



Asymmetry and Discrimination in Internet Peering Evidence from the LINX

Alessio D'Ignazio¹ and Emanuele Giovannetti²

Working Paper Number 102

¹ University of Cambridge

² University of Cape Town and University of Verona

Asymmetry and Discrimination in Internet Peering Evidence from the LINX[◇]

Alessio D'Ignazio^a and Emanuele Giovannetti^{b,*}

^a*University of Cambridge*

^b *University of Cape Town and University of Verona*

Abstract

Is the quality of interconnection between Internet operators affected by their asymmetry? While recent game theoretic literature provides contrasting answers to this question, there is a lack of empirical research. We introduce a novel dataset based on Internet routing policies, and study the interconnection decisions amongst the Internet Service Providers (ISPs) members of the London Internet Exchange Point (LINX).

Our results show that interconnection quality degradation can be significantly explained by asymmetry between providers. We also show that Competition Authorities should focus more on the role played by the “centrality of an operator”, rather than on its market share.

Keywords: Internet Peering, Two-sided Markets, Network Industries, Antitrust, Net Neutrality

JEL Classification: L14, L86, L96, C81, L40

[◇] The authors would like to thank Bipasa Datta, Joerge Lepler, Tim Griffin, Randy Bush, John Souter (LINX), Chris Fletcher (LINX), Valeria Rossi (MIX), and Daniel Akerberg for their help, useful comments and data. The usual disclaimer applies. Alessio D'Ignazio thankfully acknowledges financial support from the Bank of Italy through the “Bonaldo Stringher” Scholarship.

* Corresponding author: Prof Emanuele Giovannetti. E-mail: giovannetti@cantab.net.

1 Introduction

Antitrust authorities are showing increasing interest in the analysis of interconnection agreements used by Internet Operators to exchange traffic packets. They focus on these arrangements to detect both actual and potential abuse of a position of significant market power¹. In this setting, a dominant position may lead to the establishment of “unfair” conditions associated to the bilateral exchanges of traffic.

A growing literature is focussing on the issue of interconnection agreements between providers in the Internet Industry (Foros, Kind, and Sørsgard 2002; Crémer *et al.* 2000; Foros and Hansen, 2001; Economides 1998, to name just a few). The actual interconnection regime between a pair of providers is clearly the result of a strategic game. In particular, the Internet operators are in a relationship of both complementarity (each network must be able to access each other in order to assure the Internet *universal connectivity*) and competitiveness (they compete over downstream customers).

Broadly speaking, each pair of providers can be interconnected in two different ways: they can exchange their traffic through a *direct link* (this agreement is known as “peering”); otherwise, they will use upstream intermediaries, called “transit providers”. These two alternative ways to exchange traffic clearly affect the quality of the interconnection between the two providers²: peering assures a better quality than transit agreements, given its “dedicated” character³. Even within the category of peering, however, providers are able to “modulate” the relative quality of the link.

¹ See for example the 1998 MCI WorldCom and the 2000 MCI-Worldcom Sprint mergers enquiries by the European Commission . Official Journal Of the European Commission (2000), Regulation (EEC) N 4046/89, Merger Procedure, Bruxelles, European Commission, DGXIII. See also Buccrossi et al. (2005).

² The quality of interconnection can be measured by different parameters such as packet loss, latency, band-width.

³ In particular, the quality of a direct peering link is generally better than the one of an indirect transit link assuring that traffic is faster, there is less packet loss and less latency (waiting time).

Theoretical research is trying to model the Internet providers' interconnections decisions, mainly by using a game theoretical approach⁴. The central question addressed is "does a provider have the incentive to degrade the quality of its interconnection with some other providers?". This question becomes interesting if we consider asymmetric networks. Indeed, having the bigger provider a larger customer base than the smaller one, the degradation of the interconnection quality is more harmful for the latter, due to the asymmetric losses in good quality connectivity (the large provider loses good connectivity to less final users than the smaller provider does). This research is particularly important from an antitrust point of view, since degrading interconnection towards smaller providers can lead to increasing market power, due to a "*market tipping*" process, which can then induce a monopolistic type of restriction in Internet supply. This preoccupation about incentives towards quality discrimination, leading to *market tipping*, was indeed the main argument in the European Commission decision to block the proposed merger between MCI-Worldcom and Sprint in 2000⁵. Moreover, understanding the real extent of this problem is particularly relevant within the "Net Neutrality" debate about the potential need for introducing interconnection regulation in the Internet⁶.

The game theoretical models, referred above, provided contrasting answers to the question at hand, motivating the need for more empirical research. This is, in fact, still scarce, and mainly anecdotic, essentially because of the confidentiality that characterizes the providers' interconnection agreements and Internet traffic data. Our work provides a contribution in this direction: this is possible thanks to a novel approach to

⁴ See for example Crémer et al. (2000); Economides (2005); Foros and Hansen (2001); Baake and Wichmann (1999); Badasyan and Chakrabarti (2003); Mah (2005); Weiss and Shin (2004); Jahn and Prüfer (2004), Ida (2005).

⁵ See footnote 1.

⁶ For a summary of the increasing body of literature on the Net Neutrality see Sidak (2006).

obtain data, which follows recent advances in the fields of Theoretical Computer Science⁷.

We investigate if asymmetry between a pair of providers is associated to interconnection quality degradation. In particular, the presence of direct peering, involving a dedicated agreement between the two providers, will be considered as a “high quality interconnection”. In the absence of peering, instead, two providers exchange traffic by using the services of upstream intermediaries; we will hence consider this as the “low quality interconnection” case. Our database consists of the interconnection decisions characterizing the Internet Service Providers (ISPs) members of the London Internet Exchange Point (LINX).

The results obtained seem to support the part of the theory claiming a positive relationship between providers’ asymmetry and quality. We also find that the bigger threat to interconnection fairness does not come from a market share-based dominance, but is mostly associated to the relative centrality of the players in the Internet. This is an interesting result, since the actual Competition Authorities’ approach usually focuses on the assessment of Internet operators’ market shares.

The rest of the paper is organised as follows. Section 2 introduces the subject and discusses some technical aspects about Internet peering, while Section 3 focuses on the game theoretic models studying interconnection agreements. Section 4 explains the process of data gathering and the criteria used to classify the Internet Operators, and section 5 provides the econometric analysis of the relevant model. Finally, section 6 concludes.

⁷ The interest of Computer Scientists for interconnection agreements focuses on the representation and analysis of the evolution of the Internet topology (Bar *et al.*, 2005), and on its efficiency from a Network optimisation point of view (Heckmann *et al.*, 2004).

