

Optis vs Apple: Can a FRAND price be inferred from “comparable” SEP licences that are not FRAND?

Pekka Sääskilahti, Andrew Tuffin, and Timo Autio¹

31 October 2025

Introduction

In May this year, the Court of Appeal (England and Wales) determined that Apple must pay Optis \$502 million for an 11-year global licence covering its patents essential to mobile-telecommunication standards – almost ten times the \$56 million awarded at first instance.² In this article, **Pekka Sääskilahti**, **Andrew Tuffin** and **Timo Autio** look at the reasons for that dramatic increase, and draw out lessons for future FRAND (fair, reasonable and non-discriminatory) determinations.

Both judgments sought to determine a FRAND price by reference to “comparable” contracts – similar licences that Optis and Apple had agreed with third parties. The rationale is that similarly situated parties will agree similar prices for similar contracts, and so the prices agreed for comparable licences will indicate the price that willing parties would agree for the disputed licence. However, the judgments use different methodologies to (a) **unpack** the evidence and (b) **evaluate** it. The gap between the determinations shows how much the specifics of a methodology matter at each stage.

For each, we draw out the following lessons:

1. Presenting the evidence: “unpacking” the prices that parties agreed

- **Lesson 1: The distinction between reliable and comparable prices.** Unpacking expresses the prices that parties *actually agreed* in a common “currency”. It should not adjust those prices to estimate what the parties might have agreed for more comparable contracts, nor should it prescribe a payment structure that comparable contracts must have.
- **Lesson 2: The benefit of using both parties’ comparables.** Where reliable, mapping both the prices of the licensor’s and the licensee’s comparables onto a common scale — the implied price for the overall SEP “stack” — helps reduce the risk of relying on contracts with a common bias.

2. Evaluating the evidence: assessing how closely comparable and reasonably priced those contracts are

- **Lesson 3: Analysing the range of unpacked rates.** A wide range of unpacked prices cannot all be comparable, FRAND, and reliable. The wider the range, the less likely it is that all prices are closely comparable, FRAND, and reliable.
- **Lesson 4: Analysing comparability between contracts.** Comparability is a matter of degree. The closest comparables to the disputed licence are the most informative of the price parties would agree for it.

- **Lesson 5: Analysing how reasonable the prices of the closest comparable are.** Even the closest comparables may not be FRAND. In this case, all five of the comparables were affected by some degree of hold up or hold out.
- **Lesson 6: Inferring the FRAND rate.** Even non-FRAND contracts may help establish the FRAND rate, provided that (a) there is a structured assessment of how closely comparable they are, and (b) there is a reliable cross-check to assess the extent to which the implied price is reasonable.

Stage 1: Presenting the evidence on “unpacked” prices agreed for comparable contracts

In the original decision, the parties put forward 23 (potentially) comparable licences: 9 were licences Optis had agreed for all or part of the same SEP portfolio it would license to Apple; 14 were licences that Apple had agreed with other licensors for their SEP portfolios, each of which were essential to the same technology standards as Optis’s SEPs.³

The starting point is to recognise that these contracts have **different prices** and **different characteristics** that determined those prices. An analysis of comparable contracts needs to understand both. The price agreed for a contract is only relevant to the extent that the characteristics that determined that price are genuinely similar to the contract being valued.

In the next section we consider how to evaluate which of the (potentially) comparable contracts offered the better guide. First, however, we look at a key concern in the original judgement: how to reliably “unpack” the prices that parties actually agreed for their (potentially) comparable contracts.

The challenge with “unpacking”

Comparing prices for similar assets is usually straightforward. With house prices, for example, the **value** of each payment varies by how desirable each house is, but the **structure** of the payments is typically the same: a one-off cash sum. That makes it easy to identify which have higher prices, which have lower prices, and to evaluate *why* their prices differ. By contrast, for SEP licences, it is common for differences in payment structure to obscure how the value of various contracts’ payment terms differ.

Unpacking is the task of normalising the differences in payment structure – translating them into a common “currency” – so that the present and cash-equivalent value of each contract’s payment terms can be compared.

The original decision identified three common differences that need to be unpacked – which it referred to as “subjectivities”:⁴

- **Different royalty structures:** some licences stipulate lump-sum payments; others lay out unit royalties, typically either (i) dollars per unit (DPU), or (ii) *ad valorem* royalties (a percentage of a royalty base, which is usually the implementing product’s average selling price).
- **Cross-licences** (or other payments in kind): where the payment terms include a transfer of goods, services, or rights – it is the cash-equivalent value of these non-monetary terms that matters.
- **Future royalties and past releases:** Where a licence covers past units that are nominally “royalty-free”, or discounted to some extent, it can obscure the underlying substance of the transaction: the total number of units covered (by the licence and/or the past release) and the total consideration the licensee must provide in return.⁵

In the original decision, Justice Smith accepted the necessity of unpacking, but noted that these subjectivities “make the objective of unpacking licences difficult, if not impossible, without importing simplifying assumptions which render the outcomes not comparable.” In this case, he considered the difficulties “insurmountable” and rejected all evidence on unpacked rates presented to him.⁶

Ground 4 of the appeal challenged that reasoning. It was upheld. The Appeal concluded that dealing with these subjectivities is unavoidable. For particular contracts they may inject enough uncertainty to render a specific estimate unreliable, but they are not insurmountable generally, nor were they insurmountable in this specific case. Indeed, the original decision necessarily had to grapple with all three issues in one way or another to make its determination.⁷

Lessons for unpacking

In discussing the original decision and setting out its own approach, the Court of Appeal provides several lessons for unpacking the prices agreed for (potentially) comparable contracts.

Lesson 1: The distinction between reliable and comparable prices

The original decision described the **purpose of unpacking** as to render incomparable licences comparable.⁸ However, as the appeal observes, that is not correct.⁹

The (limited) purpose of unpacking

Unpacking **should** reveal the value of the payment terms that parties *actually agreed* for a contract as it was; it **should not** estimate the price they *might have agreed* had the characteristics that determined the price of that contract been more closely comparable to those of the disputed licence.

