Price Tracker. `https://ember-climate.org/data/data-tools/carbon-price-viewer/` (last visited Jan 10, 2023).

[16] Lixin Gao and Jennifer Rexford. 2001. Stable Internet routing without global coordination. *IEEE/ACM Transactions on networking* 9, 6 (2001), 681–692.

[17] Ward Heddeghem, Filip Idzikowski, Willem Vereecken, Didier Colle, Mario Pickavet, and Piet Demeester. 2012. Power Consumption Modeling in Optical Multilayer Networks. *Photonic Network Communications* 24 (Oct 2012). `https://doi.org/10.1007/s11107-011-0370-7`

[18] Geoff Huston. 2021. BGP in 2020 – The BGP Table. `https://labs.apnic.net/?p=1395` (last visited Jan 10, 2023), archived at `https://perma.cc/UMJ9-Q4QA`.

[19] International Energy Agency (IEA). 2022. Average CO2 intensity of power generation from coal power plants, 2000-2020. `https://www.iea.org/data-and-statistics/charts/average-co2-intensity-of-power-generation-from-coal-power-plants-2000-2020` (last visited Jan 10, 2023).

[20] Romain Jacob and Laurent Vanbever. 2022. The Internet of tomorrow must sleep more and grow old. In *HotCarbon 2022*.

[21] Lazard. 2020. Levelized Cost of Energy Analysis: Version 14.0. *Report* (2020).

[22] Yu-Hsin Liu, Jeffrey Prince, and Scott Wallsten. 2018. Distinguishing bandwidth and latency in households' willingness-to-pay for broadband internet speed. *Information Economics and Policy* 45 (2018), 1–15.

[23] Zhenhua Liu, Minghong Lin, Adam Wierman, Steven H Low, and Lachlan LH Andrew. 2011. Greening geographical load balancing. *ACM SIGMETRICS Performance Evaluation Review* 39, 1 (2011), 193–204.

[24] Richard TB Ma, Dah Ming Chiu, John CS Lui, Vishal Misra, and Dan Rubenstein. 2007. Internet Economics: The use of Shapley value for ISP settlement. In *Proceedings of the 2007 ACM CoNEXT conference*. 1–12.

[25] Andres Marentes, Thiago Teixeira, and Tilman Wolf. 2015. Exploring economic dynamics in an internet with service choices. In *2015 IEEE International Conference on Communications (ICC)*. IEEE, 5272–5277.

[26] Eli A Meirom, Shie Mannor, and Ariel Orda. 2014. Network formation games with heterogeneous players and the internet structure. In *Proceedings of the fifteenth ACM conference on Economics and computation*. 735–752.

[27] Jayne Miller. 2021. 2021 Global Internet Map Tracks Global Capacity, Traffic, and Cloud Infrastructure. `https://blog.telegeography.com/2021-global-internet-map-tracks-global-capacity-traffic-and-cloud-infrastructure` (last visited Jan 10, 2023).

[28] Mahdi Motalleb, Pierluigi Siano, and Reza Ghorbani. 2019. Networked Stackelberg Competition in a Demand Response Market. *Applied energy* 239 (2019), 680–691.

[29] Anna Nagurney and Tilman Wolf. 2014. A Cournot–Nash–Bertrand game theory model of a service-oriented Internet with price and quality competition among network transport providers. *Computational Management Science* 11, 4 (2014), 475–502.

[30] John F Nash Jr. 1950. The bargaining problem. *Econometrica: Journal of the econometric society* (1950), 155–162.

[31] OpenVault. 2021. Broadband Insights Report (OVBI), Q4/2021. `https://openvault.com/wp-content/uploads/2022/03/OVBI_4Q21_Report_FINAL-1.pdf` (last visited Jan 10, 2023).

[32] Chutiphon Pukdeboon. 2011. A review of fundamentals of Lyapunov theory. *The Journal of Applied Science* 10, 2 (2011), 55–61.
[33] Anatol Rapoport. 1989. Prisoner's dilemma. In *Game theory*. Springer, 199–204.
[34] Matthew Roughan. 2005. Simplifying the synthesis of Internet traffic matrices. *ACM SIGCOMM Computer Communication Review* 35, 5 (2005), 93–96.
[35] Simon Scherrer, Markus Legner, Adrian Perrig, and Stefan Schmid. 2021. Enabling Novel Interconnection Agreements with Path-Aware Networking Architectures. In *2021 51st Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*. IEEE, 116–128.
[36] Srinivas Shakkottai and Rayadurgam Srikant. 2006. Economics of network pricing with multiple ISPs. *IEEE/ACM Transactions On Networking* 14, 6 (2006), 1233–1245.
[37] Herbert A Simon. 1990. Bounded rationality. In *Utility and probability*. Springer, 15–18.
[38] Lakshminarayanan Subramanian, Venkata N Padmanabhan, and Randy H Katz. 2002. Geographic Properties of Internet Routing.. In *Proceedings of the USENIX Annual Technical Conference (ATC)*. 243–259.
[39] Seyedali Tabaeiaghdaei, Simon Scherrer, Jonghoon Kwon, and Adrian Perrig. 2023. Carbon Intelligent Global Routing in Path-Aware Networks. In *Proceedings of the 14th ACM International Conference on Future Energy Systems (e-Energy '23)*. ACM.
[40] Renata Teixeira, Aman Shaikh, Tim Griffin, and Jennifer Rexford. 2004. Dynamics of hot-potato routing in IP networks. In *Proceedings of the Joint International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS)*. 307–319.
[41] Thiago Teixeira, Andres Marentes, and Tilman Wolf. 2017. Economic incentives in virtualized access networks. In *2017 IEEE International Conference on Communications (ICC)*. IEEE, 1–7.
[42] Vytautas Valancius, Cristian Lumezanu, Nick Feamster, Ramesh Johari, and Vijay V Vazirani. 2011. How many tiers? Pricing in the Internet transit market. In *Proceedings of the ACM SIGCOMM 2011 Conference*. 194–205.
[43] Hal R Varian and JK MacKie Mason. 1994. Some FAQs about usage-based pricing. *Okerson [Oke95]* (1994), 77–86.
[44] Heinrich von Stackelberg. 1934. *Marktform und Gleichgewicht*. Springer Science & Business Media.
[45] Cindy Whelan. 2020. Wholesale Forecast: Backbone Services Prove Resilient, Offsetting Voice Declines. `https://apacblog.lumen.com/wholesale-forecast-backbone-services-prove-resilient-offsetting-voice-declines/` (last visited Jan 13, 2023).
[46] Tilman Wolf, James Griffioen, Kenneth L Calvert, Rudra Dutta, George N Rouskas, Ilya Baldin, and Anna Nagurney. 2014. ChoiceNet: toward an economy plane for the Internet. *ACM SIGCOMM Computer Communication Review* 44, 3 (2014), 58–65.
[47] Zhi-Li Zhang, Papak Nabipay, Andrew Odlyzko, and Roch Guerin. 2010. Interactions, competition and innovation in a service-oriented Internet: An economic model. In *2010 Proceedings IEEE INFOCOM*. IEEE, 1–5.
[48] Weihua Zhou, Yaqi Pu, Hongyan Dai, and Qingwei Jin. 2017. Cooperative interconnection settlement among ISPs through NAP. *European Journal of Operational Research* 256, 3 (2017), 991–1003.
[49] Noa Zilberman, Eve M Schooler, Uri Cummings, Rajit Manohar, Dawn Nafus, Robert Soulé, and Rick Taylor. 2022. Toward Carbon-Aware Networking. In *HotCarbon 2022: 1st Workshop on Sustainable Computer Systems Design and Implementation*.

Appendix A. Proof of Theorem 1

Appendix A.1. Unrestricted best response in Eq. (11)

For an individual source-destination pair (n_1, n_2), we can simplify Eq. (5) as follows:

$$\pi_n = \frac{dv_{r(n)}}{1+\sum_{r'\in R} v_{r'}} \cdot \left(\rho_n - \sum_{k\in K} \phi_{nk} a_{nk} - \phi_{n0}\right) - \sum_{k\in K} \gamma_{nk} a_{nk} - \gamma_{n0} \tag{A.1}$$

where the argument $\mathbf{A}$ has been omitted and

$$v_{r(n)} = \sum_{r\in R.\ n\in r} v_r. \tag{A.2}$$

Differentiating Eq. (A.1) with respect to attribute a_{nk} of ISP n yields:

$$\frac{\partial \pi_n}{\partial a_{nk}} = \frac{d\alpha_{nk}\left(1+v_{-r(n)}\right)}{\left(1+\sum_{r'\in R} v_{r'}\right)^2} \cdot \left(\rho_n - \sum_{k\in K} \phi_{nk} a_{nk} - \phi_{n0}\right) - \frac{d\phi_{nk} v_{r(n)}}{1+\sum_{r'\in R} v_{r'}} - \gamma_{nk} \tag{A.3}$$

where the abbreviations from Eq. (13) have been used. Setting Eq. (A.3) to 0 allows the following rewriting:

$$d\alpha_{nk}\left(1+v_{-r(n)}\right) \cdot \left(\rho_n - \sum_{k\in K} \phi_{nk} a_{nk} - \phi_{n0}\right) - d\phi_{nk} v_{r(n)} \cdot \left(1 + \sum_{r'\in R} v_{r'}\right) - \gamma_{nk} \cdot \left(1 + \sum_{r'\in R} v_{r'}\right)^2 = 0 \tag{A.4}$$

In the LHS, we can substitute

$$\begin{aligned} v_{r(n)} + v_{-r(n)} &= \sum_{r'\in R} v_{r'} = \alpha_{nk} a_{nk} + v_{-nk} = \alpha_{nk} a_{nk} + v_{r(n),-nk} + v_{-r(n)} \\ \sum_{k'\in K} \phi_{nk'} a_{nk'} + \phi_{n0} &= \phi_{nk} a_{nk} + \Phi_{-nk} \end{aligned} \tag{A.5}$$

and obtain:

$$\begin{aligned}
& - d\alpha_{nk}\phi_{nk}\left(1+v_{-r(n)}\right) \cdot a_{nk} + d\alpha_{nk}\left(\rho_n - \Phi_{-nk}\right)\left(1+v_{-r(n)}\right) && \text{(A.6)}\\
& - d\phi_{nk}\left(\alpha_{nk} a_{nk} + v_{r(n),-nk}\right)\left(\alpha_{nk} a_{nk} + 1 + v_{-nk}\right) \\
& - \gamma_{nk} \cdot \left(\alpha_{nk} a_{nk} + 1 + v_{-nk}\right)^2 \\
\Longleftrightarrow\ & - d\alpha_{nk}\phi_{nk}\left(1+v_{-r(n)}\right) \cdot a_{nk} + d\alpha_{nk}\left(\rho_n - \Phi_{-nk}\right)\left(1+v_{-r(n)}\right) && \text{(A.7)}\\
& - d\phi_{nk} \cdot \alpha_{nk}^2 a_{nk}^2 - d\phi_{nk} \cdot \alpha_{nk} \cdot \left(v_{r(n),-nk} + \left(1+v_{-nk}\right)\right) a_{nk} - d\phi_{nk} \cdot v_{r(n),-nk}\left(1+v_{-nk}\right) \\
& - \gamma_{nk} \cdot \left(\alpha_{nk}^2 a_{nk}^2 + 2\alpha_{nk}\left(1+v_{-nk}\right) a_{nk} + \left(1+v_{-nk}\right)^2\right) \\
\Longleftrightarrow\ & \left(-d\phi_{nk}\alpha_{nk}^2 - \gamma_{nk}\alpha_{nk}^2\right) \cdot a_{nk}^2 && \text{(A.8)}\\
& + \left(-d\alpha_{nk}\phi_{nk}\left(1+v_{-r(n)}\right) - d\alpha_{nk}\phi_{nk}\left(1+v_{r(n),-nk}+v_{-nk}\right)\right) \cdot a_{nk} \\
& + \left(-2\alpha_{nk}\gamma_{nk}\left(1+v_{-nk}\right)\right) \cdot a_{nk} \\
& + d\alpha_{nk}\left(\rho_n - \Phi_{-nk}\right)\left(1+v_{-r(n)}\right) - d\phi_{nk} \cdot v_{r(n),-nk} \cdot \left(1+v_{-nk}\right) - \gamma_{nk} \cdot \left(1+v_{-nk}\right)^2 \\
\Longleftrightarrow\ & - \alpha_{nk}^2\left(d\phi_{nk} + \gamma_{nk}\right) \cdot a_{nk}^2 && \text{(A.9)}\\
& - 2d\alpha_{nk}\phi_{nk}\left(1+v_{-nk}\right) \cdot a_{nk} - 2\alpha_{nk}\gamma_{nk}\left(1+v_{-nk}\right) \cdot a_{nk} \\
& + d\alpha_{nk}\left(\rho_n - \Phi_{-nk}\right) - d\phi_{nk} \cdot v_{r(n),-nk} \cdot \left(1+v_{-nk}\right) - \gamma_{nk} \cdot \left(1+v_{-nk}\right)^2 \\
\Longleftrightarrow\ & - \alpha_{nk}^2\left(d\phi_{nk} + \gamma_{nk}\right) \cdot a_{nk}^2 && \text{(A.10)}\\
& - 2\alpha_{nk}\left(d\phi_{nk} + \gamma_{nk}\right)\left(1+v_{-nk}\right) \cdot a_{nk} \\
& + d\alpha_{nk}\left(\rho_n - \Phi_{-nk}\right)\left(1+v_{-r(n)}\right) - d\phi_{nk} \cdot v_{r(n),-nk} \cdot \left(1+v_{-nk}\right) - \gamma_{nk} \cdot \left(1+v_{-nk}\right)^2 \\
\Longleftrightarrow\ & T_1 a_{nk}^2 + T_2 a_{nk} + T_3 = 0 && \text{(A.11)}
\end{aligned}$$

Hence, we obtain a quadratic equation in a_{nk}. This quadratic equation has the solutions:

$$
\begin{aligned}
a_{nk} &= \frac{-T_2 \pm \sqrt{T_2^2 + 4T_1T_3}}{2T_1} && \text{(A.12)}\\
&= -\frac{2\alpha_{nk}\left(d\phi_{nk} + \gamma_{nk}\right)\left(1 + v_{-nk}\right)}{2\alpha_{nk}^2\left(d\phi_{nk} + \gamma_{nk}\right)} && \text{(A.13)}\\
&\quad \pm \frac{\sqrt{4\alpha_{nk}^2\left(d\phi_{nk} + \gamma_{nk}\right)^2\left(1 + v_{-nk}\right)^2 + 4\alpha_{nk}^2(d\phi_{nk} + \gamma_{nk}) \cdot T_3}}{2\alpha_{nk}^2\left(d\phi_{nk} + \gamma_{nk}\right)} && \text{(A.14)}\\
&= -\frac{1 + v_{-nk}}{\alpha_{nk}} \pm \frac{\sqrt{\left(d\phi_{nk} + \gamma_{nk}\right)\left(1 + v_{-nk}\right)^2 + T_3}}{\alpha_{nk}\sqrt{d\phi_{nk} + \gamma_{nk}}}, && \text{(A.15)}
\end{aligned}
$$

where we expand the term under the root as follows:

$$
\begin{aligned}
&\left(d\phi_{nk} + \gamma_{nk}\right)\left(1 + v_{-nk}\right)^2 + T_3 && \text{(A.16)}\\
&= d\phi_{nk} \cdot \left(1 + v_{-nk}\right)^2 + \underline{\gamma_{nk}\left(1 + v_{-nk}\right)^2} && \text{(A.17)}\\
&\quad + d\alpha_{nk}\left(\rho_n - \Phi_{-nk}\right)\left(1 + v_{-r(n)}\right)\\
&\quad - d\phi_{nk} \cdot v_{r(n),-nk} \cdot \left(1 + v_{-nk}\right) \underline{- \gamma_{nk} \cdot \left(1 + v_{-nk}\right)^2}\\
&= d\phi_{nk} \cdot \left(1 + v_{-nk}\right)^2 - d\phi_{nk} \cdot v_{r(n),-nk} \cdot \left(1 + v_{-nk}\right) && \text{(A.18)}\\
&\quad + d\alpha_{nk}\left(\rho_n - \Phi_{-nk}\right)\left(1 + v_{-r(n)}\right)\\
&\overset{\text{(A.20)}}{=} d\phi_{nk}\left(1 + v_{-r(n)}\right)\left(1 + v_{-nk}\right) + d\alpha_{nk}\left(\rho_n - \Phi_{-nk}\right)\left(1 + v_{-r(n)}\right), && \text{(A.19)}
\end{aligned}
$$

where we have made use of the following equality in the last step:

$$
\begin{aligned}
&\left(1 + v_{-nk}\right)^2 - v_{r(n),-nk}\left(1 + v_{-nk}\right)\\
&= \left(1 + v_{r(n),-nk} + v_{-r(n)}\right)\left(1 + v_{-nk}\right) - v_{r(n),-nk}\left(1 + v_{-nk}\right) && \text{(A.20)}\\
&= \left(1 + v_{-r(n)}\right)\left(1 + v_{-nk}\right)
\end{aligned}
$$

By reinserting Eq. (A.19) in Eq. (A.15), we obtain:

$$
a_{nk} = \frac{1}{\alpha_{nk}}\left(\pm\sqrt{\frac{d\left(1 + v_{-r(n)}\right)}{d\phi_{nk} + \gamma_{nk}}\left(\phi_{nk}(1 + v_{-nk}) + \alpha_{nk}(\rho_n - \Phi_{-nk})\right)} - \left(1 + v_{-nk}\right)\right), \tag{A.21}
$$

where only the upper solution (i.e., with the positive coefficient of the square-root term) is potentially valid given the non-negativity of attributes. Hence, we arrive at $\hat{a}_{nk}^*(\mathbf{A}_{-nk})$ as in Eq. (11).

