

“Unfair” Discrimination in two-sided Peering? Evidence from LINX[◇]

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

^a*University of Cambridge*

^b *University of Rome “La Sapienza” and School of Economics UCT*

Abstract

Does asymmetry between Internet Providers affect the “fairness” of their interconnection contracts? While recent game theoretic literature provides contrasting answers to this question, there is a lack of empirical research. We introduce a novel dataset on micro-interconnection policies and provide an econometric analysis of the determinants of *peering* decisions amongst the Internet Service Providers interconnecting at the London Internet Exchange Point (LINX).

Our key result shows that two different metrics, introduced to capture asymmetry, exert opposite effects. Asymmetry in “market size” enhances the quality of the link, while asymmetry in “network centrality” induces quality degradation, hence “unfairer” interconnection conditions.

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

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) for their help, useful comments and data. The usual disclaimer applies. The authors would also like to acknowledge EU financing through the 6th Framework Project IST-2004-2012 and from the Isaac Newton Trust, Trinity College, University of Cambridge.

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 two-sided interconnection decisions 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” clauses associated to the bilateral exchanges of traffic.

A growing literature is focussing on the issue of access pricing and interconnection quality 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 (over downstream customers). One stream of the theoretical debate² focussed on the main trade off underlying the interconnection choice: a better quality of the interconnection between a pair of providers has a positive effect on the networks’ performance (network externality effect) and may justify higher prices for the service offered; however, it also reduces the networks’ differentiation in terms of quality, which may result in a loss of customers for the “higher quality” network to the advantage of the “smaller quality network”(business stealing effect). The relative size of

¹ During the 1998 MCI WorldCom merger analysis the European Commission included in the backbone market all the providers which were able to obtain global connectivity either through private or public peering, needing no transit contracts. This definition was subsequently modified: only the providers reaching global connectivity exclusively via private peering were included in the backbone market. See Official Journal Of the European Commission (2000), Regulation (EEC) N 4046/89, Merger Procedure, Bruxelles, European Commission, DGXIII. See also Buccirosi *et al.* (2005).

² 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), to name just a few.

these two effects determines the eventual extent of quality degradation in interconnection, decided by the larger (better quality) providers.

The theoretical literature on interconnection games provided contrasting results on this trade-off, motivating the need for more empirical research. This is however still scarce, and mainly anecdotic, essentially because of the confidentiality that characterizes the two-sided interconnection agreements and Internet traffic data. Our work provides a contribution in this direction: this is possible thanks to a novel approach to obtain data, which follows recent advances in the fields of Theoretical Computer Science³.

We assess the claim that asymmetry between Providers affects the “fairness” of their interconnection contracts. In other words, we investigate if asymmetry is associated to quality degradation, expressed by the refusal of peering⁴. We focus on the interconnection decisions involving providers that compete in the same market, considering, in particular, the large number of *two by two* interconnection decisions amongst the Internet Service Providers (ISPs) members of the London Internet Exchange Point (LINX).

The peculiar nature of the Internet asks for the analysis of different metrics⁵ to assess the competitor’s asymmetry. Our key result, based on the econometric analysis of the ISPs interconnection patterns, shows that the two asymmetry’s metrics introduced have opposite effects on the bilateral peering decisions. Asymmetry in “*market size*” facilitates peering, while asymmetry in “*network centrality*” makes peering less likely and hence

³ The interest of Computer Scientists and Mathematicians focus for the actual interconnection agreements is linked to the representation and analysis of the evolution of the Internet topology (Bar *et al.*, 2005), and its efficiency from a Network optimisation point of view (Heckmann *et al.*, 2004) .

⁴ Peering and Transit are the two main categories of interconnection agreements. See below for the definitions.

⁵ We will indeed focus on two different notions of size, both explained below, one, the customer cone, providing a proxy for “market share” and the second, *betweenness*, expressing the market power of any given players, by showing how unavoidable a given operator is in the Internet traffic flows, given the set of existing Interconnection policies.

interconnection conditions “unfairer”. This clear result should not be surprising when thinking about the nature of network industries and specifically of the Internet, where, given the richness and redundancy of possible interconnection modalities, necessarily one has to consider more than one dimension for the measurement of size, and therefore of asymmetry.

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 results on interconnection agreements. Section 4 explains the process of data gathering and the criteria used to classify the Internet Operators. Section 5 provides the econometric analysis of the relevant model. Section 6 concludes.

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*⁶, 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: each user is able to access each other; this is only possible thanks to the system of bilateral interconnections between the

⁶ 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).