Unpacking converts a contract’s particular payment terms into common metrics – “currencies” – for comparison. Those metrics should not change the fundamental **value** of the payment terms; they simply change how that value is **expressed**. By analogy, converting prices into USD to compare assets paid for in different currencies does not affect either how comparable an asset is, nor the amount that was paid; it simply normalises how the price for each asset is expressed.

For example, the FRAND determination for the Optis-Apple licence was a **lump sum**, but that amount can be presented in alternative, value-equivalent, “currencies” for comparison:

- **The lump sum awarded:** \$502 million.
- **The equivalent DPU:** \$0.15 per unit.
- **The equivalent ad valorem:** 0.032% using a representative handset ASP of \$470 (or 0.024% using a representative Apple ASP of \$625).

Calculating unpacked rates is straightforward. For instance, unpacking a lump sum into a DPU simply divides the total payment agreed by the volume of sales expected over the life of the contracts in present value terms.

Contracts can be unpacked into any common metric – and it often helps to unpack them into more than one, particularly where the parties disagree on which should be used. The Appeal notes that, although it was frustrating, ultimately the difference in the parties’ preferred metrics was immaterial: “by the end of the trial, as one might expect, both experts had expressed their unpacked conclusions in their rival’s “currency”.¹⁰

The distinction between the reliability and comparability of unpacked prices

In its consideration of the evidence on unpacked prices, the Appeal notes the “useful distinction” between “the **comparability** of a licence and the **reliability** of the evidence arising from it” (our *emphasis*):¹¹

“90. ... **Comparability** in the sense of a licence being a good parallel, as referred to by Lloyd LJ in *Cimetidine*, is primarily concerned with the situation of the parties and the subject matter being licensed... However this kind of comparability has nothing to do with whether the licence terms involve lump sums or ad valorem or DPU royalties... Unpacking does not alter that sort of comparability at all.

91. On the other hand **reliability** is concerned with the quality of the information derived from a given licence, perhaps involving unpacking in various respects. *With this distinction in mind, unpacking is not concerned with making a licence more or less comparable in the sense I have described it, it is about improving the reliability of the comparison.*”

The difficulty with unpacked rates typically comes from one of these two sources, but they need to be assessed and dealt with separately.

Assessing the reliability of the unpacked rates

The price unpacked should provide the same total present value as the payment terms the parties actually agreed. However, if some of the inputs to the calculation are uncertain, that unpacked price may not be reliably equivalent.

For example, if parties agree a \$300 million lump sum for a licence expected to cover 1.0 billion units, then the unpacked price is \$0.30 per unit. However, uncertainty about the expected sales affects its reliability. If the parties expected 1.2 billion units, the equivalent unit price would be \$0.25. The same issue arises with other key inputs, such as the discount rate used to convert future payments to present value, and the cash-equivalent value of any payments in kind. Uncertainty in those inputs reduces confidence that the unpacked price is a true like-for-like expression of what was agreed.

Uncertainty is a matter of degree. While an unpacked price may be sufficiently reliable, others may not be. That is an issue that must be assessed and taken into account.

Separating unpacking from the adjustments to make prices more comparable

In our experience, a common mistake in SEP cases is to ask too much from unpacking. Instead of converting payment terms into a common metric that reflects the value the parties actually agreed – given their real circumstances and the actual content of the licence – analysts start adjusting that number to account for differences between contracts.

That kind of adjustment can be necessary and important, but it is a **different** step. It does not give evidence of the price parties actually agreed; it interprets that evidence and then estimates a different price that they might have agreed under different circumstances. For transparency – and to reduce error – the two exercises should be kept separate.

To illustrate:

- **Houses.** When valuing a four-bedroom house, it is relevant evidence that a broadly comparable three-bedroom house next door sold for \$500,000. One might credibly judge that the four-bedroom equivalent would have sold for \$600,000. But the unpacked price of the three-bedroom

house is not \$600,000; that is an estimate that adjusts for the difference in their circumstances. The actual evidence is \$500,000 for the three-bedroom house.

- **SEP licences.** When valuing a disputed SEP licence, the same principle applies. It is relevant evidence that a broadly comparable portfolio was licensed for, say, \$3 per unit. However, if the disputed licence covers greater volumes, includes a larger share of past sales, or applies to products with a lower share of the latest standard, then none of those differences in substance change the fact that the unpacked price that was actually agreed for the potential comparable contracts is \$3 per unit. They may mean the contract is not closely comparable — but that is a separate question from unpacking.

Differences in payment structure do not necessarily mean that contracts are not comparable

A second mistake that can occur in SEP cases is to conflate the metric being used to express the unpacked prices with a contract's actual payment terms; in principle, there does not need to be any relationship between the two.

The Appeal highlights a “confusion” in the original judgment between:¹²

- **The “mode of overall payment”:** this refers to the actual payment structure in each contract. Apple sought a lump sum, which reflected the payment structure of its comparables. Optis favoured an ad valorem rate, which reflected the rate in most of its licences. By the time of the Appeal, both parties agreed the Optis–Apple licence would be paid as a lump sum.¹³
- **The “royalty case”:** this refers to the per-unit metric used to compare contracts. The main purpose was to (a) state prices in like-for-like terms, and then (b) *calculate* the amount of money that would be paid for the Optis–Apple licence. Apple used a DPU rate; Optis used ad valorem.

The first point to highlight is that the FRAND rate is simply “an amount of money”. It does not matter whether that amount is expressed as a gross lump sum, a per unit rate, an ad valorem rate, or another much more complicated set of terms. As the Appeal states: “either approach to a rate is manifestly capable of being FRAND” if it is adopted as the final payment structure. Furthermore, the metric used for comparison does not determine the final payment structure; in this case, it was “common ground that the FRAND licence will in fact involve a lump sum payment so that either ad valorem or DPU is just a tool on the way to producing that lump sum” (our emphasis).¹⁴

The second point to stress is that, if the parties chose different payment structures, then that may indicate a difference in their circumstance that makes the contracts less comparable. However, by itself, the structure of the payment terms is not the relevant difference. By analogy, two houses may be comparable, even if one was paid for in dollars and the other in euros; but they may cease to be comparable because one is in New York and the other in Berlin. What matters are the circumstances that affect their prices, not the difference in “currency” in which the parties expressed that price.