Appendix A.2. Confirmation of maximum

This solution is a maximum of π_n if π_n is concave in a_{nk}, which can be demonstrated by means of the second derivative:

$$
\frac{\partial^2 \pi_n}{\partial a_{nk}^2} = \frac{-2d\alpha_{kn}^2\left(1 + v_{-r(n)}\right)}{\left(1 + \sum_{r' \in R} v_{r'}\right)^3} \cdot \left(\rho_n - \sum_{k \in K} \phi_{nk} a_{nk} - \phi_{n0}\right) - \frac{2d\alpha_n \phi_{nk}\left(1 + v_{-r(n)}\right)}{\left(1 + \sum_{r' \in R} v_{r'}\right)^2} \tag{A.22}
$$

Clearly, π_n may only be non-concave under the following condition:

$$
\frac{\partial^2 \pi_n}{\partial a_{nk}^2} > 0 \implies \rho_n - \sum_{k \in K} \phi_{nk} a_{nk} - \phi_{n0} < 0 \tag{A.23}
$$

However, if the condition in Eq. (A.23) is true for some $\mathbf{a}_n \in \mathbb{R}_{\geq 0}^K$, then the profit function has a negative slope at that point (cf. Eq. (A.3)). Hence, the profit function has no extrema in the non-concave regions, and thus any extremum, in particular $\hat{a}_{nk}^*(\mathbf{A}_{-nk})$, is guaranteed to be located in the concave regions and to be a maximum.

Appendix A.3. Restricted best response in Eq. (10)

Since attribute values must be non-negative, we now investigate how the unrestricted best-response $\hat{a}^*_{nk}$ informs the best-response a^*_{nk} on the restricted domain $\mathbb{R}_{\geq 0}$. Clearly, if $\hat{a}^*_{nk} \geq 0$, then $a^*_{nk} = \hat{a}^*_{nk}$. Otherwise, if $\hat{a}^*_{nk} < 0$, the boundary point $a_{nk} = 0$ constitutes a local maximum on the restricted domain $\mathbb{R}_{\geq 0}$. To confirm this statement, we have to distinguish the cases $\phi_{nk} > 0$ and $\phi_{nk} = 0$.

If $\phi_{nk} > 0$, we note that the profit function can only be non-concave for high enough a_{nk}:

$$\rho_n - \sum_{k \in K} \phi_{nk} a_{nk} - \phi_{n0} < 0 \iff a_{nk} > \frac{\rho_n - \Phi_{-nk}}{\phi_{nk}} =: \overline{a}_{nk} \tag{A.24}$$

Hence, π_n is guaranteed to be concave for all $a_{nk} \leq \overline{a}_{nk}$. Furthermore, as argued in the previous paragraph, π_n is strictly decreasing for all $a_{nk} > \overline{a}_{nk}$. Hence, $a_{nk} = 0$ is the maximum on the restricted domain independent of $\overline{a}_{nk}$, given that $\hat{a}^*_{nk} < 0$: Either $a_{nk} = 0$ is in the concave region (if $\overline{a}_n k \leq 0$) or in the decreasing region (if $\overline{a}_{nk} > 0$).

If $\phi_{nk} = 0$, π_n can only be non-concave if $\rho_n - \Phi_{-nk} < 0$, independent of a_{nk}. Hence, π_n is either guaranteed to be consistently concave in a_{nk} (if $\rho_n - \Phi_{-nk} \geq 0$), or guaranteed to be decreasing for all $a_{nk} \geq 0$. In both cases, the boundary point $a_{nk} = 0$ constitutes a local maximum.

Finally, we investigate the case where the unrestricted best response $\hat{a}^*_{nk}$ is *undefined* on $\mathbb{R}$, i.e., if the term under the square root is negative. A necessary condition for this negativity is that $\rho_n - \Phi_{-nk}$ is negative, which according to Eq. (A.24) implies that $\overline{a}_{nk} < 0$. Hence, in this case π_n is decreasing for $a_{nk} \geq 0$, which again makes $a_{nk} = 0$ a local maximum on the restricted domain $\mathbb{R}_{\geq 0}$.

Appendix B. Proof of Theorem 2

Appendix B.1. Homogeneous profit function

In a homogeneous network, all attributes are equally valuable and costly, i.e., α_{nk}, ϕ_{nk} and γ_{nk} are equal across all attributes $k \in K$. This homogeneity allows to understand the profit function π_n of a ISP n as a function of the attribute sums $a_n = \sum_{k \in K} a_{nk}$:

$$\pi_n(\mathbf{A}) = \frac{\alpha_1 a_n + \alpha_1 \sum_{n' \in r(n) \setminus \{n\}} a_{n'} + \alpha_0}{1 + \alpha_1 a_n + \alpha_1 \sum_{n' \in N \setminus \{n\}} a_{n'} + Q\alpha_0} d(\rho - \phi_1 a_n - \phi_0) - \gamma_1 a_n - \gamma_0 \tag{B.1}$$

In the following, we thus treat the attribute sum a_n like a single attribute of ISP n.

Appendix B.2. Unrestricted equilibrium in Eq. (16)

The equilibrium conditions in Eq. (15) suggest that the equilibrium for a homogeneous parallel-path network satisfies the following equation for every $n \in N$:

$$a^+_n = \max\left(0, \frac{\sqrt{\frac{d\left(1+v_{-r(n)}(\mathbf{A}^+_{-n})\right)}{d\phi_1+\gamma_1}\left(\phi_1(1+v_{-n}(\mathbf{A}^+_{-n})) + \alpha_1(\rho - \phi_0)\right)} - \left(1 + v_{-n}(\mathbf{A}^+_{-n})\right)}{\alpha}\right) \tag{B.2}$$

where

$$v_{-r(n)}(\mathbf{A}^+_{-n}) = \alpha_1 \sum_{n' \in N \setminus r(n)} a^+_{n'} + (Q-1)\alpha_0 \tag{B.3}$$

$$v_{-n}(\mathbf{A}^+_{-n}) = v_{-r(n)}(\mathbf{A}^+_{-n}) + \alpha_1 \sum_{n' \in r(n) \setminus \{n\}} a^+_{n'} + \alpha_0 \tag{B.4}$$

Note that we effectively consider a single attribute in the style of the attribute sum, which simplifies $\Phi_{-nk} = \phi_0 \leq \rho$. Hence, the undefined case from Theorem 1 does not arise because the term under the square root in Eq. (B.2) is always non-negative.

To solve the equation system Eq. (B.2) $\forall n \in N$, we first consider the unrestricted system (i.e., without required non-negativity of solutions) in $\hat{\mathbf{A}}^+$:

$$\forall n \in N. \quad \hat{a}_n^+ = \frac{\sqrt{\frac{d\left(1+v_{-r(n)}(\hat{\mathbf{A}}_{-n}^+)\right)}{d\phi_1+\gamma_1}\left(\phi_1(1+v_{-n}(\hat{\mathbf{A}}_{-n}^+))+\alpha_1(\rho-\phi_0)\right)\Big)} - \left(1+v_{-n}(\hat{\mathbf{A}}_{-n}^+)\right)}{\alpha} \tag{B.5}$$

This equation system can be transformed such that the LHS is constant across all equations:

$$\forall n \in N. \quad 1+\alpha_1 \sum_{r\in R}\sum_{n'\in r} \hat{a}_{n'}^+ + Q\alpha_0 = \sqrt{\frac{d\left(1+v_{-r(n)}(\hat{\mathbf{A}}_{-n}^+)\right)}{d\phi_1+\gamma_1}\left(\phi_1(1+v_{-n}(\hat{\mathbf{A}}_{-n}^+))+\alpha_1(\rho-\phi_0)\right)\Big)} \tag{B.6}$$

which implies that all a_n^+ $\forall n \in N$ are equal to a value $\hat{a}^+$. This value $\hat{a}^+$ can be found by solving the following single equation:

$$1+QI\alpha_1\hat{a}^+ + Q\alpha_0 = \sqrt{\frac{d\left(1+(Q-1)(I\alpha_1\hat{a}^+ + \alpha_0)\right)}{d\phi_1+\gamma_1}\left(\phi_1(1+(QI-1)\alpha_1\hat{a}^+ + Q\alpha_0)+\alpha_1(\rho-\phi_0)\right)\Big)} \tag{B.7}$$

This equation is solved by $\hat{a}^+$ as defined in Theorem 2.

Appendix B.3. Restricted equilibrium

It remains to show how the solution $\hat{a}^+$ of the unrestricted system can be used to derive the actual solution a^+ of the restricted system. If $\hat{a}^+ \geq 0$, the unrestricted-system solution $\hat{a}^+$ is clearly also a solution to the restricted system, i.e., $a^+ = \hat{a}^+$. However, if $\hat{a}^+ < 0$, the attribute values as suggested by the unrestricted system are negative, which is invalid for the restricted system. In this case, we can show that a solution of the restricted system is given by $a^+ = 0$, i.e., $0 = a_n^*(\mathbf{0})$ (where a_n^* is the optimal choice in single-market networks according to Theorem 1). To show this property, we first observe that the condition for $\hat{a}^+ < 0$ can be simplified to a condition on sub-term T_3:

$$\hat{a}^+ < 0 \implies \frac{\sqrt{T_2^2 - 4T_1T_3} - T_2}{2T_1} < 0 \tag{B.8}$$

$$\text{For } T_1 > 0\text{:} \quad T_2^2 - 4T_1T_3 < T_2^2 \implies -4T_1T_3 < 0 \implies T_3 > 0 \tag{B.9}$$

$$\text{For } T_1 < 0\text{:} \quad T_2^2 - 4T_1T_3 > T_2^2 \implies -4T_1T_3 > 0 \implies T_3 > 0 \tag{B.10}$$

$$\implies T_3 = (1+Q\alpha_0)^2 - \frac{d(1+(Q-1)\alpha_0)}{d\phi_1+\gamma_1}\left(\phi_1(1+Q\alpha_0)+\alpha_1(\rho-\phi_0)\right) > 0 \tag{B.11}$$

Furthermore, the inequality on T_3 allows the following conclusion:

$$(1+Q\alpha_0)^2 - \frac{d(1+(Q-1)\alpha_0)}{d\phi_1+\gamma_1}\left(\phi_1(1+Q\alpha_0)+\alpha_1(\rho-\phi_0)\right) > 0 \tag{B.12}$$

$$\iff (1+Q\alpha_0) - \sqrt{\frac{d(1+(Q-1)\alpha_0)}{d\phi_1+\gamma_1}\left(\phi_1(1+Q\alpha_0)+\alpha_1(\rho-\phi_0)\right)} > 0 \tag{B.13}$$

due to the equivalence $x^2 - y > 0 \iff x^2 > y \iff x > \sqrt{y} \iff x - \sqrt{y} > 0$ (if $x, y \geq 0$).

Eq. (B.13) has a striking similarity to $\hat{a}_n^*(\mathbf{0})$ for a homogeneous parallel-path network:

$$\hat{a}_n^*(\mathbf{0}) = \frac{\sqrt{\frac{d\left(1+(Q-1)\alpha_0\right)}{d\phi_1+\gamma_1}\left(\phi_1(1+Q\alpha_0)+\alpha_1(\rho-\phi_0)\right)\Big)} - \left(1+Q\alpha_0\right)}{\alpha} \tag{B.14}$$

More precisely, Eq. (B.13) implies that $\hat{a}_n^*(\mathbf{0})$ is always below 0 if $\hat{a}^+ < 0$, and that $a_n^*(\mathbf{0})$ thus always equals 0 for $\hat{a}^+ < 0$, i.e., an attribute choice of 0 is the best response to competitor attributes being 0, making $\mathbf{A}^+ = \mathbf{0}$ an equilibrium in this case. This insight concludes the proof.

Appendix C. Proof of Theorem 3

Appendix C.1. Linearization of dynamic system

In order to prove asymptotic stability of the given Nash equilibrium, we leverage the indirect Lyapunov method [32]. This method requires that the equilibrium of an ODE system is asymptotically stable if the Jacobian matrix of the ODE system, evaluated at the equilibrium point, has exclusively negative eigenvalues. More formally, given the Jacobian matrix $\mathbf{J}(\mathbf{A}^+) \in \mathbb{R}^{|N|\times|N|}$, it must hold that $\forall \lambda \in \mathbb{R}$ that $\exists \mathbf{x} \in \mathbb{R}^{|N|}, \mathbf{x} \neq \mathbf{0}.\ \mathbf{J}(\mathbf{A}^+)\mathbf{x} = \lambda\mathbf{x} \implies \lambda < 0$. This matrix $\mathbf{J}(\mathbf{A}^+)$ is defined as follows for the dynamic system from Eq. (20):

$$J_{nn} = \frac{\partial \dot{a}_n}{\partial a_n}(\mathbf{A}^+) = -1 \tag{C.1}$$

$$n \neq m,\ r(n) = r(m): \quad J_{nm} = \frac{\partial \dot{a}_n}{\partial a_m}(\mathbf{A}^+) = \begin{cases} \frac{T_4}{T_5} - 1 & \text{if } \hat{a}^+ \geq 0 \\ 0 & \text{otherwise} \end{cases} \tag{C.2}$$

$$n \neq m,\ r(n) \neq r(m): \quad J_{nm} = \frac{\partial \dot{a}_n}{\partial a_m}(\mathbf{A}^+) = \begin{cases} \frac{T_4+T_6}{T_5} - 1 & \text{if } \hat{a}^+ \geq 0 \\ 0 & \text{otherwise} \end{cases} \tag{C.3}$$

where $\hat{a}^+$ is the unrestricted equilibrium attribute value according to Theorem 2, and

$$T_4 = d\phi_1\big(1 + v_{-r(n)}(\hat{\mathbf{A}}^+)\big), \tag{C.4}$$

$$T_5 = 2(d\phi_1 + \gamma_1)\sqrt{\frac{d\big(1 + v_{-r(n)}(\hat{\mathbf{A}}^+_{-n})\big)}{d\phi_1 + \gamma_1}\big(\phi_1(1 + v_{-n}(\hat{\mathbf{A}}^+_{-n})) + \alpha_1(\rho - \phi_0)\big)}, \text{ and} \tag{C.5}$$

$$T_6 = d\big(\phi_1(1 + v_{-n}(\hat{\mathbf{A}}^+_{-n})) + \alpha_1(\rho - \phi_0)\big). \tag{C.6}$$

Appendix C.2. Case 1: Non-negative unrestricted equilibrium ($\hat{a}^+ \geq 0$)

We first consider the case of a non-negative unrestricted equilibrium value $\hat{a}^+$ such that $\mathbf{A}^+ = \hat{\mathbf{A}}^+$. In that case, the eigenvalue condition induces the following system of equations:

$$\forall n \in N. \quad (-\lambda - 1)x_n + \left(\frac{T_4}{T_5} - 1\right) \sum_{n' \in r(n)\setminus\{n\}} x_{n'} + \left(\frac{T_4 + T_6}{T_5} - 1\right) \sum_{n' \in N\setminus r(n)} x_{n'} = 0 \tag{C.7}$$

This system has a number of solutions $(\lambda, \mathbf{x})$.

