Robert Kenny

Are traffic charges needed to avert a coming capex catastrophe?

A review of the AT Kearney paper *A Viable Future Model for the Internet*

14 August 2011

COMMUNICATIONS CHAMBERS

Contents

About the author

Robert Kenny (rob@commcham.com) is a partner with Communications Chambers, a firm of TMT strategy and policy advisors. Before joining Communications Chambers he was MD of Human Capital, a consulting firm. Past roles include heading strategy and/or M&A for Hongkong Telecom, Reach and Level 3 (all multi-billion dollar telcos). He was also a founder of IncubASIA, a Hong Kong based venture capital firm investing in online businesses.

The opinions in this report are entirely his own.

Introduction

For a number of years the US has had a lively debate on 'net neutrality'. While definitions vary significantly, in broad terms net neutrality refers to the idea that internet traffic from different sources should be treated equally by carriers. Proponents of net neutrality argue that it is key to the open nature of the internet, and that without it ISPs could become choke-points, controlling the content their users see and imposing unfair charges. Opponents argue that any regulation for net neutrality is unnecessary, might stifle innovation and could threaten the ability of carriers to fund the expansion of their networks to meet growing demand.[1]

The net neutrality debate is now gathering steam in Europe, both at the Commission level[2] and in member states. Against this background, four European telcos[3] commissioned a report from AT Kearney [ATK], to support their opposition to net neutrality regulation. This report, *A Viable Future Model for the Internet,*[4] claims that carriers are facing ballooning capex requirements to fund the growth of internet traffic and that the best way to address this structural problem is via traffic charges to online service providers [OSPs].

If massive capex is required, and this needs to be recovered from OSPs, that would be a significant argument against net neutrality regulation, since it would necessarily end the principal that consumers could access any (legal) site they wished – ISPs would block access to sites that had not paid the charges the ISPs had chosen to impose.

Broadly the logic of ATK's report as follows:

- Telco investors are already seeing lower returns than investors in other players in the internet value chain
- Telcos face ballooning capex
- This capex is unsustainable
- OSPs are not contributing to the costs of traffic
- In a two-sided market, both sides pay
- Traffic charges are necessary because otherwise OSPs have no incentive to constrain traffic costs
- OSPs can easily afford increased charges
- Increasing retail prices will be challenging
- It is practical to implement traffic charges to OSPs
- Enhanced quality services can be introduced without degrading the basic internet

However I believe both its starting assumptions and its logic are open to significant challenge. This paper reviews the ATK report, from technical, economic and regulatory perspectives, and makes the case that

[1] This is obviously a very brief summary of a highly complex issue
[2] See for instance EC, *The open internet and net neutrality in Europe*, 19 April 2011
[3] Deutsche Telekom, France Telecom-Orange, Telecom Italia and Telefónica
[4] AT Kearney, *A Viable Future Model for the Internet*, 21 December 2010

ATK's conclusion (that the best way forward is traffic charges to OSPs) is not at all well-founded. I consider in turn each of the logical steps above.

Note that the focus of the economic analysis in this paper is primarily on fixed networks, though the qualitative arguments apply equally to both fixed and mobile networks.

"Telco investors are seeing lower returns than investors in other players in the internet value chain"

ATK claims that telco investors are receiving lower returns than other players in the internet value chain, particularly OSPs. There are two problems with this – firstly, even if returns are different it is far from clear why this is in some way inappropriate, and secondly ATK's analysis doesn't actually prove its claim that returns are indeed different.

Returns for different players in any value chain are almost always different – consider airlines and the oil industry (as jet fuel suppliers) for instance. Returns are a result of range of factors that may materially differ for different links in a value chain, most obviously risk – businesses facing higher risk on average pay higher returns than those facing lower risk.

There is of course no reason to expect different parts of the value chain to have the same level of risk. The provision of basic telecommunications services by incumbents is much lower risk than the provision of online services. Twenty years ago NTT, BT, France Telecom and Deutsche Telekom were all the leading players in their markets and they still are today. By contrast, online services were essentially non-existent twenty years ago, and in that period many players have come and gone. As ATK themselves note in another report, "of the top 15 websites in the United States in 1999 … only four remained in this league table by 2009".[5] Even in more recent years, highly successful companies have faltered and failed – think of Bebo, Myspace and Lycos (the last was the 12th largest site in the world in 2006[6] - today it is not in the top thousand[7]).

Thus different returns for different links in the internet value chain don't inherently suggest a problem. Moreover, it isn't clear from ATK's analysis that the returns *are* in fact that different. ATK offers the following chart to support its view that OSP returns have outstripped connectivity (carrier) returns:

[5] AT Kearney, *Internet Value Chain Economics*, 27 April 2010
[6] Morgan Stanley, *The State of the Internet, Part 3*, 8 November 2006
[7] Doubleclick Adplanner, April 2011

