What (Exactly) Are Patents Worth at Trial?: The "Smartphone War" Example

Jonathan D. Putnam*

Charles River Associates

April 6, 2012

* Vice President, Charles River Associates, Boston, MA. Presented at the Spring meetings of the American Intellectual Property Lawyers Association, Austin, TX, May 11, 2012. This paper is derived from "Patent Portfolios, Apportionment and the Adding-Up Constraint" (Putnam 2012). I have benefited from helpful conversations with Peter Boberg, Peter Schwechheimer, Andrew Tepperman and Craig Tyler. Any errors are mine.

Introduction

With the demise of the 25 Percent Rule in *Uniloc*,[1] the Federal Circuit has removed from patent damages experts a widespread rule of thumb for allocating the profits of an infringing product between the patentee and the infringer in a hypothetical negotiation between them. For better or worse, the Rule influenced both actual awards in litigation, and licenses negotiated in anticipation of those awards, for decades. Apart from the principles for determining an "established royalty" (which predates even the *Georgia-Pacific* framework), there perhaps has been no more consistent fixture of reasonable royalty jurisprudence during the past half-century.

However unjustified its use – and the criticisms have been broad and longstanding[2] – the 25 Percent Rule filled an important conceptual niche in the *Georgia-Pacific* framework. Under *Georgia-Pacific* factor 13, the trier of fact is directed to consider "the portion of the realizable profit that should be credited to the invention." By giving an operational definition to "profit" (the Rule's author claimed that "operating profit" was the correct concept) and by establishing a "starting point" for negotiations over the "portion" – 25 percent[3] – the Rule gave experts a simple, easily implemented means of claiming some share of the gains created by the infringer's product.

The problem, of course, is that there is no reason to think that 25 percent is the right starting point, given economic conditions that vary widely across different industries,

---

[1] *Uniloc USA, Inc. v. Microsoft Corp.*, 632 F.3d 1292, 1315 (Fed. Cir. 2011).

[2] *See e.g.* R. Epstein, "The '25% Rule' For Patent Infringement Damages After *Uniloc*," 2012 *Duke L. & Tech. Rev.* 001 (Jan. 29, 2012). Available at: http://www.law.duke.edu/journals/dltr/articles/2012dltr001.

[3] In later iterations, the "Rule" was changed to the "25 to 33 Percent Rule."

technologies and other relevant circumstances.[4] Perhaps the most important source of variation – again, frequently cited – is that some complex inventions embody dozens, hundreds or even thousands of inventions, while others may embody only one. One highly simplified form of this critique pointed out that, if profit must be divided among five or more inventions, they could not all contribute 25% of profit, because the profit shares would then add to more than 100% of profit. More generally, as the number of pieces of intellectual property (or other inputs) among which profit must be divided increases, the less likely that any one piece accounts for as much as 25% of the whole.

Absent the 25 Percent Rule, however, there is no generally accepted method of allocating profit between the patentee's asserted patent and the accused infringer's patents.

This conceptual vacuum created by the Federal Circuit has coincided with the rise of the "smartphone wars." The longstanding "patent peace" among the major telecom handset vendors has broken down, having been replaced by large-scale global litigation over both the infringement of individual patents, and the licensing of entire patent portfolios. Even more intriguingly, the defenses to these claims have expanded to include antitrust and other competition-related allegations that patentees have failed to satisfy standards-related obligations to license subsets of their portfolios – certain "essential" patents – on fair, reasonable and non-discriminatory (FRAND) terms. But it is hard to identify these patents objectively, and even if they could be identified, determining the portion of the total value of the portfolio attributable to the "essential" members – or even whether such apportionment is required – is sharply disputed.

[4] Some defenders of the Rule responded that, because 25% is only a starting point, the Rule properly applied was flexible enough to reach a reasonable result in diverse circumstances. Of course, given sufficient "flexibility," *any* starting point could be used, so this justification only underscores the arbitrary nature of the starting point.

