



# **Interconnection in the Internet: the policy challenge**

David Clark

MIT CFP

May, 2011



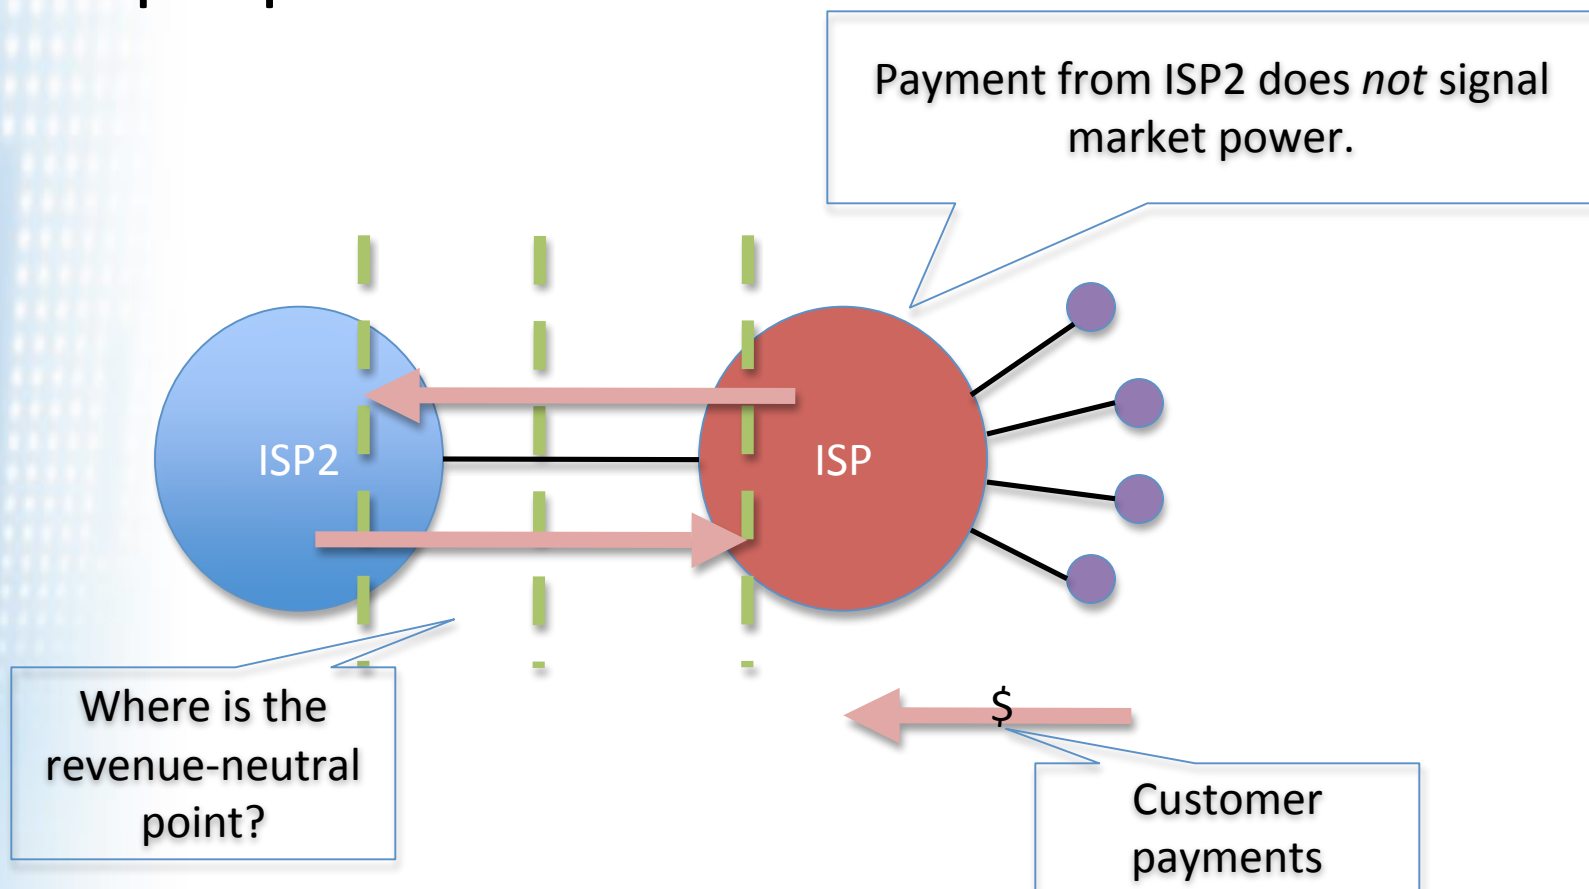
## Open interconnection?