Figure 1 ATK analysis of market capitalisation

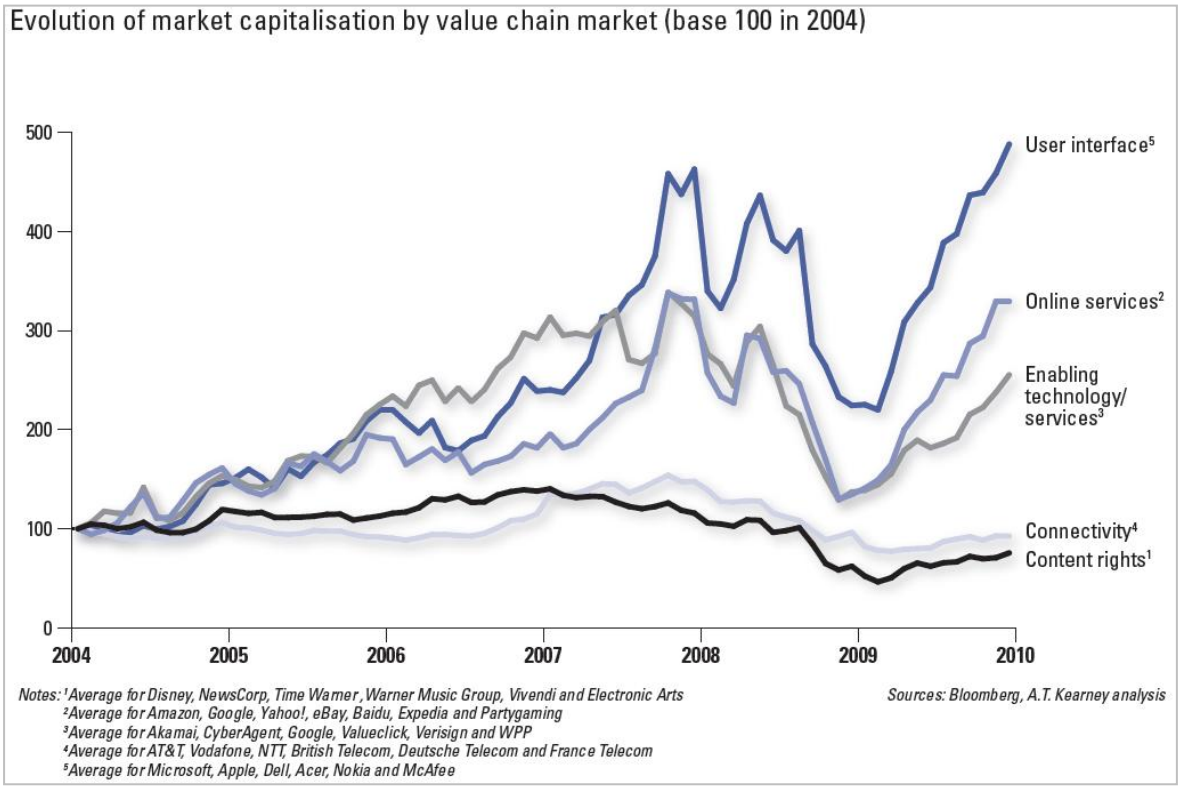

ATK has selected sample companies in each sector, and created an index of market capitalization for each. Unfortunately there are numerous problems with this approach:

- *It appears to suffer from 'survivor bias'* – representatives of each sector seem to have been chosen based on their importance today, not in 2004. This doesn't matter much for the telcos, precisely because of the stability in that sector. However, for the more volatile OSPs, it leads to significant overstatement, since by definition it will tend to select companies that have enjoyed growth, and omit those that have struggled. For instance, why is IAC (the owner of Ask), the 4th largest OSPs by market cap in 2004[8] (but far less valuable today) not included in the list? Why has Baidu, a stellar performer, been included even though it didn't list until 2005, whereas Shanda, China's most valuable internet company in 2004, been excluded?
- *It double-counts.* Google, another stellar performer, has been included in both 'online' and 'enabling technology' sectors
- *It only considers part of shareholder returns*. It looks only at market capitalization over time. However, an investor's return from a telco will include dividends as well as share price growth. Tech stocks very often do not pay dividends, so using market cap significantly overstates the return advantage of tech stocks

[8] Morgan Stanley, *Ten Questions Internet Execs Should Ask & Answer*, 16 November 2010 (2004 market cap data for 11th November that year)

- *It omits as a category another key part of the value chain*, namely telecoms network equipment suppliers (such as Alcatel-Lucent and Cisco) which have materially underperformed telcos since 2004. (Presumably ATK would not argue that because telco shareholders have done better than network equipment suppliers, there should be a value transfer from the former to the latter)

Thus we cannot reach a conclusion from ATK's analysis as to whether different links in the internet value chain have experienced different returns (even if we were to think that might be a problem).

"Telcos face ballooning capex"

ATK's next contention is that carriers will face a requirement for substantially increased capex, as internet traffic growth requires upgraded backhaul networks. For fixed networks, my focus here, ATK's capex projections are, at a high level[9], based on:

- Estimating a base level of European telco capex for backbone expansion
- Imputing from this a unit cost for capacity expansion
- Applying an annual 15% cost improvement to this figure
- Multiplying the improved cost figure for a given year by the forecast traffic growth in that year.[10] (Annual growth over the period is forecast at 35%).

A critical assumption here is the 15% cost improvement. Crudely put, if traffic growth is greater than the cost improvement, then capex will rise. Conversely, if traffic growth is lower than the cost improvement, then capex will fall. Unfortunately, ATK do not give a source for their 15%. However, costs fall rapidly for telecoms equipment, and a number of sources suggest that this figure may be appreciably too low:

- Transit costs, a good proxy for traffic carriage costs, have fallen 37% annually since 1998[11]
- Level 3 estimated an annual fall of 46% in optical and IP transmission technology costs[12]
- France Telecom cut its IP backbone unit costs by 50% simply by sourcing new IP routers[13]
- Cisco report an annual 23% decline in the cost per Gbps for their routers (a key component of backbone cost)[14]
- Network Strategy Partners estimate that bandwidth unit costs fall 20% per year[15]
- Ovum figures (cited by Deutsche Telekom) for transmission costs over time suggest a 40% annual decline[16]

[9] See pages 16 and 45 of the ATK report for a more detailed discussion
[10] Derived from the widely used Cisco VNI traffic forecasts. However, it is worth noting that these (like any forecast) are not infallible. The 2008 forecast overestimated 2010 traffic for Western Europe by 20%
[11] DrPeering International, *Internet Transit Prices - Historical and Projected*, August 2010
[12] Level 3, *Analyst and investor conference 2006*. Note that Level 3 refers to an 86% improvement in performance per dollar, equivalent to a 46% decline in cost per unit performance
[13] France Telecom, *investor day conquests 2015*, 31 May 2011
[14] Cisco, *IP NGN Backbone Routers for the Next Decade*,2011
[15] Network Strategy Partners, *Cloud based OTT video services : A business case analysis*, January 2011
[16] Deutsche Telekom, *Surfing the Gbps Wave*, 8 March 2011

Plugging any of these figures into the ATK analysis (instead of their own 15%) would give very different results. For instance, using the transit cost decline (37%), the Level 3 estimate (46%) or the Ovum estimate (40%) would imply no increase in capex at all to meet traffic growth.[17] In that case, the whole premise of ATK's report - that there is a problem of ballooning capex - simply falls away, at least as far as fixed networks are concerned.

Figure 2 West European incumbent telcos' capex (€bn)[18]

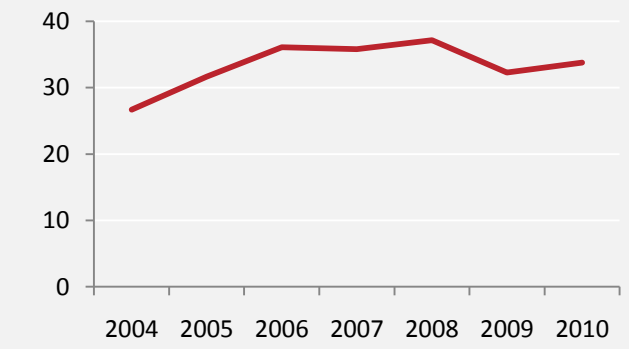

Note: Total for Belgacom, BT, DT, FT, KPN, PT, TI & Telefonica

Another way to look at this is to consider that rapid growth of internet traffic is nothing new. Western European IP traffic growth was 52% annually between 2007 and 2010, compared to the forecast figure of 35% for 2010 to 2013 that concerns ATK. [19] If the historic growth was higher (albeit from a lower base), then we might expect to see some early signs of the supposed ballooning capex that ATK predicts. In fact, West European incumbent telco's capex actually *fell* by 6% in this period (see Figure 2). If historic growth in IP traffic did not lead to an explosion in capex, it is not clear why a future growth (at a lower rate) should do so.