2 The Internet hierarchy, peering and transit

Internet operators may be classified into different categories, depending upon their position in the Internet hierarchy. At the top level there are the *Tier-1 Transit Providers* and the *Internet Backbones* (IBPs): they constitute the upstream industry (Kende, 2000) providing universal connectivity to the downstream industry, constituted by operators of smaller dimensions (*Internet Service Providers*⁸, or ISPs). At a further lower level in the Internet hierarchy there are the so-called Internet Access Providers, or IAPs, which usually obtain connectivity through a single connection to an ISP.

The dominant feature of the Internet, the Network of networks, is the universal connectivity: users are able to access to each other, whatever the provider they subscribe to. This is only possible thanks to the system of bilateral interconnections between the Internet Operators: there are a variety of commercial agreements, but these can be essentially divided into two main categories: *transit* and *peering*.

- The transit agreement leads to a unilateral *provider-to-customer* relationship: the Internet Operator “customer” is provided with connectivity to the entire Internet by the intermediary Internet Operator acting as an “upstream provider”; for this service, the customer pays a settlement fee to the provider.
- The peering agreement leads to a bilateral direct and high quality *peer-to-peer* relationship: each peer provides the other connectivity to its own network, usually without any settlement fee⁹.

One of the main advantages from a peering agreement is the minimisation of traffic costs: Internet Operators do not have to pay for the traffic routed

⁸ This term has now fallen into a general looser usage, but it is properly used to describe regional providers that typically connect to multiple backbone providers (Woodcock, 2002).

⁹ This is known as *Sender Keeps All* (SKA) peering, or *Bill and Keep* peering.

through peering networks. While peering also involves the sunk cost of interconnecting, borne by the two providers (see Norton, 2002), these costs fell sharply in the recent years, after the development and growth of the *Internet Exchange Points*¹⁰ (IXPs). IXPs are organizations that provide a centralised interconnection infrastructure to the members ISPs, so that they can exchange bilateral traffic without the need to build dedicated extra circuits. A second advantage enjoyed by peering providers, with respect to being connected through transit agreements with upstream providers, is the better performance of the traffic flows between them: this is due to the direct nature of peering and it is technically expressed through a lower latency in the transmission of packets, and a greater reliability¹¹.

On the other hand, transit also has some advantages to peering. A well known one is that, contrarily to peering, transit agreements include Service Level Agreements (SLAs) that guarantees rapid repair if problems on the interconnection link occur, while if a peering link experiences troubles, it is up to the peers to fix these. This is one of the reasons for which mutual knowledge and reputation effects between peers seem crucial (this element is strengthened by Titley, 1997) to decide about a potential peering relationship.

Another advantage of transit over peering is its feature of “certainty”; since it is too costly to have a reliable measurement of the traffic volume bilaterally exchanged, it is difficult to establish the actual advantages in terms of traffic costs saving from peering. In this regard, it is argued that the mutual presence at several different IXPs, for a pair of potential peers, enhances the likelihood of peering. In this case, in fact, the peers are “on

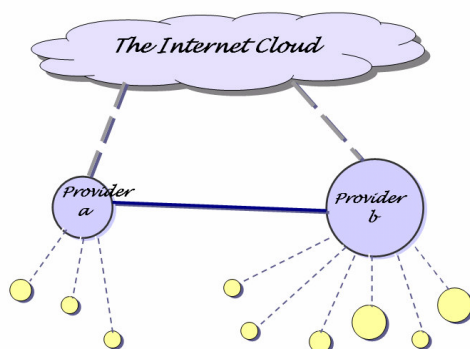
¹⁰ Xu *et al.* (2004) find that the percentage of peering agreements between ISPs participating at IXPs is significantly higher than the percentage characterising the whole Internet, providing evidence that IXPs plays an important role in shaping the relationships between Internet operators.

¹¹ For a description of the relevant Internet interconnection quality parameters in an economic framework see Giovannetti et al (2005).

average” capable of routing the free traffic to the peering network relatively soon, without thus bearing much of the cost associated to carrying traffic packets. This incentive to deliver the traffic packets to the destination network as soon as possible is commonly known with the name of “hot potato routing”¹².

Our paper empirically investigates the relationship between providers asymmetry and interconnection quality degradation. For the sake of tractability, we consider a binary case, where the presence of a peering contract represents the “high quality interconnection case”, while otherwise operators exchange traffic through their upstream providers (“low quality interconnection case”).

Figure 1: Two modalities of Interconnection



The stylized figure above shows the two modalities a pair of providers can use to exchange traffic. The thick line represents a peering agreement,

¹²Hot potato routing is crucial in peering, and it involves technological aspects of traffic routing. Since carrying traffic is costly, when a packet has to be delivered from a network A to a network B, the network A has the incentive to deliver the packet to B following the shortest path. If the networks are connected at many exchange points, each network is able to route relatively soon the traffic to the destination network. Hence, mutual presence at more exchange points is argued to positively affect peering.

used to exchange their traffic directly; the dotted lines represent one or more transit agreements with upstream providers, in the Internet Cloud.

The interconnection decision problem at hand has been addressed by several papers; many authors argued that peering is negatively affected by providers asymmetry¹³; in particular, two commonly argued reasons seem to induce a large provider to refuse peering with a small operator: they are the so-called *backbone free riding* and the *business stealing effect*.

To understand the *backbone free riding* problem, we have to notice that, in any peering agreement, the smaller network gets the bigger benefit. Since the sunk and maintenance costs associated with the peering link are equally shared by the providers, the smaller network free rides on the bigger one.

The *business stealing* refers to quality differentiation. Due to a network externality effect, a big provider is able to offer a better quality service to its customers than a small provider. This quality differentiation is hence relevant to gain more customers. If two networks of different sizes peer, however, this quality differentiation is dramatically reduced thanks to the new peering link (we can think about the new link as joining the two providers into one big network); as a consequence, the larger network may lose customers to the advantage of the, usually cheaper, smaller network. For instance, let us consider the figure 1 above, and assume that the peering link is initially absent. If peering is realised, it might be possible that some customers of provider *b* decide to leave and join *a*, whose quality is now improved relatively to *b*.

¹³ See for example Norton (2002), Kende, (2000) and Filstrup, (2001). According to Filstrup , who reports the selective peering criteria released by WorldCom, the symmetry in size is expressed in terms of a balance in the geographic scope , traffic across the peering point, capacity and traffic volume.