- We have said (for some time) that issues around interconnection will be the next phase in the debates over “network neutrality”, and the future of the “open” Internet.
- It is time...

# Background

- Dispute between Comcast and L3 attracted attention.
  - CDNs raise specific issues.
  - See our TPRC papers, etc., for more background.
- Erosion of old models
  - Revenue neutral peering no longer obvious efficient outcome.
  - Who pays may vary even without recourse to market power argument.
- Policy implications
  - Points up problems with NN antidiscrimination rules.
  - No new regulation now, but that may not remain true.
    - Concern is valid
  - Better data and/or disclosure requirements may help.

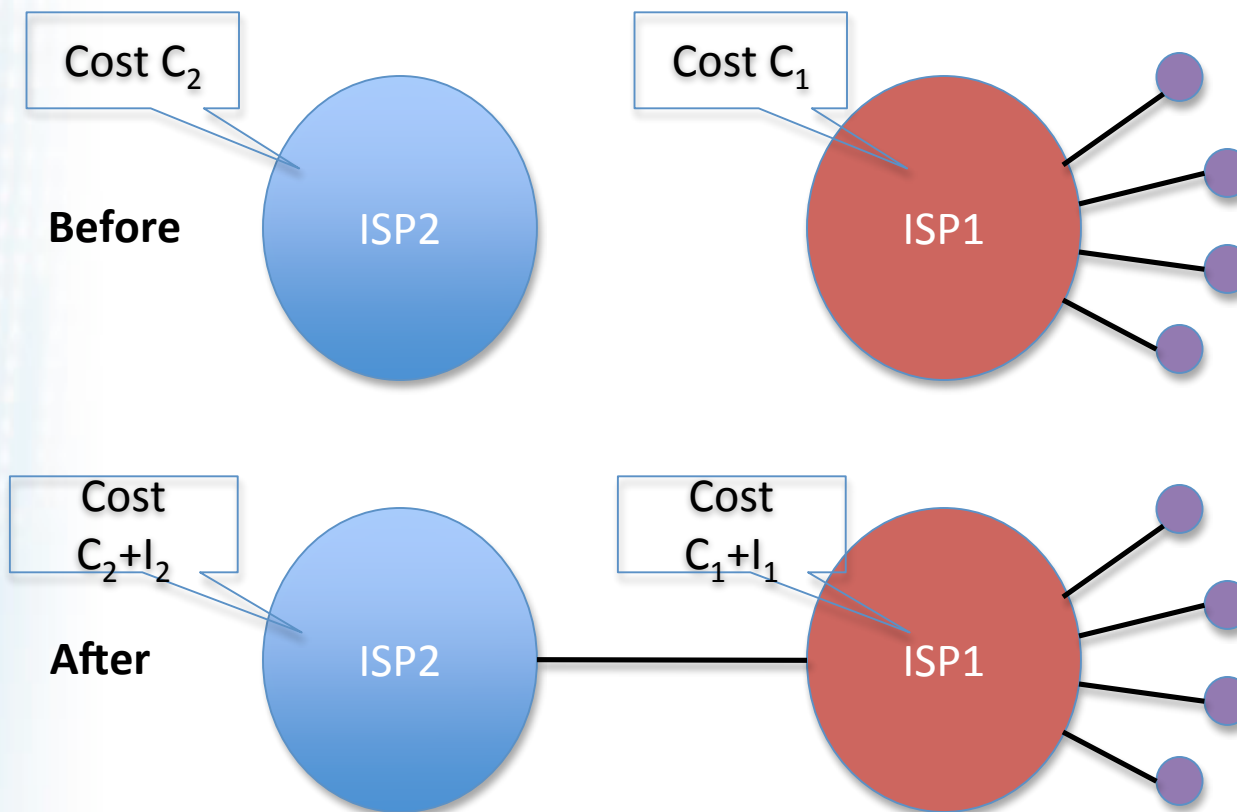
# Simple picture



Money flows in from the edges.

- So there must be a point where the flows meet.


## Some terminology



Several cases:

- $I_1 < 0$ ;  $I_2 < 0$  (Traditional peering—both sides save money.)
- $I_1 > 0$ ;  $I_2 < 0$  (Many CDNs)
- $I_1 < 0$ ;  $I_2 > 0$  (CDN for rural ISP)
- $I_1 > 0$ ;  $I_2 > 0$  (No interconnection will occur unless some other party pays)

# Cost-based argument



Negotiation to set transfer payment  $W$  is bounded by cost savings.  
(But actual values for  $I$  are hard to determine and not public...)

## Traditional peering

- When two similar ISPs (e.g. tier 1 ISPs) propose to peer:
  - Probably both have similar internal cost structures.
  - Incremental cost  $I$  is negative for both of them
    - Both save money (cost of transit).
  - No new traffic is associated with the agreement.
  - Complexity of negotiation leads to setting transfer payment  $W$  to 0.
    - Bill would say this is often an efficient outcome.



If they are not similar?

- A more common situation today.
  - Negotiation between ISPs of different size.
  - Negotiation between ISPs with different internal cost structures.
  - Negotiation between ISPs with different classes of customers.
    - Obvious current case: ISP1 is broadband access network; ISP2 is CDN.



## Different sizes?

- Paper by Odlyzko and Tilly: “A refutation of Metcalf’s law...”