$\lambda = -T_4/T_5$. First, for $\lambda_1 = -T_4/T_5$, the first two terms in Eq. (C.7) obtain the same coefficient, and the equation system is reduced from $|N|$ ISP-specific to $|R|$ path-specific equations:

$$\forall r \in R. \quad \left(\frac{T_4}{T_5} - 1\right) X_r + \left(\frac{T_4 + T_6}{T_5} - 1\right) \sum_{r' \in R\setminus\{r\}} X_{r'} = 0 \quad \text{where} \quad X_r = \sum_{n' \in r} x_{n'} \tag{C.8}$$

Equation systems of this form may have three types of solutions in $\mathbf{x}$. For $T_6 = 0$ and $T_4 = T_5$, any $\mathbf{x}$ is a solution, as the coefficients of the variables X_r are 0. For $T_6 = 0$ and $T_4 \neq T_5$, any $\mathbf{x}$ with entries summing up to 0 is a solution, as the sum of all X_r has a single non-zero coefficient. For $T_6 \neq 0$, any $\mathbf{x}$ with $X_r = 0\ \forall r \in R$ is a solution. More importantly for the proof, $\lambda_1 = -T_4/T_5$ is consistently negative given the parameter ranges, except for the case where $\phi_1 = 0$ and hence $\lambda_1 = 0$. The case of $\phi_1 = 0$ is indeed an interesting special case for which the equilibrium is not unique, as we will show in §3.3; for this special case, $\lambda_1 = 0$ describes the fact that the dynamics do not converge to a specific equilibrium if they have already converged onto another equilibrium.

$\lambda \neq -T_4/T_5$. After discovering the first eigenvalue $\lambda_1 = -T_4/T_5$, we now consider the case where $\lambda \neq -T_4/T_5$. In this case, the symmetric structure of the equation system implies that all eigenvector entries x_n associated with the same path r are equal, and thus $x_n = X_r/I$. Hence, a reduction of the equation system to $|R|$ equations is possible again:

$$\forall r \in R. \quad \left(\frac{-\lambda - 1}{I} + \frac{I-1}{I}\left(\frac{T_4}{T_5} - 1\right)\right) X_r + \left(\frac{T_4 + T_6}{T_5} - 1\right) \sum_{r' \in R \setminus \{r\}} X_{r'} = 0 \tag{C.9}$$

Again, this equation system admits different types of solutions.

The first type is obtainable by assuming $X_r = -\sum_{r' \in R \setminus \{r\}} X_{r'}$, and is associated with the following eigenvalue:

$$\frac{-\lambda_2 - 1}{I} + \frac{I-1}{I}\left(\frac{T_4}{T_5} - 1\right) = \frac{T_4 + T_6}{T_5} - 1 \quad \Longrightarrow \quad \lambda_2 = -\frac{T_4 + IT_6}{T_5} \tag{C.10}$$

Importantly in this case, λ_2 is consistently negative.

The second type of solution for the equation system in Eq. (C.9) is obtainable by assuming equal X_r across all paths $r \in R$, and isassociated with the following eigenvalue:

$$\frac{-\lambda_3 - 1}{I} + \frac{I-1}{I}\left(\frac{T_4}{T_5} - 1\right) = -(Q-1)\left(\frac{T_4 + T_6}{T_5} - 1\right) \quad \Longrightarrow \quad \lambda_3 = QI\left(\frac{T_4 + T_6}{T_5} - 1\right) - \frac{T_4 + IT_6}{T_5} \tag{C.11}$$

By inspection of λ_3, we confirm that the maximum λ_3 is negative:

$$\max_{\substack{\alpha_1,\alpha_0,\phi_1,\phi_0,\\ \gamma_1,\rho,d,Q,I}} \lambda_3 = \max_{\phi_1,Q} \lim_{\substack{\alpha_1,\alpha_0,\rho,d\\ \to\infty}} \lim_{\substack{\phi_0,\gamma_1\\ \to 0}} \lim_{I \to 1} \lambda_3 < -\frac{1}{2} \tag{C.12}$$

Hence, all eigenvalues of $\mathbf{J}(\mathbf{A}^+)$ for $\hat{a}^+ > 0$ are negative, i.e., the equilibrium is asymptotically stable in this case.

Appendix C.3. Case 2: Negative unrestricted equilibrium ($\hat{a}^+ < 0$)

It remains to show that the equilibrium $\mathbf{A}^+$ is also asymptotically stable for the case $\hat{a}^+ < 0$ such that $\mathbf{A}^+ = \mathbf{0}$. This part of the proof is trivial: For $\hat{a}^+ < 0$, the Jacobian $\mathbf{J}(\mathbf{A}^+)$ corresponds to the negative identity matrix, which only has the negative eigenvalue $\lambda = -1$. Hence, the proof is concluded.

Appendix D. Proof of Theorem 4

Appendix D.1. NBS attribute

To characterize the NBS attribute a°, we first require an understanding of the aggregate profit of ISPs on the path:

$$\sum_{n \in r} \pi_n(\mathbf{A}) = d \frac{\alpha_1 \sum_{n \in r} a_n + \alpha_0}{1 + \alpha_1 \sum_{n \in r} a_n + \alpha_0} (I(\rho - \phi_0) - \phi_1 \sum_{n \in r} a_n) - \gamma_1 \sum_{n \in r} a_n - I\gamma_0 \tag{D.1}$$

The aggregate profit can thus be considered a function of the sum a_r of ISP attributes on path, i.e., $a_r := \sum_{n \in r} a_n$. By Theorem 1, the unrestricted optimal attribute sum a_r° is $a_r^\circ = \max(0, \hat{a}_r^\circ)$, where:

$$\hat{a}_r^\circ = \frac{\sqrt{\frac{d}{d\phi_1 + \gamma_1} \left(\phi_1(1 + \alpha_0) + I\alpha_1(\rho - \phi_0)\right)} - (1 + \alpha_0)}{\alpha_1}. \tag{D.2}$$

Clearly, the Nash bargaining attributes $\{a_n^\circ\}_{n \in r}$ must sum to a_r° in order to be optimal. Moreover, the Nash bargaining solution is fair for the cooperating entities, requiring equal profit for all ISPs in our context. As a result, the Nash bargaining solution stipulates a single attribute value a°, which is adopted by all ISPs. This NBS attribute a° is $a^\circ = \max(0, \hat{a}^\circ)$, where $\hat{a}^\circ = \hat{a}_r^\circ/I$.

Appendix D.2. Equilibrium attribute

The equilibrium attribute a^+ is defined as in Theorem 2, but can be considerably simplified for the case $Q = 1$. In particular, the equilibrium attribute a^+ for $Q = 1$ is $a^+ = \max(0, \hat{a}^+)$, where:

$$\hat{a}^+ = \frac{\sqrt{T_2^2 - 4T_1T_3} - T_2}{2T_1}, \tag{D.3}$$

$$T_1 = I^2\alpha_1^2, \tag{D.4}$$

$$T_2 = 2I\alpha_1(1+\alpha_0) - \frac{d}{d\phi_1+\gamma_1}\alpha_1\phi_1(I-1), \text{ and}$$

$$T_3 = (1+\alpha_0)^2 - \frac{d}{d\phi_1+\gamma_1}\big(\phi_1(1+\alpha_0) + \alpha_1(\rho-\phi_0)\big). \tag{D.5}$$

Appendix D.3. Comparison of attributes

We show that $a^+ \leq a^\circ$ by showing that $\hat{a}^+ \leq \hat{a}^\circ$. This inequality can be rewritten as

$$\frac{\sqrt{T_2^2 - 4T_1T_3} - T_2}{2T_1} \leq a^\circ \iff T_2\hat{a}^\circ \geq -(T_1 a^{\circ 2} + T_3) \tag{D.6}$$

The two sides of the second inequality in Eq. (D.6) expand to:

$$\begin{aligned} T_2\hat{a}^\circ = & -2(1+\alpha_0)^2 + 2(1+\alpha_0)\sqrt{\frac{d}{d\phi_1+g_1}(\phi_1(1+\alpha_0) + I\alpha_1(r-\phi_0))} \\ & - \frac{d}{d\phi_1+g_1}(I-1)\alpha_1\phi_1\hat{a}^\circ \end{aligned} \tag{D.7}$$

$$\begin{aligned} -(T_1a^{\circ 2} + T_3) = & -2(1+\alpha_0)^2 + 2(1+\alpha_0)\sqrt{\frac{d}{d\phi_1+g_1}(\phi_1(1+\alpha_0) + I\alpha_1(r-\phi_0))} \\ & - \frac{d}{d\phi_1+g_1}(I-1)\alpha_1(r-\phi_0) \end{aligned} \tag{D.8}$$

Since these terms lend themselves to considerable simplification, Eq. (D.6) reduces to:

$$\rho - \phi_1\hat{a}^\circ - \phi_0 \geq 0 \tag{D.9}$$

Interestingly, $\hat{a}^\circ$ is guaranteed to satisfy this inequality. To see why, assume the opposite for the sake of contradiction: $\rho - \phi_1\hat{a}^\circ - \phi_0 < 0$. If $\phi_1 = 0$, this inequality conflicts with the model assumption $\rho - \phi_0 \geq 0$. If $\phi_1 > 0$, the same model assumption indicates that $\hat{a}^\circ > (\rho - \phi_0)/\phi_1 \geq 0$. Hence, the profit function of any ISP n is negative:

$$\pi_n(\hat{\mathbf{A}}^\circ) = \underbrace{d\frac{\alpha_1 I\hat{a}^\circ + \alpha_0}{1 + \alpha_1 I\hat{a}^\circ + \alpha_0}}_{>0}\underbrace{(\rho - \phi_1\hat{a}^\circ - \phi_0)}_{<0}\underbrace{-\gamma_1\hat{a}^\circ - \gamma_0}_{\leq 0} \tag{D.10}$$

However, this negative profit could be strictly improved by choosing the lower attribute value $a' = (\rho - \phi_0)/\phi_1 < \hat{a}^\circ$. This profit improvement is a contradiction to the character of $\hat{a}^\circ$ as the profit-optimizing attribute value. Hence, Eq. (D.9) holds, and therefore also the proposition $\hat{a}^+ \leq \hat{a}^\circ$ holds. This insight concludes the proof.

Appendix E. Proof of Theorem 5

Appendix E.1. Equilibrium for competitive network $\mathcal{N}_2$

We begin the proof by characterizing the equilibrium for the competitive network $\mathcal{N}_2$, in which every ISP n optimizes the following profit function π_n:

$$\pi_n(a_n) = d' \left(\sum_{q=1}^{Q} \frac{v_{r(m_{q1},m_{q2},n)}}{1 + \sum_{r' \in R(m_{q1},m_{q2})} v_{r'}} \right) (\rho - \phi_1 a_n - \phi 0) - \gamma_1 a_n - \gamma_0 \tag{E.1}$$

where $r(m_{q1}, m_{q2}, n)$ denotes the unique path connecting (m_{q1}, m_{q2}) via ISP n. In the unrestricted equilibrium $\hat{\mathbf{A}}^+$, every ISP n has the optimal attribute value $\hat{a}_n^+$ given competitor attributes $\hat{\mathbf{A}}_{-n}$, which can be found by setting $\partial \pi_n / \partial a_n = 0$:

$$d' \left(\sum_{q=1}^{Q} \frac{\alpha_1 \left(1 + v_{-r(m_{q1},m_{q2},n)}\right)}{\left(1 + \sum_{r' \in R(m_{q1},m_{q2})} v_{r'}\right)^2} \right) (\rho - \phi_1 a_n - \phi 0) - d'\phi_1 \left(\sum_{q=1}^{Q} \frac{v_{r(m_{q1},m_{q2},n)}}{1 + \sum_{r' \in R(m_{q1},m_{q2})} v_{r'}} \right) - \gamma_1 = 0 \tag{E.2}$$

Since this equation is equivalent for every ISP n, the equilibrium $\hat{a}_n^+$ is identical for all ISPs n, i.e., $\hat{a}_n^+ = \hat{a}^+$. This simplification allows the following transformation of Eq. (E.3):

$$d'Q \frac{\alpha_1(1 + (Q-1)(I\alpha_1\hat{a}^+ + \alpha_0))}{(1 + Q(I\alpha_1\hat{a}^+ + \alpha_0))^2} (\rho - \phi_1 a_n - \phi 0) - d'Q\phi_1 \frac{I\alpha_1\hat{a}^+ + \alpha_0}{1 + Q(I\alpha_1\hat{a}^+ + \alpha_0)} - \gamma_1 = 0 \tag{E.3}$$

This equilibrium condition is identical to the equilibrium condition for a homogeneous parallel-path network with a single origin-destination pair and demand limit $d = d'Q$. Hence, the unrestricted equilibrium value $\hat{a}^+$ from Theorem 2 (with $d'Q$ substituted for d) also applies to the competitive network $\mathcal{N}_2$.

Appendix E.2. Equilibrium for competition-free network $\mathcal{N}_1$

Moreover, we note that a single sub-network (for one origin-destination pair) of the competition-free network $\mathcal{N}_1$ is equivalent to the network $\mathcal{N}_2$ for $Q = 1$. Since the identical, isolated sub-networks of the competition-free network $\mathcal{N}_1$ do not influence each other, the equilibrium attribute value $\hat{a}^+$ is thus equal in that whole network.

Appendix E.3. Comparison of equilibria

Hence, if $\hat{a}^+(Q)$ is considered the equilibrium attribute for the competitive network $\mathcal{N}_2$, we can prove the proposition $a^+(\mathcal{N}_2) \geq a^+(\mathcal{N}_1)$ for $Q \geq 1$ by showing $\hat{a}^+(Q) \geq \hat{a}^+(1)$ for $Q \geq 1$. To show this property, we solve the following inequality:

$$\hat{a}^+(Q) - \hat{a}^+(1) \geq 0 \tag{E.4}$$

$$\iff \frac{\sqrt{T_2(Q)^2 - 4T_1(Q)T_3(Q)} - T_2(Q)}{2T_1(Q)} - \hat{a}^+(1) \geq 0 \tag{E.5}$$

$$\iff \sqrt{T_2(Q)^2 - 4T_1(Q)T_3(Q)} \geq T_2(Q) + 2T_1(Q)\hat{a}^+(1) \tag{E.6}$$

$$\iff T_2(Q)^2 - 4T_1(Q)T_3(Q) \geq \left(T_2(Q) + 2T_1(Q)\hat{a}^+(1)\right)^2 \tag{E.7}$$

In Eq. (E.5), the equilibrium constituent terms T_1, T_2, and T_3 from Theorem 2 are considered functions of Q. In Eq. (E.6), the transformation is possible by the fact that $T_1(Q) > 0$ for $Q \geq 1$:

$$T_1(Q) = Q^2 I^2 \alpha_1^2 - \frac{Qd'}{Qd'\phi_1 + \gamma_1} (QI - 1)(Q-1) I\alpha_1^2\phi_1 > 0 \iff Q > \frac{d'\phi_1}{(d'\phi_1(I+1) + I\gamma_1)} \tag{E.8}$$

where the RHS in the last inequality is consistently below 1, and $T_1(Q) > 0$ thus holds for all $Q \geq 1$. Using lengthy rewriting, the inequality in Eq. (E.7) can then be transformed into the following inequality containing a quadratic equation of Q:

$$T_7 Q^2 + T_8 Q + T_9 \leq 0 \tag{E.9}$$

where

$$\begin{aligned} T_7 = {} & \hat{a}^+(1)^2 \alpha_1^2 I \left(I\gamma_1 + d\phi_1 (I+1)\right) + \hat{a}^+(1)\alpha_1 \left(2I\alpha_0\gamma_1 + (2I+1)\, d\phi_1\alpha_0 - d\,(\rho - \phi_0)\, I\alpha_1\right) && \text{(E.10)} \\ & + \alpha_0^2 \left(d\phi_1 + \gamma_1\right) - d\alpha_0\alpha_1 \left(\rho - \phi_0\right), \\ T_8 = {} & -\hat{a}^+(1)^2 dI\alpha_1^2\phi_1 + \hat{a}^+(1)\alpha_1 \left(2I\gamma_1 + d\phi_1(I+1) - d\phi_1\alpha_0 + d\,(\rho - \phi_0)\, I\alpha_1\right) && \text{(E.11)} \\ & + 2\alpha_0\gamma_1 + (\alpha_0 - 1)\, d\alpha_1 \left(\rho - \phi_0\right) + d\alpha_0\phi_1, \text{ and} \\ T_9 = {} & \gamma_1. && \text{(E.12)} \end{aligned}$$

To solve Eq. (E.9), we make use of the following two properties.