Although some SEP FRAND decisions have included the payment structure itself as a relevant aspect of assessing comparability or offering FRAND terms, in our view, it is the *value* of the payment terms, and the factors that affect that value, that matters.¹⁵ There are reasons why some parties choose running royalties and some choose lump sums. In some cases, those reasons may be substantial enough to mean that a contract is not comparable to another, but not in all cases. As set out in *Optis v Apple*, the payment structure is “just a tool” for calculating that value.

Lesson 2: The benefit of using both parties' comparables

The Appeal notes that the original decision sought to include both the licensor's and licensee's comparables in the evidence base, and was right to do so.¹⁶

The main benefit is that it reduces risk if all comparables from one party are tainted by a common factor. In this case, for example, Optis argued all Apple's comparables were depressed by hold out, while Apple argued that all of Optis's were inflated by hold up. Placing the prices agreed for each set of comparables on a common scale helps evaluate the plausibility of those claims – as we explore in section 2 below.

The common scale required to compare the licensor's comparables with the licensee's comparables is the Implied ARR (Aggregate Royalty Rate) an implementer would pay to license all SEPs, if (a) all SEP holders licensed their portfolios, and (b) did so on the same basis. The judgment refers to this as the price of the "stack". An implied stack price is needed because the licensor's contracts and *each* of the licensee's contracts cover different patents. Unless the relative contribution of each portfolio is taken into account, their prices can't be compared meaningfully.

Comparing the licensor's contracts with *any* of the licensee's contracts comes with a large and important caveat: the estimated price for the stack needs to be sufficiently reliable to be informative.

The Appeal recognises that it is difficult to reliably estimate a portfolio's contribution to the stack, but also that it can be done to some extent. A portfolio's contribution depends not only on the share of declared SEPs, but also on the relative proportion that are valid, essential, and infringed, and on the relative importance of the patented technology. Errors in these inputs can materially distort the implied stack price – especially for small portfolios such as Optis's.¹⁷

The practical consequence is that – other things equal – comparing the prices unpacked from the SEP holder's licences will always be more reliable than comparing stack prices implied by licences for different SEP portfolios. However, provided the risks are identified and managed where possible, unpacking both parties' comparables will better inform the valuation. The Appeal emphasises this point in paragraph 95, which we quote in full.

*"95. Before leaving this topic I will say a bit more about starting from the SEP holder's own licences. I maintain that this is the place to start but that is all it is. Factors like hold out and hold up may well render licences of the same portfolio less good as comparables. In this case there are also licences to the putative licensee (Apple). Such licences are capable of being useful comparables, again subject to hold up and hold out, but using them also involves a further dimension which is why, although they may well be useful in the end, they are not the best place to start. Their comparability (not reliability) also depends on the relationship between the patent portfolio being licensed and the SEP holder's portfolio. Not only does one need a view about stack shares, the issue of portfolio quality arises. It is not enough to render them comparable to say that the SEP holder's portfolio is average. The other licensed portfolios also have to be examined."*¹⁸

Stage 2: Evaluating the evidence to assess how closely comparable and reasonably priced those contracts are

Having unpacked the prices actually agreed for the (potentially) comparable contracts, the substance of an analysis of comparables lies in **evaluating** those unpacked prices to infer a reasonable price for the disputed licence.

By the time of the Appeal – after the concessions noted at paragraph 99 – only 5 licences remained for evaluation: Optis's licence with Google; and 4 of Apple's licences, all from licensors with

relatively substantial portfolios (Ericsson, Nokia, InterDigital, Sisvel). In principle, however, all 23 licences could have been compared on the same basis.

For illustrative purposes, **Figure 1** shows the range of unpacked prices for both Optis and Apple comparables.

- **Optis's comparables.** The exact figures are confidential, but the judgment indicates Optis's Google licence implied an aggregate stack rate above 15% (\$71 DPU), and other Optis comparables exceeded 30%.¹⁹ Purely for the sake of illustration, in **Figure 1** we highlight 15% to represent Optis's Google comparable.
- **Apple's comparables.** The exact figures are also confidential, however, the judgment indicates Apple's 14 (potentially) comparable contracts imply stack prices that range from more than 3% (about \$14 DPU on a representative \$470 handset) up to 8% (about \$38 DPU).²⁰ The highest of those had been rejected in the original decision as too high to be reasonable, and the 4 comparables used were not at the bottom of the range.²¹ Purely for the sake of illustration, in **Figure 1** we highlight 5% to represent the 4 Apple comparables.

Figure 1: Illustration of 4G SEP rates for smartphones assessed in Optis v Apple



Source: Compass Lexecon analysis based on data from [2025] EWCA Civ 552, [2017] EWHC 711 (Pat), [2019] *TCL Communication Technology Holdings Ltd v Telefonaktiebolaget LM Ericsson*, USCA Fed. Cir. No. 18-1363, and a translation of [2023] *OPPO v Nokia*, The People's Republic of China Chongqing First Intermediate Court (2021 Yu 01 Minchu no. 1232), page 100 ('Reasoning' Section V. (I) 1.)). Note: Implied Aggregate Royalty Rate (ARR) as % of a representative \$470 phone..

For context, **Figure 1** also shows:

- **The original determination:** with an implied stack price below 1% (or \$4 per unit), it is less than a third of the price implied by Apple's lowest priced licence, and almost an order of magnitude less than the rate implied by its most expensive.
- **The three potential FRAND rates the Court of Appeal considered in its determination;**²² and
- **A range of other potential reference points:** *Unwired Planet v Huawei* determined a FRAND stack price for 4G SEPs equivalent to 8.8%; *Nokia v Oppo* (Chongqing) and *TCL v Ericsson* (US) both determined FRAND stack prices equivalent to 6-8%, based on the range of stack prices major licensors had said they expected before LTE became the global 4G standard.²³

The challenge with evaluating (potentially) comparable SEP contracts

In *Optis–Apple*, the prices unpacked from the 23 (potentially) comparable licences were not equally informative. The breadth of unpacked prices was wide. Intuitively, some contracts are better indicators of a reasonable price for the Optis–Apple licence than others. Each party emphasised the contracts it considered the most informative.