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 customer is provided connectivity to the entire Internet, in front of a payment of a settlement fee, by the provider. In other words, the customer can send to the provider traffic that is destined anywhere and the provider will get the traffic there.
- The peering agreement leads to a bilateral *peer-to-peer* relationship: each peer provides the other connectivity to its own network (customers) only, usually without any settlement fee⁷.

A relevant incentive to the realisation of peering is that it leads to a reduction in traffic costs, since the Internet Operators do not have to pay for the traffic routed through peering networks⁸. Norton (2002a) notices that, although that argument is given for granted, this is not always true, since the costs incurred to connect the networks and maintaining the peering point may offset the advantages in terms of less transit costs. However, the costs of peering fell sharply in the recent years, after the development of the *Internet Exchange Points*⁹ (IXPs). Peering also leads to a better performance in the traffic between the two networks involved, through a lower latency in the transmission of packets.

Although the advantages from peering, outlined above, are relevant, many Internet Operators still exchange traffic through a transit relationship.

⁷ This is known as *Sender Keeps All* (SKA) peering, or *Bill and Keep* peering.

⁸ This is a general feature of the peering agreement. However, recently some networks reached paid peering interconnection (Miller, 2002).

⁹ Internet Exchange Points (IXPs) are meeting points that provide switches to the members ISPs, so that they can exchange traffic without the need to build dedicated extra circuits. However, this does not necessarily mean that private peering is more costly than peering within the IXP, since in the latter case the providers bear the cost of getting to the exchange point, the membership fees, and the cost for using the switches. 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,

This reflects the fact that peering is a bilateral agreement, and in order for peering to be realised the two networks must both find it profitable. A very consistent argument is that peering is based on the equality of size¹⁰ between the two providers (Norton, 2002a; Kende, 2000; Filstrup, 2001 to name just a few). Indeed, two commonly argued reasons seem to prevent peering between a large operator and a small operator: they are the so-called *backbone free riding* and the *business stealing effect*.

The *free riding* refers to the fact that the smaller network uses the capacity of the largest network for free. The *business stealing* refers to the following: larger size implies better performance of the network; if two networks of different sizes peer, then this barrier falls and the larger network may lose customers to the advantage of the, usually cheaper, smaller network.

Another element that is said to negatively affect peering is the fact that, contrary to transit, peering agreements do not have Service Level Agreements (SLAs) that guarantees rapid repair if problems occur. Norton (2003) argues that denying peering may lead to a transit relationship, from which the upstream provider benefits: this may prevent large upstream providers to engage in peering with “potential downstream customers”.

A further element involves the technological aspect of traffic routing. Since carrying traffic is costly, when a packet has to be delivered from a network *a* to a network *b* (either because there is a peering or transit relationship between them), 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

providing evidence that IXPs plays an important role in shaping the relationships between Internet operators.

¹⁰ According to Filstrup (2001), 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.

traffic to the destination network, which thus bears the large part of the transmission packet's costs. Hence, mutual presence at more exchange points is argued to positively affect peering. Interestingly, this conclusion is strengthened by Titley (1997): he maintains that the pairs of providers having several peering relationships enjoy a positive reputation effect, which enhances the likelihood of further peering being realised.

Other elements make the understanding of peering even more complex. Indeed, Norton (2002b) argues that also the way in which peering negotiations are conducted is likely to affect the final decision. The theoretical models of peering games all assume the so-called *direct approach*: the *initiator network* asks the *target network* for peering; then, the *target network* accepts only if he finds peering profitable. In reality, there are many different tactics. Norton presents nineteen manoeuvres used by the interviewed Peering Coordinators, showing that the choice of the right tactic is crucial, since it can lead to peering even when the direct approach would not work¹¹.

3 Game theoretic models of Internet peering

One of the earliest theoretical works on the interconnection strategies between competing 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 unattached consumers. They consider a two stage game. In the first

¹¹ Two tactics that proved to be very effective are the "transit with peering migration" and the "end run". The first manoeuvre consists in reaching a transit agreement to be muted into a peering agreement if the peering prerequisites fixed by the target network are reached. The second tactic consists of negotiating peering directly with the largest customers of the target network, in order to reduce in this way the costs of transit. "Less honest", but apparently quite effective, manoeuvres are the "traffic manipulation" and the "bluff". The first has the aim to the instigating ISP forces its traffic over the potential peers' transit services, to maximize the target ISP's cost of accessing its traffic; concerning the "bluff", sometimes the assertion of strength is not subject to proof by analysis of traffic, hence can work".

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 externality¹². The larger backbone gains competitive advantage on 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 he can offer a better quality service of the other¹³).