- $Q = 1$ is a root of of $\hat{a}^+(Q) - \hat{a}^+(1)$, which implies:

$$T_7 + T_8 + T_9 = 0 \iff T_7 + T_8 = -T_9. \tag{E.13}$$

- The inspection of T_7 yields the following insight, which we derived by means of the symbolic algebra system in MATLAB:

$$T_7 \leq \lim_{\substack{d,\alpha_0 \\ \to 0}} T_7 = \gamma_1 = T_9 \tag{E.14}$$

Given the lower root $\underline{Q}$ and the higher root $\overline{Q}$ of the quadratic function in Eq. (E.9) (which are guaranteed to exist at least at $Q = 1$ and are identical if $T_7 = 0$), the inequality is solved by the following Q:

$$Q \in \begin{cases} [\underline{Q},\ \overline{Q}] & \text{① if } T_7 > 0 \\ (-\infty,\ \underline{Q}] \cup [\overline{Q},\ \infty) & \text{② if } T_7 < 0 \\ (-\infty,\ \underline{Q}] & \text{③ if } T_7 = 0 \wedge T_8 > 0 \\ [\overline{Q},\ \infty) & \text{④ if } T_7 = 0 \wedge T_8 < 0 \\ (-\infty,\ \infty) & \text{⑤ if } T_7 = 0 \wedge T_8 = 0 \wedge T_9 \leq 0 \\ \emptyset & \text{⑥ if } T_7 = 0 \wedge T_8 = 0 \wedge T_9 > 0 \end{cases} \tag{E.15}$$

This area of Q (leading to non-positive values of the quadratic function in Eq. (E.9)) includes $[1, \infty)$ in all cases:

1. $T_7 \neq 0$. (Eq. (E.15)① and ②): For $T_7 \neq 0$, the property in Eq. (E.13) facilitates finding the solutions $(\underline{Q},\ \overline{Q})$:

$$(\underline{Q},\ \overline{Q}) = \frac{-T_8 \pm \sqrt{T_8^2 - 4T_7T_9}}{2T_7} = \frac{-T_8 \pm \sqrt{(2T_7 + T_8)^2}}{2T_7} = \frac{-T_8 \pm (2T_7 + T_8)}{2T_7} \tag{E.16}$$

   (a) $T_7 > 0$ (Eq. (E.15)①): For $T_7 > 0$, we note that

$$2T_7 + T_8 \overset{\text{Eq. (E.13)}}{=} T_7 - \gamma_1 \overset{\text{Eq. (E.14)}}{\leq} 0. \tag{E.17}$$

   Hence, the solutions from Eq. (E.16) are:

$$\underline{Q} = \frac{-T_8 + (2T_7 + T_8)}{2T_7} = 1 \qquad \overline{Q} = \frac{-T_8 - (2T_7 + T_8)}{2T_7} = \frac{\gamma_1}{T_7} \geq 1 \tag{E.18}$$

   where the higher solution $\overline{Q}$ is spurious and has been introduced by the squaring operation in Eq. (E.7).

(b) $T_7 < 0$: (Eq. (E.15)②): For $T_7 < 0$ (Eq. (E.15)②), the solutions are as follows:

$$\underline{Q} = \frac{-T_8 - (2T_7 + T_8)}{2T_7} = \frac{\gamma_1}{T_7} < 0 \qquad \overline{Q} = \frac{-T_8 + (2T_7 + T_8)}{2T_7} = 1 \tag{E.19}$$

2. $T_7 = 0$ (Eq. (E.15)③–⑥): For $T_7 = 0$, it holds that $T_8 = -\gamma_1 - T_7 = -\gamma_1 \leq 0$ and $\overline{Q} = -T_9/T_8 = (-\gamma_1)/(-\gamma_1) = 1$.

   (a) $T_8 > 0$ (Eq. (E.15)③): The case $T_8 > 0$ thus cannot arise.

   (b) $T_8 < 0$ (Eq. (E.15)④): For $T_8 < 0$, the proposition clearly holds.

   (c) $T_8 = 0$ (Eq. (E.15)⑤ and ⑥): For $T_8 = 0$, the equality $T_7 + T_8 = -T_9$ from Eq. (E.13) implies $T_9 = 0$. Hence, the case in Eq. (E.15)⑤ always arises if $T_7 = T_8 = 0$, whereas the case in Eq. (E.15)⑥ never arises.

Since Eq. (E.4) thus always holds for $Q \geq 1$, the proposition is proven.

Appendix F. Proof of Theorem 6

To start the proof, we note that both the equilibrium attribute sum $a^+(\mathcal{N}_1)$ and the NBS attribute sum $a^\circ(\mathcal{N}_1)$ for the competition-free network are found by analyzing a single path, since the isolated sub-paths in the competition-free network do not influence each other. Hence, $a^+(\mathcal{N}_1)$ and $a^\circ(\mathcal{N}_1)$ are as in Theorem 4, which relates to the single-path context and thus states that $a^+(\mathcal{N}_1) \leq a^\circ(\mathcal{N}_1)$. Therefore, the interval $[a^+(\mathcal{N}_1), a^\circ(\mathcal{N}_1)]$ is never empty.

From the proof of Theorem 5, we know that the proposition $\pi^+(\mathcal{N}_2) \geq \pi^+(\mathcal{N}_1)$ is equivalent to the proposition $\Delta\pi = \pi(Q, a^+(\mathcal{N}_2)) - \pi(1, a^+(\mathcal{N}_1)) \geq 0$, where $\pi(Q, a)$ is defined as follows:

$$\pi(Q, a) = Qd' \frac{I\alpha_1 a + \alpha_0}{1 + Q(I\alpha_1 a + \alpha_0)} (\rho - \phi_1 a - \phi_0) - \gamma_1 a - \gamma_0. \tag{F.1}$$

Clearly, $\pi(Q, a^\circ(\mathcal{N}_1))$ is optimal for $Q = 1$, i.e., the NBS attribute sum is optimal in the competition-free network. Hence, it also holds that $\pi(1, a^+(\mathcal{N}_1)) \leq \pi(1, a^\circ(\mathcal{N}_1))$, i.e., the equilibrium profit in the competition-free network is generally sub-optimal. Moreover, since $\pi(Q, a)$ is consistently concave in a in the relevant regions, the assumption $a^+(\mathcal{N}_2) \in [a^+(\mathcal{N}_1), a^\circ(\mathcal{N}_1)]$ implies

$$\pi(1, a^+(\mathcal{N}_1)) \leq \pi(1, a^+(\mathcal{N}_2)). \tag{F.2}$$

Given Eq. (F.2), we can lower bound the profit difference:

$$\Delta\pi = \pi(Q, a^+(\mathcal{N}_2)) - \pi(1, a^+(\mathcal{N}_1)) \geq \pi(Q, a^+(\mathcal{N}_2)) - \pi(1, a^+(\mathcal{N}_2)) =: \underline{\Delta\pi} \tag{F.3}$$

Hence, if $\underline{\Delta\pi} \geq 0$ holds, the proof proposition $\Delta\pi \geq 0$ follows. At this point, we also note that $a^+(\mathcal{N}_2) \in [a^+(\mathcal{N}_1), a^\circ(\mathcal{N}_1)]$ is only a sufficient, but not a necessary condition for $\underline{\Delta\pi} \geq 0$; hence, profit increases might also happen if $a^+(\mathcal{N}_2) \notin [a^+(\mathcal{N}_1), a^\circ(\mathcal{N}_1)]$.

We can reformulate the lower bound $\underline{\Delta\pi}$ on the profit difference as follows:

$$\begin{aligned}
\underline{\Delta\pi} &= \pi(Q, a^+(\mathcal{N}_2)) - \pi(1, a^+(\mathcal{N}_2)) \\
&= d' \left(\frac{Q(I\alpha_1 a^+(\mathcal{N}_2) + \alpha_0)}{1 + Q(I\alpha_1 a^+(\mathcal{N}_2) + \alpha_0)} - \frac{I\alpha_1 a^+(\mathcal{N}_2) + \alpha_0}{1 + I\alpha_1 a^+(\mathcal{N}_2) + \alpha_0} \right) \left(\rho - \phi_1 a^+(\mathcal{N}_2) - \phi_0\right) \\
&= d' \left(\frac{(Q-1)\left(I\alpha_1 a^+(\mathcal{N}_2) + \alpha_0\right)}{(1 + Q(I\alpha_1 a^+(\mathcal{N}_2) + \alpha_0))(1 + I\alpha_1 a^+(\mathcal{N}_2) + \alpha_0)} \right) \left(\rho - \phi_1 a^+(\mathcal{N}_2) - \phi_0\right)
\end{aligned} \tag{F.4}$$

Given $d' > 0$ and $Q \geq 1$, the first and second factor of $\underline{\Delta\pi}$ in Eq. (F.4) are non-negative. Hence, $\underline{\Delta\pi} \geq 0$ is equivalent to $\rho - \phi_1 a^+(\mathcal{N}_2) - \phi_0 \geq 0$. This latter condition also always holds, which is demonstrable by contradiction. Let $\rho - \phi_1 a^+(\mathcal{N}_2) - \phi_0 < 0 \iff a^+(\mathcal{N}_2) > (\rho - \phi_0)/\phi_1$, which makes the minuend in $\pi(1, a^+(\mathcal{N}_2))$ negative (cf. Eq. (F.1)). In that case, all $a' > a^+(\mathcal{N}_2)$ would lead to lower profit $\pi(1, a')$. This observation contradicts the optimality of the NBS attribute sum $a^\circ(\mathcal{N}_1)$ regarding $\pi(1, a')$, as $a^\circ(\mathcal{N}_1) \geq a^+(\mathcal{N}_2)$.

Hence, since $\underline{\Delta\pi} \geq 0$, it holds that $\Delta\pi \geq 0$ and the theorem proposition follows.

Appendix G. Proof of Theorem 7

In order to be a Nash bargaining solution, the attribute values $\mathbf{A}^\circ$ should both optimize the aggregate profit function $\pi(\mathbf{A}) = \sum_{n\in r} \pi_n(\mathbf{A})$, and create a maximally equitable profit distribution across the ISPs $n \in r$. This maximum fairness is achieved by optimizing the Nash bargaining product, i.e.:

$$\mathbf{A}^\circ = \arg\max_{\mathbf{A}\in\mathbb{R}_{\geq 0}} \Pi_{n\in r} \pi_n(\mathbf{A}) \tag{G.1}$$

This optimization of the Nash bargaining product must be performed subject to the constraints in Theorem 7 that are associated with optimal aggregate profit. In the following, we characterize this aggregate-profit function, and show that the conditions stated in Theorem 7 are both sufficient and necessary in order for $\mathbf{A}^\circ$ to satisfy aggregate-profit optimality.

Appendix G.1. Aggregate-profit function

The aggregate profit $\pi(\mathbf{A})$ in our setting is:

$$\pi(\mathbf{A}) = \sum_{n\in r} \pi_n(\mathbf{A}) = d\frac{v_r(\mathbf{A})}{1 + v_r(\mathbf{A})} \left(\sum_{n\in r} \rho_n - \phi_{n0}\right) - \sum_{n\in r} \left(\sum_{k\in K} \gamma_{nk} a_{nk} + \gamma_{n0}\right) \tag{G.2}$$

This aggregate-profit function has the following first and second derivative in any a_{nk}:

$$\frac{\partial}{\partial a_{nk}} \pi(\mathbf{A}) = d\frac{\alpha_{nk}}{(1 + v_r(\mathbf{A}))^2} \left(\sum_{n\in r} \rho_n - \phi_{n0}\right) - \gamma_{nk} \tag{G.3}$$

$$\frac{\partial^2}{\partial a_{nk}^2} \pi(\mathbf{A}) = -d\frac{2\alpha_{nk}^2}{(1 + v_r(\mathbf{A}))^3} \left(\sum_{n\in r} \rho_n - \phi_{n0}\right) \tag{G.4}$$

As the second derivative is non-positive for all $\mathbf{A} \in \mathbb{R}_{\geq 0}$, the aggregate-profit function is consistently concave in any a_{nk} on the valid domain $\mathbb{R}_{\geq 0}$. Therefore, if the first derivative $\partial/\partial a_{nk}\, \pi(\mathbf{A})$ is negative for any a_{nk}, all reductions of a_{nk} increase aggregate profit, and all increases of a_{nk} reduce the aggregate profit (The reverse holds for a positive first derivative). This condition on the first derivative is equivalent to the following condition, which is central for the proof:

$$\forall n \in r, k \in K. \quad \frac{\partial}{\partial a_{nk}} \pi(\mathbf{A}) < 0 \iff v_r(\mathbf{A}) > \sqrt{\frac{\alpha_{nk}}{\gamma_{nk}}} \sqrt{d \sum_{n\in r} (\rho_n - \phi_{n0})} - 1 \tag{G.5}$$

Appendix G.2. Sufficiency of conditions

After this characterization of the aggregate-profit function, we now demonstrate that the conditions in Theorem 7 are *sufficient*, i.e., any $\mathbf{A}^\circ$ with the conditions optimizes the aggregate profit. Sufficiency can be demonstrated by performing the following case distinction:

1. $\forall (n, k) \in r \times K.\ \alpha_{r0} > \sqrt{\frac{\alpha_{nk}}{\gamma_{nk}}} \sqrt{d \sum_{n\in r} (\rho_n - \phi_{n0})} - 1$
   According to Theorem 7, all optimal attribute values $\mathbf{A}^\circ$ must be 0 in this case:

$$v_r(\mathbf{A}^\circ) = v_r^\circ \overset{\text{Eq. (22)}}{=} \alpha_{r0} \iff \mathbf{A}^\circ = \mathbf{0} \tag{G.6}$$

Moreover, the first derivative $\partial/\partial a_{nk}\ \pi(\mathbf{A}^\circ)$ for all $n \in r$, $k \in K$ must be negative:

$$\forall n \in r, k \in K.\ v_r(\mathbf{A}^\circ) = \alpha_{r0} > \sqrt{\frac{\alpha_{nk}}{\gamma_{nk}}}\sqrt{d\sum_{n\in r}(\rho_n - \phi_{n0})} - 1 \overset{\text{Eq. (G.5)}}{\iff} \frac{\partial}{\partial a_{nk}}\pi(\mathbf{A}) < 0 \tag{G.7}$$

Hence, the only way to further increase the aggregate profit π would be by reductions in any a_{nk}. However, since every $a_{nk} = 0$, such reductions are not possible given the restricted domain $\mathbb{R}_{\geq 0}$. Hence, $\mathbf{A}^\circ = \mathbf{0}$ is optimal.