  - Metcalf’s law: all users value each other equally  
-> value of net goes up as  $N^2$ .
  - Consequence: if two networks connect (e.g. peer), the gain in value to each is **independent of relative size**. (e.g. revenue neutral peering is valid).
  - Their view: users do *not* value each other equally.
  - Consequence: small network get more value from peering, and thus should pay.

## Value-based negotiation

- A very dangerous topic.
- Under what circumstances should delivery fees be conditioned on the value of the exchange?
  - To over-simplify: are all bytes equal?
- What is relation of “value” to cost and cost recovery?

# An important distinction

- With commercial content, there is a payment from consumer to the producer/programmer: the content payment.
  - This is separate from the delivery (conduit) fee.
  - (Content fees take all sorts of forms. Heavily studied by economists.)
- In other circumstances (e.g. telephony) there is no content fee.
  - We just talk to each other—peer production of content.
  - If the context of negotiation is access to commercial content, then discussions of value can signal two different things:
    - Attempts to tap into “content payment”.
    - Negotiation over allocation/recovery of the incremental costs.
- Telephone example.
  - “800” numbers.

## History: tapping the content payment

- Railroads used to have value-based pricing.
  - Sustained by regulation.
- Trucking undercut high-value pricing with “bytes is bytes” (a ton is a ton) pricing.
  - Contributed to collapse of railroads.
- Lessons:
  - Bad game plan.
  - Presumption: value pricing implies either regulatory intervention or market power.
    - Or that distortion from a cost basis is minimal.

## Avoiding the danger zone

- When is it safe(r) to discuss value pricing?
  - One answer: *two-sided markets*
- Term in economics:
  - Provider with two customer classes, *which depend on each other*.
- Classic example: singles club.
  - Charge men and women different prices: “ladies’ night”.
- Lots of economic theory on rationale to set prices for classes.
  - The railroad situation was *not* a two-sided market.
    - The different customers were not dependent.