3 Game theoretic models of Internet peering

One of the earliest theoretical works on the interconnection strategies between competing Internet operators is due to Crémer, Rey, and Tirole (2000). They study the interconnection decision between two backbones, with one having a larger installed base of consumers; the backbones compete à la Cournot over the remaining part of still unattached consumers. They consider a two stage game. In the first stage each backbone i chooses a quality θ_i for the interconnection; the effective quality of interconnection is then $\min\{\theta_1, \theta_2\}$. Given the interconnection quality, the backbones choose their capacities and prices. The solution of the game relies on the comparison between two effects of degrading interconnection quality. If the connectivity between the two networks is degraded, both backbones face a demand reduction (their customers' access to each others deteriorates). However, the degradation of the connectivity leads to a greater quality differentiation between the two networks, which increases with the extent of network externality¹⁴. The larger backbone gains competitive advantage over the smaller one. Hence, Crémer *et al.* show that the largest network has incentives to degrade interconnection with the smaller networks to further increase its market share (it attracts customers because it can offer a better quality service of the other¹⁵).

On the same line are the results of Jahn and Prüfer (2004), and Weiss and Shin (2004). Jahn and Prüfer (2004) consider two Internet Operators that have a fixed base of customers, while they compete in prices over consumers

¹⁴ Indeed, in the model of Crémer *et al.*, the quality of the service of the backbone i is given by $s_i = v[(\beta_i + q_i) + \theta(\beta_j + q_j)]$, where β_i is the installed base of customers of the backbone i , q_i is the number of unattached customers enrolled by backbone i , $\theta \in [0,1]$ is the quality of interconnection, and v a parameter that reflects the importance of connectivity.

¹⁵ We referred before to this as the *business stealing* effect.

located in a *battlezone*¹⁶. They show that sufficiently symmetric in size (represented by the number of customers locked) networks reach a peering agreement; otherwise an upstream intermediary is used to exchange traffic. Weiss and Shin (2004) argue that the choice of the interconnection regime is based on the traffic volume, which, in turn, is linked to market share. Their model shows that symmetry in traffic positively affects peering¹⁷.

Although the result that difference in size negatively affects peering is commonly accepted, there are some situations where it does not seem to work. First of all, peering does not necessarily imply business stealing if the networks are sufficiently differentiated. Secondly, the negative effects of business stealing and free riding may be offset by other positive effects caused by network externalities. We now briefly point at these issues.

Since Internet Operators compete for downstream customers (either end users or other Internet providers), their interconnection strategy depends upon the preferences of these customers. Courcoubetis and Weber (2003) argue that *“the decision as to whether or not peering is beneficial depends on the way the networks are differentiated and on the importance that their customers place on the differentiating parameters, such as size and location.”* In this direction, Foros and Hansen (2001) consider horizontal differentiation

¹⁶ The two networks are ex ante connected through an intermediary, defined as the cheapest Tier-1 provider. In the first stage of the game, the two networks decide non cooperatively about the interconnection regime: if they do not reach a peering regime (either bill and keep or paid), then they remain connected through the intermediary. In the subsequent stage the two networks set prices, competing à la Hotelling over the consumers on the battlezone. Finally, consumers choose the network to subscribe with. Hence, while in Crémer et al. the strategic variable is the interconnection quality, here the strategic variable is the interconnection regime.

¹⁷ In their model there is one IBP in the upstream market and two ISPs in the downstream market. The realisation of peering between the two ISPs occurs where both of them take advantage from the reduction in the transit costs. Given the assumptions of the model, where traffic is associated with the market share, this occurs when the difference in the traffic volume of the two ISPs does not exceed a certain value k . Indeed, when the traffic generated differs significantly, the larger provider mainly routes its traffic within its network, and the fees paid to the upstream IBP are minimal. Hence, the large provider's dominant strategy is not to peer, while the small provider would be better off in case of peering.

between two Internet Service Providers that compete á la Bertrand¹⁸, obtaining the opposite result as Crémer *et al.* (2000). They present a two stage game: in the first stage, the two Internet Operators choose the interconnection quality, while in the second the two firms compete over end customers. In this setting, where also the assumption of the Operators having an installed customer base is removed, the network externality effect is the driving force that leads the firms to increase the interconnection quality. Mason (1999) studies competition between ISPs that are both horizontally and vertically differentiated, obtaining results in line with Foros and Hansen (2001).

The network externality effect is also relevant in Baake and Wichmann (1999). In their model two Internet Operators competing a la Cournot are interconnected through a backbone, and the interconnection quality can be improved by direct peering. Baake and Wichmann show that the peering decision may be profitable even if leads to a lower market share (because of the business stealing effect) for one of the networks; indeed, both networks may charge higher prices for the increased quality of the service offered after that peering is realised¹⁹. On the same line, Economides (2005) shows that, *“with the same assumptions as Crémer et al. (2000) except now allowing for customer migration, the market equilibrium shows no (size) dominance by any firm and no network has incentive to degrade interconnection”*. Indeed, when customers can migrate, the interconnection degradation becomes unprofitable, and the possibility to exploit network externalities between

¹⁸ Preference for variety due to differentiation is driving the incentives for ISPs of interconnection in Giovannetti (2002).

¹⁹ The effect of an increase in the interconnection quality on Operator *i*'s profit can be divided into three main components: a direct effect, an indirect effect and the *business stealing* effect. The direct effect is positive given the assumptions in the model, and its value depends on both cost and network effects: an increase in the interconnection quality lowers the cost paid for transit, and also it increases the perceived network size for *i*'s customers, and hence the price they are willing to pay. The indirect effect, which also depends on both a cost and a price component, is negative. This effect is

operators leads to an increase in interconnection. This result is particularly relevant given the development of ISP multihoming²⁰, since it allows greater customer migration between different upstream providers. Hence, while in Crémer *et al.* (2000) even a slightly larger network will refuse to interconnect with other networks, in Economides (2005) network externalities and demand for universal connectivity will force networks to interconnect. In this setting, other strategies, such as increase in the prices of the service offered, are more profitable than degrading interconnection. The role of network externalities (modelled by the weight that consumers attribute to congestion and connection failure when choosing the provider) is present also in Badasyan and Chakrabarti (2003). They study the incentives of Internet providers, already connected through a National Access Point (NAP), to engage in private peering. Contrarily to the other models, in this work the peering decisions are endogenous, following the theory of endogenous network formation²¹.

4 Gathering the data and classification of Internet Operators

4.1 *Inferring the commercial agreements*

Obtaining data from Internet Operators is a particularly difficult task; almost everything that is relevant to the Economic Research is labelled “confidential”: prices, traffic flows, commercial agreements, and so on. Our

strengthened by the business stealing. The combination of the effects illustrated above makes it possible that peering might still be profitable despite losing market share.