It is no accident that with the rise of the "smartphone wars" has come the systematic pursuit of large-scale patent portfolios.[5] This shift in demand, and the actual and prospective transactions it has produced, have led some observers to conclude that there is a "patent bubble," by which is meant a temporary overvaluation of patent rights that is not grounded in market fundamentals. But while the identification of "bubbles" first requires that one specify those fundamental valuations, there is no agreed-upon method of relating patent value to profits.

A particularly prominent example of the need for reliable apportionment methods has arisen in ongoing patent and copyright litigation between Oracle and Google over Google's alleged use of Java-related technology in its Android mobile handset operating system. Following an initial claim of more than $6 billion in damages, Oracle has been forced repeatedly to trim its demands, through a series of pre-trial damages motions and hearings that are virtually unprecedented in scope and mathematical sophistication. The final pre-trial skirmish prominently featured arguments over techniques by which to determine the share of value accounted for by the top 3.9% of patents in a portfolio of more than 500 patents.[6]

These events highlight the need for better methods by which to determine the share of total value – of either a portfolio, or a profit stream – that can be reliably attributed to a single patent. Any such method should have several properties:

[5] Nortel's creditors recently sold Nortel's portfolio of 4,000 US patents to a consortium (including Apple, Microsoft and Research in Motion) for $4.5 billion, or about $1.1 million per patent. Other benchmarks include Google's recent acquisition of Motorola Mobility and its 17,000 US patents for $12.5 billion (about $735,000 per patent), and its rumored acquisition of InterDigital and its 1,500 US patents ($2 million per patent) at the time of the rumor. Because these latter two transactions include all the assets of the target company, as well as equivalent foreign patents, they overstate the value of the US patent portfolio.

[6] http://www.groklaw.net/article.php?story=20120307180750431; see also ORDER GRANTING IN PART AND DENYING IN PART GOOGLE'S *DAUBERT* MOTION TO EXCLUDE DR. COCKBURN'S THIRD REPORT, March 13, 2012, available at: http://www.groklaw.net/article.php?story=20120313221158509.

- It should satisfy the Federal Circuit's objections to the 25 Percent Rule; in particular, it should be "tied to the facts of the case."
- It should satisfy the indicia of scientific testimony set forth in *Daubert* and related cases: peer review, known error rate, etc.
- It should produce valuations that are self-consistent, for example, by ensuring that the individual patents thus valued add up to the value of the entire portfolio or (relevant portion of the) profit stream.

This paper describes one method that meets these desirable criteria. In addition, it has been accepted as reliable by trial courts and employed as the basis for actual damages awards. It also has the virtue of ruling out damages claims that are inconsistent with these criteria.

The remainder of this paper explains the method and shows how to apply it. The technical details, and the data on which the method is based, are found in a related paper (Putnam 2012).[7]

Legal context and economic assumptions

U.S. patent law has long required that a patentee whose patent has been infringed bears the burden of proving the portion of the infringer's profit that his invention has caused[8] and, if that portion is less than the whole, to apportion the infringer's profit between that which is legally attributable to the use of the patentee's invention and that which is attributable to the

[7] "Patent Portfolios, Apportionment, and the Adding-Up Constraint," unpublished manuscript.
[8] 35 U.S.C. §284 awards damages "adequate to compensate for the infringement, but in no event less than a reasonable royalty for the use made of the invention." Because damages are compensatory, disgorgement of the infringer's profit is not a remedy available to the patent owner (as it may be for, say, copyright or trade secret owners), so the infringer's profit (or a portion thereof) is not a direct measure of damages.

infringer's own inputs.[9] Among the *Georgia-Pacific* factors, factor 13 requires the trial court to examine "the portion of the realizable profit that should be credited to the invention, as distinguished from non-patented elements, the manufacturing process, business risks, or significant features or improvements added by the infringer."

In *Uniloc,* the Federal Circuit found that Uniloc's expert applied the 25 Percent Rule without justifying its use as a starting point, or linking it to the portion of profit caused by the invention, or to the portion of profit attributable to other elements added by the infringer.[10] Other than citing to *Daubert,*[11] however, the Federal Circuit did not affirmatively state the attributes of an acceptable apportionment method.