Nor do the telcos themselves predict dramatic capex growth. For instance, France Telecom expects its capex for the period 2011-2013 to be €5.9bn annually, compared to €5.5bn in 2010 (excluding FTTH), an increase only slightly ahead of inflation. Thereafter, it expects capex to fall, so as to be 10% of revenues in 2014-15, compared to 12.2% in 2010.[20] Telecom Italia predicts its Italian capex will drop from €3.1bn in 2010 to an average of €2.9bn in 2011-13.[21] Indeed, ATK itself acknowledges that "fixed operators are … forecasting flat capex plans".[22] Public companies take their financial forecasts very seriously indeed, and these flat capex forecasts presumably represent their best belief about the future. Thus they must disagree with one or more of ATK's assumptions that lead to an expectation of rapid capex growth.

"This capex is unsustainable"

ATK's next step is to use its forecast of incremental capex resulting from IP growth (€10bn for fixed networks in 2010-14) to project the additional revenues required to cover this cost. Their methodology is to assume a fixed 34% ratio of capex to revenue, and assume this also applies to the incremental

[17] Strictly, it would be appropriate to use a blend of these figures, since the ATK assumption is for the decline in the cost of capacity, which is – capacity being a combination of transport, routers and so on. However, the transit cost decline is just such a blend of inputs

[18] Berenberg Bank, *Will the real cash flow please stand up… again?*, 9 June 2011

[19] Communications Chambers analysis of Cisco VNI forecasts of 2008 and 2011 for W Europe traffic

[20] France Telecom, *investor day conquests 2015*, 31 May 2011

[21] Telecom Italia, *FY10 Results & Plan Update*, 25 February 2011

[22] Page 17

capex. This suggests that €3 of revenue are required for each €1 of capex, and results in an additional revenue requirement for fixed networks that rises to €9bn in 2014.

Assuming a fixed ratio of capex to revenue is an extremely generous way to estimate required revenue. For the company as a whole, revenue will need to cover capex, opex, tax, returns to shareholders and so on. Thus, total capex will be relatively low fraction of total revenue. However, that does not mean this fraction can simply be applied to incremental capex to assess incremental revenue. In this particular case, an upgrade to capacity to handle incremental traffic, there will be minimal costs beyond the capex itself. No extra customer care is required, no additional sales and marketing expense, no further corporate overhead and so on. The same number of engineers is needed to install a high capacity router as a low capacity one. ATK suggest additional maintenance fees and power may be required, but these are likely to be trivial.

Thus virtually all of the incremental revenue beyond the amount of capex is pure profit. While it is legitimate for shareholders to expect returns, a €3: €1 ratio of revenue to capex effectively implies a four month payback period on this investment. This is very generous indeed – four *years* might be more typical. Consequently, ATK's estimate of required revenue (even based on their excessive capex forecasts) is substantially too large.

Figure 3 UK Advertised broadband speed mix[23]

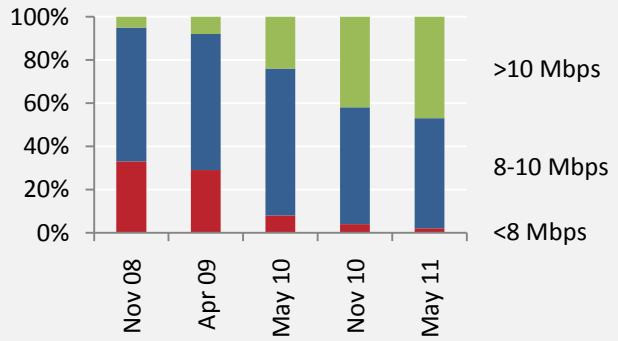

Note: Speeds are 'up to', not necessarily achieved speeds

ATK also appear to assume that by default this extra traffic (and capex) will bring no extra income, since "future traffic growth is more driven by an increase in traffic per customer which, with the current pricing structure, does not drive much incremental revenue."[24] If all the traffic growth came from consumers using their existing connection more intensely, this might be true. But in reality consumers have been upgrading to higher bandwidth connections, precisely because they are making heavier use of the internet (see Figure 3 for the change in mix of UK broadband, for example). These higher speed connections are more expensive – for instance, a UK consumer upgrading from Virgin Media's basic 'L' package (up to 10 Mbps) to the mid-tier 'XL' package (up to 30 Mbps) will spend an additional 37% per month.[25] Thus without any change in prices, nor any new revenue streams, carriers can expect appreciable revenue growth from traffic growth, as a result of consumers upgrading their connection speeds (though in some cases some of this incremental revenue will be required to cover the cost of upgrading the last mile connection).

[23] Ofcom, *UK fixed-line broadband performance, May 2011*, 27 July 2011 and *UK fixed broadband speeds, November/December 2010*, 2 March 2011

[24] Page 18

[25] £18.50 vs 13.50 per month. Broadband only. Prices per Virgin's website as at 29 July 2011

Indeed, just such an increase is one of the four key factors underpinning the Cisco VNI traffic forecasts that are the basis of ATK's analysis – Cisco says it expects that "[t]he average global residential Internet connection download speed will grow fourfold from 2010 to 2015, from 7.0 Mbps to 28 Mbps".[26]

Further evidence that ATK may have not allowed for the change in mix in broadband connections comes from their claim that "Between 2007 and 2009 the price of the average fixed standalone broadband connection declined by 44 percent". In fact, the source they cite for this[27] says that the price of a *2-4 Mbps connection* dropped by this amount[28]. The average connection would undoubtedly have had a much lower price drop (or possibly none at all), since the trend to ever-faster and more pricy connections continued during that period.

Thus ATK have both overestimated the incremental revenue required to support their projected capex (itself an overestimate), and have ignored incremental revenue that carriers will anyway receive from consumers upgrading connection speeds.