2. $\exists (n,k) \in r \times K.\ \alpha_{r0} \leq \sqrt{\frac{\alpha_{nk}}{\gamma_{nk}}}\sqrt{d\sum_{n\in r}(\rho_n - \phi_{n0})} - 1$
In this case, the attributes $(n^\circ, k^\circ) \in r \times K$ with the maximal ratio $\alpha_{n^\circ k^\circ}/\gamma_{n^\circ k^\circ}$ play a special role according to Theorem 7. We denote the set of these attributes by K_r°:

$$K_r^\circ = \left\{ (n^\circ, k^\circ) \mid (n^\circ, k^\circ) = \arg\max_{(n,k)\in r\times K} \frac{\alpha_{nk}}{\gamma_{nk}} \right\}. \tag{G.8}$$

This maximal ratio also determines the optimal path valuation $v_r(\mathbf{A}^\circ)$ according to Theorem 7:

$$\forall (n^\circ, k^\circ) \in K_r^\circ.\ v_r(\mathbf{A}^\circ) = \sqrt{\frac{\alpha_{n^\circ k^\circ}}{\gamma_{n^\circ k^\circ}}}\sqrt{d\sum_{n\in r}(\rho_n - \phi_{n0})} - 1 \tag{G.9}$$

In contrast, for all attributes $(n^\oslash, k^\oslash) \notin K_r^\circ$, the following condition holds:

$$\begin{aligned}\forall (n^\oslash, k^\oslash) \notin K_r^\circ.\ v_r(\mathbf{A}^\circ) = \sqrt{\frac{\alpha_{n^\circ k^\circ}}{\gamma_{n^\circ k^\circ}}}\sqrt{d\sum_{n\in r}(\rho_n - \phi_{n0})} - 1 > \sqrt{\frac{\alpha_{n^\oslash k^\oslash}}{\gamma_{n^\oslash k^\oslash}}}\sqrt{d\sum_{n\in r}(\rho_n - \phi_{n0})} - 1 \\ \iff \frac{\partial}{\partial a_{n^\oslash k^\oslash}}\pi(\mathbf{A}) < 0\end{aligned} \tag{G.10}$$

Hence, the only way to increase the aggregate profit π would be by reductions in any $a_{n^\oslash k^\oslash}$ for $(n^\oslash, k^\oslash) \notin K_r^\circ$. However, since $a_{n^\oslash k^\oslash} = 0\ \forall (n^\oslash, k^\oslash) \notin K_r^\circ$ by Theorem 7, such reductions are not possible, and hence $\mathbf{A}^\circ$ is optimal.

Appendix G.3. Necessity of conditions

After demonstrating that the conditions in Theorem 7 are sufficient for optimal aggregate profit, we now demonstrate that the conditions are also *necessary*, i.e., no choice of attribute values $\mathbf{A}^\circ$ that violates these conditions can be optimal. For the sake of contradiction, we assume that some attribute values $\mathbf{A}^\circ$ are optimal while satisfying the following conditions:

$$\exists (n^\oslash, k^\oslash) \notin K_r^\circ.\ a^\circ_{n^\oslash k^\oslash} > 0. \tag{G.11}$$

A contradiction can be produced in all cases of the following case distinction:

1. $\forall (n,k) \in r \times K.\ \alpha_{r0} > \sqrt{\frac{\alpha_{nk}}{\gamma_{nk}}}\sqrt{d\sum_{n\in r}(\rho_n - \phi_{n0})} - 1$
Since $a^\circ_{n^\oslash k^\oslash} > 0$ for the fixed attribute $(n^\oslash, k^\oslash)$, the optimal path valuation $v_r(\mathbf{A}^\circ)$ exceeds α_{r0}, and hence:

$$\forall (n,k) \in r \times K. \quad v_r(\mathbf{A}^\circ) > \alpha_{r0} > \sqrt{\frac{\alpha_{nk}}{\gamma_{nk}}}\sqrt{d\sum_{n\in r}(\rho_n - \phi_{n0})} - 1 \implies \frac{\partial}{\partial a_{n^\oslash k^\oslash}}\pi(\mathbf{A}) < 0 \tag{G.12}$$

The aggregate profit can thus be increased by reducing $a^\circ_{n^\oslash k^\oslash}$, and such a reduction is also possible since $a_{n^\oslash k^\oslash} > 0$. Hence, the attribute values $\mathbf{A}^\circ$ are not optimal, which causes a contradiction.

2. $\exists (n,k) \in r \times K.\ \alpha_{r0} \leq \sqrt{\frac{\alpha_{nk}}{\gamma_{nk}}}\sqrt{d\sum_{n\in r}(\rho_n - \phi_{n0})} - 1$
   In that case, we perform a sub-case distinction on the value of $v_r(\mathbf{A}^\circ)$:
   (a) $v_r(\mathbf{A}^\circ) \leq \sqrt{\frac{\alpha_{n^\oslash k^\oslash}}{\gamma_{n^\oslash k^\oslash}}}\sqrt{d\sum_{n\in r}(\rho_n - \phi_{n0})} - 1$
   Since $a_{n^\oslash k^\oslash}/\gamma_{n^\oslash k^\oslash} < a_{n^\circ k^\circ}/\gamma_{n^\circ k^\circ}$, it follows that:

$$v_r(\mathbf{A}^\circ) < \sqrt{\frac{\alpha_{n^\circ k^\circ}}{\gamma_{n^\circ k^\circ}}}\sqrt{d\sum_{n\in r}(\rho_n - \phi_{n0})} - 1 \iff \frac{\partial}{\partial a_{n^\circ k^\circ}}\pi(\mathbf{A}) > 0, \tag{G.13}$$

   which implies that the profit can be increased by increasing the value of $a_{n^\circ k^\circ}$ $\forall (n,k) \in K_r^\circ$, which contradicts the assumption that $\mathbf{A}^\circ$ is optimal.
   (b) $v_r(\mathbf{A}^\circ) > \sqrt{\frac{\alpha_{n^\oslash k^\oslash}}{\gamma_{n^\oslash k^\oslash}}}\sqrt{d\sum_{n\in r}(\rho_n - \phi_{n0})} - 1$
   This condition implies:

$$\frac{\partial}{\partial a_{n^\oslash k^\oslash}}\pi(\mathbf{A}) < 0. \tag{G.14}$$

   Hence, the profit can be increased by reducing $a_{n^\oslash k^\oslash}$, which is possible given $a_{n^\oslash k^\oslash} > 0$. Therefore, we again produce a contradiction to the optimality of $\mathbf{A}^\circ$.

We have thus identified the conditions on $\mathbf{A}^\circ$ that are sufficient and necessary for optimal aggregate profit. Thereby, the theorem is proven.

Appendix H. Proof of Theorem 8

Since the equilibrium conditions in Theorem 8 are highly similar to the optimality conditions in Theorem 7, the proof of Theorem 8 is analogous to the proof of Theorem 7. The proof is analogous because the derivatives of the individual profit functions π_n have equivalent properties to the derivatives of the aggregate-profit function π from Eq. (G.2). In particular, every first derivative satisfies:

$$\forall n \in r, k \in K. \quad \frac{\partial \pi_n(\mathbf{A})}{\partial a_{nk}} < 0 \iff v_r(\mathbf{A}) > \sqrt{\frac{\alpha_{nk}}{\gamma_{nk}}}\sqrt{d(\rho_n - \phi_{n0})} - 1 \tag{H.1}$$

Moreover, every individual profit function π_n is consistently concave in any relevant attribute a_{nk}:

$$\forall n \in r, k \in K.\ \frac{\partial^2}{\partial a_{nk}^2}\pi_n(\mathbf{A}) = -d\frac{2\alpha_{nk}^2}{(1 + v_r(\mathbf{A}))^3}(\rho_n - \phi_{n0}) \tag{H.2}$$

Building on these properties, the equilibrium conditions can be shown to be sufficient and necessary analogously to Theorem 7.

Appendix I. Proof of Theorem 9

In order to show that $v_r(\mathbf{A}^+) \leq v_r(\mathbf{A}^\circ)$, it is enough to show that:

$$\sqrt{\frac{\alpha_{n^+k^+}}{\gamma_{n^+k^+}}}\sqrt{d(\rho_{n^+} - \phi_{n^+0})} - 1 \leq \sqrt{\frac{\alpha_{n^\circ k^\circ}}{\gamma_{n^\circ k^\circ}}}\sqrt{d\sum_{n\in r}(\rho_n - \phi_{n0})} - 1 \tag{I.1}$$

where $(n^+, k^+) \in K_r^+$ and $(n^\circ, k^\circ) \in K_r^\circ$. This inequality can be transformed into the following form:

$$\sqrt{\frac{\alpha_{n^+k^+}}{\gamma_{n^+k^+}}}\sqrt{\frac{d(\rho_{n^+} - \phi_{n^+0})}{d\sum_{n\in r}(\rho_n - \phi_{n0})}} \leq \sqrt{\frac{\alpha_{n^\circ k^\circ}}{\gamma_{n^\circ k^\circ}}}. \tag{I.2}$$

Thanks to the following two insights, this inequality is always satisfied:

$$\sqrt{\frac{d(\rho_{n^+} - \phi_{n^+0})}{d\sum_{n\in r}(\rho_n - \phi_{n0})}} \leq 1 \qquad \frac{\alpha_{n^+k^+}}{\gamma_{n^+k^+}} \leq \frac{\alpha_{n^\circ k^\circ}}{\gamma_{n^\circ k^\circ}} = \max_{(n,k)\in r\times K} \frac{\alpha_{nk}}{\gamma_{nk}} \tag{I.3}$$

Hence, the theorem holds.

Appendix J. Proof of Theorem 10

Appendix J.1. Analogy to Theorem 8

To start off, we once more characterize the derivatives of the individual-profit functions π_n for any attribute value a_{nk}:

$$\forall n \in N, k \in K, r = r(n). \quad \frac{\partial \pi_n(\mathbf{A})}{\partial a_{nk}} = d\frac{\alpha_{nk}\,(1 + v_s(\mathbf{A}))}{(1 + v_r(\mathbf{A}) + v_s(\mathbf{A}))^2}(\rho_n - \phi_{n0}) - \gamma_{nk} \tag{J.1}$$

$$\frac{\partial^2 \pi_n(\mathbf{A})}{\partial a_{nk}^2} = -d\frac{2\alpha_{nk}^2\,(1 + v_s(\mathbf{A}))}{(1 + v_r(\mathbf{A}) + v_s(\mathbf{A}))^3}(\rho_n - \phi_{n0}) \tag{J.2}$$

For better readability, this proof denotes the alternative path $\overline{r}$ to path r by s.

Since the second derivative is never positive, every profit function π_n is consistently concave in the attribute values controlled by ISP n. Hence, a negative first derivative $\partial\pi(\mathbf{A})/\partial a_{nk} < 0$ indicates that a_{nk} must be reduced if the profit is to be increased. The case of a negative first derivative in a_{nk} can be expressed as follows (for $r = r(n)$):

$$\frac{\partial \pi_n(\mathbf{A})}{\partial a_{nk}} < 0 \iff v_r(\mathbf{A}) > \sqrt{\frac{\alpha_{nk}}{\gamma_{nk}}}\sqrt{d\,(\rho_n - \phi_{n0})}\sqrt{1 + v_s(\mathbf{A})} - (1 + v_s(\mathbf{A})) \tag{J.3}$$

When thinking of Eq. (J.3) as an extension of Eq. (G.5) with v_s as a fixed quantity, an analogous proof to the proof of Theorem 7 can be performed. The extension by fixed v_s does not change the finding that only attributes $(n^\circ, k^\circ) \in K_r^\circ$ might have non-zero values in equilibrium. However, the extension by v_s changes the equilibrium path valuation v_r^+ from Eq. (25) to:

$$v_r^+ = \max\left(\alpha_{r0},\ \sqrt{\frac{\alpha_{n^\circ k^\circ}}{\gamma_{n^\circ k^\circ}}}\sqrt{d\,(\rho_{n^\circ} - \phi_{n^\circ 0})}\sqrt{1 + v_s} - (1 + v_s)\right). \tag{J.4}$$

Crucially, this condition simultaneously holds for both paths r and s in a two-path scenario, creating an interdependence of the equilibrium path valuations:

$$\begin{aligned}\forall r \in R. \quad v_r^+ &= \max\left(\alpha_{r0},\ \sqrt{\frac{\alpha_{n^\circ k^\circ}}{\gamma_{n^\circ k^\circ}}}\sqrt{d\,(\rho_{n^\circ} - \phi_{n^\circ 0})}\sqrt{1 + v_s^+} - \left(1 + v_s^+\right)\right) \\ &= \max\left(\alpha_{r0},\ \psi_r\sqrt{d}\sqrt{1 + v_s^+} - \left(1 + v_s^+\right)\right) = \max\left(\alpha_{r0},\ \hat{v}_r^*\left(v_s^+\right)\right),\end{aligned} \tag{J.5}$$

where $\hat{v}_r^*(v_s)$ is the *unrestricted best-response valuation* for path r given competing-path valuation v_s. The *characteristic ratio* ψ_r is reflected in Eq. (29).

The remainder of the proof illustrates how to derive the equilibrium path valuations v_r^+ and v_s^+.

Appendix J.2. Unrestricted equilibrium $\hat{v}_r^+$

Considering a relaxed setting in which the constraint $\mathbf{A}^+ \in \mathbb{R}_{\geq 0}^{|N|\times|K|} \iff v_r^+ \geq \alpha_{r0}$ is ignored, the unrestricted equilibrium path valuations $\hat{v}_r^+$ satisfy the following system of equations:

$$\forall r \in R. \quad \hat{v}_r^+ = \hat{v}_r^*(\hat{v}_s^+). \tag{J.6}$$

In this relaxed setting, this system of two equations can be conventionally solved for the unrestricted equilibrium path valuations $\hat{v}_r^+$, $r \in R$, resulting in the unique solution denoted in Eq. (28).

Moreover, we make the following important observation:

$$\hat{v}_r^*\left(\hat{v}_s^*(v_r)\right) \leq v_r \iff v_r \geq \hat{v}_r^+. \tag{J.7}$$

Appendix J.3. Restricted equilibrium v_r^+

We now rely on the equilibrium gained by relaxation to characterize the equilibrium under the re-introduced constraint $v_r^+ \geq \alpha_{r0}$ $\forall r \in R$. In particular, we want to show that the calculation provided in Theorem 10 is correct:

$$\forall r \in R. \quad v_r^+ = \max\left(\alpha_{r0},\ \hat{v}_r^*\left(\max\left(\alpha_{s0}, \hat{v}_s^+\right)\right)\right) \tag{J.8}$$

In other words, v_r^+ as calculated by Eq. (J.8) should satisfy the equilibrium conditions on v_r^+ in Eq. (J.5). To satisfy these conditions, we consider all cases regarding $\hat{v}_r^+$ and $\hat{v}_s^+$:

1. $\hat{v}_r^+ \geq \alpha_{r0}$:

   (a) $\hat{v}_s^+ \geq \alpha_{s0}$: In that case, Eq. (J.8) suggests that

$$\begin{aligned} v_r^+ &\overset{1.(a)}{=} \max\left(\alpha_{r0}, \hat{v}_r^*\left(\hat{v}_s^+\right)\right) \overset{(J.6)}{=} \max\left(\alpha_{r0}, \hat{v}_r^+\right) \overset{1.}{=} \hat{v}_r^+, \text{ and} \\ v_s^+ &\overset{1.}{=} \max\left(\alpha_{s0}, \hat{v}_s^*\left(\hat{v}_r^+\right)\right) \overset{(J.6)}{=} \max\left(\alpha_{s0}, \hat{v}_s^+\right) \overset{1.(a)}{=} \hat{v}_s^+. \end{aligned} \tag{J.9}$$

   These values satisfy the equilibrium conditions in Eq. (J.5):

$$v_r^+ \overset{(J.5)}{=} \max\left(\alpha_{r0},\ \hat{v}_r^*\left(v_s^+\right)\right) \overset{(J.9)}{=} \max\left(\alpha_{r0}, \hat{v}_r^*\left(\hat{v}_s^+\right)\right) \overset{(J.6)}{=} \max\left(\alpha_{r0}, \hat{v}_r^+\right) \overset{1.}{=} \hat{v}_r^+, \tag{J.10}$$

   and symmetrically for v_s^+.

