The valuation of two-sided platforms in arbitration

Ryan Williams, Cornerstone Research

CORNERSTONE RESEARCH



Ryan Williams is a Principal in Cornerstone Research's London office and a Research Fellow at Imperial College London.

The views expressed in this article are solely those of the author, who is responsible for the content, and do not necessarily represent the views of Cornerstone Research.

An increasing proportion of the global economy is represented by platforms: for example, marketplaces, app stores, payment networks, ride-hailing and delivery apps, advertising exchanges, freight, and accommodation booking sites. What these businesses have in common is that they do not only sell a product of their own. A marketplace sells buyers to sellers and sellers to buyers; an advertising exchange sells audiences to advertisers and advertisers to publishers. Economists refer to them as 'two-sided markets' and have studied them for more than two decades.[1] The framework is standard in competition economics, where it shapes market definition and the assessment of market power, and in the strategy literature on subsidising one side of a market to attract the other.[2] However, disputes involving two-sided platforms have not to date been a significant feature of the arbitration landscape. That is likely to change, in which case the valuation of platforms may well take centre stage. The purpose of this article is to introduce arbitration practitioners to some of the factors that may need to be taken into account when performing such valuations. The objective is not to be prescriptive but rather to elucidate certain key differences relative to valuations that a reader may have previously encountered.

The discounted cash flow (DCF) method values any business by projecting the cash flows it is expected to generate and discounting them at a rate that reflects their risk. Nothing about a platform changes the second half of that exercise. The first half is another matter. A conventional projection forecasts revenues and costs line by line, on the assumption that the lines can be forecast independently. For a platform they cannot, because participation on each side of the market depends on participation on the other. In the remainder of this article, I explain what that dependence does to a projection of cash flows, why it matters for the measurement of loss, and what questions a reader should be asking when a platform valuation appears in an expert's report.

The basic premises are simple. First, sellers join a platform because buyers are there, and buyers join because sellers are there. Second, an event that removes participants from one side therefore removes participants from the other, even though nothing has happened to the other side directly. Third, the size of that indirect effect is a matter of evidence rather than judgement, and it can be a multiple of the direct effect. Operationalising these premises in a way that a tribunal can follow and an opposing expert can test is the subject of what follows.

The one-sided projection

Consider first a business that is not a platform: a manufacturer that buys inputs from suppliers, transforms them and sells the output to customers. Its annual cash flow is revenue from customers, less payments to suppliers, less fixed costs. To project it, an expert forecasts the number of customers and the price each pays, the quantity of inputs and the cost of each, and the fixed costs, and takes the difference year by year.

The two lines are forecast independently because they are independent. A supplier of steel does not decide whether to sell to the manufacturer on the basis of how many customers the manufacturer has, and a customer does not decide whether to buy on the basis of how many suppliers there are. If a supplier is lost, the manufacturer sources elsewhere, perhaps at a higher cost, and the customer line is untouched. If customers are lost, the manufacturer buys fewer inputs, and the supplier line adjusts mechanically.

Most damages projections are built this way, and for most businesses that is correct. The difficulty arises when the same template is applied to a platform.

Two sides, two sets of cash flows

The first thing a platform changes is the direction of the cash flows. In the manufacturer's case, money comes in from customers and goes out to suppliers. On a platform, either side may pay and either side may be paid. A newspaper charges advertisers and subsidises readers; a marketplace charges sellers listing fees and commissions and may subsidise buyers with free delivery or discounts; a payment network charges merchants and rewards cardholders. The projection therefore needs, for each side, a net contribution per participant: what the platform receives from a participant on that side, less what it pays out in subsidies and servicing costs. Annual cash flow is the sum across both sides, less fixed costs.

That is a small change, and on its own it is not enough. If the number of participants on each side is still forecast independently, the two-sided projection collapses into the one-sided one with the labels rearranged. The substantive change is the next one.