²⁰ An ISP is multihomed when it has two or more upstream providers (large backbones or regional backbones). The main reason to multihome is that it permits to maintain full connectivity even if one of the upstream providers has huge problems. The rationale behind ISP customers multihoming is exactly the same.

²¹ Badasyan and Chakrabarti (2003) consider both the Bala and Goyal (2000) fully non-cooperative approach, where Internet Operators signal their willingness to engage in peering, and peering is realised when a reciprocal will is found, and the Jackson and Wolinsky (1996) approach, where mutual consent is needed for the peering to be reached.

interest lies in the study of commercial agreements. A source of information is available on the websites of some Internet Exchange Points; in particular, these websites provide a symmetric matrix (the peering matrix) with entries 0 or 1, where 1 indicates the presence of interconnection (through peering or transit). The major drawback associated to these data is that it is not possible to analyse the strategic decisions between peering and transit. In the present work we overcome this problem, following recent developments in the field of *Theoretical Computer Sciences*. Indeed, we apply recently developed algorithms in order to infer the actual bilateral business relationship between any given pair of Internet providers from publicly available data.

The algorithms used to infer the business relationships can be grouped into two main categories, depending upon the source of data on which the inference is based upon:

- Inference from Border Gateway Protocol (*BGP*)
- Inference from the Internet Routing Registry (*IRR*)

The Border Gateway Protocol is a series of “instructions” that govern the transmission of packets over the Internet through connected independent networks. These instructions govern the micro-specification of the interconnection policies established between Internet Operators²². These policies, specified in the BGP data set, represent a “*information treasure*” for our research .

Our second source of data is derived from Internet Routing Registries. These *IRRs* are large databases where Internet Service Providers willingly publish their routing policies²³. More specifically, the data we used were obtained mainly by using the algorithm devised by Huber *et al.* (2004), based

²² Technically known as *Interdomain Routing*.

²³ Routing policies mainly consists of two elements: route preferences and filtering policies; route preferences indicate, when multiple routes to the same destination are available, which one is

on the Internet Routing Registry²⁴; this information was complemented with inference based on the *BGP* tables²⁵ (Gao, 2001; Subramanian *et al.*, 2002; Di Battista *et al.*, 2003)²⁶.

4.2 Units of Analysis: IBPs and ISPs

Following Filstrup (2001) and Weiss and Shin (2004), we differentiate between three classes of providers among Internet Operators, according to their “size”: Tier-1, Internet Backbones (IBPs) and Internet Service Providers

preferred; filtering policies are instead used in order to hide some of the exported routes, or to filter some of the routes imported from Internet Operators.

²⁴ The fact that the information provided in the IRR by the Internet Operator is merely voluntary led to the beliefs that the IRR is poorly maintained, with obvious consequences on the actual reliability of the inferred relationships. However, Siganos and Faloutsos (2004) were able to derive a relatively large subset of data from IRR that were up to date and consistent with the observed *BGP* tables. .

²⁵ This approach starts by using the *BGP* table paths to derive an undirected graph that connects providers, (Autonomous Systems, ASes). Then it infers the existing the commercial relationships from these paths. A central assumption for this inference is that valid paths are *valley free*; in other words, in any path there can be only one consecutive chain of upstream relationships and one consecutive of downstream relationships: the path starts with an AS, which is customer of the next upstream provider, and so on until the path reaches a peak, where it starts to descend. The economic logic of the *valley free* assumption is straightforward, nobody would like to act as the valley AS, paying two upstream providers just to transfer traffic neither originated nor terminated at this node. The inference of the commercial relationships can be seen as a two step process. In the first stage, given the undirected graph obtained from the *BGP* tables, the following *Type of Relationship* problem is solved: “Given an undirected graph G , a set of paths, and an integer k , find an orientation to *all the edges* of G such that the number of invalid paths is at most k ”. In the second step, the directed graph obtained as the solution of the previous problem is refined to introduce peering relationships. The problem to be solved is the following: “Given an undirected graph G , a set of paths, and an integer k , find an orientation to *some of the edges* of G such that the number of invalid paths is at most k ”. The first attempt in this direction is due to Gao (2001). The algorithm used by Gao bases the inference on the degree of each node (the degree of a node is defined as the number of edges that touch that node), considered an indicator of the AS’s size. Subramanian *et al.* (2002) analyse the *BGP* tables-related graph from different vantage points, and base the inference on a probability measure attached to each edge orientation. Di Battista *et al.* (2003) introduce a new algorithm that reduces the number of invalid paths estimated with the approach of Subramanian *et al.* (2002). Dimitropoulos *et al.* (2005) provide some arguments against the approach of Subramanian *et al.* (2002) and Di Battista *et al.* (2003), showing that other approaches that are not devoted to minimise the number of invalid paths produce more realistic results. An evaluation of the inference methods is provided by Xia and Gao (2004). They find that both the Gao approach and the Subramanian *et al.* approach are very effective in detecting transit relationships, while the accuracy for peering is significantly lower.

²⁶ The drawbacks characterising the *BGP* approach depends instead on the assumptions made to translate paths into commercial relationships. Xia and Gao (2004) evaluated several *BGP*-based inference approaches, showing that about 98% of the relationships inferred as transit are correct, while about 70% of the relationships inferred as peering are correct. Huber *et al.* (2004) find that the algorithm based on the IRR produces good inference with respect to the *BGP*-based inference.

(ISPs)²⁷. We follow a two-step process: firstly we classify the providers into the above categories, and then we perform our econometric analysis on the inferred interconnection patterns among ISPs.

The population of Internet Operators considered is given by the members of the London Internet Exchange Point (*LINX*), one of the most important Internet Exchange Points in Europe according to both number of members and traffic routed. Although it is not possible to find a clear cut point to separate Internet Operators into the categories of IBPs and ISPs, it is indeed feasible to approximately accomplish this task by looking at some “size” metrics. We use the *customer cone*, introduced by the Cooperative Association for Internet Data Analysis (*CAIDA*). Broadly speaking, the customer cone of an Internet provider is given by the number of the provider’s customers (i.e., the providers that buy transit services from the first), plus the providers’ customers’ customers, and so on²⁸. This metrics, which is the closest possible empirical estimate of “market share” is also used to rank the providers. We consider both the *customer cone* and the rank measure to separate the providers into IBPs and ISPs.

The original list of *LINX* members is given by 179 Internet Operators. 49 providers were deleted. We firstly deleted the smallest Operators, with a very low customer cone²⁹, classified as Internet Access Providers (IAPs); as

²⁷ Today there are less than 10 Tier-1 providers and over 40 Internet backbones, and their number is increasing. Tiers-1 are characterised by the fact that they exchange traffic between them through peering, while they have generally only transit agreements with ISPs. There are more than 10,000 ISPs; they obtain universal connectivity through multiple interconnections with Tier-1 and or backbone providers (through transit or peering).