In short form, then, the objective of the present paper is to describe a method of implementing factor 13 that satisfies the concerns the Federal Circuit expressed in *Lucent* and *Uniloc,* and that is supportable under *Daubert*. For these purposes, I assume the following:

---

[9] See *Lucent v. Microsoft*, 580 F.3d 1301, tracing the history of profit apportionment to *Garretson v. Clark* (1884), which held:

> When a patent is for an improvement, and not for an entirely new machine or contrivance, the patentee must show in what particulars his improvement has added to the usefulness of the machine or contrivance. He must separate its results distinctly from those of the other parts, so that the benefits derived from it may be distinctly seen and appreciated. . . . The patentee . . . must in every case give evidence tending to separate or apportion the defendant's profits and the patentee's damages between the patented feature and the unpatented features, and such evidence must be reliable and tangible, and not conjectural or speculative; or he must show, by equally reliable and satisfactory evidence, that the profits and damages are to be calculated on the whole machine, for the reason that the entire value of the whole machine, as a marketable article, is properly and legally attributable to the patented feature.

[10] "In short, [the expert's] starting point of a 25 percent royalty had no relation to the facts of the case, and as such, was arbitrary, unreliable, and irrelevant. The use of such a rule fails to pass muster under *Daubert* and taints the jury's damages calculation." *Uniloc USA, Inc. v. Microsoft Corp*., 632 F.3d 1292, 1315 (Fed. Cir. 2011) (slip op. at 47).

[11] *Daubert* and its progeny look to certain indicia of reliability, including:

- Empirical testing: is the theory or technique falsifiable, refutable, and testable.
- Whether or not the theory been subjected to peer review and publication.
- The existence of a known or potential error rate.
- The existence and maintenance of standards and controls concerning the method's operation.
- The degree to which the theory and technique is generally accepted by a relevant scientific community.

Daubert v. Merrell Dow Pharmaceuticals, 509 U.S. 579 (1993).

- Actual (or forecast) profit is the correct measure of "realizable" profit
- The profit to be apportioned is defined properly and undisputed
- The allocation of profit to non-patent causes (or at least a lower bound on that allocation) has been determined by a reliable method
- Within the portion of profit that is attributable to patented technology ("patent profit"), the number of patents (contributed both by the infringer and by others) is known
- Among these contributions to patent profit, the patents can be ranked in order by the size of their contributions
- The value of any two patents, taken together, is equal to the sum of their individual values
- "The parts sum to the whole" – or put differently, every patent's individual share of patent profit adds to 100%

These assumptions are plausible, or at least close approximations to the truth, in many situations involving large numbers of patents. For example, it is often relatively straightforward to identify the number of patents that contribute to the profit of a given device, either through company records or technical expert testimony.[12] Perhaps more importantly, they can be justified on a variety of economic and legal grounds. But their justification is beyond the scope of this short paper, so for the purposes of this paper I simply take them as having been established.

[12] Note that, because a patent conveys the right to exclude others, it may contribute to a product's profit even though the product does not embody the claims of the patent.

The “count, rank and divide” method

The method described below can be analogized to carving a turkey at Thanksgiving, or dividing a pie: the number of guests is known, and the relative sizes of their appetites are known. The objective is to allocate portions in such a way that every guest receives a portion commensurate with his or her appetite, no turkey is left over, and the allocation rule used is consistent with the scientific literature concerning the rate at which guests consume turkey. In other words, the procedure is to *count* the guests, *rank* them by their appetites, and *divide* the turkey among them.

Under the maintained assumptions, the number of “relevant” patents is known, and their rank is known. Thus, the first two steps of this method are taken as given. These two steps are sufficient to construct what is known as a “Lorenz graph.”[13] A Lorenz graph plots the percentage of portfolio value (the vertical axis) accounted for by the bottom *X* percent of patents. An example Lorenz graph is shown in Figure 1.

For example, in the sample graph labeled “0.5,” the bottom 50% of all patents account for about 30% of the portfolio’s value, while the top 50% of patents account for the remaining 70% of value. In the graph labeled “1.0,” the bottom 50% of patents account for only about 15% of the portfolio’s value. Colloquially speaking, the relationship between the bottom *X* patents, and the share of total value they account for, is known as the “value distribution.”