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 located in a *battlezone*¹⁴. 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 interconnection regime is based on the

¹² 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.

¹³ We referred before to this as the *business stealing* effect.

¹⁴ 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.

traffic volume, which, in turn, is linked to market share. Their model shows that symmetry in traffic positively affects peering¹⁵.

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 between two Internet Service Providers that compete à la Bertrand¹⁶, 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

¹⁵ 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.

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 that are Cournot competitors 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 realised¹⁷. 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 operators leads to an increase in interconnection. This result is particularly relevant given the development of ISP multihoming and ISP customers multihoming¹⁸, 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

¹⁶ Preference for variety due to differentiation is driving the incentives for ISPs of interconnection in Giovannetti (2002).

¹⁷ 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 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.

¹⁸ 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

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 formation¹⁹.

4 Gathering the data and classification of Internet Operators

4.1 *Inferring the commercial agreements*

Obtaining data on two-sided interconnection agreements from firms is a very difficult task. This is particularly true for the Internet, where almost everything that is relevant to the Economic Research is labelled “confidential”: prices, traffic flows, commercial agreements, and so on. Our 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 with entries 0 or 1 (the peering matrix), where 1 indicates the presence of interconnection (through peering or transit). The major drawback with these data is that it is not possible to analyse the strategic decisions between peering and transit. The dataset used

of the upstream providers has huge problems. The rationale behind ISP customers multihoming is exactly the same.

¹⁹ 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.

in the present work permits to overcome this problem. The strategy followed to create the dataset is linked to 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 data²⁰.

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. This protocol governs what is known as *Interdomain Routing*, the micro-specification of the two-sided interconnection policies established between competing providers in need of Universal connectivity. These policies, specified in the BGP data set, are indeed the empirical implementations of the interconnection strategies analysed in the game theoretic literature discuss above.

Our second source of data is derived from Internet Routing Registries. These *IRRs* are large databases where Internet Service Providers willingly publish their routing policies²¹. More specifically the data used in our econometric estimation of the drives of the bilateral interconnection decisions among Internet Operators where obtained mainly by using the algorithm devised by Huber *et al.* (2004), based on the Internet Routing

²⁰ The interest behind the inference of business relationships between Internet Operators in the field of Computer Science and Mathematics is linked to the analysis of the Internet Protocols used to transfer packets in the Internet.

²¹ 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

Registry; this information was complemented with inference based on the *BGP* tables²² (Gao, 2001; Subramanian *et al.*, 2002; Di Battista *et al.*, 2003). The fact that the information provided in the IRR by the Internet Operator is merely voluntary²³ 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* tables²⁴.

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.

²² This approach uses the *BGP* table paths to derive an undirected graph that connects Autonomous Systems. Then it makes use of some central assumptions to infer the commercial relationships from these paths. A central assumption 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 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-D-Simple* 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.

²³ There seem not to be any economic incentive or regulatory duty for Internet Operators to update the Internet Routing Registries.

²⁴ 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.

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 Operators²⁵: Tier-1, IBPs and ISPs. 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 the Internet structure is continuous and it is not possible to find a clear cut point to separate Internet Operators into the categories of IBPs and ISPs, it is possible to roughly accomplish this preliminary task by looking at the estimated *customer cone* of the Internet Operators. This concept, the closest possible empirical estimate of “Market Size” for an Internet Operator has been introduced by the Cooperative Association for Internet Data Analysis (*CAIDA*). They estimated the *customer cone* of the Internet Providers, and then ranked the providers according to this measure. We used both the *customer cone* and the rank measure to separate the providers into IBPs and ISPs, where the *customer cone* measure is given by the number of /24 address spaces contained in all the customers of each provider²⁶.


²⁵ 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 to multiple interconnections with Tier-1 and or backbone providers (through transit or peering). IAPs are typically tied to the Internet mainly through a single connection to an ISP, and do not generally participate in any peering points.

²⁶ 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). 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

The original list of *LINX* members is given by 179 Internet Operators. 49 providers were deleted since there were no estimated interconnection agreements²⁷. 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 4,000,000/24s. 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. 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.

The following **figure 1** represents the inferred commercial agreements for the class of Internet Service Providers. The table shows the symmetric peering matrix characterising ISPs. The Internet Operators are sorted according to their increasing rank²⁸ in the Internet hierarchy.

Figure 1: Inferred Interconnection Agreements



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.

²⁷ This is essentially due to lack of data in the Internet Routing Registry about the providers involved.

²⁸ Notice that a larger rank corresponds to a lower customer cone.

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 *LINX*.