Endogenous participation and the network multiplier

The second thing a platform changes is that the number of participants on each side is not an independent forecast. It is a response to the number on the other side. The simplest way to capture this is to treat each side as having a base of participants who would join regardless, plus a component that responds to the size of the other side:[3]

Buyers = base buyers + δ × Sellers

Sellers = base sellers + β × Buyers

Here δ is the number of additional buyers that each seller attracts and β is the number of additional sellers that each buyer attracts. 'Participants' here means units of participation, not heads. Most platform users multihome: a seller lists on rival marketplaces at once, a buyer splits spend across apps, and an event rarely makes either leave outright; it shifts their activity. Measured on listings, transactions or volumes, that shift is exactly what δ and β capture, while a count of registered accounts would miss it. Both are cross-side responses. Neither is a matter of assumption: both can be estimated from the platform's own history across regions and cohorts, from its past changes in fees and subsidies, or from the experience of comparable platforms.[4] Estimating δ and β is not straightforward, and an opposing expert may well challenge the method. The two sides move together, so the estimates must rely on events that hit one side directly and reach the other only through participation: a change in seller fees, a buyer-side promotion, a staggered regional launch. The techniques for doing this are well established.[5] Difficulty is no reason to skip the estimate. An expert who skips it has still made an assumption, usually that buyers stay put when sellers leave, or the reverse. An estimate at least puts its evidence where the opposing expert and the tribunal can examine it.

Solving the two equations together gives the equilibrium size of each side:

Buyers = (base buyers + δ × base sellers) / (1 - δβ)

with a symmetric expression for sellers. The term 1/(1 - δβ) is the network multiplier. It summarises the whole feedback loop: a seller attracts buyers, who attract sellers, who attract further buyers, and so on. When δ and β are both zero the multiplier is one and the projection is the conventional one. As the product δβ approaches one the multiplier becomes very large, and the platform's value comes to rest on a knife-edge.[6]

The projection then proceeds in five steps:

1. Estimate the base level of participation on each side, that is, the participants who would be present regardless of the other side;
2. Estimate the cross-side responses δ and β;
3. Solve the two equations for the equilibrium number of participants on each side;
4. Apply the net contribution per participant on each side, and deduct fixed costs, to obtain annual cash flow; and
5. Discount the resulting cash flows at a rate that reflects their risk, in the usual way.

To measure a loss, the same five steps are carried out with and without the event complained of, and the difference in value is taken. The critical point is that the event enters at step 1 or step 2, on the side it affects, and its consequences for the other side emerge at step 3. They are not added by hand.

A worked example

Table 1 sets out the assumptions for an illustrative marketplace. Each buyer contributes 10 per year net of subsidies and servicing costs, each seller contributes 4, and fixed costs are 5. Ten buyers and ten sellers would participate regardless of the other side, each seller attracts a further 1.8 buyers and each buyer attracts a further 0.2 sellers. The product δβ is 0.36 and the multiplier is 1.56.

TABLE 1 Assumptions for the illustrative marketplace

Input	Value
Net contribution per buyer, per year	10
Net contribution per seller, per year	4
Fixed costs, per year	5
Base buyers (present regardless of sellers)	10
Base sellers (present regardless of buyers)	10
Buyers attracted by each seller, δ	1.8
Sellers attracted by each buyer, β	0.2
Product, δβ	0.36
Network multiplier, 1/(1 - δβ)	1.56

Table 2 shows the feedback loop at work. The ten base sellers attract 18 buyers, so 28 are present; those 28 buyers attract 5.6 further sellers; the new sellers attract 10.08 further buyers; and so on. From the third round onwards each round adds 36 per cent (δβ) of what the previous round added, so the additions shrink geometrically and sum to a finite total: 43.75 buyers and 18.75 sellers. The multiplier reaches the same total in one step, 1.56 × 28 = 43.75 buyers, which is all it does: it sums the rounds.

TABLE 2 **The feedback loop, round by round**

Round	Buyers added	Buyers on the platform	Sellers added	Sellers on the platform
Base participants	10.00	10.00	10.00	10.00
1	18.00	28.00	5.60	15.60
2	10.08	38.08	2.02	17.62
3	3.63	41.71	0.73	18.34
4	1.31	43.02	0.26	18.60
5	0.47	43.49	0.09	18.70
All later rounds	0.26		0.05	
Equilibrium		43.75		18.75

Annual cash flow at the equilibrium is 507.5, as shown in the first column of Table 3. Now suppose an event removes two of the ten base sellers. It does not matter for present purposes whether the event is a breach of contract, a regulatory measure or the conduct of a competitor; what matters is that it lands on the seller side only.
The second and third columns of Table 3 show what the two projections make of it.