We can characterize value distributions in general, and Lorenz graphs in particular, by how “unequal” their allocations are. If every patent contributed the same amount to total profits, then their shares would be equal. In that case, the bottom 10% of patents would account for 10%

[13] Lorenz graphs are used to study a wide variety of outcomes in which the nature of the distribution’s inequality is relevant. For example, the distribution of income is often analyzed and compared using Lorenz graphs.

of portfolio value, the bottom 20% for 20% of value, etc. The graph that corresponds to complete equality in the distribution runs in a straight line from the southwest to the northeast corners of Figure 1 (sometimes called the “45-degree line”). As the distribution becomes increasingly unequal, the Lorenz graph describing that distribution shifts further and further away from the 45 degree line towards the southeast. Thus, the graph labeled “3.0” describes a distribution in which the bottom 90% of all patents account for only 5% of portfolio value, while the top 10% account for 95% of portfolio value – a very unequal distribution.

These characteristics are helpful for understanding the general properties of a Lorenz graph. In actual litigation, the damages issues typically concern one or a handful of patents; in particular, how to *divide* portfolio value between the asserted patents and the remainder. Lorenz graphs are useful for this purpose at well. It can be shown that the ratio of the value of any given patent to the value of the average patent in the portfolio is equal to the slope of a Lorenz graph at the point corresponding to that patent. I call this ratio K. When K refers to a particular patent ranked in the n^{th} percentile, I refer to it as K_n. The bottom line of the paper is this: if you know (*a*) the correct Lorenz graph, (*b*) the percentile ranking n of the patent in question, and (*c*) the value of the average patent (or, equivalently, the total portfolio value and the number of patents in it), then you know K_n, which (when multiplied by the average value) gives you the expected value of the n^{th} patent.

Example: Suppose that a product embodies 1,000 patents and generates $1 billion in (patent) profit. Then, on average, each patent contributes $1 million to profit. Most patents contribute relatively little: for example, for the median (50^{th} percentile) patent, the ratio is $K_n = 0.2$, which means that this patent is expected to account for only about 0.2 x $1,000,000 = $200,000 of the $1 billion total. And by definition, half the portfolio’s patents are worth even less than that. But as you move to the right of Figure 1, this ratio increases, at an increasing rate. At the 90th percentile, a patent's expected contribution is about $2.1 million; at the 95th percentile, about $3.7 million, and at the 99th percentile, about $12 million. And, if you add up all of these values, you get exactly $1 billion. This is the "adding-up constraint." Expressed in shares, the values are 0.21%, 0.37% and 1.2%, respectively, which must, of course, add up to 100%.

Figure 2 shows these results, and the intuition behind them, graphically. As a baseline, consider the 45-degree line. By definition, its slope must be one: each 10 percentage point increase in the share of patents (the horizontal axis) is accompanied by a 10 percentage point increase in the share of portfolio value (the vertical axis). A slope of one implies that every patent's value has a ratio of one to the value of the average patent, *i.e.*, every patent is equal to the average patent in value.

In contrast, the slope of a Lorenz graph varies markedly, from nearly flat for the lowest-value patents to nearly vertical for the highest-value patents. At the 90th percentile (*i.e.*, the patent ranked 101st out of 1,000), the slope is about 2.1, so a 90th-percentile patent is worth a little more than twice the average.

Figure 2 also illustrates other calculations that may prove useful in litigation. For example, suppose one knew that a given patent ranked in the top 5% of patents, but the precise ranking within that group was unavailable.[14] Figure 2 shows that the average value of patents in

[14] A determination analogous to this was made by Judge Alsup after a variety of pre-trial submissions by the parties in the ongoing Oracle-Google matter. Judge Alsup ruled that Oracle's expert could testify that some of Oracle's patents ranked in the top 22 out of 569 relevant patents (*i.e.*, in the top 3.9%), but that the evidence did not support a more precise ranking. *See* ORDER GRANTING IN PART AND DENYING IN PART GOOGLE'S *DAUBERT* MOTION TO EXCLUDE DR. COCKBURN'S THIRD REPORT, March 13, 2012, available at: http://www.groklaw.net/article.php?story=20120313221158509.

this group is about 10 times the value of the average patent overall or, in our example, about $10 million. I refer to the ratio of the mean value of the top 1 – *n* percent of patents to the mean patent overall as M_n; in this example, $M_{95} = 10.0$.