Broadly, the reasons that explain why unpacked prices differ fall into three categories:

- **The contracts' characteristics are not equally comparable.** The contracts differed in various ways and to varying degrees – in particular: the contracts differ by (a) the licensed patents they cover, (b) the licensed products they cover, (c) the other contract terms in the licence, and (d) the characteristics of the parties at time of the negotiation. Each of these differences may affect the prices parties agree for a particular licence. While some disparities can be adjusted for reliability, that is not always possible – and relying heavily on adjustments dilutes the rationale for using a comparables approach: that market prices indicate the value of a disputed licence's value, not an analyst's estimate.
- **The unpacked prices are not equally reliable.** The Appeal accepts that uncertainties can undermine the quality of the information that (potentially) comparable contracts provide. For example, unpacked prices may be sensitive to: differences in licensed product volumes, the discount factors used to calculate present value, estimates of the cash-equivalent gross royalty rate inferred from net royalty rates, the treatment and scale of volumes subject to past-release, and the patent data used to estimate a portfolio's contribution to the stack.²⁴ Any of these issues may render an unpacked estimate unreliable.
- **The unpacked prices are not equally reasonable.** In SEP licence negotiation, neither side truly has an "outside option", so there is little if any competitive pressure on the prices they agree to. Therefore, the unpacked prices may be too high to be FRAND (due to licensor hold up), or too low (due to licensee hold out).

The original decision criticised the parties for wrongly adopting what it called an "exclusionary" approach to selecting the most informative comparable contracts.²⁵ Instead, the judgement used an "inclusive approach" – selecting all 20 licences that had agreed a lump-sum payment structure, calculating the lump sum that each one implied for the entire stack, excluded outliers, and then took a simple average to the remaining licences in order to award Optis a proportional share of the implied stack price.

The Appeal overturned that approach. Even on its own terms, the method contained errors that would have required correction (ground 8). The Appeal also rejected the use of averaging (ground

7). More fundamentally, the Appeal determined that “the right approach was to adopt a comparables based approach (ground 2) in the sense of being one based on identifying the best comparable or comparables, excluding others and working from there.”²⁶

Having rejected the original determination, the Appeal performed its own evaluation of the evidence on the comparables’ unpacked prices – which it had restored (ground 4).²⁷

Lessons for evaluating comparable contract prices

The evaluation of the unpacked prices in *Optis v Apple* provides useful lessons for FRAND determinations using comparable contracts.

Lesson 3: Analysing the range of unpacked rates.

In our experience, courts are often presented with a wide range of unpacked prices. In those cases, it is not unusual for the following dubious propositions to be claimed:

- that the evidence presents a range within which *any* price is reasonable; or
- that a simple (or weighted) average provides a FRAND price.

The Appeal on *Optis-Apple* concludes that both propositions are wrong.²⁸

“109. In truth there is a simple explanation for the spread, which is not solved by taking any kind of average. Either these licences are not all equally good comparables for various reasons or, putting it another way, not all these rates are FRAND.”

A wide range cannot be FRAND

The central premise of the comparables approach is that similarly situated parties will pay similar prices for similar licences. On that basis, if all contracts are perfectly comparable and their unpacked prices are reasonable and reliably calculated, then they will converge on a narrow range of prices. If they do not converge, then at least some of them are not comparable, not FRAND, or unreliable.

In paragraph 108, the Appeal comments on the range of implied prices for the stack considered in the original decision. We quote in full, but add emphasis:²⁹

*“108. The problem is apparent on the face of Table 13 itself. Bear in mind Table 13 puts all the licences on a common scale. The most striking thing about the table is that the range of values is very wide even when Blackberry and Qualcomm are excluded. Assuming each licensee has obtained a share of the overall “pie” in proportion to their stack share, it follows that the pie from which the licensees at the top of the list have obtained a slice is far larger than the pie used for the licensee at the bottom. **The obvious conclusion is that either this approach does not work or that these licences cannot all be useful comparables.** The approach puts them too far apart from each other. The judge recognised the problem at [483](i) (“**There is no convergence on a single price. That does undermine confidence in the figures...**”). It led the judge to make the exclusions mentioned already but in my judgment he ought to have realised that the spread in the table after he had made these adjustments was still too large to take forward. The problem identified has not been solved.”*

Averaging prices that are not comparable, reliable, and reasonable does not solve the problem

To deal with the wide range in unpacked prices, the original judgment took a simple average.³⁰ That approach was appealed (ground 7), and upheld, as it: “had no precedent or basis in the evidence before him nor can it be justified in principle.”³¹

A weighted average was proposed, but the appeal notes the problem with averaging is a matter of principle, not the specific approach. “Both forms of averaging share the same weakness in giving weight to all licences in the table as comparables, albeit that weight is different.”³² The problem with giving weight to all licences, is that they are not equally informative, as we explore in the next lesson.

Lesson 4: Analysing comparability between contracts

Having identified an implausibly wide range of unpacked prices, the task was not to accept the spread – or average it away. The task was to **analyse, explain, and interpret** the reasons for the differences. The first issue is how closely comparable the circumstances that determine their prices are.

The Appeal concluded that the conventional comparables approach – termed the “exclusionary approach” – had wrongly been rejected. “The right approach was to adopt a comparables based approach (ground 2) in the sense of being one based on identifying the best comparable or comparables, excluding others and working from there.”³³

The Appeal sets out the rationale for the comparables approach – it refers to English Law, but the principles apply more broadly.