TABLE 3 One-sided and two-sided projections of the same event

	Before the event	One-sided projection	Two-sided projection
Sellers removed by the event		2.00	2.00
Buyers who leave because sellers have left			5.63
Sellers who leave because buyers have left			1.13
Sellers on the platform	18.75	16.75	15.63
Buyers on the platform	43.75	43.75	38.13
Cash flow from buyers (10 each)	437.50	437.50	381.25
Cash flow from sellers (4 each)	75.00	67.00	62.50
Less fixed costs	5.00	5.00	5.00
Annual cash flow	507.50	499.50	438.75
Annual loss		8.00	68.75
Loss as a proportion of cash flow		1.6%	13.5%

The one-sided projection deducts two sellers and their contribution and leaves the buyers where they were. The loss is 8 per year, about 1.6 per cent of cash flow. The two-sided projection lets the buyers respond. Each base seller removed takes with it 2.8 buyers (1.8 × 1.56), and those departing buyers take a further 0.56 sellers, so the event costs the platform 3.1 sellers and 5.6 buyers in all. Annual cash flow falls to 438.75, a loss of 68.75 per year, about 13.5 per cent of cash flow and more than eight times the one-sided figure. Nothing happened to any buyer directly. The buyers left because the sellers did.

Although this is an illustrative example, the intuition carries through that in any two-sided business, an event that lands on one side is transmitted to the other through the cross-side responses, and the size of the transmission is a matter of evidence about δ and β. The arithmetic runs equally well in reverse: an event that removes buyers reaches the sellers, and a projection that stops at the buyer line will miss it.

What moves the answer

Three features of the projection deserve particular attention, and each cuts in both directions.

The first is the multiplier itself. Table 4 repeats the exercise with β raised from 0.2 to 0.5, which moves $\delta\beta$ to 0.9 and the multiplier from 1.56 to 10. The platform's undamaged annual cash flow rises nearly seven times with no change to the economics of any individual participant, and the same two lost base sellers now cost 20 sellers and 36 buyers rather than 3 and 6.

TABLE 4 The same event at a higher multiplier

	$\delta\beta = 0.36$ ($\beta = 0.2$)	$\delta\beta = 0.9$ ($\beta = 0.5$)
Network multiplier	1.56	10.00
Buyers before the event	43.75	280.00
Sellers before the event	18.75	150.00
Annual cash flow before the event	507.50	3,395.00
Sellers lost to the event	3.13	20.00
Buyers lost to the event	5.63	36.00
Annual loss	68.75	440.00
Loss as a proportion of cash flow	13.5%	13.0%

Small differences in the estimated cross-side responses therefore produce large differences in the but-for valuation, which is a reason to examine the evidence for δ and β with the rigour normally reserved for the discount rate. A claimant model that embeds a high multiplier is asserting that the platform was worth a great deal and was extraordinarily fragile before the event; both halves of that claim deserve scrutiny. A respondent model that sets the cross-side responses to zero is asserting that the platform is not a platform at all.

The second is the discount rate. One of the standard corporate finance textbooks warns against 'fudge factors', by which it means the practice of loading a project-specific risk into the discount rate rather than into the projected cash flows.[7] Platform valuations attract fudge factors readily, because network effects feel risky and the temptation is to add a few points to the rate and move on. The two-sided projection removes the excuse. The exposure of a platform to an event that hits one side is a cash flow question: what probability attaches to the event, and what does the participation system do when it occurs. That is where the risk belongs. The transmission of a particular event through a particular platform's participation system is a feature of that platform, not of the market as a whole, and it has no place in the discount rate.

The third is timing. The equilibrium in the example is a resting point, and participants do not move to it overnight. Sellers on annual contracts leave when the contract ends, buyers drift away as choice thins, and some on both sides never leave. A multihoming participant adjusts fastest of all, shifting share without ever leaving. The speed of adjustment governs when the losses arrive, which matters for the damages period, for discounting and for mitigation. A platform that raises subsidies to hold the unaffected side in place is spending money to interrupt the transmission, and that expenditure is itself a loss the projection can size. The five steps can be carried out period by period with an adjustment lag, but the lag is a further parameter that needs evidence, and it is one the platform's own churn data will usually supply.