Similar calculations can be made for other intervals of the Lorenz graph. Figure 2 shows, for example, that the average value of a patent falling within the "inter-quartile range" (between the 25th and 75th percentiles) is about 0.3 times the overall mean, or about $300,000.

<u>Identification and robustness</u>

The basic relationship linking the slope of a Lorenz graph to the mean is valid regardless of which particular graph is chosen. But, as Figure 1 shows, the slope varies depending on the shape of the particular graph. So the first order of business is to identify the best Lorenz graph.

For the past 25 years, economists have investigated patent value distributions, using a wide variety of techniques, models and data. Some of these studies are surveyed in "How to Count Patents and Value Intellectual Property,"[15] which describes some of this variety. I have since updated that survey to include more recent studies. I have also shown that, for present purposes, all of their results regarding the distribution of patent values can be boiled down to a single parameter,[16] "sigma," which is one of the two parameters of a log-normal distribution.[17] When sigma equals zero, all patents are equal. As sigma increases, the distribution becomes increasingly unequal. Figure 1 plots Lorenz graphs for values of sigma ranging from 0 to 3.

[15] J. Lanjouw, A. Pakes and J. Putnam, "How to Count Patents and Value Intellectual Property: Uses of Patent Renewal and Application Data," *Journal of Industrial Economics 46*(4), December 1998: 405-32.
[16] A parameter is a number that governs the relationship between any particular outcome and the likelihood of that outcome, for a given distribution. For example, a coin flip is an example of a binomial distribution, which is governed by the parameter *p* (the probability of heads). By definition, if a coin is fair, $p = 0.5$.
[17] A normal (Gaussian) distribution is governed by two parameters, mu (the mean of the distribution) and sigma (the standard deviation of the distribution). In a log-normal distribution (*i.e.*, the logarithms of the distribution's values follow a normal distribution), sigma is sometimes called the "scale parameter."

The value of sigma varies somewhat from study to study. Most of the studies report a value for sigma in the range of 1.5 – 2.0. In Figure 1, the heavier line plots the Lorenz graph generated by the median value for sigma across all of the studies surveyed in Putnam (2012) (about 1.67 for the upper tail). This is the graph used to generate the values reported in the text and in Figure 2.

Having determined a value for sigma, one should then ask a couple of follow-up questions: Should you choose different sigmas under different circumstances (depending on, say, the country or the technology)? And, how sensitive are the results to the choice of sigma?

Choice of sigma. The economic literature does not indicate that sigma varies systematically by industry, country, technology field or other observable variables that might be thought to influence patent values. As summarized in Putnam (2012), the variation that is observed appears to be related to different modeling techniques and estimation methods, rather than to observable invention characteristics. Thus, while the estimates produced by "count, rank and divide" method do vary with the choice of sigma, there appears to be no reason, at present, to select a sigma based on the industry, technology or country of the litigated patent. But, even if one thought these characteristics ought to be taken into account, for some reason,[18] this choice turns out not to matter very much for most purposes, as I explain below.

[18] Recently, the Federal Circuit has emphasized the importance of "comparability" in the selection of license agreements used to produce a reasonable royalty analysis, as a means of reducing abuses by damages experts who tried to import the terms of unrelated agreements into the hypothetical negotiations. Among characteristics thought to render a license "comparable" is the similarity of the technology. As an economic matter, similarity of technology is neither necessary nor sufficient to establish that one license is comparable to another – what matters, as in all determinations of price, are the factors underlying supply and demand for the patented technology, which may be different for similar technologies and similar for different technologies. In any event, a credible damages analysis should prove rather than assume the reliability of the relationship between a candidate license (or other "comparable" data point) and the claimed damages figure.