- **The “closest possible parallel” principle (from *Cimetidine*).** “The object of the comparability exercise, in this as in any other branch of the law, is to find the closest possible parallel. If there is an exact parallel, there is no point in looking any further. If there are slight differences, an allowance may be made. But once you have found your comparables, whether one or more, which enable you to arrive at the appropriate figure, it would surely be erroneous to modify that figure by reference to other cases which are not truly comparable at all, so as to bring the case into line with a predetermined range. This was, with great respect, the mistake which the hearing officer made.”³⁴
- **The “as relevant, but not the same” clarification (from *Unwired Planet*).** “In my judgment, if a group of comparables are at least potentially as relevant as each other and are not the same, it is not right to elevate a small subset above the others. That is also not what Lloyd LJ in *Cimetidine* said one must do; instead, he said that, assuming there is no exact parallel, once true comparables have been determined one should be careful not to dilute them by reference to other cases which are not truly comparable at all. [...] If a group of good comparables corroborate one another then no doubt that is a factor to take into account but equally if apparently good comparables, when properly understood, contain different rates that is also relevant too.”³⁵

The importance of the first principle should be straight forward. By analogy, if the semi-detached house next door is sold for \$500,000, with the same footprint, layout, and condition as the house being valued, that likely indicates its reasonable price; it is close to an exact parallel. Adding the prices of bungalows, flats, and tenements that are located many miles away only dilutes the evidence base.

The same principle applies to licences for intellectual property. Those covering patents, products, and with price-relevant characteristics that closely resemble those of the disputed contract, will provide a better indication of its price than licences that differ to a greater extent.

However, in SEP cases, very close parallels are rare. Furthermore, the factors that affect the prices parties agree for contracts may be harder to identify, or have their significance disputed. In these circumstances, the clarification in *Unwired Planet* applies. Several contracts may broadly be as close as each other to the licence being valued, but in different respects and with different prices. In that situation there may be no objective basis to prefer one subset's prices over another's; they should be treated as equally informative.

Nonetheless, that does not mean demonstrably weaker parallels can be placed on the same footing as closer ones. Nor does it mean that *any degree* of price dispersion between notionally "comparable" contracts can be tolerated. If prices differ, there are reasons for it. Sometimes they cannot be determined. But the wider the disparity, the more effort should be made to identify and explain the drivers of difference – and, where possible, address them.

In principle, the differences between contracts' characteristics can be **adjusted for**. And where it is **reliable** to do so, they should be adjusted for. However, that is a distinct exercise from identifying the closest comparables, and it is not always possible to adjust contract prices reliably. The best approach is to identify the contracts that most closely compare to the licence being valued – as they are less in need of adjustment to start with.

Lesson 5: Analysing how reasonable the prices agreed for the closest comparables are

The second issue to consider when looking at the range of unpacked prices is how reasonable the prices that parties agreed for those contracts are. In other words, to what extent they might be affected by a degree of hold up or hold out.

A striking point from *Optis–Apple* is that none of the comparables were considered reasonable. The price implied by the Optis–Google licence was judged to be "too high". The prices implied by Apple's licences (with Ericsson, Nokia, InterDigital, Sisvel) were judged to be "too low".³⁶

We discuss the reasons below, but first, it is important to recognise what is meant – and what is *not* meant – by the finding that none of the comparable contacts had FRAND prices.

The FRAND rate as an idealised standard

In his judgment, quoting the Supreme Court, Lord Justice Birss describes the purpose of the FRAND rate as to protect *both* parties from the "mischiefs" of hold up and hold out.³⁷ He then made three important clarifications.³⁸

- **The FRAND rate is an "idealised legal standard"**. As a matter of principle, it is the rate which a willing licensee and willing licensor would agree, free of hold up or hold out.
- **Hold up and hold out are "plainly a matter of degree"**. The judgment notes that "willing licensees do not engage in hold out to any degree and willing licensors do not engage in hold up to any degree either. Neither party needs to be protected from the behaviour of the other. By contrast real parties negotiate as hard as they can."
- **A rate that is not FRAND is not "intrinsically unlawful or illegitimate"**. The judgment explains that "a finding that a degree of hold up or hold out was involved in a given real negotiation is not a finding that either party has acted in an unlawful manner. It is simply a finding that that outcome cannot be taken as the FRAND rate. It may be close to it or far away, and if a view can be taken about the degree of hold up or hold out involved that might shed some light, but these are different issues."

The Optis-Google licence would set a price that is too high to be FRAND

The Appeal determines that a price reflecting the one agreed for the Optis-Google licence would be “too high” for two reasons. First, it is too far from the stack prices implied by Apple’s own comparables to be plausibly FRAND. The second reason comes from a “top-down sense check” against the average selling price of a phone. At over 15%, the Optis comparable implies a stack price that is not “consistent with a sensible business model” – in other words, the stack price seems unaffordable on its face.³⁹

Given the scope of the Appeal, the judgment does not explore the reasons why the contract’s implied stack price is so high. However, from the limited information provided in the judgment, our own view is that there are three potential reasons (although there may be others).

- **On reliability: the implied stack price may be artificially high.** Optis has a very small portfolio; estimated to constitute 0.38% or 0.61% of the stack depending on the licensee. The first problem with portfolios of this scale is that the implied stack price is very sensitive to their estimated share. A minor change in the estimate would substantially change the implied stack price. Second, the implied stack price assumes a linear relationship between the portfolio share and the price. That may break down for small and large portfolios – the original judgment draws a useful analogy with taxi fares that may apply in this context too: while most trips may reflect a fixed rate per mile or minute, very short or long journeys don’t. For short trips, a fixed “hail fare” would distort the notional per mile/minute rate upward; and for long trips, a negotiated “fixed price” might be agreed that suppresses the notional per mile/minute rate. That does not make the pricing structure illegitimate; it just means that simple extrapolations based on the majority of trips may not be reliable when applied to cases at the extremes – and vice versa.⁴⁰
- **On comparability: low-volume and high-volume licences are not necessarily comparable.** The Google licence covers far fewer licensed units than any of Apple’s comparables. This may affect their comparability in two respects. First, the transaction cost of negotiations does not scale linearly with volumes. So, for a small volume contract like Google’s, the transaction costs would constitute a larger proportion of the per unit rate, as they are spread less thinly than they are on large volume contracts. Secondly, a small-volume contract provides less benefit to the licensor, as it reduces its exposure to unlicensed sales to a lesser extent than a large volume contract does. Therefore, a licensor might recognise the additional benefit of a large volume contract by accepting a lower *rate*. In *InterDigital v Lenovo*’s discussion of volume discounts, the court accepts the *relevance* of these issues but considers they could not explain very large differences between the prices small volume and large volume licensees agree.⁴¹
- **On reasonableness: the marginal benefit of a lower royalty may be outweighed by the marginal cost of pursuing it.** Apple argued that the lump sum paid was comparable to the cost of litigation with Optis, and so Google were in effect paying Optis off based on the “nuisance value” of the dispute. The logic here is that, when both the licensed portfolio and licensed product volumes are small, the expected royalty revenue is relatively small. In that circumstance, it is possible that the cost of pursuing a lower royalty *rate* may be offset by the total cost of negotiating it – meaning it is cheaper to settle for a high rate. In this specific case, however, the original judgment dismissed this general concern on the basis that a company of Google’s stature would have scrutinised the sum with some care.⁴²