²⁸ CAIDA provides three alternative measures of the customer cone of a given Autonomous System (an Autonomous System, or AS, is a network that is administered by a single set of management rules that are controlled by one person, group or organization). The simplest measure of the customer cone of a certain AS is given by the number of its customers (other ASes), its customers’ customers, and so on. A more precise measure considers instead not the number of customers in the cone but the total number of prefixes that they advertise. Each prefix consists of several /24-address-space-segments, hence the most precise measure of customer cone of a certain AS considers the total number of /24-network-segments contained in all its customers. We use the #/24-network-segments metric to rank the ISPs, since this is the metric that promises the least number of inaccuracies.

²⁹ Measuring between 0 and 16 units. Twelve providers were classified as IAPs.

seen in Section 2, IAPs are small providers below the category of ISPs in the Internet hierarchy. The other providers were deleted since we could only estimate a few interconnection agreements for them; where these operators constitute a relevant proportion of the LINX members, there not seems to be any selection issue involved with their exclusion. Among these providers, in fact, there are several non commercial Operators belonging to organizations such as APNIC (Asia Pacific Network Information Centre), ARIN (American Registry for Internet Numbers)³⁰ and Réseaux IP Européens (RIPE). Moreover, another issue to be considered is the presence in our sample of mirror providers and replica ASes³¹. Finally, the reasons for excluding other providers seem to be due essentially to their poor maintenance of the IRR database³².

Among the remaining 130 providers, we individuated 5 *Top Tier-1* Operators (Level3, Global Crossing, CWA, UUNet, NTT/Verio); these providers have *customer cone* greater than 4,000,000 units³³. The group of IBPs (18) is given by the providers with rank below 50; these providers are all characterised by *customer cone* between 3,600,000 and 3,500,000 units. Finally, the set of ISPs (98) consists of the providers having rank greater than 50 and *customer cone* lower by at least one order of magnitude with respect to the IBPs; this category is very heterogeneous, containing providers with *customer cone* between 380,000 and 16 units.

The following **figure 2** represents the inferred commercial agreements for the class of Internet Service Providers at LINX. The Internet Operators

³⁰ APNIC and ARIN are present at LINX with the Operators AS2914, AS2828, AS4788, AS13768, AS22822.

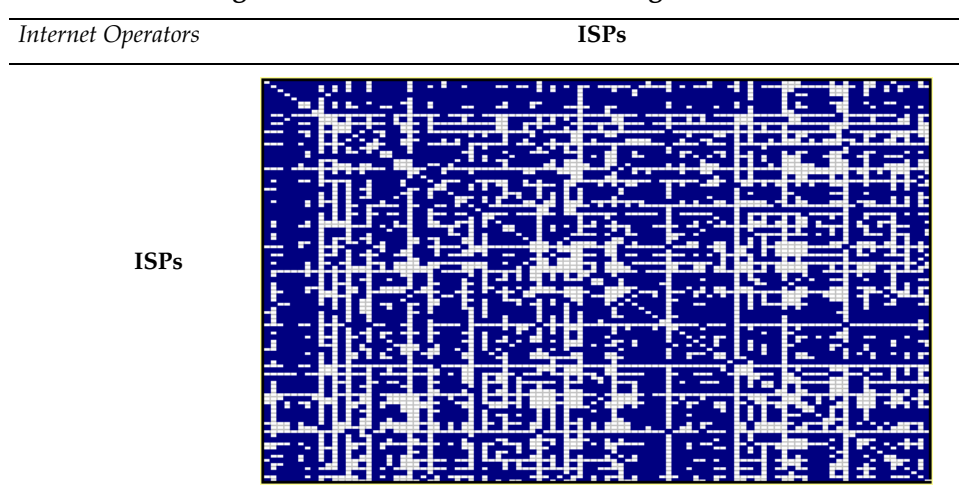
³¹ AS3741 is for instance a Mirror AS created by AS27822 to express its routing policy within the RIPE database. AS25310 is a “replica AS” for Cable and Wireless, already present in the LINX with the main AS3561.

³² Again, it does not seem to exist a possible selection issue, since these latter providers have very different sizes and market power.

³³ The units of measurement employed, described in the previous footnote, is /24s.

are sorted according to their increasing rank³⁴ in the Internet hierarchy. Each square of the symmetric matrix shows the inferred agreement between the pair of providers indicated by the corresponding row and column. A dark dot indicates a peering relationship, while a white dot indicates that the two providers exchange traffic using their transit agreements with upstream providers.

Figure 2: Inferred Interconnection Agreements



5 ISPs interconnection model

This section is devoted to the econometric analysis of the interconnection relationships among competing Internet Operators. As we argued before, we focussed on the class of Internet Service Providers that are members of the London Internet Exchange Point.

5.1 Empirical specification

The interconnection patterns between ISPs are expressed by a binary model, with the two possible outcomes given by *peering* and *indirect interconnection*;

³⁴ Notice that a larger rank corresponds to a lower customer cone.

in the latter case, the providers will exchange traffic by using their upstream providers as intermediaries. 98 ISPs were considered, giving rise to 4753 pairs; among these, 2674 were inferred as connected through peering, while 2079 were inferred as connected through upstream providers.

The dependent variable is the peering decision, assuming value 1 when peering between the pair of providers occurs and 0 otherwise. The explanatory variables are devised to model the competitors' asymmetry, the geographical differentiation (in terms of both headquarters location and IXPs coverage) and some technical elements, such as the *hot potato* routing. As we have seen before, *hot potato* routing refers to the fact that carrying traffic in the Internet is costly, and providers have the incentive to deliver traffic following the shortest way to the destination network³⁵.

The peculiar nature of the Internet asks for the utilisation of different metrics to assess the asymmetry between any pair of providers. The first measure we considered is the difference in the providers' *customer cones*. In particular, the *customer cone* is used as a proxy for market shares: for any pair of providers, the difference in their *customer cones* (*diff_base*) gives a market share-based measure of asymmetry.

The second measure introduced involves instead a market power-based measure of asymmetry, given by the difference in the providers' *betweenness* (*diff_centrality*). This metrics is derived from BGP paths. Each one of these paths provides the instructions indicating the sequence of different providers that a given traffic unit (called information packet) should follow, starting from the originator provider to reach its final destination³⁶.

³⁵ While assessing their incentives towards peering providers will take into account the possibility of delivering traffic to the peer's network as soon as possible; this means that mutual presence at more exchange points is thought to positively affect their incentives to do peering.