To understand why the choice of sigma is likely to be unrelated to technology, geography, etc., it is important to specify the factors that influence a patent's value, and how they do so. Most of the factors that are known to influence a patent's value do so through the *size* of the pie to be divided, not through the distribution of *shares*. For example, a U.S. patent is likely to be more valuable than an otherwise equivalent French patent because the U.S. economy (and so the level of expected sales) is larger than the French economy. A pharmaceutical compound patent is more valuable than a mechanical arts patent, other things equal, for many reasons related both to technology and to industry: the scope of its claims is easier to define, there may be fewer ways to invent around the claims, these claims and other regulation create more effective barriers to competitive entry, market demand is less elastic, consumers do not face market prices, etc. All of these factors imply that pharmaceutical profits, and therefore the size of the relevant pie, should be higher (again, other things equal).

But when comparing the individual slices of a pie, it is critically important to distinguish differences in the overall size of the pie from differences in shares within pies of a given size. In other words, 10% of a larger pie always produces a bigger piece than 10% of a smaller pie, though the share is the same. Thus, if one knows the distribution rule used to allocate shares in Pie *A*, and one has reason to believe that that allocation rule is similar for Pie *B*, it is entirely irrelevant that *A* and *B* have different *sizes*, if one's only objective is to compare the difference in *shares*. If the allocation rule is the same, the shares are the same also, regardless of size.

Sensitivity to the choice of sigma. For whatever reason, sigma does vary, so it is important to determine the sensitivity of the overall results to the choice of sigma. As it turns out, in the vast majority of cases, value shares do not change much, even when sigma changes.

One can see this stability graphically in Figure 1, by comparing (for example) the slope of the Lorenz graph for "sigma = 1.5" to the slope of the graph for "sigma = 2.0" at any given point. Over most of the range (up to about the 93rd percentile), the "sigma = 1.5" graph is slightly steeper at any given point. Thus, for most patents, one would obtain slightly higher estimates of value if one chose sigma = 1.5 than sigma = 2.0. The intuitive reason is that, because the sigma = 2.0 distribution is more unequal than the sigma = 1.5 distribution, its low-value patents must have smaller shares.

For the same reason, in the upper right tail where the most valuable patents are located, the sigma = 2.0 graph is steeper, which means that patents in this part of the distribution have larger shares than they would using the sigma = 1.5 distribution.

In general, the choice of sigma matters most for the highest-value patents. But even here, the variability is relatively small (at least relative to typical differences of opinion between damages experts). For example, assuming sigma = 1.5, a 99th percentile patent is worth about 10.6 times the mean, or about $10.6 million in our example. For the median sigma = 1.67, the figure is $11.9 million, and for sigma = 2.0, the expected value is about $14.2 million. More formally speaking, we say that the expected shares are robust to the choice of sigma.

<u>The "known error rate"</u>

As with the calculation of other expected values, it is important to understand the potential variability of the results in a sample of any given size. For example, when 10 fair coins are tossed, the expected number of heads is 5, and the probability of obtaining all heads is about 1/1000, or 10^{-3}. When 100 fair coins are tossed, the expected number of heads is 50, but the probability of obtaining all heads is less than 10^{-31}. In words, as the sample size increases, the variability of the results relative to the expectation decreases.

This general proposition is illustrated in Figure 3, which shows the 95% confidence interval for the share attributable to the 99th percentile patent, in samples of different sizes. In very small samples, this range is large, and even in a portfolio of 100 patents, the share of the 99th percentile (2nd-ranked) patent can range from about 5% to about 20% of the total. But as the portfolio size increases to 1,000, the confidence interval around the expected share is small: from 1% to 2%.

Ruling out inconsistent claims

By completely specifying the distribution of patent values as a function of their rank, this model provides another salutary benefit: one can distinguish damages claims that are merely unlikely from those that are logically inconsistent, and rule out the latter.

A logically inconsistent damages claim is one that claims a share of profit that is inconsistent with the asserted patent's rank among other patents. To be logically inconsistent is to claim that the sum of the individual parts is greater than the whole.