5.1 Empirical specification

The interconnection patterns between ISPs are expressed by a binary model, with the two possible outcomes given by *peering* and *no interconnection*; indeed, the traffic agreement typically involves relationships between ISPs and IBPs only. The inferred relationships are consistent with this theoretical argument: 98 ISPs were considered, giving rise to 4753 pairs; among these relationships, 2674 were inferred as peering, while 2079 were inferred as no relationship.

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 (coverage), and some technical elements, such as the *hot potato* routing effect.

As we argued before, 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 considered is the difference in their *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_24s*) gives a market share-focussed measure of asymmetry.

The second measure introduced involves instead a market power-focussed measure of asymmetry, given by the difference in the providers' *betweenness* (*diff_bet*). Following Shimbel (1953) we calculated the following measure of *betweenness* centrality for each Internet Operator v ²⁹ :

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

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 the market power of any given player, by showing how unavoidable, a given operator is, in the Internet traffic flows, given the set of existing Interconnection policies.

The geographical differentiation is modelled by using two variables. A first variable, *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 rule³⁰. A second variable, *diff_ixp*, takes into account the different geographical coverage: for any pair of providers, this variable represents the difference in the number of memberships among the most important Internet Exchange Points all over the world³¹ that they have.

²⁹ D'Ignazio and Giovannetti (2006) have used this metric to assess HHI market concentration indexes, we focus instead on the micro bilateral interconnection choices.

³⁰ $dist(Operator1 - Operator 2) = RadiusEarth * ArcCos(Cos(Radians(90-Lat1)) * Cos(Radians(90-Lat2)) + Sin(Radians(90-Lat1)) * Sin(Radians(90-Lat2)) * Cos(Radians(Long1-Long2)))$

³¹ 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, Xchangeoint

In order to model the technical element behind the *hot potato* routing effect, we constructed a variable, *both_ixp*, indicating, for each pair of providers, the number of IXPs at which they are both present³². The variable *both_ixp* could also be interpreted as expressing a reputation effect, following Titley (1997).

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_24s</i>	Difference in the number of /24s for any pair of providers in units of thousands
<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_bet</i>	Difference in the <i>betweenness</i> measure in thousands of units.

5.2 Estimation results

We estimated both a probit and a logit model by maximum likelihood.

London, Equinix 7 locations USA, Jpnep 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).

³² 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.

Table 2: ISPs binary model results

Dependent Variable: <i>peering</i>		
Independent Variable	<i>Probit</i>	<i>Logit</i>
<i>diff_24s</i>	0.004 (10.26)	0.008 (9.26)
<i>dist_hq</i>	0.025 (2.83)	0.044 (3.01)
<i>both_ixp</i>	0.281 (10.45)	0.463 (9.91)
<i>diff_ixp</i>	0.043 (3.67)	0.070 (3.62)
<i>diff_bet</i>	-0.128 (23.87)	-0.220 (20.90)
<i>constant</i>	-0.096 (1.80)	-0.170 (1.90)
Number of Observations	4753	4753
Percent of correctly predicted	0.686	0.687
Pseudo R-Square	0.1638	0.1637
Log-likelihood value	-2723.59	-2723.89
Absolute value of z statistics in parentheses		

Probit and logit produce very consistent results. All the variables introduced are statistically significant. Interestingly, the two variables representing the competitors' asymmetry seem to have opposite effects. Indeed, the difference in the *betweenness*, which has also the highest z statistic, is negatively related to peering, while the difference in the *customer cone* positively affects peering.

A possible interpretation relies upon the fact that *customer cone* expresses asymmetry in "size", while the *betweenness* expresses asymmetry in "market power associated to *unavoidability*". The asymmetry in size can be seen as a differentiating element, which lowers the extent of competition between providers, hence reducing business stealing and positively affecting peering. 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. In this case, hence, both the *backbone free riding* and the *business stealing effects* play a decisive role against peering.

Our results also show 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. 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. The estimated partial effects (see Appendix) provide some evidence about the magnitude of the covariates' effects on peering.

An effective measure of the estimated model's goodness of fit is related to the percent of correctly predicted pairs. Among the 4753 pairs, 2674 peering agreements were inferred using the algorithms described in Section 4; hence, the peering ratio is 0.562. If we were asked to predict the interconnection regime based on this information only, we would always predict peering, being right in the 56.2% of the cases. The use of the estimated model significantly improves the predictions: the percent of correct predictions is 0.686 for both probit and logit, with an improvement by 22% relatively to the rule "*always predict peering*". Moreover, in this case we are able to predict both peering and no interconnection. The Pseudo R-Square is 0.164 for both probit and logit. A *RESET* test was conducted to test the hypothesis of no misspecification of the econometric model. Results are in the direction of no misspecification for both probit and logit (see Appendix).