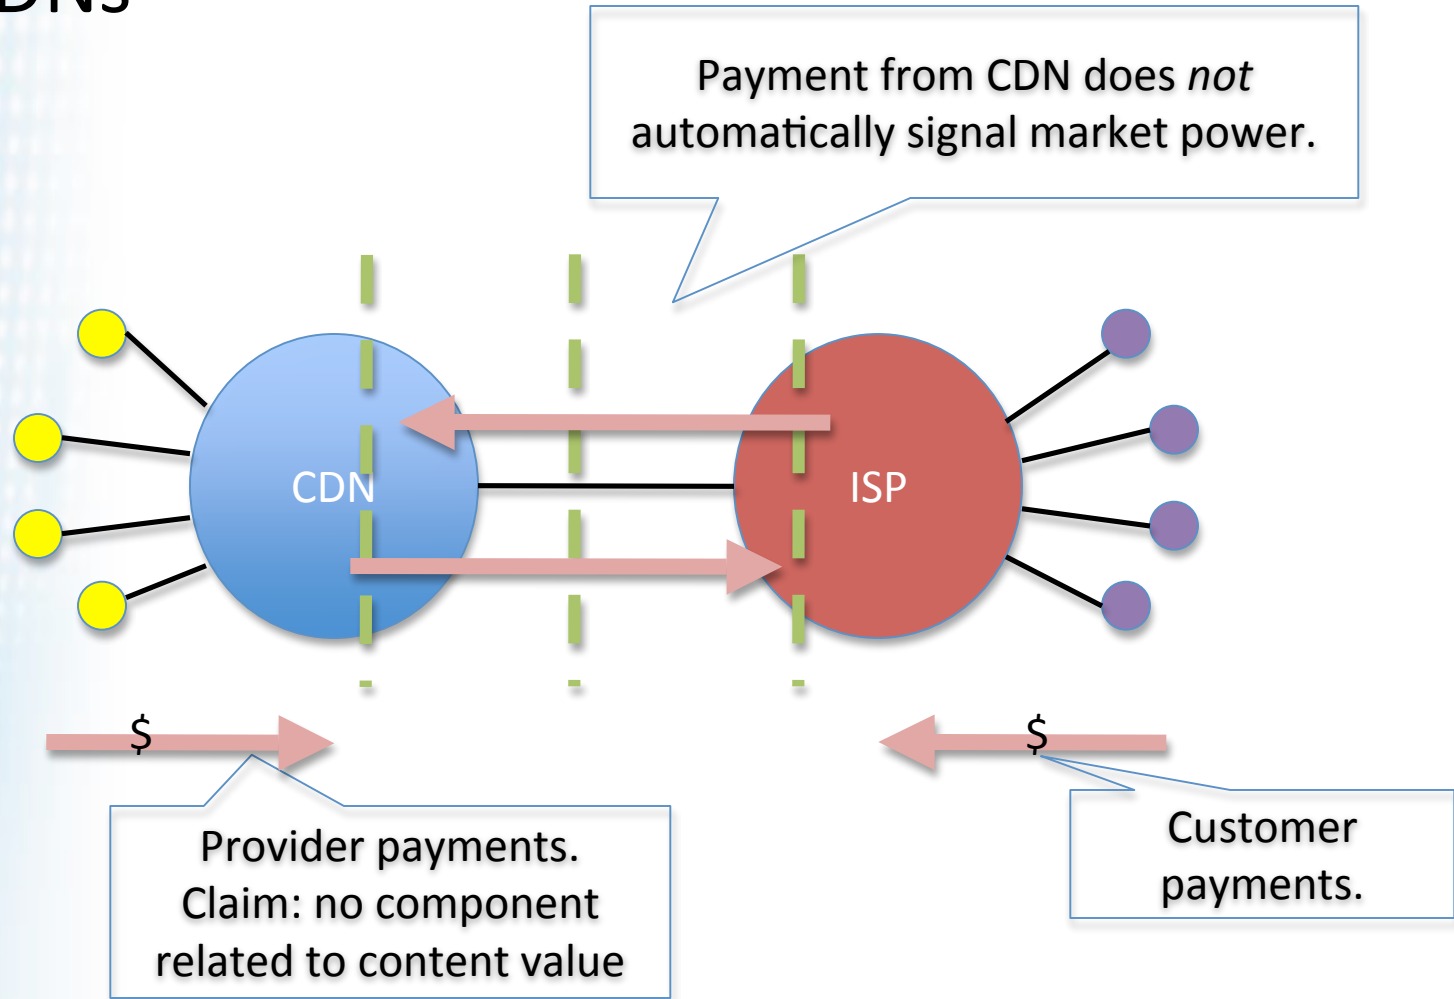
## Consider the ISP

- Is an ISP a provider in a two-sided market?
  - Should they charge all of their customers the same?
  - Are there mutually dependent customer classes?
    - CDN and customer?
  - In a two-sided market, one *must* discuss the relative value of the different customer classes, even if the discipline of the market drives total pricing to recovering short-run costs.
    - Singles clubs are highly competitive.
- What is the true signal of market power?
  - Rents (e.g. high profits).

## Specific case: CDNs

- Are CDNs a “class” of customer?
  - Lots of different content payment models.
    - Netflix: consumer pays Netflix pays CDN (pays ISP?).
    - ESPN3: consumer pays ISP pays ESPN (pays MLBAM) pays CDN (pays ISP?).
    - Ad-based: Advertiser pays programmer pays CDN (pays ISP?).
      - Only significant case with additional infusion of money.
  - But all the content seems to be “commercial”.
- To my knowledge, CDNs do not/cannot demand a “value-based component” in their pricing.
  - Highly competitive and commodity.
  - Would seem to suggest that all CDN traffic could validly be put into one “value class” in two-sided analysis.