"OSPs are not contributing to the costs of traffic"

ATK not only believes that consumers will not contribute to increased traffic costs, they also seem to believe that OSPs also do not contribute. They give the impression that OSPs do all they can to avoid carrying costs for traffic, suggesting they simply connect "their hosting infrastructure to the nearest Internet exchange". However, just because OSPs do not pay the last mile connectivity provider does not mean they do not invest substantially in carrying traffic.

For instance, many larger OSPs make use of content distribution networks (CDNs) such as Akamai, Limelight and Level 3[29]. CDNs host content on behalf of OSPs in multiple locations, in part to ensure that consumers are more likely to be able to access content from 'nearby' servers. Akamai, for example, has a worldwide network of over 80,000 servers.[30] One effect of delivering traffic from these local servers is to reduce traffic on the internet backbone, reducing costs for carriers.

CDNs are heavily used. To use Akamai as an example again, all 30 of the top media and entertainment companies are amongst its customers.[31] These are of course key generators of the video traffic that is such an important driver of the growth that concerns ATK. Total spend on CDNs is substantial – it is forecast to be a $2.6bn market in 2011, and growing rapidly.[32]

[26] Cisco, *Entering the Zettabyte Era*, 1 June 2011
[27] EC, *Europe's Digital Competitiveness Report, Volume 1: i2010 — Annual Information Society Report 2009 Benchmarking i2010: Trends and main achievements*, 4 August 2009
[28] From €52 to €29 per month, April 2007 to April 2009
[29] Level 3 is primarily an internet backbone provider, but is also a significant player in the CDN market
[30] Akamai, *The Akamai network*, 2011
[31] Akamai, *Investor summit*, 2010
[32] IDATE, *Evolution of the CDN market*, 9 November 2010

CDNs are of course just one way in which OSPs invest in traffic distribution. Larger players have substantial national and international networks of their own. For instance, Google is a 20% participant[33] in the Unity trans-Pacific cable, which has a cost of $300m and a design capacity of 7.68 Tbps.[34] This is just one component of Google's much larger network, connecting 65 public and 61 private peering points[35] around the world at which it exchanges traffic with other operators. This obviously goes far beyond simply connecting "their hosting infrastructure to the nearest Internet exchange".

"In a two-sided market, both sides pay"

While many OSPs do invest substantially in traffic carriage cost, it is fair to say that they generally do not contribute to the cost of carrying traffic on the last leg between an interconnect point with a given ISP and the homes (or businesses) of that ISP's customers.

As ATK notes, ISPs operate in a two-sided market, bringing together consumers and OSPs, to the benefit of both. According to ATK "[a]ll successful providers of two-sided market services must find a pricing balance between the two sides. This balance accounts for the relative value derived by each side."[36] The example they offer of a two sided market is a trade-fair, which charges both attendees and exhibitors.

ATK's language throughout the document seems to suggest that it is a natural outcome of two-sided markets that both sides contribute to the cost of the platform. However, this is certainly not a requirement of the economic theory[38], and while both parties pay in *some* two sided markets, there are many examples where they do not (see Figure 4). To take a very parallel case to the trade-fair, while customers pay for entry here, they do not pay for entry to a shopping mall. In a mall, the merchants cover the entire cost. The 'right' market-driven outcome is highly situational, and frequently involves a cross-subsidy from one set of participants to the other.

Figure 4 Examples of two-sided markets with single payer[37]

Market	Charged participant	Uncharged participant
Credit cards	Merchant	Card holder
Job websites	Employer	Job seeker
Shopping mall	Merchant	Shopper
Night club's 'ladies night'	Men	Women
PDF documents	Document creators	Document readers
Academic journal	Readers	Authors
Free newspaper	Advertiser	Reader

[33] Terabit Consulting, *International Capacity Markets in Asia*, 16 March 2010
[34] Google, *Global Consortium to Construct New Cable System Linking US and Japan to Meet Increasing Bandwidth Demands*, 26 February 2008
[35] Per *peeringdb.com*, as at 29 July 2011
[36] Page 8
[37] Note that while within each of these markets generally only one participant is charged, there are exceptions. Examples include consumer charges for 'gold' credit cards, and author charges for academic papers containing very high resolution images
[38] Jean-Charles Rochet & Jean Tiroley, *Platform Competition in Two-Sided Markets*, May 29, 2001

Indeed, there are other two-sided markets within telecoms where pricing structures vary considerably. Telephone calls are a two-sided market – generally, both the calling and the called party benefit from the call. However, the costs can be charged entirely to the caller (a standard call), to the called party (an 800 number) or shared between the two (non-geographic numbers[39]). Notably, calls where the entire cost is met by one party or the other are far more common. This is parallel to the situation for ISPs, where the cost of connectivity is borne by the consumer not the OSP.

Figure 5 Reasons for no home internet access (EU)[40]

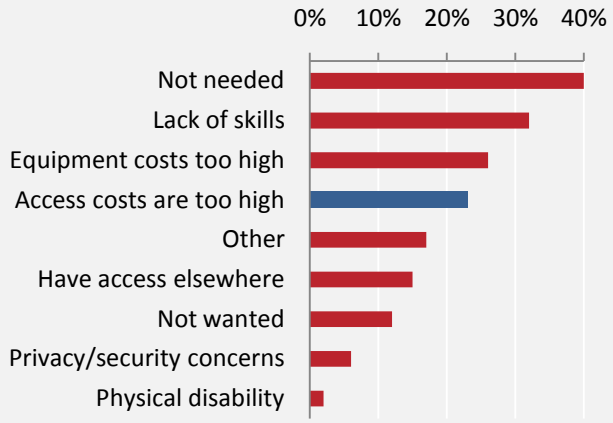

ATK note, rightly, that in a two-sided market it is possible for parties on one side of the market to *benefit* from paying increased charges, if this enables lower charges to those on the other side and thereby greater adoption of the platform. In the particular case of ISPs, ATK suggests that if OSPs pay more, consumers could pay less and this would lead to increased adoption of broadband, to the benefit of OSPs.[41] However, this argument is much weakened by the reality that access charges are not actually a particularly important barrier to internet adoption amongst those currently not online – of those without internet access at home, they are cited by only 23% as a reason (see Figure 5).

Moreover, the argument can be reversed – consumers may benefit from increased charges. If the alternative is increased charges for OSPs, which persuades many of them to stop providing service, then consumers may suffer more harm than if they'd paid higher prices.

Figure 6 shows an illustrative example. If the connectivity provider is making a loss (a notional '20'), then it could seek to recover it from the OSP. However, since the OSP only has a surplus of 15, its response would be to cease offering service, at which point consumers would lose their surplus of 100, for an overall loss of 95. Conversely, if consumers are charged, then they will still happily continue to consume the service (since their surplus will still be 80), and there will be no overall economic loss.