   (b) $\hat{v}_s^+ < \alpha_{s0}$: In that case, Eq. (J.8) suggests that

$$\begin{aligned} v_r^+ &\overset{1.(b)}{=} \max\left(\alpha_{r0}, \hat{v}_r^*\left(\alpha_{s0}\right)\right) \\ v_s^+ &\overset{1.}{=} \max\left(\alpha_{s0}, \hat{v}_s^*\left(\hat{v}_r^+\right)\right) \overset{(J.6)}{=} \max\left(\alpha_{s0}, \hat{v}_s^+\right) \overset{1.(b)}{=} \alpha_{s0}. \end{aligned} \tag{J.11}$$

   For that case, we perform another level of sub-case distinctions:

   i. $\hat{v}_r^*(\alpha_{s0}) \geq \alpha_{r0}$: In that case, Eq. (J.11) is simplified to

$$\begin{aligned} v_r^+ &\overset{(J.11)}{=} \max\left(\alpha_{r0}, \hat{v}_r^*\left(\alpha_{s0}\right)\right) \overset{1.(b).i}{=} \hat{v}_r^*(\alpha_{s0}) \\ v_s^+ &\overset{(J.11)}{=} \alpha_{s0} \end{aligned} \tag{J.12}$$

   Using Eq. (J.7) and case condition 1.(b), we can also deduce:

$$\hat{v}_s^*\left(\hat{v}_r^*(\alpha_{s0})\right) < \alpha_{s0}. \tag{J.13}$$

   Then, we can again verify that these values satisfy the equilibrium conditions from Eq. (J.5):

$$\begin{aligned} v_r^+ &\overset{(J.5)}{=} \max\left(\alpha_{r0},\ \hat{v}_r^*\left(v_s^+\right)\right) \overset{(J.12)}{=} \max\left(\alpha_{r0}, \hat{v}_r^*\left(\alpha_{s0}\right)\right) \overset{1.b.i}{=} \hat{v}_r^*(\alpha_{s0}) \\ v_s^+ &\overset{(J.5)}{=} \max\left(\alpha_{s0},\ \hat{v}_s^*\left(v_r^+\right)\right) \overset{(J.12)}{=} \max\left(\alpha_{s0}, \hat{v}_s^*\left(\hat{v}_r^*\left(\alpha_{s0}\right)\right)\right) \overset{(J.13)}{=} \alpha_{s0}. \end{aligned} \tag{J.14}$$

   ii. $\hat{v}_r^*(\alpha_{s0}) < \alpha_{r0}$: In that case, Eq. (J.11) is simplified to

$$\begin{aligned} v_r^+ &\overset{(J.11)}{=} \max\left(\alpha_{r0}, \hat{v}_r^*\left(\alpha_{s0}\right)\right) \overset{1.(b).ii}{=} \alpha_{r0} \\ v_s^+ &\overset{(J.11)}{=} \alpha_{s0} \end{aligned} \tag{J.15}$$

   Moreover, as proven in Appendix J.3.1, the current case implies

$$\hat{v}_s^*(\alpha_{r0}) \leq \alpha_{s0}. \tag{J.16}$$

   Based on these findings, the equilibrium conditions from Eq. (J.5) are satisfied:

$$\begin{aligned} v_r^+ &\overset{(J.5)}{=} \max\left(\alpha_{r0},\ \hat{v}_r^*\left(v_s^+\right)\right) \overset{(J.15)}{=} \max\left(\alpha_{r0}, \alpha_{r0}\right) = \alpha_{r0} \\ v_s^+ &\overset{(J.5)}{=} \max\left(\alpha_{s0},\ \hat{v}_s^*\left(v_r^+\right)\right) \overset{(J.15)}{=} \max\left(\alpha_{s0}, \hat{v}_s^*\left(\alpha_{r0}\right)\right) \overset{(J.16)}{=} \alpha_{s0}. \end{aligned} \tag{J.17}$$

2. $\hat{v}_r^+ < \alpha_{r0}$:

(a) $\hat{v}_s^+ \geq \alpha_{s0}$: This case is symmetric to case 1.(b).

(b) $\hat{v}_s^+ < \alpha_{s0}$: In that case, Eq. (J.8) suggests that

$$\begin{aligned} v_r^+ &\overset{2.(b)}{=} \max\left(\alpha_{r0}, \hat{v}_r^*\left(\alpha_{s0}\right)\right) \\ v_s^+ &\overset{2.}{=} \max\left(\alpha_{s0}, \hat{v}_s^*\left(\alpha_{r0}\right)\right) \end{aligned} \tag{J.18}$$

Hence, we again need to perform another level of sub-case distinctions:

i. $\hat{v}_r^*(\alpha_{s0}) \geq \alpha_{r0}$: For that case, we show in Appendix J.3.2 that

$$\hat{v}_s^*(\alpha_{r0}) \leq \alpha_{s0}. \tag{J.19}$$

Hence, Eq. (J.18) simplifies to:

$$\begin{aligned} v_r^+ &\overset{2.(b)}{=} \max\left(\alpha_{r0}, \hat{v}_r^*\left(\alpha_{s0}\right)\right) \overset{2.(b).i}{=} \hat{v}_r^*\left(\alpha_{s0}\right) \\ v_s^+ &\overset{2.}{=} \max\left(\alpha_{s0}, \hat{v}_s^*\left(\alpha_{r0}\right)\right) \overset{(J.19)}{=} \alpha_{s0} \end{aligned} \tag{J.20}$$

Moreover, using Eq. (J.7) and case condition 2.(b), we can again deduce:

$$\hat{v}_s^*\left(\hat{v}_s^*(\alpha_{r0})\right) < \alpha_{s0}. \tag{J.21}$$

Based on these findings, the equilibrium conditions from Eq. (J.5) are satisfied:

$$\begin{aligned} v_r^+ &\overset{(J.5)}{=} \max\left(\alpha_{r0},\ \hat{v}_r^*\left(v_s^+\right)\right) \overset{(J.20)}{=} \max\left(\alpha_{r0},\ \hat{v}_r^*\left(\alpha_{s0}\right)\right) \overset{2.(b).i}{=} \hat{v}_r^*\left(\alpha_{s0}\right) \\ v_s^+ &\overset{(J.5)}{=} \max\left(\alpha_{s0},\ \hat{v}_s^*\left(v_r^+\right)\right) \overset{(J.20)}{=} \max\left(\alpha_{s0}, \hat{v}_s^*\left(\hat{v}_r^*\left(\alpha_{s0}\right)\right)\right) \overset{(J.21)}{=} \alpha_{s0} \end{aligned} \tag{J.22}$$

ii. $\hat{v}_r^*(\alpha_{s0}) < \alpha_{r0}$: In that case, Eq. (J.18) simplifies to:

$$\begin{aligned} v_r^+ &\overset{2.(b)}{=} \max\left(\alpha_{r0}, \hat{v}_r^*\left(\alpha_{s0}\right)\right) \overset{2.(b).ii}{=} \alpha_{r0} \\ v_s^+ &\overset{2.}{=} \max\left(\alpha_{s0}, \hat{v}_s^*\left(\alpha_{r0}\right)\right) \end{aligned} \tag{J.23}$$

To further simplify Eq. (J.23), we perform another sub-case distinction:

A. $\hat{v}_s^*(\alpha_{r0}) \geq \alpha_{s0}$: This case is symmetric to case 2.(b).i.

B. $\hat{v}_s^*(\alpha_{r0}) < \alpha_{s0}$: In that case, Eq. (J.23) directly simplifies to:

$$\begin{aligned} v_r^+ &\overset{2.(b)}{=} \max\left(\alpha_{r0}, \hat{v}_r^*\left(\alpha_{s0}\right)\right) \overset{2.(b).ii}{=} \alpha_{r0} \\ v_s^+ &\overset{2.}{=} \max\left(\alpha_{s0}, \hat{v}_s^*\left(\alpha_{r0}\right)\right) \overset{2.(b).ii.B}{=} \alpha_{s0} \end{aligned} \tag{J.24}$$

Clearly, the equilibrium conditions from Eq. (J.5) are satisfied by these values:

$$\begin{aligned} v_r^+ &\overset{(J.5)}{=} \max\left(\alpha_{r0},\ \hat{v}_r^*\left(v_s^+\right)\right) \overset{(J.24)}{=} \max\left(\alpha_{r0}, \hat{v}_r^*\left(\alpha_{s0}\right)\right) \overset{2.(b).ii}{=} \alpha_{r0} \\ v_s^+ &\overset{(J.5)}{=} \max\left(\alpha_{s0}, \hat{v}_s^*\left(v_r^+\right)\right) \overset{(J.24)}{=} \max\left(\alpha_{s0}, \hat{v}_s^*\left(\alpha_{r0}\right)\right) \overset{2.(b).ii.B}{=} \alpha_{s0} \end{aligned} \tag{J.25}$$

Since the values calculated according to Eq. (J.8) always satisfy the restricted-equilibrium conditions from Eq. (J.5), the theorem is proven.

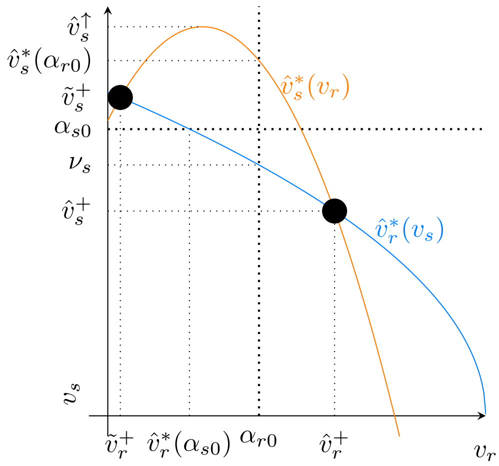

Figure J.6: Visualization of non-unique equilibrium $(\hat{v}_r^+, \hat{v}_s^+)$ in Appendix J.3.1.

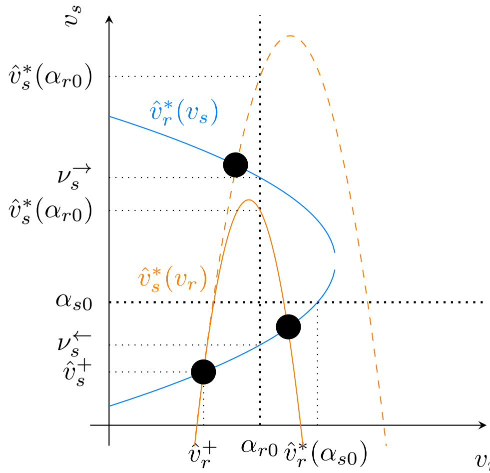

Figure J.7: Visualization of non-unique equilibrium $(\hat{v}_r^+, \hat{v}_s^+)$ in Appendix J.3.2.

Appendix J.3.1. Upper bound on $\hat{v}_s^*(\alpha_{r0})$ for case 1.(b).ii

We can show that case 1.(b).ii implies

$$\hat{v}_s^*(\alpha_{r0}) \leq \alpha_{s0}. \tag{J.26}$$

In particular, let us assume the opposite for the sake of contradiction, i.e., we assume

$$\hat{v}_s^*(\alpha_{r0}) > \alpha_{s0}. \tag{J.27}$$

For this proof, we first investigate the functions $\hat{v}_r^*$ and $\hat{v}_s^*$ more thoroughly, and then produce a contradiction by showing the existence of a second unrestricted equilibrium $(\tilde{v}_r^+, \tilde{v}_s^+) \neq (\hat{v}_r^+, \hat{v}_s^+)$. The proof idea is visualized in Fig. J.6.

$\hat{v}_r^*$. In case 1.(b).ii, we can more precisely characterize the function $\hat{v}_r^*$ based on the case conditions. In particular, we know that $\hat{v}_r^*$ evolves from value $\hat{v}_r^*(\hat{v}_s^+) = \hat{v}_r^+ \geq \alpha_{r0}$ (1.) at argument $\hat{v}_s^+$ down to value $\hat{v}_r^*(\alpha_{s0}) < \alpha_{r0}$ (1.(b).ii) at argument $\alpha_{s0} > \hat{v}_s^+$ (1.(b)). Hence, the intermediate-value theorem and the concavity of $\hat{v}_r^*$ stipulate that

$$\exists \nu_s \in [\hat{v}_s^+, \alpha_{s0}).\quad \hat{v}_r^*(\nu_s) = \alpha_{r0} \quad \text{and} \quad \forall v_s > \nu_s.\ \hat{v}_r^*(v_s) < \alpha_{r0}. \tag{J.28}$$

$\hat{v}_s^*$. The assumption in Eq. (J.27) suggests that $\hat{v}_s^*$ reaches a value above α_{s0} at argument α_{r0}. Hence, the maximum of $\hat{v}_s^*$ is also at least α_{s0}:

$$\hat{v}_s^{\uparrow} = \max_{v_r} \hat{v}_s^*(v_r) \geq \alpha_{s0} \overset{(J.28)}{>} \nu_s. \tag{J.29}$$

Contradiction. Based on these properties of $\hat{v}_r^*$ and $\hat{v}_s^*$, we now show that there exist $\tilde{v}_r^+ < \alpha_{r0}$ and $\tilde{v}_s^+ > \nu_s$, with the unrestricted-equilibrium properties $\tilde{v}_r^+ = \hat{v}_r^*(\tilde{v}_s^+)$ and $\tilde{v}_s^+ = \hat{v}_s^*(\tilde{v}_r^+)$. To verify the existence of these valuations, note that the value $\tilde{v}_s^+$ satisfies the condition:

$$\tilde{v}_s^+ = \hat{v}_s^*(\tilde{v}_r^+) = \hat{v}_s^*(\hat{v}_r^*(\tilde{v}_s^+)) \iff \hat{v}_s^{**}(\tilde{v}_s^+) := \tilde{v}_s^+ - \hat{v}_s^*(\hat{v}_r^*(\tilde{v}_s^+)) = 0. \tag{J.30}$$

We now evaluate the 'reflector' function $\hat{v}_s^{**}$ at two arguments v_s, namely ν_s and $\hat{v}_s^{\uparrow}$. For $v_s = \nu_s$, it holds that

$$\hat{v}_s^{**}(\nu_s) \overset{(J.30)}{=} \nu_s - \hat{v}_s^*(\hat{v}_r^*(\nu_s)) \overset{(J.28)}{=} \nu_s - \hat{v}_s^*(\alpha_{r0}) \overset{(J.27)}{<} \nu_s - \alpha_{s0} \overset{(J.28)}{<} 0 \tag{J.31}$$

For $v_s = \hat{v}_s^{\uparrow} > \nu_s$, it holds that

$$\hat{v}_s^{**}(\hat{v}_s^{\uparrow}) \overset{(J.30)}{=} \hat{v}_s^{\uparrow} - \hat{v}_s^*(\hat{v}_r^*(\hat{v}_s^{\uparrow})) \overset{(J.29)}{=} 0. \tag{J.32}$$

Since $\hat{v}_s^{**}$ is continuous, the intermediate-value theorem stipulates that a $\tilde{v}_s^+ \in (\nu_s, v_s^{\uparrow}]$ exists that satisfies $\hat{v}_s^{**}(\tilde{v}_s^+) = 0$. Then, since $\tilde{v}_s^+ > \nu_s$, it follows from Eq. (J.28) that

$$\tilde{v}_r^+ = \hat{v}_r^*(\tilde{v}_s^+) < \alpha_{r0}. \tag{J.33}$$

Since the unrestricted equilibrium valuations $(\hat{v}_r^+, \hat{v}_s^+)$ are unique, it must hold that $\hat{v}_r^+ = \tilde{v}_r^+$. However, the case condition $\hat{v}_r^+ \geq \alpha_{r0}$ conflicts with the derived condition $\tilde{v}_r^+ < \alpha_{r0}$ from Eq. (J.33). We thus arrive at a contradiction, which invalidates Eq. (J.27) and confirms Eq. (J.26).

Appendix J.3.2. Upper bound on $\hat{v}_s^*(\alpha_{r0})$ for case 2.(b).i

The following proof is similar in structure and goal as the proof in Appendix J.3.1, namely to prove $\hat{v}_s^*(\alpha_{r0}) \leq \alpha_{s0}$ by assuming

$$\hat{v}_s^*(\alpha_{r0}) > \alpha_{s0}. \tag{J.34}$$

The arguments of the proof are visualized in Fig. J.7.