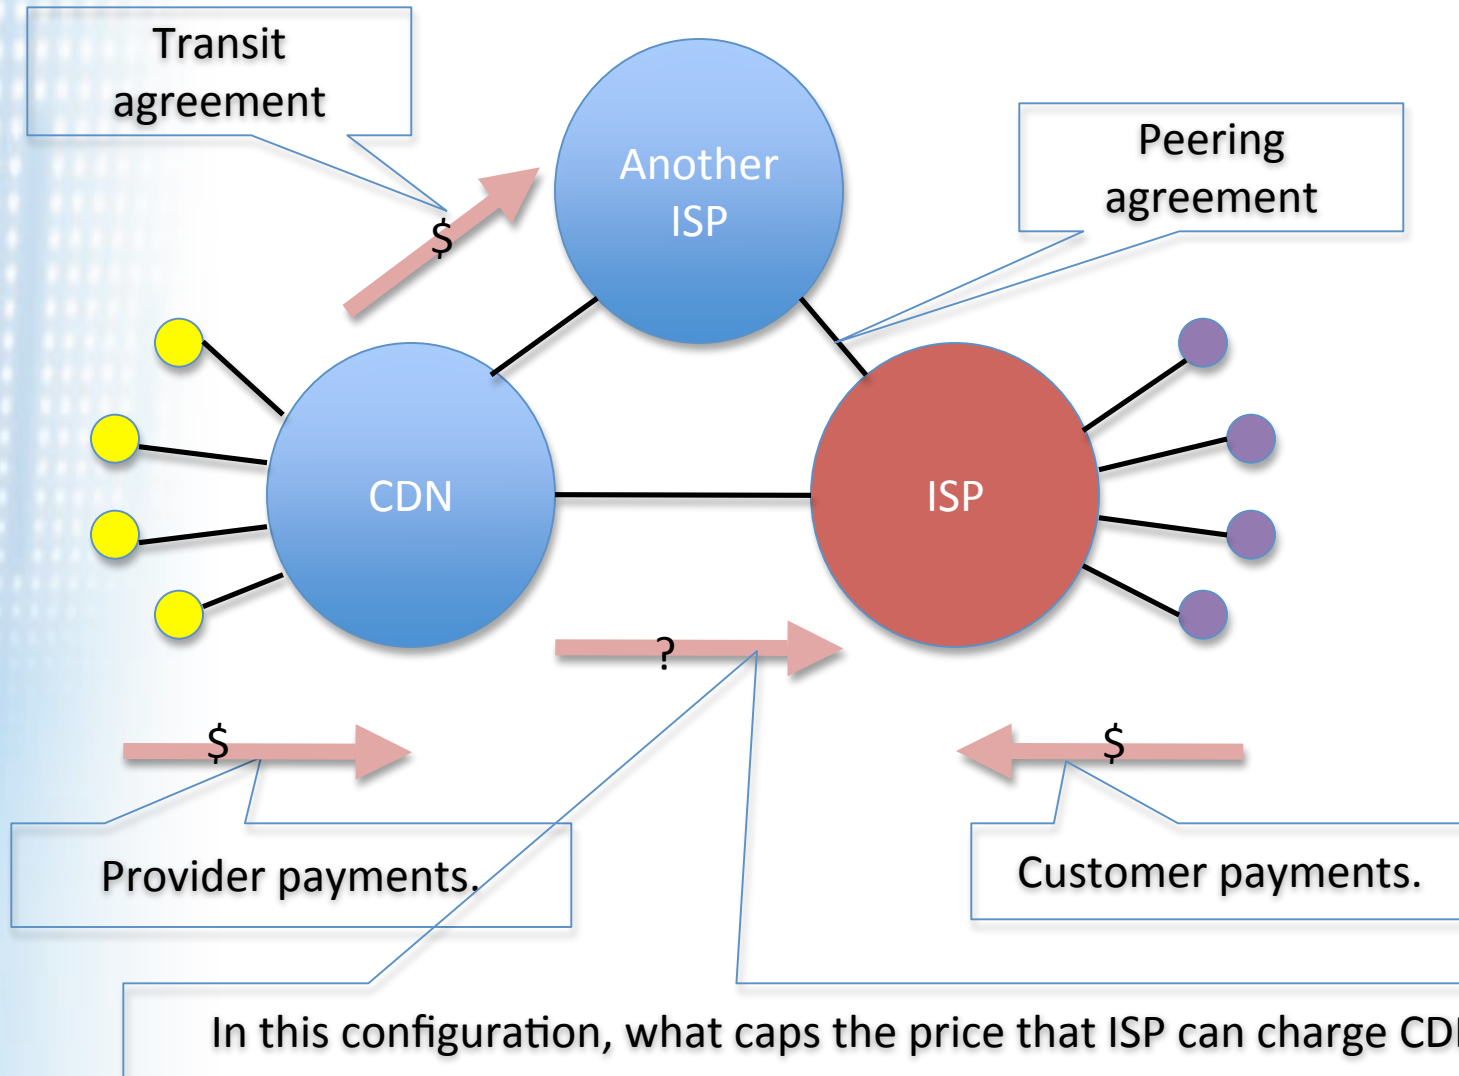
# CDNs



Hypothesis:  $I_{\text{CDN}} < 0$  (costs go down);  $I_{\text{ISP}} > 0$  (costs go up).



# An argument based on topology



## Topological limits

- In the previous picture, why would CDN ever agree to pay ISP more than the cost of transit to CDN if it reaches ISP by its peering partner?
  - Better performance
    - Yes, but probably not sufficient to justify a big distortion.
  - ISP blocks CDN traffic coming from “Another ISP” peer.
    - That would be pretty blatant discrimination.
  - ISP de-peers “Another ISP”, or demands paid peering.
    - Back where we started.
    - But with a tougher job of value discrimination.

# Finding the bright line

- Payment from CDN to ISP could be:
  - Extraction of a rent based on the value of the content.
  - Resolution of the “two-sided market” value-based cost allocation.
- Where is the line?
  - Claim: the line is not at “zero payment”, but at a point that is a function of customary transit costs.
    - Perhaps some discount (due to routing restrictions).
    - Perhaps a slight premium (for enhanced service)?
  - Even though specific agreements and incremental costs are NDA stuff, could a “customary function” emerge?
    - Customary function might be easier to explicate than cost models for incremental costs  $I_1$  and  $I_2$ .
- But nature of negotiation still undefined.
  - Who pays whom how much?