Figure 6 Threat to economic surplus from inappropriate charges (illustrative)

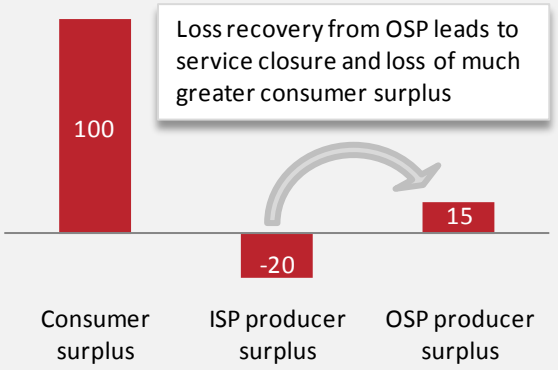

[39] Such as 0845 numbers in the UK, where the caller pays only the cost of a local call, and the called party covers the remainder of the cost

[40] Eurostat

[41] See pages 8 and 28 of the ATK report

Given that there are numerous online services provided at a loss, or 'freemium' services with loss leading components, this is a highly plausible scenario. (Note that in the example above, the OSP could in theory recover the connectivity charge from consumers, but for very many sites, which are non-commercial or purely ad-funded, this is not in reality a practical option.)

While it is not inherent to the economics of two-sided markets, from a regulatory and policy perspective, it is better for the revenue to be garnered where there is the potential for price competition. This will naturally keep prices low and operators efficient. However, for ISPs the potential for competition is on the consumer side, *not* on the OSP side, since the delivery of inbound traffic is one of the last great telecoms monopolies. If Ms Smith requests a video or webpage, there is no alternative for an OSP but to deliver it through her ISP, and thus the ISP has the potential to hold that OSP 'to ransom'. Conversely, if they try to overcharge Ms Smith, she will simply take her business elsewhere.

"Traffic charges are necessary because otherwise OSPs have no incentive to constrain traffic costs"

ATK make much of the economic inefficiency that could arise since "the absence of price signals means end users and Online Service Providers have no incentive to manage demand or to optimise the traffic they send or receive".[42] As noted above, one of the findings from the economic theory of two sided markets is that it can in fact be value maximising for one set of participants to cross-subsidise another, and thus we should not assume there is some form of market failure (and economic inefficiency) simply because OSPs are not paying ISPs for connectivity.

ATK also significantly overstate the reality when they say OSPs "have no incentive to manage demand or to optimise the traffic they send or receive". As we have seen, many OSPs spend substantially on CDNs, global networks and multiple peering points and so on. Even if this isn't the full end-to-end cost of traffic, it gives them a substantial incentive to limit their traffic volumes.

Moreover, the need to give a good customer experience also encourages OSPs to minimize the amount of traffic required to carry a page or provide a service. Their users' broadband connections vary widely in their bandwidth, and OSPs that make their content usable on lower speed connections will have a much larger addressable market. Further, if page weights[43] are lower then all users will experience quicker page loads. This in turn has real impact on revenues – for instance, Microsoft found that 1 second delay in page loads resulted in a 2.8% drop in revenue per user.[44]

[42] Page 14
[43] The total size of the files sent to the client browser to render a page; note that this is one of a number of factors influencing load times
[44] Eric Schurman & Jake Brutlag, *Performance related changes and their user impact*, 23 June 2009

"OSPs can easily afford increased charges"

Based on its estimate of needed incremental revenue for ISPs and its view that it is appropriate for OSPs to provide this revenue, ATK calculate a charge of €0.05 per GB (for fixed networks). It calculates the cost-per-download this would generate for various forms of paid content, to seek to show that such charges would be affordable.[45] However, these are far from a representative examples. Many forms of internet traffic bring no incremental revenue to the OSP. For instance, neither YouTube nor flickr charge users, and much of their traffic carries no advertising. Indeed, the relationship between traffic and profit is relatively weak (see Figure 7).

Services in the bottom right of this matrix (high bandwidth, low profitability) will be particularly threatened by traffic charges. To take a practical example, Skype represents 2.44% of all European traffic[46]. Were it to pick up a pro-rata portion of the €9bn that ATK claim ISPs need to raise from traffic charges for fixed networks, Skype would face an additional cost of €220m ($316m) in Europe alone.[47] Globally, Skype made an operating loss of $21m on revenues of $860m. Clearly such charges would be a substantial threat to the viability of the service.

Figure 7 Profitability and bandwidth intensity for various applications (illustrative)

This disconnect between bandwidth intensity and profitability exists even within companies. Google search is a highly profitable service which (per user) uses very little bandwidth. Conversely Google's YouTube service is highly bandwidth intense, but is much less profitable – despite its global success, it is reported to have only just reached break-even.[48] Its *global* revenues are estimated at $450m[49], but the ATK charges would result in *European* fixed network data charges of $1,541m.[50]

This analysis demonstrates how misleading is ATK's claim that its proposed charges would only increase Google's service delivery cost from 3% to 6% of revenues.[51] Google search would see a minimal increase in this percentage, but YouTube would see a catastrophic increase, such that service delivery costs

[45] Figure 20, page 31
[46] Sandvine, *Global Internet Phenomena Spotlight Europe, Fixed Access, Spring 2011*, 13 May 2011
[47] Note however that Skype is primarily P2P traffic, which has implications for which party might incur the traffic charges – see below for more discussion of P2P issues
[48] New York Times, *YouTube Ads Turn Videos Into Revenue*, 2 September 2010
[49] ibid
[50] Source of traffic data and methodology as for Skype example above
[51] Figure 21, page 32

became a *high multiple* of revenues. It is not obvious why Google should continue to cross subsidise YouTube in such a situation, and the result could well be closure. (Equally, there's no reason that carriers should cross-subsidise YouTube – the question is whether consumers or OSPs should cover any structural ISP loss, should one exist).

ATK's suggestion is that those OSPs facing increased data charges should charge "higher advertising fees or subscriptions".[52] This will be of cold comfort for the many many non-commercial sites on the web. Equally, given that advertising is a highly competitive, cross-media market, it is not clear that higher advertising fees would be possible. For businesses such as Google search, ad rates are already set by sophisticated auctions to extract maximum value from advertisers – increasing rates might not be possible at all, and is certainly far more challenging than simply issuing a new rate card.