The Apple licences would have set a price that is too low to be FRAND

The Appeal determines that a price reflecting Apple’s licences would be “too low” to be FRAND. The first reason provided is that the price they imply would sit too far from the stack price implied by Optis’s comparable to be plausibly FRAND.⁴³ This is significant. Although the Optis comparable

is not FRAND, that does not render it irrelevant – the *degree* to which it is too high matters. Given that, and the very large difference between its price and the prices agreed for the Apple comparables, the simplest explanation is that both sets of contracts are outside the reasonable range on different sides.

The second reason provided is that Apple's contracts were judged to have been suppressed by hold out. This issue was a focus of the Appeal (ground 3). The appeal was upheld, concluding (emphasis added):⁴⁴

*"123. Therefore on this ground, and expressed in deliberately qualitative terms, in my judgment the judge was wrong to place weight on the values derived from the Apple licences as a whole. **There is hold out involved**, particularly as an explanation for the spread of values of these licences whenever they are put on a common scale by stack share. The difference is not explicable by the disproportionate effect of errors in the stack share when they are grossed up."*

The conclusion has two foundations.

The first is that Apple's licensing framework contains what the judgment calls "indefensible elements", such as patent by patent licensing and the SSPPU (smallest saleable patent practicing unit) approach.⁴⁵ The former is considered a practice that no willing party would adopt when seeking to agree terms for a global licence.⁴⁶ The latter was dismissed in the original judgment for being "economically and rationally indefensible" – an "artificial attempt" to reduce the prices which "in no way assists in deriving a price for the technology licence here in issue".⁴⁷ That conclusion was not appealed.⁴⁸

The second comes from an inspection of the range of stack prices implied by Apple's licences. The judgment considers that the range – from 3% to 8% – is too wide to be explained by any factors that willing licensors and licensees would agree to. Further, with some exceptions, the licensors with larger portfolios imply higher stack prices than those with smaller portfolios – which is opposite of what the royalty stacking theory and "nuisance value" of small licensors would predict. The judgement reasons that: *"there is only one, hardly surprising conclusion. When it can do so, Apple's significant negotiating strength leads some parties to agree lower rates than would be agreed between a willing licensor/willing licensee. There is a degree of hold out involved."*⁴⁹

That comparable SEP contracts may reveal unreasonable prices should not surprise anyone. They are all negotiated in a situation where neither side has an outside option that would otherwise apply competitive pressure on the price they negotiate. No implementer can opt for a rival technology. No licensor can opt to implement its technology exclusively – nor can it license to the highest bidder, leveraging the rivalry between implementers seeking to secure the benefits that its technology provides.

Lesson 6: Inferring the FRAND rate

The Appeal's determination was inferred from the comparables, but it was not a price that any of the comparables actually had. This is normal when analysing comparables, as there is rarely a perfect parallel. Some allowance must be made to recognise the impact that differences in circumstances or uncertainty may have. However, this inevitably introduces risk into the analysis. What started out as a valuation based on the prices that real parties agreed to, may become reliant on estimates that are untested in the real world – which is how subjectivity, bias, and error may enter the process.

Broadly, there are two (complementary) approaches:

- **Adjusting the unpacked prices for specific factors that affect their value.** This approach identifies and explains the extent to which a contract is (a) not closely comparable, and if possible (b) not reasonable — and then adjusts for those differences. This can improve the valuation, but it is risky unless there is clear evidence of how the factor being adjusted for actually affects the prices parties pay. These two types of adjustment should always be kept distinct. Comparability adjustments are empirical: they correct for factual differences that really do affect the prices parties agree. These adjustments, therefore, have more hope of being reliable. Reasonableness adjustments are different. They correct for the effect of hold up or hold out — i.e. behaviour that should not drive the price agreed between willing parties. This is difficult to do with granular adjustments. It is safer using a cross-check, that helps assess whether a price is reasonable.
- **Triangulating from “cross-checks”, including other contracts and benchmarks of value.** This approach does not try to “fix” each contract step-by-step. Instead, it compares candidate FRAND rates against two type of reference point: (a) the unpacked prices in other contracts, taking into account *the reasons why* they might not be closely comparable, reliable, and reasonable; and (b) a top-down cross-check, that indicates the value of the functionality the technology provides, or helps assess how affordable the proposed royalty would be.

In effect, the Court of Appeal followed this second approach. It considered three possible FRAND rates — \$0.10, \$0.15, and \$0.20 per unit — and evaluated each against (a) the non-FRAND comparables, and (b) the share of a handset’s average selling price that the implied total stack would represent.

Note that the judgment does not determine a FRAND range, and it does not draw a boundary between rates that are “FRAND” and “not FRAND.” Instead, it concludes that — for this licence, on this evidence — a rate at \$0.20 or above would not be FRAND, a rate at \$0.10 or below would not be FRAND, and a rate at \$0.15 is FRAND. That is sufficient for the determination.

The absence of a defined boundary is also underpinned by the judgment’s logic. FRAND is an idealised rate: it is the price that willing parties would agree for a perfectly comparable licence, in negotiations without any degree of hold up or hold out. As a contract becomes less closely comparable, and/or its price is increasingly affected by hold up or hold out, then the price it implies for the disputed licence increasingly departs from FRAND. It is a matter of degree.