$\hat{v}_r^*$. In case 2.(b).i, the strict concavity of $\hat{v}_r^*$, together with the knowledge of $\hat{v}_r^*(\alpha_{s0}) \geq \alpha_{r0}$ (2.(b).i), imply:

$$\begin{aligned} \exists\, \nu_s^{\leftarrow}, \nu_s^{\rightarrow}.\quad & \alpha_{s0} \in [\nu_s^{\leftarrow}, \nu_s^{\rightarrow}] \text{ and } \hat{v}_r^*(\nu_s^{\leftarrow}) = \alpha_{r0} \text{ and } \hat{v}_r^*(\nu_s^{\rightarrow}) = \alpha_{r0} \text{ and} \\ & \forall v_s \in [\nu_s^{\leftarrow}, \nu_s^{\rightarrow}].\ \hat{v}_r^*(v_s) \geq \alpha_{r0} \text{ and } \forall v_s \notin [\nu_s^{\leftarrow}, \nu_s^{\rightarrow}].\ \hat{v}_r^*(v_s) < \alpha_{r0}. \end{aligned} \tag{J.35}$$

$\hat{v}_s^*$. Given Eq. (J.34), we find the maximum of $\hat{v}_s^*$:

$$\hat{v}_s^{\uparrow} = \max_{v_r} \hat{v}_s^*(v_r) \geq \hat{v}_s^*(\alpha_{r0}) \overset{(J.34)}{>} \alpha_{s0}. \tag{J.36}$$

Contradiction. Similar as in Appendix J.3.1, we show the existence of $(\tilde{v}_r^+, \tilde{v}_s^+) \neq (\hat{v}_r^+, \hat{v}_s^+)$, which satisfy the unrestricted-equilibrium properties $\tilde{v}_r^+ = \hat{v}_r^*(\tilde{v}_s^+)$ and $\tilde{v}_s^+ = \hat{v}_s^*(\tilde{v}_r^+)$, and thus contradict the uniqueness of the unrestricted equilibrium valuations $(\hat{v}_r^+, \hat{v}_s^+)$. To that end, we again introduce a reflector function $\hat{v}_s^{**}$ with

$$\hat{v}_s^{**}(\tilde{v}_s^+) = \tilde{v}_s^+ - \hat{v}_s^*(\hat{v}_r^*(\tilde{v}_s^+)) = 0. \tag{J.37}$$

However, unlike in Appendix J.3.1, we additionally have to consider the relative position of $\nu_s^{\rightarrow}$ from Eq. (J.35) and $\hat{v}_s^*(\alpha_{r0})$.

- $\hat{v}_s^*(\alpha_{r0}) \geq \nu_s^{\rightarrow}$: For that case, we evaluate the reflector function $\hat{v}_s^{**}$ at arguments $\nu_s^{\rightarrow}$ from Eq. (J.35) and $\hat{v}_s^{\uparrow} \geq \nu_s^{\rightarrow}$:

$$\begin{aligned} \hat{v}_s^{**}(\nu_s^{\rightarrow}) &\overset{(J.37)}{=} \nu_s^{\rightarrow} - \hat{v}_s^*(\hat{v}_r^*(\nu_s^{\rightarrow})) \overset{(J.35)}{=} \nu_s^{\rightarrow} - \hat{v}_s^*(\alpha_{r0}) \leq 0 \\ \hat{v}_s^{**}(\hat{v}_s^{\uparrow}) &\overset{(J.37)}{=} \hat{v}_s^{\uparrow} - \hat{v}_s^*(\hat{v}_r^*(\hat{v}_s^{\uparrow})) \overset{(J.36)}{\geq} 0 \end{aligned} \tag{J.38}$$

  By the intermediate value theorem, there must thus exist a $\tilde{v}_s^+ \in [\nu_s^{\rightarrow}, \hat{v}_s^{\uparrow}]$ with $\hat{v}_s^{**}(\tilde{v}_s^+) = 0$. Since $\tilde{v}_s^+ \in [\nu_s^{\rightarrow}, \hat{v}_s^{\uparrow}]$, Eq. (J.35) implies that $\tilde{v}_s^+ \geq \nu_s^{\rightarrow} \geq \alpha_{s0}$, which conflicts with $\hat{v}_s^+ < \alpha_{s0}$ from case condition 2.(b) and the uniqueness of the unrestricted equilibrium.

- $\hat{v}_s^*(\alpha_{r0}) < \nu_s^{\rightarrow}$: For that case, we evaluate the reflector function $\hat{v}_s^{**}$ at arguments $\nu_s^{\leftarrow}$ from Eq. (J.35) and $\hat{v}_s^{\uparrow}$ from Eq. (J.36):

$$\begin{aligned} \hat{v}_s^{**}(\nu_s^{\leftarrow}) &\overset{(J.37)}{=} \nu_s^{\leftarrow} - \hat{v}_s^*(\hat{v}_r^*(\nu_s^{\leftarrow})) \overset{(J.35)}{=} \nu_s^{\leftarrow} - \hat{v}_s^*(\alpha_{r0}) \overset{(J.34)}{<} \nu_s^{\leftarrow} - \alpha_{s0} \overset{(J.35)}{\leq} 0 \\ \hat{v}_s^{**}(\hat{v}_s^{\uparrow}) &\overset{(J.37)}{=} \hat{v}_s^{\uparrow} - \hat{v}_s^*(\hat{v}_r^*(\hat{v}_s^{\uparrow})) \overset{(J.36)}{\geq} 0 \end{aligned} \tag{J.39}$$

  By the intermediate value theorem, there must thus exist a $\tilde{v}_s^+ \in [\nu_s^{\leftarrow}, \hat{v}_s^{\uparrow}]$ with $\hat{v}_s^{**}(\tilde{v}_s^+) = 0$. Since $\tilde{v}_s^+ \in [\nu_s^{\leftarrow}, \hat{v}_s^{\uparrow}] \subset [\nu_s^{\leftarrow}, \nu_s^{\rightarrow}]$, Eq. (J.35) implies that $\tilde{v}_r^+ = \hat{v}_r^+(\tilde{v}_s^+) \geq \alpha_{r0}$, which conflicts with $\hat{v}_r^+ < \alpha_{r0}$ from case condition 2 and the uniqueness of the unrestricted equilibrium.

Appendix K. Proof of Theorem 11

Appendix K.1. Proof idea

To confirm the asymptotic stability of an equilibrium $\mathbf{A}^+$, we demonstrate that the Jacobian matrix $\mathbf{J}(\mathbf{A})$ of the process in Eq. (30) is negative definite when evaluated at the equilibrium $\mathbf{A}^+$. On a high level, the Jacobian matrix $\mathbf{J}(\mathbf{A})$ is defined as follows:

$$\forall(n,k),\ (n',k') \in N \times K. \quad J_{I(n,k),I(n',k')}(\mathbf{A}) = \frac{\partial}{\partial a_{n'k'}} \left(a^*_{nk}(\mathbf{A}_{-nk}) - a_{nk}\right) \tag{K.1}$$

where $I(n,k)$ is an index corresponding to attribute (n,k). The derivatives of the restricted best-response a^*_{nk} in any attribute prevalence $a_{n'k'}$ are as follows:

$$\frac{\partial}{\partial a_{n'k'}} a^*_{nk}(\mathbf{A}) = \frac{\partial}{\partial a_{n'k'}} \max(0, \hat{a}^*_{nk}(\mathbf{A})) = \begin{cases} \frac{\partial}{\partial a_{n'k'}} \hat{a}^*_{nk}(\mathbf{A}) & \text{if } \hat{a}^*_{nk}(\mathbf{A}) \geq 0, \\ 0 & \text{otherwise.} \end{cases} \tag{K.2}$$

To show this negative definiteness of $\mathbf{J}^+ = \mathbf{J}(\mathbf{A}^+)$, we demonstrate that every eigenvalue λ of $\mathbf{J}^+$ has a negative real part $\mathrm{Re}(\lambda)$. To find the eigenvalues λ of $\mathbf{J}^+$, we solve the equation system $\mathbf{J}^+\mathbf{x} = \lambda\mathbf{x}$ for $\lambda \in \mathbb{C}$, $\lambda \neq 0$, and $\mathbf{x} \in \mathbb{C}^{|N||K|}$, $\mathbf{x} \neq \mathbf{0}$. This equation system can be represented in the following form:

$$\forall(n,k) \in N \times K. \quad \left(J^+_{I(n,k),I(n,k)} - \lambda\right) x_{I(n,k)} + \sum_{\substack{(n',k') \in N \times K. \\ (n',k') \neq (n,k)}} J^+_{I(n,k),I(n',k')} x_{I(n',k')} = 0 \tag{K.3}$$

Appendix K.2. Simplification of Eq. (K.3)

To concretize Eq. (K.3), we instantiate the Jacobian matrix $\mathbf{J}^+$. As a^*_{nk} is independent of a_{nk}, the diagonal entries of $\mathbf{J}^+$ are:

$$\forall(n,k) \in N \times K. \quad J^+_{I(n,k),I(n,k)} \overset{(K.1)}{=} \frac{\partial}{\partial a_{n'k'}} \left(a^*_{nk}(\mathbf{A}_{-nk}) - a_{nk}\right) = -1, \tag{K.4}$$

Now, we consider the entries not on the diagonal of $\mathbf{J}^+$, i.e., $J^+_{I(n,k),I(n',k')}$ for all $(n,k) \neq (n',k')$. First, we specifically consider the rows of $\mathbf{J}^+$ associated with attributes $(n,k) \in L^+$, where L^+ is the set of attributes which must have zero prevalence $a^+_{nk} = 0$ in equilibrium:

$$L^+ = N \times K \setminus (K^+_r \cup K^+_{\bar{r}}). \tag{K.5}$$

The following inequality holds on the equilibrium path valuation v^+_r for each path r (cf. Theorems 8 and 10):

$$\forall(n,k) \in L^+. \quad v^+_{r(n)} \geq \psi_r \sqrt{d} \sqrt{1 + v^+_{\bar{r}(n)}} - (1 + v^+_{\bar{r}(n)}) > \sqrt{\frac{\alpha_{nk}}{\gamma_{nk}} d(\rho_n - \phi_{n0})} \sqrt{1 + v^+_{\bar{r}(n)}} - (1 + v^+_{\bar{r}(n)}) \tag{K.6}$$

Then, remember the following equivalence from Theorem 8 for any attribute (n,k):

$$v^+_{r(n)} > \sqrt{\frac{\alpha_{nk}}{\gamma_{nk}} d(\rho_n - \phi_{n0})} \sqrt{1 + v^+_{\bar{r}(n)}} - (1 + v^+_{\bar{r}(n)}) \iff \frac{\partial \pi_n(\mathbf{A}^+)}{\partial a_{nk}} < 0. \tag{K.7}$$

Together with the concavity of π_n, we thus note that the attribute value a_{nk} needs to be decreased to optimize the profit π_n in a_{nk}. Given $a^+_{nk} = 0$, we note that the unrestricted best-response attribute $\hat{a}^*_{nk}$ for $(n,k) \in L^+$ is thus negative in the equilibrium:

$$\forall(n,k) \in L^+. \quad \hat{a}^*(\mathbf{A}^+_{-nk}) < 0 \tag{K.8}$$

Given the definition of the Jacobian entries in Eq. (K.1) and the derivative of the restricted best response a^*_{nk} in Eq. (K.2), we note that:

$$\forall (n,k) \in L^+.\ \forall (n',k') \neq (n,k). \quad J^+_{I(n,k),I(n',k')} = 0. \tag{K.9}$$

The eigenvalue equation system in Eq. (K.3) can thus be written as:

$$\forall (n,k) \in K^+_r \cup K^+_{\overline{r}}. \quad -(\lambda+1)x_{I(n,k)} + \sum_{\substack{(n',k') \in N \times K.\\ (n',k') \neq (n,k)}} J^+_{I(n,k),I(n',k')} x_{I(n',k')} = 0 \tag{K.10}$$

$$\forall (n,k) \in L^+. \quad -(\lambda+1)x_{I(n,k)} = 0 \tag{K.11}$$

Interestingly, the equation system in Eqs. (K.10) and (K.11) can be considerably simplified in our proof, which can be shown by a case distinction on $L^+ = (N \times K) \setminus (K^+_r \cup K^+_{\overline{r}})$, i.e., the set of attributes that certainly have zero prevalence in the equilibrium.

- $\boldsymbol{L^+ = \emptyset}$ **:** In this case, the equation system can be simplified in two respects. First, we note that $L^+ = \emptyset$ implies that no equations in the form of Eq. (K.11) exist in the equation system. Second, we note that we only investigate networks with a unique equilibrium, i.e., non-zero equilibrium prevalence is possible for only one attribute on each path ($|K^+_r| = |K^+_{\overline{r}}| = 1$). Hence, if $L^+ = \emptyset$, we know that $K^+_r \cup K^+_{\overline{r}}$ covers both of the two attributes of the network, one on each path:

$$K^+_r \cup K^+_{\overline{r}} = N \times K = \{(n(r), k(r)), (n(\overline{r}), k(\overline{r}))\} \tag{K.12}$$

  where $(n(r), k(r))$ is the single attribute with possibly non-zero equilibrium prevalence on path r. These insights allow to reduce the equation system in Eq. (K.10) to only two equations (No equations like Eq. (K.11) exist):

$$-(\lambda+1)x_r + J^+_{\overline{r}} x_{\overline{r}} = 0 \qquad\qquad -(\lambda+1)x_{\overline{r}} + J^+_r x_r = 0 \tag{K.13}$$

  where we have abbreviated:

$$J^+_r = J^+_{I(n(r),k(r)),I(n(\overline{r}),k(\overline{r}))} \qquad\qquad x_r = x_{I(n(r),k(r))} \tag{K.14}$$

  Since the eigenvector $\mathbf{x}$ in the current case only has the two entries x_r and $x_{\overline{r}}$, we require $\mathbf{x} = (x_r, x_{\overline{r}})^\top \neq \mathbf{0}$. We find that $\lambda = -1$ is an eigenvalue of the system in Eq. (K.13) if and only if $J^+_r = 0$ or $J^+_{\overline{r}} = 0$, i.e., at least one of the two relevant Jacobian entries is zero. Since $\lambda = -1$ would preserve negative definiteness of $\mathbf{J}^+$, we do not need to consider this case further.

- $\boldsymbol{L^+ \neq \emptyset}$ **:** If $L^+ \neq \emptyset$, the equation system contains equations in the form of Eq. (K.11). Then, $\lambda = -1$ may be an eigenvalue of $\mathbf{J}^+$, which would preserve negative definiteness of $\mathbf{J}^+$; hence, this case is not further considered. Conversely, if $\lambda = -1$ is not a solution of the system, the equations in the form of Eq. (K.11) imply that $x_{I(n,k)} = 0$ for all $(n,k) \in L^+$. This insight then again allows the simplification to the equation system in Eq. (K.13). Crucially, since $x_{I(n,k)} = 0$ for all $(n,k) \in L^+$, it must hold that $(x_r, x_{\overline{r}})^\top \neq \mathbf{0}$ such that $\mathbf{x} \neq \mathbf{0}$, i.e., such that $\mathbf{x}$ is a valid eigenvector.

Appendix K.3. Solution of Eq. (K.13)

In summary, we only need to consider the equation system in Eq. (K.13) and the case $\lambda \neq -1$. Furthermore, we require $(x_r, x_{\overline{r}})^\top \neq \mathbf{0}$. Without loss of generality, let r be the path with $x_r \neq 0$. Then, we can perform the following transformation:

$$-(\lambda+1)x_r + J^+_{\overline{r}} x_{\overline{r}} = 0 \implies \lambda + 1 = J_{\overline{r}} \frac{x_{\overline{r}}}{x_r} \tag{K.15}$$

$$-(\lambda+1)x_{\overline{r}} + J^+_r x_r = 0 \implies x_{\overline{r}} = \frac{J_r}{\lambda+1} x_r \tag{K.16}$$

Inserting Eq. (K.16) into Eq. (K.15) yields a quadratic equation in λ:

$$(\lambda+1)^2 - J_r^+ J_{\overline{r}}^+ = 0 \implies \lambda_{1,2} = -1 \pm \sqrt{J_r^+ J_{\overline{r}}^+} \tag{K.17}$$

If $J_r^+ J_{\overline{r}}^+ = 0$, then $\lambda_{1,2} = -1$ produces a contradiction to the assumption $\lambda \neq -1$, which implies that no eigenvalue $\lambda \neq -1$ exists.