Of course, even if Europe went down the path of charging OSPs substantial traffic fees, it doesn't follow that other countries would follow suit. In such circumstances services such as YouTube and its successors might simply choose not to offer their product in Europe, to the detriment of European consumers. If Europeans were not able to download from such services, they would be much less likely to upload their content to such services, reducing European voices online. Finally, if Europeans couldn't use such services, Europe would be an unlikely place for companies to develop and launch new bandwidth intensive services.

"Increasing retail prices will be challenging"

One of the reasons that ATK focuses on increasing charges to OSPs is that it believes that garnering increased revenues from consumers will be difficult. They claim "the biggest challenge would be achieving the average price increases needed in highly competitive markets".[53]

However, it is inconsistent with basic economics that in a competitive market players would not be able to increase prices if *all* were facing increased input costs, as ATK presumably believes all ISPs are. This is exactly what happened in response to the 2010 UK VAT increase for instance – all the leading players increased their prices to reflect this rise in their costs.[54] In fact, increases in prices for broadband services are not that unusual. OECD data shows that of the 30 incumbent telcos they track, more than half raised broadband prices at least once between 2006 and 2010, and several did so twice.[55]

ATK also errs in assuming that increases in retail pricing would be required to increase revenues. ATK say "fixed access charges would need to rise by an average €6 per month by 2014"[56]. However, to cover the

[52] Page 30
[53] Page 26
[54] HomePhone Choices, *Home phone price hike after VAT increase*, 5 January 2010
[55] Communications Chambers analysis of OECD, *OECD Communications Outlook 2011*, 8 July 2011. Data is available here. Price increases included only for years when service characteristics unchanged – lack of data for other years will lead to an undercount of price increases
[56] Page 27

notional increase in capex, an increase in average revenue per user (ARPU) is required,[57] not necessarily prices. As discussed above,[58] increasing adoption of higher speed broadband will increase ARPU without any change in tariffs.

ATK also has a practical concern with usage based tariffs, since "end users often do not have full control or awareness of the actual traffic they are downloading". While it is true that users have little idea what might be the page weight of a web page they are reading, this has not been a barrier to usage charges for mobile networks, which are far more common than unlimited plans. Indeed, as more consumers make use of mobile data, they are almost certainly getting a better understanding of what online applications are bandwidth intense (if only because they are less likely to work consistently on a mobile connection).

Even on fixed networks, some element of usage based pricing is not unusual. In 2010 27% of broadband tariffs in OECD countries include an explicit limit on monthly traffic (a bandwidth cap).[59] Interestingly this figure is *down* from 35.7% in 2008. If carriers truly believe traffic growth seriously threatens their investment case for broadband services, it is puzzling that they are moving away from bandwidth caps.

The relevance of bandwidth caps can be seen from the distribution of fixed network usage (Figure 8). The top 10% of consumers use approximately 55% of the traffic. Thereafter, traffic is comparatively evenly distributed across the remaining 90% of users. To the extent to which traffic growth is a problem, it is one that is materially driven by a small number of users. This distribution also suggests that concerns over consumers' knowledge of their data usage may be unnecessary. Those in the top 10% of users, who may run up against caps or material usage charges, have very different consumption from the majority (likely significant high quality video and/or file sharing), and almost certainly are well aware of their heavy usage.

Figure 8 European fixed network usage distribution[60]

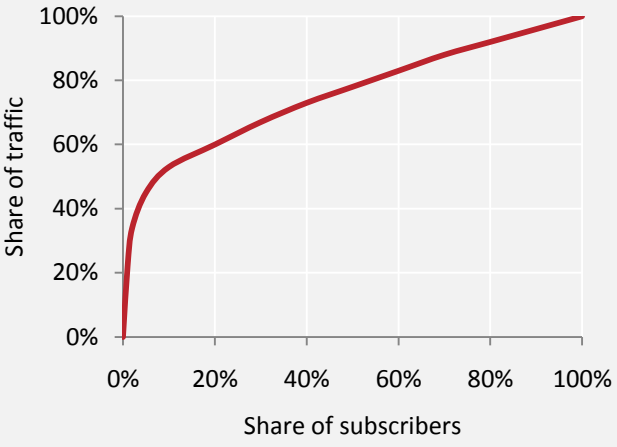

[57] ATK does in fact refer to an increase in ARPU in its figure at the bottom of Page 28
[58] See page 7
[59] OECD, *OECD Communications Outlook 2011*, 8 July 2011
[60] Adapted from Sandvine, *Global Internet Phenomena Spotlight Europe, Fixed Access, Spring 2011*, 13 May 2011

This is all the more important because a substantial volume of traffic actually comes from users, not OSPs. ATK offers a simplified model of the internet as a communications platform linking OSPs to consumers,[62] and this is certainly valid for reasonable portion of traffic (likely a majority). This model accurately describes web browsing and most real time entertainment, both of which represent significant portions of total traffic (see Figure 9). However over 30% of traffic is peer-to-peer (P2P) filesharing, the great majority of which is BitTorrent. BitTorrent is a protocol for consumers to share files with each other, in practice primarily pirated video content and software. This traffic does not originate from a centralised OSP – rather it travels between consumers, since those downloading files will generally be uploading them also.

Figure 9 Share of European traffic, fixed access [61]

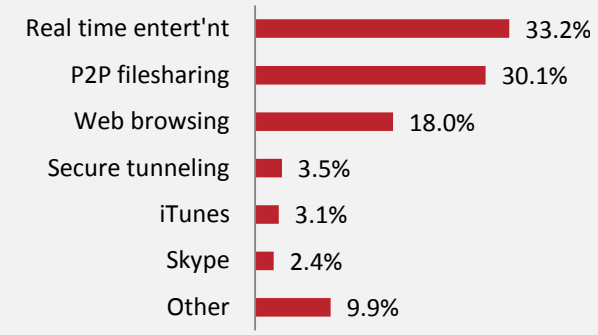

Any charging mechanism seeking to create appropriate economic incentives regarding traffic (quite apart from any concerns regarding the legality of the content travelling in this traffic) will be very incomplete if it does not address this substantial portion of traffic flows. Clearly this can only happen via consumer charges, since there is no OSP involved in this traffic.

Moreover, P2P filesharing is only part of the total P2P traffic. Increasingly legitimate media OSPs (included in the 'real time entertainment' category in Figure 9) are using P2P to deliver much of their content. The BBC does so for its iPlayer service, and Spotify for its music service. For Spotify, for example, 80% of bits received by a user come from other users.[63] If the traffic for these services is to be charged on a usage basis, then the charges will have to be levied on consumers. Further, if consumer traffic is not charged for, but OSP traffic is, then this will simply create a powerful incentive for media companies not currently using P2P to start doing so.