Refining the approach

The Court of Appeal acknowledged that its exercise was “a broad axe,” and that “the nature of the evidence here does not justify fine[r] distinctions” than the three candidate rates.⁵⁰ But the underlying principles are useful, and their application in future cases could be refined in two ways.

First, both approaches — adjustments and triangulation — benefit from an explicit and thorough understanding of how closely comparable contracts are. In particular, an explicit framework helps (i) identify the factors that determine comparability, (ii) articulate the extent to which contracts may be closely or distantly comparables with respect to each factor, and (iii) explain the relationship those differences have with contract prices. The better that understanding, the more reliably one can (a) identify contracts that are closely comparable to start with, and (b) adjust for their deficiencies or triangulate between them. As the quality of that understanding improves, the two approaches start to converge.

Second, the better the benchmark used to “cross-check” the various rates, the more accurately one can assess whether they are reasonable. This case demonstrates that the rates parties agree are

not necessarily reasonable. In principle, the *entire* range of unpacked rates, whether wide or narrow, may not be FRAND if they are all biased. An independent cross-check helps assess that risk.

The Appeal used the average selling price of a phone as that benchmark: a representative handset at \$470, and a representative Apple handset at \$625. Such an approach allows a broad consideration of whether the implied stack price is “plainly untenable” or not. But otherwise, it is less than ideal due to two limitations.

- **Limitation 1: There is no direct relationship between the price of the implementing product, and contribution that the licensed technology makes to that price.** Clearly, an implementer would not be able to charge as much for the phone if it lacked the functionality that the licensed technology enabled. But in abstract, we cannot say to what extent. The licensed technology may contribute a lot to that \$470, or it might not. There are many technologies in a smartphone that contribute to its value: cellular technology is one of them, but there are many others too. Some of those contribute to the phone independently of the licensed technology. Others overlap, adding more value if the quality of the cellular technology is greater – for instance, large high-resolution screens and high-speed processors may be less valuable to consumers if they lack the download speeds to use them effectively, and more valuable if they do.
- **Limitation 2: Prices can be updated to reflect reasonable costs.** Implementers set prices based on their expectations of the reasonable costs they will pay. In principle, those expectations might be optimistic. Using the ASP as a benchmark implies that the price of a phone is fixed *regardless* of the royalty that is determined to be FRAND, but strictly speaking, that is not true. The ASP is not a *cap* on reasonable costs; if an implementer underestimates reasonable licensing costs, it should update its prices to reflect them, as it would for other input costs.

To illustrate, consider how an implementer may adjust its pricing to reflect licensing costs before and after a higher than expected FRAND determination. Before, the implementer has contracts for 50% of the stack, and an average implied stack price of 5% (\$23.50 per unit, using the lower ASP).⁵¹ So, the implementer might reasonably anticipate that it must allow for \$11.75 per unit in licensing costs – which it might pay as a lump sum. Now, after a determination that the FRAND price is in fact 8.4%, there are two considerations. First, if all existing contracts were renewed on equivalent terms, its expected costs would rise to \$19.50. Second, if all *other* SEP holders came forward – which in principle they would be entitled to do – the expected cost would rise to \$39.00 per unit: the full stack price. The original ASP, however, is not a constraint. If needed, the implementer could adjust that price in light of the reasonable costs. Alternatively, if it has sufficient pricing power that it does not need to alter its price, then it would simply pay the reasonable costs as a lump sum.

We have discussed the benefits of other benchmarks in a previous paper, so we don’t repeat them here.⁵² A benchmark will better inform an evaluation of reasonable prices, the extent to which it isolates the value of the implementing product’s functionality *depends* on the licensed technology, and the extent to which it can account for the pass-through of changes in the incremental licensing costs expected – if required.

The broad point, however, is the same: whether the **price** proposed for a contract is reasonable or not depends on how it compares to the **value** of the benefits that the licensed technology actually provides.

Pekka Sääskilahti is an Executive Vice President, Andrew Tuffin is a Vice President, and Timo Autio is a Vice President at Compass Lexecon. We have benefitted from comments by Ben Dubowitz and Janne Sirnio. Adam Sanderson provided research assistance. The views expressed in this article are the views of the authors only and do not necessarily represent the views of Compass Lexecon, its management, its subsidiaries, its affiliates, its employees or its clients.

For the Appeal, see [2025] EWCA Civ 552. For the original decision, see [2023] EWHC 1095 (Ch). Apple sought permission to appeal to the UK Supreme Court against the orders giving effect to the Court of Appeal's ruling. UKSC/2025/0145. <https://www.supremecourt.uk/cases/uksc-2025-0145>. Updated 6 November: On 4 November, the UK Supreme Court granted Apple permission to appeal.

[2025] EWCA Civ 552, paragraphs 50-52.

"50. Optis's comparables approach is addressed at [397] – [398]. As [397] explains, Optis had identified 19 licences which covered all or a sub-set of the relevant patent portfolio and from that pool Optis had identified 9 (the Optis Comparables) as constituting suitably reliable and useful data points. Many of these licences were based on ad valorem rates but for those which were not, such as the Google licence, the lump sum features were unpacked."

"52. [...] Apple had identified 19 licences in which it was a licensee and excluded 5 altogether making 14 comparables....]"

[2023] EWHC 1095 (Ch), paragraphs 301-317.

[2023] EWHC 539 (Pat) paragraphs 520-545, including: "528. Having given this issue considerable thought, I have reached the conclusion that limitation periods do not have a role in the relationship between willing licensor and willing licensee and, indeed, that they are inconsistent with that relationship. As I have explained above, a willing licensee will, notionally or otherwise, set aside funds to pay for its licence. If, for some reason, those willing parties are not able to reach a deal for some time (assuming the negotiations last for longer than 6 years), I do not believe that a willing licensee would refuse to pay whatever licence fees were eventually determined to be applicable in respect of units produced and sold more than 6 years prior to the determination. A licensee who did that would no longer qualify as 'willing'."

[2023] EWHC 1095 (Ch), paragraphs 297, 302, and 314-317.

[2025] EWCA Civ 552, paragraphs 86-97.