For example, in a portfolio of 100 patents, it is logically possible (though, according to Figure 3, very unlikely) that the 99th percentile (2nd-ranked) patent is worth 20 times the average patent. But this valuation is logically possible only if the patent in question is ranked correctly. To see this, suppose that the patent actually ranks 6th (*i.e.*, 95th percentile) rather than 2nd in the portfolio. Then each of the 5 patents ranked above it must also be worth at least 20 times the average patent. If each of these patents is worth 20 or more times the average, then their sum must be 100 or more times the value of the average patent. But, in a portfolio of 100 patents, the sum is (by definition) equal to exactly 100 times the average patent. The claims as to the patent's ranking (6th) and valuation (20x) are mutually inconsistent, and should be rejected.

Applications and experience

The Federal Circuit found that one of the essential shortcomings of the 25 Percent Rule was its failure to take account of the facts of a particular case. When it comes to apportionment among a group of patents, perhaps the most important of those facts is the rank of the asserted patent among its peers. The determination of that ranking is itself likely to be a fact-intensive inquiry. Thus, the present method – which assumes that this rank is known, or can at least be placed in a range – is not a "rule of thumb" divorced from case-specific facts, but rather a means of translating a fact-based ranking into a logically consistent valuation.

In a case that consolidated infringement claims brought by both parties, the "count, rank and divide" method was accepted by a trial judge who had tried more than 100 patent cases.[19] The eventual damages award used the result produced by this method.[20] Testimony based on this method has been accepted, over objection, elsewhere.[21]

A prominent case in which this method could have provided a helpful check on the parties' claims is *Lucent v. Microsoft.*[22] A jury awarded Lucent damages of about $358 million

---

[19] *LG Display v. AU Optronics et al. v. LG Display,* Civ. Ac. No. 06-726 (JJF) (D.Del. 2009).
[20] "… the Court finds [the] methodology to be credible and consistent with Federal Circuit case law and the *Georgia Pacific* factors, despite LGD's assertions to the contrary…. With AUO's aggregate claim against LGD assessed, Dr. Putnam then used a method described as "count, rank, and divide" to determine the portion of the claim attributable to the four asserted patents. This method takes into account *Georgia Pacific* factors 9-11. Based on the value share of each patent in AUO's portfolio and based on the assumption that these patents are in the top 5% of AUO's portfolio, Dr. Putnam determined that AUO's damages for infringement of all four patents would total $305,399 …" *See* MEMORANDUM OPINION, JULY 8, 2010, available at: http://www.scribd.com/doc/34218584/LG-Display-CO-Ltd-V-AU-Optronics-Corporation-et-al-C-A-No-06-726-JJF-D-Del-July-8-2010 .
[21] *Energy Transportation Group, Inc. v. Sonic Innovations, Inc. et al.,* Civil Action No. 05-CV-422 (GMS) (D.Del. 2008).
[22] *Lucent Techs., Inc. v. Gateway et al.,* 580 F.3d 1301 (Fed. Cir. 2009).

for Microsoft's infringement of a patent on a drop-down menu for selecting a date (the so-called "date-picker" feature), used in Microsoft's Outlook personal organizer software. Lucent's expert testified that damages of about $562 million (out of $8 billion in sales) were appropriate; Microsoft's expert testified that damages should be $6.5 million, or about 1% of Lucent's demand.[23]

On appeal, the Federal Circuit found that, "The evidence can support only a finding that the infringing feature contained in Microsoft Outlook is but a tiny feature of one part of a much larger software program."[24] While this observation might seem to be nearly tautological—and therefore to imply a necessarily small share of Microsoft's profit—it is routine to see opposing patent damages experts differ by two or more orders of magnitude in their assessment of the share of profit attributable to an individual patent, particularly when the invention represents one input into a complex device generating sales in the billions of dollars. Even the qualitative skepticism expressed by the court as to the price Lucent claimed still leaves room for substantial quantitative disagreement, perhaps amounting to hundreds of millions of dollars over a large royalty base (like Windows), as occurred in *Uniloc*.