6 Conclusions

In recent years, many game theoretic models have analysed interconnection agreements between Internet Operators, at the same time there is an increasing interest in this issue among Competition Authorities. A relevant question is whether or not the asymmetry between Providers affects the “fairness” of their interconnection contracts, eventually leading to interconnection quality degradation. While theoretical models provide contrasting results, there is a lack of empirical analysis on this issue. This work fills this gap: we provide some empirical evidence 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 the 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 refusal of peering by the larger providers.

We introduced two distinct metrics to model the providers’ asymmetry: the *customer cone*, providing a proxy for “market share” and the *betweenness*, expressing the market power of any given player, by showing its degree of *unavoidability* in the Internet traffic routing.

The binary model introduced showed a particularly interesting result: the two measures used to represent asymmetry have opposite effects on the interconnection quality among pairs of providers. Hence we suggest that the contrasting results achieved in the game theoretic literature can be seen in the light of which type of asymmetry prevails: Cr  mer, Rey and Tirole, (2000), findings reflect a dominant role of asymmetry in terms of network

centrality, expressing relative market power, on the other hand, the results obtained by Foros and Hansen (2001), and Economides (2005), pointing to a major role played by differentiation and network externalities in driving the peering decision³³ can be seen as linked to a notion of asymmetry in terms of customer shares .

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. "Antitrust Analysis for the Internet Upstream Market: A Border Getaway Protocol Approach", forthcoming in the *Journal of Competition Law & Economics*.
- 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]

³³ 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 Economics* 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
- 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. (2002a) "Internet Service Providers and Peering", Equinix Inc.
- NORTON, W.B. (2002b) "The Art of Peering: The Peering Playbook", Equinix Inc.
- NORTON, W. B. (2003) "The Evolution of the U.S. Internet Peering Ecosystem."

- SHIMBEL, A. (1953) "Structural parameters of communication networks". *Bulletin of Mathematical Biophysics* 15, 501-507.
- 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.1: ISPs binary models, variables summary statistics

Variable	Mean	Std. Dev.	Min	Max
<i>peering</i> = 1 (2674 obs)				
<i>diff_24s</i>	24.73925	64.24227	0	380.151
<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_bet</i>	1.628163	2.388645	0	16.45
<i>peering</i> = 0 (2079 obs)				
<i>diff_24s</i>	14.12309	39.74619	0	380.151
<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_bet</i>	5.674763	6.65285	0	16.414

Table A.2: RESET test for probit and logit

Dependent Variable: <i>peering</i>		
Independent Variable	<i>Probit</i>	<i>Logit</i>
<i>diff_24s</i>	0.005 (6.32)	0.009 (6.12)
<i>dist_hq</i>	0.028 (3.02)	0.048 (3.03)
<i>both_ixp</i>	0.345 (5.21)	0.534 (4.88)
<i>diff_ixp</i>	0.050 (3.71)	0.078 (3.49)
<i>diff_bet</i>	-0.144 (8.78)	-0.239 (8.29)
<i>fitvar2</i>	-0.522 (1.06)	-0.575 (0.71)
<i>constant</i>	0.013 (0.11)	-0.050 (0.26)
Number of Observations	4753	4753
Pseudo R-Square	0.1640	1638
Log-likelihood value	-2723.0329	-2723.6346
RESET test	<i>fitvar2</i> = 0 <i>chi2</i> (1) = 1.11 <i>Prob</i> > <i>chi2</i> = 0.2914	<i>fitvar2</i> = 0 <i>chi2</i> (1) = 0.50 <i>Prob</i> > <i>chi2</i> = 0.4777
Absolute value of z statistics in parentheses		

Table A.3: ISPs binary model partial effects

Dependent Variable: $P(y=1 \mathbf{x})$			
Independent Variable	<i>Probit</i>	<i>Logit</i>	\bar{x}
<i>diff_24s</i>	0.001748 (10.26)	0.002031 (9.26)	20.0957
<i>dist_hq</i>	0.00974 (2.83)	0.01082 (3.01)	2.16791
<i>both_ixp</i>	0.111385 (10.45)	0.115027 (9.91)	1.46118
<i>diff_ixp</i>	0.017183 (3.67)	0.017395 (3.62)	2.1843
<i>diff_bet</i>	-0.0507 (23.87)	-0.05469 (20.90)	3.39818
Number of Observations	4753	4753	
$y = \text{Pr}(\text{peering}) \text{ (predict)}$.54626609	.54251271	
Absolute value of z statistics in parentheses			