³⁶ Indeed, each path specifies with which other networks one provider should interconnect to deliver its off-net traffic.

Typically there are multiple paths available to reach the same *off-net* destination for traffic with the same origin. In this case we focus on the shortest path: given that carrying traffic is costly, the shortest paths are often preferred to others. It is clearly an advantage, for a provider, to appear in as many shortest paths as possible, in the sense it becomes an *almost unavoidable* step for Internet traffic going from and to other providers. We capture this notion of *network centrality* by using a simple measure: the number of shortest paths an operator can be found in. We calculated this metrics, known in the literature as *betweenness centrality*,³⁷ for each Internet Operator v ³⁸:

$$B_s(v) = \sum_{s \neq v \neq t \in V} \sigma_{st}(v)$$

where $\sigma_{st}(v) = \sigma_{ts}(v)$ is the number of shortest *BGP* paths from the Internet Operator s to the Operator t on which the v lies on. Hence, *betweenness* expresses, from a *network's topology* aspect, the market power of any given provider by showing how unavoidable it is, in the Internet traffic flow paths, given the set of existing interconnection policies³⁹.

In order to take into account also possible size effects, we introduce in the estimation two further variables for each pair of providers: customer cone of

³⁷ Introduced by Shimbel (1953).

³⁸ D'Ignazio and Giovannetti (2006) have used this metric to assess HHI market concentration indexes, we focus instead on the micro bilateral interconnection choices.

³⁹ The use of the difference in the betweenness measure could raise some endogeneity issues, in the sense that betweenness inevitably depends on the actual peering relationships pattern. A large number of peering agreements is normally reflected in high betweenness. In order to tackle this problem we introduced another regressor, aimed to capture this "size effect", given by the "maximum betweenness" for each pair of providers. Once we control for this effect, the difference in the betweenness between any pair of providers does not seem to depend much on their eventual peering relationship; indeed, if this were the case, then the peering link would increase both the providers' betweenness in a similar measure, with very little effect on the difference.

the largest provider (*max_base*), and betweenness of the largest provider (*max_centrality*)⁴⁰.

We also focus on the possible role that geographical differentiation can play in the peering decision. It is often argued that proximity of the operators will facilitate mutual knowledge and trust. On the other hand, peering with a provider located further away will provide high quality interconnection with a differentiated customer base. Geographical differentiation can exert a positive impact on peering if two providers, located further away, perceive themselves more as complements than as substitutes.

Geographical differentiation is captured by two independent variables. The first, *dist_hq*, expresses the distance (in thousands of miles) between the headquarters of the Internet Operators. The distance was calculated following a two steps process: first, we located each Internet Operator by considering the latitude and longitude of its headquarter; then we estimated the distance between headquarters using the great circle distance rule⁴¹. The second variable, *diff_ixp*, takes into account the different geographical coverage: for any pair of providers, it represents the difference in the number of memberships among the most important Internet Exchange Points all over the world⁴² that they have.

⁴⁰ We are thankful to Daniel Akerberg for precious suggestions on the econometric specification of our model.

⁴¹ $dist(Operator1 - Operator 2) = RadiusEarth * ArcCos(Cos(Radians(90 - Lat1)) * Cos(Radians(90 - Lat2)) + Sin(Radians(90 - Lat1)) * Sin(Radians(90 - Lat2)) * Cos(Radians(Long1 - Long2)))$

⁴² We considered 45 IXPs. All the 35 members of Euro-IX were included (Aix Athens, Ams-ix Amsterdam, Bcix Berlin, Bix Budapest, Bnix Brussels, Catnix Barcelona, Cixp Geneva, De-cix Frankfurt, Espanix Madrid, Ficix Helsinki, Gigapix Lisbon, Gn-ix Groningen, In-ex Dublin, Lix Luxembourg, Mix Milan, Msk-ix Moscow, Namex Rome, Ndix Enschede, Netnod Stockholm, Nix Oslo, Nix.cx Prague, Nota Miami, Parix Paris, Ronix Bucharest, Six Ljubljana, Tix Zurich, Topix Torino, Vix Vienna, Linx London, Lipex London, Lonap London, Manap Manchester, Xchange point London, Equinix 7 locations USA, Jpnap Tokyo). Other European IXPs were included (Free-ix Paris, Inxs Munich, NI-ix Amsterdam, Swiss-ix Zurich) and Extra-European IXPs (Ape Auckland, Hk-ix Hong Kong, Jp-ix Tokyo, Nyi-ix New York, Six Seattle, Tor-ix Toronto).

In order to model the technical elements behind the *hot potato* routing effect, discussed before, we constructed a variable, *both_ixp*, indicating, for each pair of providers, the number of IXPs at which they are both present⁴³. This variable could also be interpreted as expressing a reputation effect, following Titley (1997). Apart from the difference in the *customer cone* measures, which was built using the March 2005 CAIDA database, all the remaining data, including the interconnection agreements inference, were gathered in July 2005.

Table 1: Description of the variables

dependent variable	
<i>peering</i> (dummy)	Assumes value 1 in case of peering between providers, 0 otherwise.
independent variables	
<i>diff_base</i>	Difference in the customer cone for any pair of providers in units of thousands
<i>max_base</i>	Customer cone of the largest among the two providers
<i>dist_hq</i>	Distance (thousands of miles) between the headquarters of the two providers
<i>both_ixp</i>	Number of IXPs in which both the providers are present.
<i>diff_ixp</i>	Difference in the number of IXPs in which both the providers are present
<i>diff_centrality</i>	Difference in the <i>betweenness</i> measure in thousands of units.
<i>max_centrality</i>	Betweenness of the largest among the two providers

⁴³ In order to generate this matrix of data we created a *visual basic* routine that cross-checked the memberships for each pair of providers among the most important IXPs all over the world. See footnote 33 for the list of IXPs considered.

5.2 Estimation results

We estimated a logit model by maximum likelihood. The presence of multiple observations for each ISP in our dataset is likely to lead to correlated residuals; we decided to tackle this problem by adding ISP fixed effects⁴⁴. The results are reported below.