If $J_r^+ J_{\overline{r}}^+ \neq 0$, then we know that

$$\mathrm{Re}(\lambda_{1,2}) < 0 \iff \mathrm{Re}\left(\sqrt{J_r^+ J_{\overline{r}}^+}\right) < 1 \iff J_r^+ J_{\overline{r}}^+ < 1. \tag{K.18}$$

Appendix K.4. Bounding of $\lambda_{1,2}$

To verify that the condition in Eq. (K.18) always holds, we first find J_r^+ for any path r:

$$J_r^+ = \begin{cases} \frac{\alpha_{n(\overline{r})k(\overline{r})}}{\alpha_{n(r)k(r)}} \left(\frac{\psi_r \sqrt{d}}{2\sqrt{1+v_{\overline{r}}^+}} - 1 \right) & \text{if } \hat{a}^*_{n(r)k(r)}\left(\mathbf{A}^+_{-n(r)k(r)}\right) \geq 0, \\ 0 & \text{otherwise.} \end{cases} \tag{K.19}$$

Given $J_r^+ J_{\overline{r}}^+ \neq 0$, stability requires:

$$J_r^+ J_{\overline{r}}^+ \overset{\text{(K.19)}}{=} \left(\frac{\psi_r \sqrt{d}}{2\sqrt{1+v_{\overline{r}}^+}} - 1 \right) \left(\frac{\psi_{\overline{r}} \sqrt{d}}{2\sqrt{1+v_r^+}} - 1 \right) \overset{\text{(K.18)}}{<} 1 \tag{K.20}$$

$$\iff \psi_r \psi_{\overline{r}} d - 2\sqrt{d} \left(\psi_r \sqrt{1+v_{\overline{r}}^+} + \psi_{\overline{r}} \sqrt{1+v_r^+} \right) < 0 \tag{K.21}$$

Moreover, $J_r^+ J_{\overline{r}}^+ \neq 0$ implies that the restricted equilibrium valuation v_r^+ for each path r corresponds to the unrestricted equilibrium valuation $\hat{v}_r^+$ from Theorem 10:

$$\begin{aligned} &\forall r \in R.\ J_r^+ J_{\overline{r}}^+ \neq 0 \implies \forall r \in R.\ \hat{a}^*_{n(r)k(r)}(\mathbf{A}^+_{-n(r)k(r)}) \geq 0 \\ \implies &\forall r \in R.\ \hat{v}_r^+ \geq \alpha_{r0} \implies \forall r \in R.\ v_r^+ = \hat{v}_r^+. \end{aligned} \tag{K.22}$$

Hence, we can expand (symmetrically for $\psi_{\overline{r}}\sqrt{1+v_{\overline{r}}^+}$):

$$\begin{aligned} \psi_r \sqrt{1+v_r^+} &\overset{\text{(K.22)}}{=} \psi_r \sqrt{1+\hat{v}_r^+} \\ &\overset{\text{Th10}}{=} \psi_r \sqrt{\frac{\psi_r^3 \psi_{\overline{r}}}{(\psi_r^2+\psi_{\overline{r}}^2)^2} \left(\sqrt{d\,(\psi_r^2+\psi_{\overline{r}}^2) + \frac{1}{4}\psi_r^2\psi_{\overline{r}}^2 d^2} + \frac{d}{2}\psi_r\psi_{\overline{r}} \right) + \frac{\psi_r^2}{\psi_r^2+\psi_{\overline{r}}^2}} \\ &= \psi_r^2 \sqrt{\frac{\psi_r \psi_{\overline{r}}}{(\psi_r^2+\psi_{\overline{r}}^2)^2} \left(\sqrt{d\,(\psi_r^2+\psi_{\overline{r}}^2) + \frac{1}{4}\psi_r^2\psi_{\overline{r}}^2 d^2} + \frac{d}{2}\psi_r\psi_{\overline{r}} \right) + \frac{1}{\psi_r^2+\psi_{\overline{r}}^2}} \end{aligned} \tag{K.23}$$

We use this equality to rewrite the inequality in Eq. (K.21):

$$\begin{aligned} &\psi_r \psi_{\overline{r}} d - 2\sqrt{d} \left(\psi_r \sqrt{1+v_{\overline{r}}^+} + \psi_{\overline{r}} \sqrt{1+v_r^+} \right) \\ \overset{\text{(K.23)}}{=} &\psi_r \psi_{\overline{r}} d - 2\sqrt{d}\left(\psi_r^2+\psi_{\overline{r}}^2\right) \sqrt{\frac{\psi_r \psi_{\overline{r}}}{(\psi_r^2+\psi_{\overline{r}}^2)^2} \left(\sqrt{d\,(\psi_r^2+\psi_{\overline{r}}^2) + \frac{1}{4}\psi_r^2\psi_{\overline{r}}^2 d^2} + \frac{d}{2}\psi_r\psi_{\overline{r}} \right) + \frac{1}{\psi_r^2+\psi_{\overline{r}}^2}} \\ = &\psi_r \psi_{\overline{r}} d - \sqrt{2d^2\psi_r^2\psi_{\overline{r}}^2 + 4d\left(\psi_r^2+\psi_{\overline{r}}^2\right)^2 \left(\frac{\psi_r \psi_{\overline{r}}}{(\psi_r^2+\psi_{\overline{r}}^2)^2} \sqrt{d\,(\psi_r^2+\psi_{\overline{r}}^2) + \frac{1}{4}\psi_r^2\psi_{\overline{r}}^2 d^2} + \frac{1}{\psi_r^2+\psi_{\overline{r}}^2} \right)} \\ < &\left(1-\sqrt{2}\right) d\psi_r\psi_{\overline{r}} < 0 \end{aligned} \tag{K.24}$$

The last upper bound holds because we can exclude $\psi_r = 0$ for any path r in the current case $v_r^+ = \hat{v}_r^+$, as $\psi_r = 0$ produces the contradiction $v_r^+ < v_r^+$:

$$\begin{aligned} v_r^+ &\overset{\text{(K.22)}}{=} \hat{v}_r^+ \overset{\text{(J.6)}}{=} \psi_r \sqrt{d}\sqrt{1+\hat{v}_{\overline{r}}^+} - (1+\hat{v}_{\overline{r}}^+) \overset{\psi_r=0}{=} -(1+\hat{v}_{\overline{r}}^+) \overset{\text{(K.22)}}{=} -(1+v_{\overline{r}}^+) \\ &\overset{\text{Th10}}{\leq} -(1+\alpha_{\overline{r}0}) \overset{\alpha_{\overline{r}0}\geq 0}{<} 0 \overset{\alpha_{r0}\geq 0}{\leq} \alpha_{r0} \overset{\text{Th10}}{\leq} v_r^+. \end{aligned} \tag{K.25}$$

In summary, we now have shown that $J_r^+ J_{\overline{r}}^+ < 1$, which ensures a negative real part $\mathrm{Re}(\lambda_{1,2}) < 0$ (Eq. (K.18)) of the eigenvalues $\lambda_{1,2}$ from Eq. (K.17), and thus confirms that $\mathbf{J}^+$ is negative definite. Since $\mathbf{J}^+$ is negative definite, the equilibrium $\mathbf{A}^+$ from Theorem 10 is asymptotically stable with respect to the process in Eq. (30), which concludes the proof.

Appendix L. Proof of Theorem 12

For the competitive network $\mathcal{N}_4$, the unrestricted equilibrium network valuation $\hat{V}^+$ is

$$\hat{V}^+(\mathcal{N}_4) = \hat{v}_r^+ + \hat{v}_{\overline{r}}^+ = \frac{\psi_r\psi_{\overline{r}}}{\psi_r^2+\psi_{\overline{r}}^2}\left(\sqrt{d\left(\psi_r^2+\psi_{\overline{r}}^2\right)+\frac{1}{4}\psi_r^2\psi_{\overline{r}}^2 d^2}+\frac{d}{2}\psi_r\psi_{\overline{r}}\right)-1, \tag{L.1}$$

where $\hat{v}_r^+$ and $\hat{v}_{\overline{r}}^+$ are as in Theorem 10. The restricted equilibrium network valuation $V^+(\mathcal{N}_4)$ is equal to $\hat{V}^+(\mathcal{N}_4)$ if $\hat{v}_r^+ \geq \alpha_{r0}$ and $\hat{v}_{\overline{r}}^+ \geq \alpha_{\overline{r}0}$. These unrestricted equilibrium path valuations are monotonically increasing in the demand limit d (cf. Theorem 10). Hence, if d is high enough, it holds that $V^+(\mathcal{N}_4) = \hat{V}^+(\mathcal{N}_4)$.

For the competition-free network $\mathcal{N}_3$, the unrestricted equilibrium network valuation $\hat{V}^+$ is

$$\hat{V}^+(\mathcal{N}_3) = \psi_r\sqrt{d_r} - 1 + \psi_{\overline{r}}\sqrt{d_{\overline{r}}} - 1. \tag{L.2}$$

Among all demand distributions $(d_r, d_{\overline{r}})$ with $d_r + d_{\overline{r}} = d$, the demand distribution maximizing $\hat{V}^+(\mathcal{N}_3)$ can be found as follows:

$$\frac{\partial}{\partial d_r}\left(\psi_r\sqrt{d_r}+\psi_{\overline{r}}\sqrt{d-d_r}-2\right)=0 \iff d_r = \frac{\psi_r^2}{\psi_r^2+\psi_{\overline{r}}^2}d, \tag{L.3}$$

where the maximum character of this value is ensured by a consistently non-positive second derivative of $\hat{V}^+(\mathcal{N}_3)$ in d_r. In the following, we thus consider only the maximum unrestricted equilibrium network valuation $\hat{V}^+(\mathcal{N}_3)$:

$$\hat{V}^+(\mathcal{N}_3) = \psi_r\sqrt{\frac{\psi_r^2}{\psi_r^2+\psi_{\overline{r}}^2}d}+\psi_{\overline{r}}\sqrt{\frac{\psi_{\overline{r}}^2}{\psi_r^2+\psi_{\overline{r}}^2}d}-2 = \sqrt{d(\psi_r^2+\psi_{\overline{r}}^2)}-2. \tag{L.4}$$

Again, for $V^+(\mathcal{N}_3) = \hat{V}^+(\mathcal{N}_3)$, d must be high enough such that

$$\psi_r\sqrt{\frac{\psi_r^2}{\psi_r^2+\psi_{\overline{r}}^2}d}-1 \geq \alpha_{r0} \quad \text{and} \quad \psi_{\overline{r}}\sqrt{\frac{\psi_{\overline{r}}^2}{\psi_r^2+\psi_{\overline{r}}^2}d}-1 \geq \alpha_{\overline{r}0}. \tag{L.5}$$

If d is high enough, the difference of the equilibrium network valuations is thus:

$$\Delta V^+ = V^+(\mathcal{N}_4) - V^+(\mathcal{N}_3) = \frac{\psi_r\psi_{\overline{r}}}{\psi_r^2+\psi_{\overline{r}}^2}\left(\sqrt{d\left(\psi_r^2+\psi_{\overline{r}}^2\right)+\frac{1}{4}\psi_r^2\psi_{\overline{r}}^2 d^2}+\frac{d}{2}\psi_r\psi_{\overline{r}}\right)-\sqrt{d(\psi_r^2+\psi_{\overline{r}}^2)}+1 \tag{L.6}$$

Clearly, this difference is eventually positive when increasing the demand limit d, meaning that $V^+(\mathcal{N}_4)$ exceeds $V^+(\mathcal{N}_3)$ for high enough d:

$$\lim_{d\to\infty}\Delta V^+ = \infty. \tag{L.7}$$

This last insight proves the theorem.

Appendix M. Proof of Theorem 13

The following proof is constructive, i.e., we demonstrate how to choose $(\psi_r, \psi_{\overline{r}})$ and $(\alpha_{r0}, \alpha_{\overline{r}0})$ such that $V^+(\mathcal{N}_4) < V^+(\mathcal{N}_3)$ holds given demand distribution $(d_r, d_{\overline{r}})$. In this construction, the goal is to create a scenario where the competitive network $\mathcal{N}_4$ will be at minimum valuation $\alpha_{r0} + \alpha_{\overline{r}0}$, but the competition-free network has a path r with an equilibrium valuation v_r^+ exceeding the minimum path valuation α_{r0}.

Regarding the path-characteristic ratios $(\psi_r, \psi_{\overline{r}})$, our first step consists of choosing the ratios such that the competitive network is at minimum valuation, i.e., such that $v_r^+ = \alpha_{r0}$ and $v_{\overline{r}}^+ = \alpha_{\overline{r}0}$. To do so, we first determine $\psi_{\overline{r}}$ such that $\hat{v}_{\overline{r}}^+ \leq \alpha_{\overline{r}0}$ for all ψ_r, which is achieved by $\psi_{\overline{r}} = 0$:

$$\lim_{\psi_{\overline{r}} \to 0} \hat{v}_{\overline{r}}^+ \overset{\text{Th10}}{=} -1 \overset{\alpha_{\overline{r}0} \geq 0}{<} \alpha_{\overline{r}0}. \tag{M.1}$$

Having selected $\psi_{\overline{r}}$ such that $\hat{v}_{\overline{r}}^+ \leq \alpha_{\overline{r}0}$, it holds that $v_{\overline{r}}^+ = \alpha_{\overline{r}0}$ by Theorem 10. As a result, the equilibrium path valuation for path r in the competitive network $\mathcal{N}_4$ is:

$$v_r^+(\mathcal{N}_4) \overset{\text{Th10}}{=} \max\left(\alpha_{r0},\ \psi_r \sqrt{d}\sqrt{1 + \alpha_{\overline{r}0}} - (1 + \alpha_{\overline{r}0})\right). \tag{M.2}$$

To ensure that v_r^+ is minimal (i.e., equals α_{r0}), the following condition must hold:

$$\alpha_{r0} \geq \psi_r \sqrt{d}\sqrt{1 + \alpha_{\overline{r}0}} - (1 + \alpha_{\overline{r}0}) \iff \psi_r \leq \frac{1 + \alpha_{r0} + \alpha_{\overline{r}0}}{\sqrt{d}\sqrt{1 + \alpha_{\overline{r}0}}} \tag{M.3}$$

If ψ_r is chosen according to Eq. (M.3), the equilibrium network valuation in the competition-free network is minimal, i.e., $V^+(\mathcal{N}_4) = \alpha_{r0} + \alpha_{\overline{r}0}$.

To let $V^+(\mathcal{N}_3)$ of the competition-free network exceed $V^+(\mathcal{N}_4)$ of the competitive network, we further need to choose ψ_r such that $\hat{v}_r^+(\mathcal{N}_3) > \alpha_{r0}$. This condition can be transformed in the following fashion:

$$\psi_r \sqrt{d_r} - 1 > \alpha_{r0} \iff \psi_r > \frac{1 + \alpha_{r0}}{\sqrt{d_r}} \tag{M.4}$$

To allow a selection of ψ_r that achieves $v_r^+(\mathcal{N}_4) = \alpha_{r0}$ but $v_r^+(\mathcal{N}_3) > \alpha_{r0}$, it must thus hold that

$$\frac{1 + \alpha_{r0}}{\sqrt{d_r}} < \frac{1 + \alpha_{r0} + \alpha_{\overline{r}0}}{\sqrt{d}\sqrt{1 + \alpha_{\overline{r}0}}} \iff \alpha_{r0} < \frac{\sqrt{d_r}\alpha_{\overline{r}0}}{\sqrt{d}\sqrt{1 + \alpha_{\overline{r}0}} - \sqrt{d_r}} - 1, \tag{M.5}$$

This condition always holds when choosing $\alpha_{r0} = 0$ and $\alpha_{\overline{r}0} > d_{\overline{r}}/d_r$:

$$\begin{aligned} &\frac{\sqrt{d_r}\alpha_{\overline{r}0}}{\sqrt{d}\sqrt{1 + \alpha_{\overline{r}0}} - \sqrt{d_r}} - 1 \overset{\alpha_{\overline{r}0} > d_{\overline{r}}/d_r}{>} \frac{\sqrt{d_r}\frac{d_{\overline{r}}}{d_r}}{\sqrt{d}\sqrt{1 + \frac{d_{\overline{r}}}{d_r}} - \sqrt{d_r}} - 1 \overset{d_r + d_{\overline{r}} = d}{=} \frac{\sqrt{d_r}\frac{d_{\overline{r}}}{d_r}}{\sqrt{d}\sqrt{\frac{d}{d_r}} - \sqrt{d_r}} - 1 \\ &\overset{\frac{\sqrt{d_r}}{d_r} = \frac{1}{\sqrt{d_r}}}{=} = \frac{\frac{d_{\overline{r}}}{\sqrt{d_r}}}{\frac{d}{\sqrt{d_r}} - \frac{d_r}{\sqrt{d_r}}} - 1 \overset{d_r + d_{\overline{r}} = d}{=} \frac{\frac{d_{\overline{r}}}{\sqrt{d_r}}}{\frac{d_{\overline{r}}}{\sqrt{d_r}}} - 1 = 0 \overset{\alpha_{r0} = 0}{=} \alpha_{r0} \end{aligned} \tag{M.6}$$

Hence, $(\psi_r, \psi_{\overline{r}})$ and $(\alpha_{r0}, \alpha_{\overline{r}0})$ can be chosen such that $V^+(\mathcal{N}_3) > V^+(\mathcal{N}_4)$, which concludes the proof.