In summary: ATK believes that it may not be possible to increase consumer prices to cover traffic costs, and in particular consumer usage charges may be difficult to implement. As we have seen, it is not at all clear that increases in charges are necessary to increase ARPUs, but even if they are required, previous consumer price increases provide reassurance that it will be possible to implement them. ATK is particularly cautious about usage based tariffs for consumers. However, given substantial existing P2P volumes and the potential for even more traffic to be delivered that way, if traffic costs are to be recovered via usage charges, then inevitably consumers will need to be charged too.

[61] Adapted from Sandvine, *Global Internet Phenomena Spotlight Europe, Fixed Access, Spring 2011*, 13 May 2011. Share is of peak period traffic, which drives capacity requirements and hence cost
[62] Page 7
[63] Gunnar Kreitz & Fredrik Niemel, *Spotify – Large Scale, Low Latency, P2P Music-on-Demand Streaming*, 27 August 2010

"It is practical to implement traffic charges to OSPs"

Thus ATK overstate the challenges of increasing consumer charges, which is after all a modification to existing structures. Conversely, they underestimate the challenges in implementing traffic charges to OSPs , which would require a reengineering of commercial arrangements across multiple participants around the world.

ATK's proposal (under its Option 2) is that ISPs should charge for inbound traffic. Connectivity providers sending this traffic would then in turn pass these charges on to those sending them traffic, and so on along the chain to the OSP originating the traffic. There are real issues about the practicality of this.

The first practical challenge is the massive billing requirement. There are over 100m active websites[64] and over 2bn internet connections[65], across approximately 10,000 ISPs[66]. At the moment, given relatively little usage charging and no variation in charges by destination, this is manageable from a billing perspective. Transit charges are usually based on sampling traffic flows every 5 minutes[67].

However, under the ATK proposal each ISP would have its own charges for terminating traffic, and possibly more than one if they had both fixed and mobile end users.[68] Thus the destination of individual packets becomes a key consideration – and European traffic runs at 409 trillion packets per day,[69] and any given packet can travel across multiple transit and peering links in its journey from source to destination. Rather than simply measuring flow, carriers will have to also record the mix of destinations for all the packets in the sample, and then apply tariff tables to this to determine charges.

To make all this work, and to create the economic incentives ATK seeks, the 10,000 ISPs would need to all introduce completely new billing systems, and absorb the initial expense and the ongoing cost of managing and storing an order of magnitude more billing data. It seems plausible these costs would consume a significant portion of the traffic charges generated. (Telcos today spend $4.9bn just on billing software[70] - introducing packet-based tariffs for IP could create a challenge of similar magnitude to telcos' current billing for calls, texts and so on).

A second practical challenge is that of financial flows and risk. The ATK proposal envisions €28bn[71] for Europe alone being transferred annually (by 2014) from OSPs to access providers. However, this would not be a simple or direct transfer. It would flow (like the associated packets) through a complex interconnected mesh of intervening carriers. These carriers would see no net gain from this – they would simply receive money on the inbound side, and pass it all on on the outbound side.

---

[64] *Netcraft*, July 2011
[65] AFP, *Number of internet users worldwide reaches 2bn : UN*, 26 January 2011
[66] Fred Baker, *Views of IPv6 Site Multihoming*, IPJ, June 2011
[67] Typically the monthly 95th percentile sample is used as the basis for charging
[68] Off peak charges would be another potential dimension
[69] Based on VNI traffic figures and an assumed average packet size of 500 bytes
[70] Analysys Mason, *Telecoms software: worldwide forecast 2010–2014*, 4 November 2010
[71] Fixed plus mobile – Page 2

This has the potential to create serious cashflow problems. All participants would have the incentive to be slower paying than they were collecting, since they would be able to receive interest income on the substantial float of cash. Indeed, this would be the only potential benefit to these intermediaries to offset the substantial billing tasks they were being asked to undertake.

Even worse, the bankruptcy of one of these intermediaries could have cascading effects. The networks downstream from the bankrupt would still be financially obliged to those downstream to them, but might only receive cents on the dollar (and that much delayed) from the bankrupt. Of course all commercial businesses run the risk that the customers might go bankrupt, but the ATK proposal essentially requires internet backbone companies to take this risk on a large, but negative margin business.[72]

To get a sense of the scale of the proposed charges relative to the business of the internet backbones (who would, amongst others, pass through these charges), consider that Level 3, the largest internet backbone provider, has *global* revenues across *all* its lines of business of €2.56bn.[73] Compare this to the access traffic charges for Europe alone of €28bn that ATK proposes to flow through business such as Level 3. In effect, the ATK proposal risks recreating in the internet the financial interconnectedness and systemic risk of the banking system – not, in light of the last few years, an encouraging precedent.

The third practical challenge is that of regulation. ATK says:[74]

> "There is a real possibility [if traffic charges are introduced] that online service providers with a strong market franchise may be able to frustrate [them] by withholding some of their content or services from some networks. In extreme cases, competition authorities would likely to be asked to review whether this is a legitimate negotiation tactic or an inappropriate use of market power".

This is a perplexing suggestion. ATK is describing a situation in which an OSP is declining to accept the offer of services (delivery of traffic) from the carrier at a given price, and suggesting that the appropriate response is that the OSP be *forced* to buy the service. Competition authorities frequently mandate companies to *sell* services at specified prices, but I'm unaware of them ever requiring a customer to *buy* services at a given price. ATK's suggestion is all the more surprising because the carrier in this scenario has an absolute monopoly over the delivery of traffic to the end-user device in question, and thus has no shortage of market power of its own.

If ATK's proposal depends (as they themselves imply) on competition authorities taking this highly novel and extreme approach, then this would seem to be a massive implementation barrier.

[72] Negative since it is zero-margin before billing costs
[73] Revenues of $3.65bn in 2010 per Google Finance. Pre Global Crossing acquisition
[74] Page 33

ATK also believe that to prevent an OSP refusing to buy traffic delivery from a single ISP introducing traffic charges "a more coordinated approach is probably required where all Retail Connectivity Providers in a market set similar charges".[75] This sounds ominously like a cartel.

The fourth practical challenge is that of fraudulent traffic. ATK note "a mechanism would have to be put in place to prevent smaller networks fraudulently initiating inbound traffic to generate revenue". In off-peak hours, inbound traffic would carry no incremental cost, but would carry incremental revenue. This would create a powerful incentive for ISPs to artificially create requests for content.