Table 2: ISPs binary model results

independent variable	Coeff.	Std. Err.	z	P>z
<i>diff centrality</i>	-.309	.019	-15.80	0.000
<i>max centrality</i>	-.037	.050	-0.73	0.464
<i>diff_ixp</i>	.152	.030	5.13	0.000
<i>both_ixp</i>	.809	.086	9.44	0.000
<i>dist_hq</i>	.060	.027	2.20	0.028
<i>diff_base</i>	-.013	.011	-1.13	0.260
<i>max_base</i>	.033	.012	2.79	0.005
Number of Observations	4753			
Pseudo R-Square	R2 = 0.3912			
Log pseudolikelihood	-1982.9028			

Table 3: ISPs binary model, partial effects

dependent Variable: $P(y=1 x)$		Std. Err.	z	$P>z$	\bar{x}
independent variable					
<i>diff_centrality</i>	-.0764	0.004810127	-15.80	0.000	3.39818
<i>max_centrality</i>	-0.00907	0.012424658	-0.73	0.464	2.22423
<i>diff_ixp</i>	.0375131	0.007312495	5.13	0.000	2.1843
<i>both_ixp</i>	.1993277	0.021115222	9.44	0.000	1.46118
<i>dist_hq</i>	.0148	0.006727273	2.20	0.028	2.16791
<i>diff_base</i>	-0.00315	0.002787611	-1.13	0.260	20.0957
<i>max_base</i>	0.00801	0.002870968	2.79	0.005	21.2218

All the variables introduced are statistically significant, but the difference in the *customer cone* and the maximum value for the *betweenness*. The two variables representing the competitors' asymmetry seem to affect peering in the same way. Indeed, both the difference in the *betweenness*, which has also the highest z statistic, and the difference in the *customer cone*, which is however not statistically significant, are negatively related to peering. This result supports the claim that the quality of interconnection degrades as the asymmetry increases.

A possible interpretation relies upon the fact that *customer cone* expresses asymmetry in "size", and the *betweenness* expresses asymmetry in "market power associated to *unavoidability*". The asymmetry in size can be seen as a "installed base of customers" element, which negatively affects peering, like in Cremer *et al.* (2000) and Jahn and Prüfer (2004). On the other hand, the asymmetry in the *betweenness* expresses difference in the bargaining power associated to the traffic routing; moreover, since high *betweenness* presumably implies a large traffic, this measure of asymmetry may also indicate traffic imbalances between pairs of providers. This result seems to support Weiss and Shin (2004); moreover, it also seems to show

⁴⁴ To do so, we introduce as many dummy variable as the number of ISPs. For each observation involving any two providers, the two relevant dummy variables are set equals to one.

that both the *backbone free riding* and the *business stealing effects* seem to play a decisive role against peering.

Our results also indicate that peering seems more likely when the geographical differentiation increases: both the distance between headquarters, as well as the geographic IXP's coverage, positively affect peering. In this sense, we can see some support also for that part of research recognising the importance of differentiation in interconnection decisions (Foros and Hansen, 2001, Economides, 2005). Finally, the mutual presence at several IXPs increases the chances of peering, following the logic of the *hot potato effect*; an alternative interpretation of this result lies instead on the importance of knowledge and reputation effects on peering decisions (Titley, 1997). The estimated partial effects (see Appendix) provide some evidence about the magnitude of the covariates' effects on peering.

6 Conclusions

In recent years, many game theoretic models have analysed the incentive structure underlying the interconnection agreements between Internet Operators. This research has also been playing an increasingly relevant role in informing recent Competition Authorities decisions in relevant Internet antitrust cases. The main issue at stake is whether or not the asymmetry between Internet Operators affects the quality of their interconnection modalities, by providing incentives to interconnection quality degradation. While theoretical models provide contrasting results, there is a lack of empirical analysis on this issue. This paper is an attempt to fill this gap: we provided an empirical analysis thanks to a novel approach to obtain data about interconnection regimes, which are otherwise usually kept confidential by the Internet Operators. In particular, we exploited some recent advances in the field of Theoretical Computer Science providing the

tools to infer aspects of the business' nature of interconnection agreements from publicly available data.

Our model focused on the interconnection patterns between competing Internet Service Providers (ISPs) at the London Internet Exchange Point (LINX). We investigated if asymmetry is associated to quality degradation, expressed by the systematic absence of peering between providers of different size. We introduced two distinct metrics to model the providers' relevance, and therefore asymmetry: the *customer cone*, providing a proxy for "market share", and the *betweenness*, expressing the market centrality of any given player, by showing its degree of *unavoidability* in the Internet traffic routing.

The binary model introduced showed that both the *customer cone* based and the *betweenness* based measures of asymmetry have a negative effect on the likelihood of establishing a peering relationship. Therefore, asymmetry seems to consistently provide incentives towards a *quality degraded* form of interconnection. With the *customer cone* picking up the installed base of customers, our results seems to show some support for Cr mer, Rey and Tirole, (2000) and Jahn and Pr fer (2004) although, in our data, there is little statistical significance for this effect. Definitely more significant is the effect associated to asymmetry measured in terms of network centrality, expressing relative market power as well as traffic imbalances. In this latter interpretation, our analysis provides empirical support to the theoretical results obtained by Weiss and Shin (2004). Hence, our results suggest that Competition Authorities should mostly be concerned about the "centrality of a player", rather than its market share, to avoid quality degradation strategies adopted by bigger providers towards smaller ones. So far, however, the Competition Authorities based their antitrust decisions on market shares analysis.

On the other hand, the results obtained by Foros and Hansen (2001), and Economides (2005), pointing to the role played by differentiation and network externalities in driving the peering decision⁴⁵ are captured in our analysis with the estimated positive effects on the likelihood of observing bilateral peering induced by geographical distance and difference in the extent of markets covered.

References

- BAAKE, P. AND WICHMANN, T. (1999) "On the Economics of Internet Peering". *Netnomics* 1 (1), 89-105.
- BADASYAN, N. AND CHAKRABARTI, S. (2003) "Private Peering among Internet Backbone Providers", WUSTL: Working Paper n. 0301002.
- BALA, V. AND GOYAL, S. (2000) "A Noncooperative Model of Network Formation", *Econometrica* 68, 1181-1229.
- BAR, S., GONEN, M. AND WOOL, A. (2005) "A Geographic Directed Preferential Internet Topology Model". *arXiv:cs.NI/0502061 v1* [<http://arxiv.org/abs/cs.NI/0502061>]
- BESEN, S., MILGROM, P., MITCHELL, B. AND SRINAGESH, P. (2001) "Advances in Routing Technologies and Internet Peering Agreements". *The American Economic Review* 91(2), 292-296.
- BUCCIROSSI, P, FERRARI BRAVO, P. AND SICILIANI, P. (2005) "Competition in the Internet Backbone Market". *World Competition* 28 (2), 235-254.
- CRÉMER, J., REY, P. AND TIROLE, J. (2000) "Connectivity in the Commercial Internet". *Journal of Industrial Economics* 48 (4), 433-472.
- COURCOUBETIS, C. AND WEBER, R. (2003) *Pricing Communication Networks: Economics, Technology and Modelling*. Wiley, England.
- DI BATTISTA, G., PATRIGNANI, M. AND PIZZONIA, M. (2003) "Computing the Types of the Relationships between Autonomous Systems". In *Proc. IEEE INFOCOM*.
- D'IGNAZIO, A. AND GIOVANNETTI, E. (2006) "Antitrust Analysis for the Internet Upstream Market: A Border Getaway Protocol Approach". *Journal of Competition Law & Economics* 20(5), 757-796
- DIMITROPOULOS, X., KRIOUKOV, D., HUFFAKER, B., CLAFFY, KC AND RILEY, G. (2005) "Inferring AS Relationships: Dead End or Lively Beginning?" *Lecture Notes in Computer Science* 3503, [http://www.caida.org/outreach/papers/2005/asrelationships/as_relationships.pdf]

⁴⁵ The importance of the network externality effect is also suggested by Baake and Wichmann (1999) and Badasyan and Chakrabarti (2003).