## If ISP had market power

- It would mean that the ISP forced the CDN to make payments that are passed through to the providers, based on the market power of the ISP and the recognized value of the commercial content?
- This question takes the CDN out of the analysis and pits the power of the ISP against the power of the content owner.
  - And now, not all content is in the same class.
  - But the CDN is an agent for all the relevant content classes.
  - Either different prices for different CDN customers (complex and blatant discrimination), or some content might be priced out of the CDN business (not desirable).

## Assume no market power

- Prices charged to CDNs reflect cost recovery in a two-sided market context.
- Still need to resolve the question of whether the CDN should pay, and how much.
  - Still assuming that the payment is capped by some function of customary transit pricing.

## Balance of flows

- The traditional basis for agreeing to revenue-neutral peering has been balance of flows.
  - A long tradition.
  - But actually no obvious basis in cost.
    - Circuits cost the same no matter which way the traffic flows. (Asymmetry leads to under-utilized capacity.)
  - Seems to be based on assumption that balance of flows signals “similarity of character”.
  - Seems to be based on rough rule that value follows the packets.
    - But this could be totally backwards.
- When ISPs are clearly not similar, no clear reason why balance of flows is a good rule of thumb.
  - Level 3 has challenged this idea.

## An aside—does any of this matter?

- Consider the Netflix case.
  - Consumer pays ISP
  - Consumer pays Netflix.
    - All the money comes from consumer.
  - So does it matter how it gets to ISP?
    - Is this whole “two-sided” discussion irrelevant?
      - Perhaps true when both of the customer classes are actually paid from the same source of money.
    - (But does not apply to other cases: e.g. advertising-based content.)

## New idea: patterns of usage

- Consider two extremes.
  - All users consume Netflix content equally.
  - Only one user consumes Netflix content.
- If the ISP-CDN connection is revenue neutral:
  - In the former case, users have equal total usage, so the Netflix-specific usage is balanced out.
  - In the latter case, the usage-related costs of this user are being spread across all the users.
  - But not all users have equal usage, as we know.



## Cross-user subsidy

- If all users pay the same for unequal usage, this is cross-user subsidy.
  - This is *not* a two-sided situation—the users are not mutually dependent.
- Why is this subsidy sustainable?
  - The amount is small and not worth thinking about.
  - All participants in the market prefer it.
  - ISPs are not subject to competition.
  - Regulators demand it.

## How to charge heavy users

- Usage caps and tiers.
  - The users pay directly.
- The “other side” pays—the provider or the CDN.
- Example: Australia
  - Low monthly caps.
  - Providers and CDNs pay for “premium service” so that their bytes do not count against the cap.

## An enhanced negotiation

- When two ISPs negotiate:
  - Specifically an access ISP and a CDN
- First try the “balance of flows” rule and see if both sides are satisfied.
- If not, consider the degree of uniformity of destinations from the CDN across the ISP
  - If uniform, ISP *might* agree that it is satisfactory to have the users pay cost directly.
  - If highly non-uniform, ISP would ask payment so that these customer are not being subsidized by all the other users.
    - Alternative is to go to usage tiers and bill users directly.

## Assumptions and ideas

- CDN market is competitive
  - CDNs do not partake of the content payment.
    - (Except perhaps as collection agent.)
- Interconnection can result in significant incremental costs, positive or negative.
  - May be many reasons.
- Transit costs may provide a customary basis to cap reasonable outcomes of negotiation.

## Some summary thoughts

- Rational discussion of interconnection cannot be separated from discussion of the carriage of commercial content.
- In the context of commercial content, consider:
  - *To encourage broadband deployment and preserve and promote the open and interconnected nature of the public Internet*, consumers are entitled to access the lawful Internet content of their choice.
- Retail pricing policy (e.g. whether to have low usage caps) will end up being tangled with bulk interconnection negotiation.