Such a scenario is entirely plausible – in the 1990s, during the disintegration of the settlement rate regime for international calls, situations arose where the cost of making an international call to certain countries was actually less than the carriers in those countries received for terminating those calls. The consequence was 'openline' calls – essentially auto-dialler calls made by the carriers from overseas to themselves, where every minute of traffic generated incremental profits.

Artificial inbound internet traffic could be very difficult to detect. The ISP concerned could scatter its traffic requests across ISPs around the world. The impact on any given website (out of the 100m total) would almost certainly be too small to notice.

In summary, ATK's proposal for mandatory traffic charges carries substantial technical, regulatory, financial and enforcement challenges, that certainly dwarf any difficulties associated with the alternative of increasing retail tariffs (even if such an increase be required).

"Enhanced quality services can be introduced without degrading the basic internet"

ATK's Options 3 and 4 both propose an overlay layer of enhanced quality, charged service for OSPs, over the public internet and via bilateral agreements respectively. The current 'best efforts' internet would stay in place. A standard objection to such approaches is that operators might deliberately degrade (or under-invest in) their best efforts offer, to effectively force OSPs to pay for the enhanced quality service. In ATK's view this would not happen, because of "the stiff competition and operators' interest in retaining overall market share in the retail market. Such a short sighted move would lead to millions of regular-service-level[76] customers becoming dissatisfied and switching to other Connectivity Providers".

While debatable, let us accept ATK's contention for the moment. If it is true that the current best-efforts internet will not be degraded (which I take to mean congestion will get no worse), then it is not clear which material group of customers will have an incentive to spend extra money on an enhanced service. The traffic growth that so concerns ATK is not coming from new services – it is coming from increasing usage of existing services. This is true for fixed networks, but is particularly true for mobile networks,

[75] Page 30

[76] It is not clear why ATK refers to 'regular-service-level' consumer customers, since it is the OSPs who choose between regular or enhanced services (and they have no choice as to which Connectivity Provider to use – they must deal with those providing access to their end users)

where bandwidth is not being consumed by bleeding-edge on-demand HD TV, but rather by extensive use of long-existing web services such as browsing, email, YouTube and so on.

The existing internet works well for a wide variety of quite demanding applications. Netflix has built a substantial video-on-demand business over the open internet, and expects to have 22m streaming customers by the end of Q3 2011.[77] The BBC iPlayer (on demand TV) served 96m programmes over the open internet in June 2011.[78] Blizzard Entertainment's World of Warcraft has 11.1m subscribers,[79] and is just one component of an online gaming market estimated to be worth $20bn in 2010 that is operating over today's best efforts internet.[80]

If the internet is not going to be degraded, then presumably none of these applications will suffer. If all this gaming and IPTV is going to work well over the best efforts internet, then presumably it is relatively small, niche applications that will feel the need for enhanced quality. Moreover, the need will have to be a serious one to convince them to pay for connectivity, when relatively capable best efforts internet is available for free. This is not to say there are no such applications, only that it seems unlikely that there enough such applications with sufficient willingness-to-pay to provide revenues that would meaningfully contribute to ATK's estimated need of €28bn.

ATK's own view is that video and gaming over the internet will be a €42bn market in Europe by 2014, and that the relevant OSPs might be prepared to pay 10-15% of this for delivery. However, unless ATK believe that a significant portion of this €42bn will come from applications that are much more bandwidth intense than those in use today, it is hard to understand why OSPs would be willing to pay 10-15% for a premium service when the best efforts internet was already working satisfactorily to deliver them this €42bn.

Finally, segmenting the network into different levels of quality of service is widely felt to be a very inefficient way of delivering quality. It is cheaper (because it is simpler) to just upgrade capacity for all traffic, rather than partitioning the traffic and improving performance for all of it. Gary Bachula of Internet2[81] testified to the US Senate:

> "For a number of years, we seriously explored various "quality of service" schemes, including having our engineers convene a Quality of Service Working Group. As it developed, though, all of our research and practical experience supported the conclusion that it was far more cost effective to simply provide more bandwidth. With enough bandwidth in the network, there is no congestion and

[77] Netflix, *Q2 2011 Letter to shareholders*, 25 July 2011
[78] BBC, *Monthly Performance Pack, June 2011*, 2 August 2011
[79] Gamasutra, *World of Warcraft Subscriptions Continue To Decline, Though More Slowly*, 3 August 2011
[80] Trend Micro, *From Virtual Worlds To Real-World Threats*, 25 October 2010
[81] Internet2 is a not-for-profit consortium focused on the future of the internet. Its members include 221 U.S. universities, 45 corporations, 66 government agencies, laboratories and other institutions of higher learning, 35 regional and state research and education networks and more than 100 national research and education networking organizations representing over 50 countries

video bits do not need preferential treatment. All of the bits arrive fast enough, even if intermingled."[82]

In other words, if the purpose of differentiated service levels is to create appropriate incentives for OSPs and thereby remove economic inefficiencies, it is a somewhat perverse way to go about it, since it creates a whole new set of inefficiencies. It will be a more expensive way to deliver the improved performance for the most demanding traffic than simply upgrading all the traffic.

Conclusions

ATK's contention is that rapid growth in carrier fixed[83] network capex necessitates traffic charges to OSPs to cover this cost and create incentives for economic efficiency for OSPs. However, as we have seen:

- It is far from clear that fixed network capex will grow rapidly –improvement in the price/performance of network equipment may offset demand growth
- Even if capex does grow, it is not obvious that major changes are needed for telcos to recover this cost. ATK's required revenue projections are almost certainly overestimates, and consumer ARPUs may anyway grow as consumers continue their switch to higher bandwidth broadband
- Even if ARPUs do not rise, increasing retail charges is entirely plausible (and is unlikely to have material adverse consequences). Moreover, any system of traffic charges must anyway involve charges for consumers, given the importance of P2P traffic
- OSPs already have substantial incentives to manage their traffic volumes, given that they invest considerably in bringing traffic to the terminating ISPs
- That OSPs do not bear 100% of the end-to-end cost of their traffic is not a sign of economically inefficient incentives – it is a completely normal feature of double-sided markets
- Implementing mandatory traffic charges for OSPs carries major technical, regulatory, financial and enforcement challenges, given the need to flow very large payments through a highly complex interconnected system of internet carriers
- Differentiated service levels are an economically inefficient approach to funding a notional capex need, and moreover are unlikely to raise significant revenue, given the capabilities of the current internet (unless of course those capabilities are degraded)

[82] Gary Bachula, *Testimony before the United States Senate Committee on Commerce, Science and Transportation Hearing on Net Neutrality*, 7 February 2006

[83] ATK's report, unlike this one, also addresses the capex growth of mobile networks