- ECONOMIDES, N. (1998) "The incentive for non-price discrimination by an input monopolist". *International Journal of Industrial Organization* 16, 271-284.
- ECONOMIDES, N. (2006) "The Economics of the Internet Backbone", forthcoming in S. Majumdar, I. Vogelsang and M. Cave (eds), *Handbook of Telecommunications Economics*, Volume 2. Amsterdam, North Holland.
- FILSTRUP, B. (2001) "Internet Interconnection Agreements", Final Project Report, S1646: Information Economics.
- FOROS, Ø. AND HANSEN, J. (2001) "Competition and compatibility among Internet Service Providers". *Information Economics and Policy* 13, 411-425.
- FOROS, Ø., KIND, H. AND SØRGARD, L. (2002) "Access Pricing, Quality Degradation and Foreclosure in the Internet". *Journal of Regulatory Economics* 22, 59-83.
- GAO, L. (2001) "On Inferring Autonomous System Relationships in the Internet." In *IEEE/ACM trans. Networking*.
- GIOVANNETTI, E., (2002). "Interconnection, differentiation and bottlenecks in the Internet". *Information Economics and Policy* 14(3): 385-404
- GIOVANNETTI, E. AND RISTUCCIA, C. (2005) "Estimating Market Power in the Internet backbone using the I.P. transit Band-X database". *Telecommunications Policy* 29, 269-284.
- HECKMANN, O., SCHMITT, J. AND STEINMETZ, R. (2004) "Optimizing Interconnection Policies". *Computer networks* 46, 19-39.
- HUBER, B., LEINEN, S., O'DELL, R. AND WATTENHOFER, R. (2004) "Inferring AS Relationships Beyond Counting Edges" D-INFK Tech Report Nr. 446.
- IDA, T. (2005) "Analysis of Internet Topology with a Three-Level Components Model". *Managerial and Decision Economics* 26, 527-534.
- JACKSON, M. O. AND WOLINSKY, A. (1996) "A Strategic Model of Social and Economic Networks". *Journal of Economic Theory* 71, 44-74.
- JAHN, E. AND PRÜFER, J. (2004) "Transit versus (Paid) Peering: Interconnection and Competition in the Internet Backbone Market",
[http://papers.ssrn.com/sol3/papers.cfm?abstract_id=590582]
- KENDE, M. (2000) "The Digital Handshake: Connecting Internet Backbones", OPP Working Paper Series No. 32, Federal Communications Commission.
- MAH, D. (2005) "Explaining Internet Connectivity". *The Information Society* 21 (5), 353-366.
- MASON, R. (1999) "Compatibility between Differentiated Networks". University of Southampton Discussion Paper in Economics and Econometrics, No. 9909.
- MILLER, R. (2002) "The Economics of Peering", Carrierhotels.com.
- NORTON, W.B. (2002) "Internet Service Providers and Peering", Equinix Inc.

- SHIMBEL, A. (1953) "Structural parameters of communication networks". *Bulletin of Mathematical Biophysics* 15, 501-507.
- SIDAK, G. (2006) "A Consumer-Welfare Approach to Network Neutrality Regulation of the Internet". *Journal of Competition Law and Economics* 2 (3), 349-474.
- SIGANOS, G. AND FALOUTSOS, M. (2004) "Analyzing BGP Policies: Methodology and Tool". In *Proc. IEEE INFOCOM*.
- SUBRAMANIAN, L.; AGARWAL, S.; REXFORD, J. AND KATZ, R. (2002) "Characterizing the Internet hierarchy from multiple vantage points." In *Proc. IEEE INFOCOM*.
- TITLEY, N. (1997) "An Analytical Model of Peering between Internet Service Providers", Release 1.
- WEISS, M. B. AND SHIN, S. J. (2004) "Internet Interconnection Economic Model and its Analysis: Peering and Settlement". *Netnomics* 6 (1), 43-57.
- WOODCOCK, B. (2000) "BGP for Bankers (White Paper on Transactions and Valuation Associated with Inter-Carrier Routing of Internet Protocol Traffic)", Packet Clearing House.
- XIA, J. AND GAO, L. (2004) "On the Evaluation of AS Relationship Inferences."
- XU, K., DUAN, Z., ZHANG, Z. AND CHANDRASHEKAR, J. (2004) "On Properties of Internet Exchange Points and Their Impact on AS Topology and Relationship", *Networking*, [<http://www.cs.fsu.edu/~duan/publications/networking04.ps>].

Appendix

Table A: ISPs binary models, variables summary statistics

Variable	Mean	Std. Dev.	Min	Max
<i>peering = 1 (2674 obs)</i>				
<i>diff_base</i>	24.73925	64.24227	0	380.151
<i>max_base</i>	26.50060	65.85295	0.032	380.167
<i>dist_hq</i>	2.272856	2.40177	0	12.20257
<i>both_ixp</i>	1.586761	0.888546	1	6
<i>diff_ixp</i>	2.280853	1.726976	0	8
<i>diff_centrality</i>	1.628163	2.388645	0	16.45
<i>max_centrality</i>	1.875564	2.092597	0.024	16.067
<i>peering = 0 (2079 obs)</i>				
<i>diff_base</i>	14.12309	39.74619	0	380.151
<i>max_base</i>	14.43220	40.28598	0.016	380.167
<i>dist_hq</i>	2.032935	2.286581	0	12.2011
<i>both_ixp</i>	1.299663	0.60885	1	5
<i>diff_ixp</i>	2.060125	1.638311	0	8
<i>diff_centrality</i>	5.674763	6.65285	0	16.414
<i>max_centrality</i>	2.762608	3.932070	0.024	16.067