Implications for damages in arbitration

The two-sided projection is relevant wherever the event complained of falls on one side of a platform's market. In commercial arbitration that includes the termination of a distribution, payments or technology agreement that supplied one side; a price parity or most-favoured-nation clause that constrained the platform's pricing to one side; and a supplier's failure to deliver that thinned the offering to the other. In investment treaty arbitration it includes measures directed at one side: a licensing requirement for foreign sellers, a data localisation rule, a ban or a divestiture order, a discriminatory tax. In post-transaction disputes it includes earn-outs and price adjustments that turn on user numbers, and warranties as to the size of one side of the market.

In each case the valuation exercise is the same: the value of the platform with and without a shock to one side, with the other side allowed to respond. When such a valuation comes before a tribunal, five questions will usually locate the point of disagreement:

1. Has the expert projected both sides of the market together, or has the side not directly affected by the event been held fixed or adjusted by judgement?
2. Where do the cross-side responses come from? Are they estimated from the platform's own data, from comparable platforms, or assumed? Is their product well below one, and if not, is the pre-event valuation credible?
3. Is the counterfactual solved as a system, with the same parameters in the actual and but-for worlds, so that the loss is the difference between two consistent equilibria rather than the sum of two independent adjustments?
4. Where has the risk of the event been placed: in probability-weighted cash flows, or in the discount rate?
5. How quickly does the projection assume each side adjusts, and does that speed match the platform's observed churn?

None of these questions requires a tribunal to solve an equation. They require only the recognition that a platform's product is the other side of its own market, and that a loss on one side does not stay there.

Endnotes

1 Jean-Charles Rochet and Jean Tirole, 'Platform Competition in Two-Sided Markets', *Journal of the European Economic Association* 1(4) (2003), 990; Mark Armstrong, 'Competition in Two-Sided Markets', RAND Journal of Economics 37(3) (2006), 668. The underlying idea of a network effect is older: Michael Katz and Carl Shapiro, 'Network Externalities, Competition, and Compatibility', American Economic Review 75(3) (1985), 424.

2 Lapo Filistrucchi, Tobias Klein and Thomas Michielsen, 'Market Definition in Two-Sided Markets: Theory and Practice', *Journal of Competition Law and Economics* 8(2) (2012), 293; Thomas Eisenmann, Geoffrey Parker and Marshall Van Alstyne, 'Strategies for Two-Sided Markets', Harvard Business Review 84(10) (2006), 92.

3 The model is set out in Ryan Williams, 'Network Leverage and the Value of Platforms', Imperial College London working paper, where the cross-side responses are themselves written as functions of the subsidy paid to one side and the fee charged to the other, so that the platform's pricing to each side can be varied within the same framework.

4 The most direct evidence is usually the platform's own. A platform that operates in several regions or launched in several waves will have observed participation on both sides at different levels, and its past changes in fees or subsidies provide a further source of variation, since a change in what one side pays alters that side's participation and, through the cross-side response, the other's. Where the platform's own data are thin, the disclosed user numbers of comparable listed platforms can be used, with the usual caveats about comparability. For a mature platform with stable prices and participation, within-platform variation may be thin. The estimation then leans more heavily on comparables, on regional differences and on the discrete events in the platform's own history, and the resulting uncertainty belongs in the valuation as a stated range rather than a single point. Such platforms are also where the dispute is usually a moderate shock to an observed equilibrium, which is exactly where locally estimated responses are most defensible.

5 For estimation of cross-side responses using pricing and entry variation, see Marc Rysman, 'Competition Between Networks: A Study of the Market for Yellow Pages', Review of Economic Studies 71(2) (2004), 483; Junhong Chu and Puneet Manchanda, 'Quantifying Cross and Direct Network Effects in Online Consumer-to-Consumer Platforms', Marketing Science 35(6) (2016), 870.

6 In the linear form set out in the text, value is unbounded as $\delta\beta$ approaches one. In practice each side is drawn from a finite population, and the cross-side responses fall as a side approaches saturation. The responses estimated at the observed equilibrium should therefore be understood as local: adequate for measuring the effect of a shock of moderate size, less so for projecting a platform from a standing start. Time is handled the same way. The responses are estimated as of the valuation date and held constant across the actual and but-for worlds, so that the measured loss reflects the event rather than drift in the parameters; over a long damages period, the sensitivity of the result to movement in the responses should be shown alongside the central estimate.

7 Richard A Brealey, Stewart C Myers and Franklin Allen, *Principles of Corporate Finance* (13th edn., 2019), 241.