[2023] EWHC 1095 (Ch), paragraph 297.

[2025] EWCA Civ 552, paragraph 90.

[2025] EWCA Civ 552, paragraph 16.

[2025] EWCA Civ 552, paragraph 90-91.

[2025] EWCA Civ 552, paragraph 41.

[2025] EWCA Civ 552, paragraph 99.

[2025] EWCA Civ 552, paragraph 89.

For instance, the Karlsruhe court in *Philips v Wiko* considered that Philips' offer was discriminatory because it did not adequately explain why it agreed a lump-sum payment with one party and not a running royalty as offered to Wiko. See *Philips v Wiko*, 6 U 183/16 Karlsruhe Higher Regional Court (30 October 2019). Similarly, the Düsseldorf Higher Regional Court considered that once an SEP licensor agreed the first licence, then the specific structure and value of the price agreed for that licence became binding, such that it must be applied to all subsequent contemporaneous licences. *Unwired Planet v Huawei*, I-2 U 31/16 Düsseldorf Higher Regional Court (22 March 2019) 413-414 and 420.

[2025] EWCA Civ 552, paragraph 129: "One of the key things the judge was trying to do was to reach a result which took into account licences both from Optis and to Apple. I believe the judge's instincts here were correct. However the flawed methodology led him to be misled into the impression that they imply similar FRAND prices. They do not."

[2025] EWCA Civ 552, paragraph 19.

[2025] EWCA Civ 552, paragraph 95.

[2025] EWCA Civ 552, paragraph 50.

[2025] EWCA Civ 552, paragraph 52.

[2025] EWCA Civ 552, paragraph 110 “That might have led the judge to wonder if Blackberry and Qualcomm were as much outliers in his own method as he thought, but in my judgment it is too late to try to adjust the judge’s approach that way by putting Blackberry and/or Qualcomm back in at full value. Neither side suggests that.”

[2025] EWCA Civ 552, paragraphs 137 and 146.

[2017] EWHC 711 (Pat), paragraphs 476 and 807.

[2019] *TCL Communication Technology Holdings Ltd v Telefonaktiebolaget LM Ericsson* USCA Fed. Cir. No. 18-1363, page 10.

A translation of [2023] *OPPO v Nokia*, The People’s Republic of China Chongqing First Intermediate Court (2021 Yu 01 Minchu no. 1232), page 100 (‘Reasoning’ Section V. (I) 1.)).

[2025] EWCA Civ 552, paragraphs 56, 92, 96, 99, and 112.

[2023] EWHC 1095 (Ch), paragraphs 290-291.

[2025] EWCA Civ 552, paragraphs 110 and 115.

[2025] EWCA Civ 552, paragraph 97.

[2025] EWCA Civ 552, paragraph 109.

[2025] EWCA Civ 552, paragraph 108.

The range was artificially narrower as (a) Apple’s Qualcomm and BlackBerry licences were removed as outliers, and (b) the Optis’ Google licence was misstated – the lump sum was assessed without properly accounting for the very small number of units the licence covered. The Appeal notes that had the Google licence been properly unpacked, “that might have led the judge to wonder if Blackberry and Qualcomm were as much outliers in his own method as he thought, but in my judgment it is too late to try to adjust the judge’s approach that way by putting Blackberry and/or Qualcomm back in at full value. Neither side suggests that.”

[2025] EWCA Civ 552, paragraph 110.

[2025] EWCA Civ 552, paragraph 106.

[2025] EWCA Civ 552, paragraph 108.

[2025] EWCA Civ 552, paragraph 115.

Smith Kline & French Laboratories Ltd’s (Cimetidine) Patents [1990] RPC 203 (Lloyd LJ), cited in [2025] EWCA Civ 552, paragraph 35.

[2017] EWHC 2988 (Pat), paragraphs 173-174, cited in [2025] EWCA Civ 552, paragraph 36.

[2025] EWCA Civ 552, paragraphs 138-139.

[2025] EWCA Civ 552, paragraph 118. “the Supreme Court describes the purpose of the FRAND undertaking given by a SEP holder as being to protect (my emphasis) implementers and ETSI from hold up, while [59] also explains there is a balance: to protect implementers from hold up and SEP holders from hold out. In other words both hold up and hold out are behaviours by one party, mischiefs, which the other party is to be protected from; and this underpins the FRAND system itself”.

[2025] EWCA Civ 552, paragraphs 118-119.

[2025] EWCA Civ 552, paragraphs 139-140

[2023] EWHC 1095 (Ch), paragraphs 431-432.

[2023] EWHC 539 (Pat) paragraphs 495-507, in particular, 498: “At best, the arguments can only justify relatively small volume discounts, where, for example the overall transactional costs (including the costs of litigation) are of the same order of magnitude as the quantified discount. As I have indicated above, this argument gets nowhere near justifying the sizes of discounts... [claimed in this case]”.

[2025] EWCA Civ 552, paragraphs 43 and 122.

[2025] EWCA Civ 552, paragraph 138.

[2025] EWCA Civ 552, paragraphs 116-123.

[2025] EWCA Civ 552, paragraph 122 and [2023] EWHC 1095 (Ch), 360.

[2017] EWHC 711 (Pat), paragraph 543.

[2023] EWHC 1095 (Ch), paragraphs 201, 213, and 212-220.

[2025] EWCA Civ 552, paragraphs 31-33.

[2025] EWCA Civ 552, paragraph 122.

[2025] EWCA Civ 552, paragraphs 128 and 143.

This is an approximate parallel to the situation the judgment describes for Apple. [2025] EWCA Civ 552, paragraph 57.

52

<https://www.compasslexecon.com/insights/publications/the-analysis-validating-that-royalties-inferred-from-comparable-sep-licences-are-frand> For a full description see: Hietamäki, Eeva and Sääskilahti, Pekka and Tuffin, Andrew and Väänänen, Lotta, Evaluating Royalties Against the Value of the Benefits that Depend on the Licensed Technology (June 04, 2025). Available at SSRN: <https://ssrn.com/abstract=5289146> or <http://dx.doi.org/10.2139/ssrn.5289146>