To assess the plaintiff's claim in *Lucent*, suppose that one assumes (charitably, it appears) that the date-picker invention corresponded to a 95th-percentile patent in a portfolio of 500

[23] Slip op., p. 31.
[24] *Id.*, p. 48. The court went on to observe:

> We find it inconceivable to conclude, based on the present record, that the use of one small feature, the date-picker, constitutes a substantial portion of the value of Outlook….
>
> . . . the only reasonable conclusion is that most of the realizable profit must be credited to non-patented [by Lucent] elements, such as "the manufacturing process, business risks, or significant features or improvements added by [Microsoft]."

Id., p. 49.

patents.[25] As I have explained, such a patent is expected to be worth about K_{95} = 3.7 times the average patent. Thus, to be internally consistent, the damages award implies that the average Microsoft patent must have been worth \$358 million / 3.7 = \$96.8 million. Since there were, by hypothesis, 500 such patents, the total profit to be apportioned must have been at least 500 × \$96.8 million = \$48.4 billion.[26] But the accused *sales* amounted to only about \$8 billion. Something doesn't "add up": either the patent's percentile ranking is too low, or the jury's valuation was too high, because the combination of the two violates the adding-up constraint.

To give some idea of the correction required: if one arbitrarily places Microsoft's Outlook operating margin at 40%, one reaches a mean profit per patent of \$8 billion × 40% / 500 patents = \$6.4 million (still probably an order of magnitude too high).[27] Thus, even under implausibly conservative assumptions, the jury's \$358 million award was more than 50 times the average, a valuation reserved only for the most extreme realizations from the distribution. A better guess would have been K_{95} = 3.7 times this generous mean, or \$23.7 million—less than 1/15 the jury's award but still greater than the Federal Circuit's finding of a "tiny" contribution.[28]

---

[25] According to the US Patent and Trademark Office, Microsoft has received more than 18,000 US patents. I assume, for the purposes of illustration, that 500 of these patents are embodied in, or otherwise advantage, Microsoft Outlook.

[26] This figure does not include whatever profit might have been attributable to Microsoft's software copyrights, trade name, network externalities, etc.

[27] There are at least three reasons why this figure is almost surely overstated:

- As I explained previously, even if one employed recent large portfolio transactions – said to evidence a "patent bubble" – as a benchmark, the mean value per invention falls in the range of \$0.7 - \$2 million, a figure that includes patent rights outside the U.S. and non-patent assets.
- This figure covers only the damages period, not the entire life of the patents.
- Again, this calculation implausibly assumes that all Microsoft's profit is attributable to patents.

Note that if one assumes that the relevant Microsoft portfolio comprises fewer than 500 patents, the implied value per patent is higher than the already unlikely \$6.4 million.

[28] In a retrial, the court repeatedly rejected various apportionment methods that Lucent proposed, through "three rounds of motions in limine." Eventually, based on a "Business Realities approach" (essentially, a 50/50 Nash division of the supposed bargaining surplus), Lucent requested and received \$70 million. Citing what it considered Lucent's remaining apportionment failures, Microsoft filed a post-trial motion to reduce the award to \$5 million. The trial court agreed in part, reducing the award to \$26.3 million. ORDER GRANTING IN PART AND DENYING IN

Conclusion

Much has been written over the past decade about outsize patent damages awards, particularly the amounts demanded and received by non-practicing entities (NPEs) in litigation over complex products. Commentators have proposed various solutions to "reform" the patent system and the litigation process. But perhaps the simplest and, in the long run, most economically consistent reform is to ensure that each piece of technology – both the patentee's, and the infringer's – that contributes to a product's profit should receive a rent consistent with that contribution. The present method (while not a cure-all or a replacement for a complete economic analysis) does – unlike the 25 Percent Rule – provide a rigorous means of implementing this requirement, as specified in *Georgia-Pacific* factor 13.

As always, there remains significant work to be done. It is no small task, either as a matter of law or as a matter of economics, to define – let alone measure – the multiple causes of a firm's profit. This method represents one step in that direction.

PART MICROSOFT'S MOTION FOR JUDGMENT AS A MATTER OF LAW AND IN THE ALTERNATIVE, A NEW TRIAL WITH REMITTITUR. November 10, 2011.

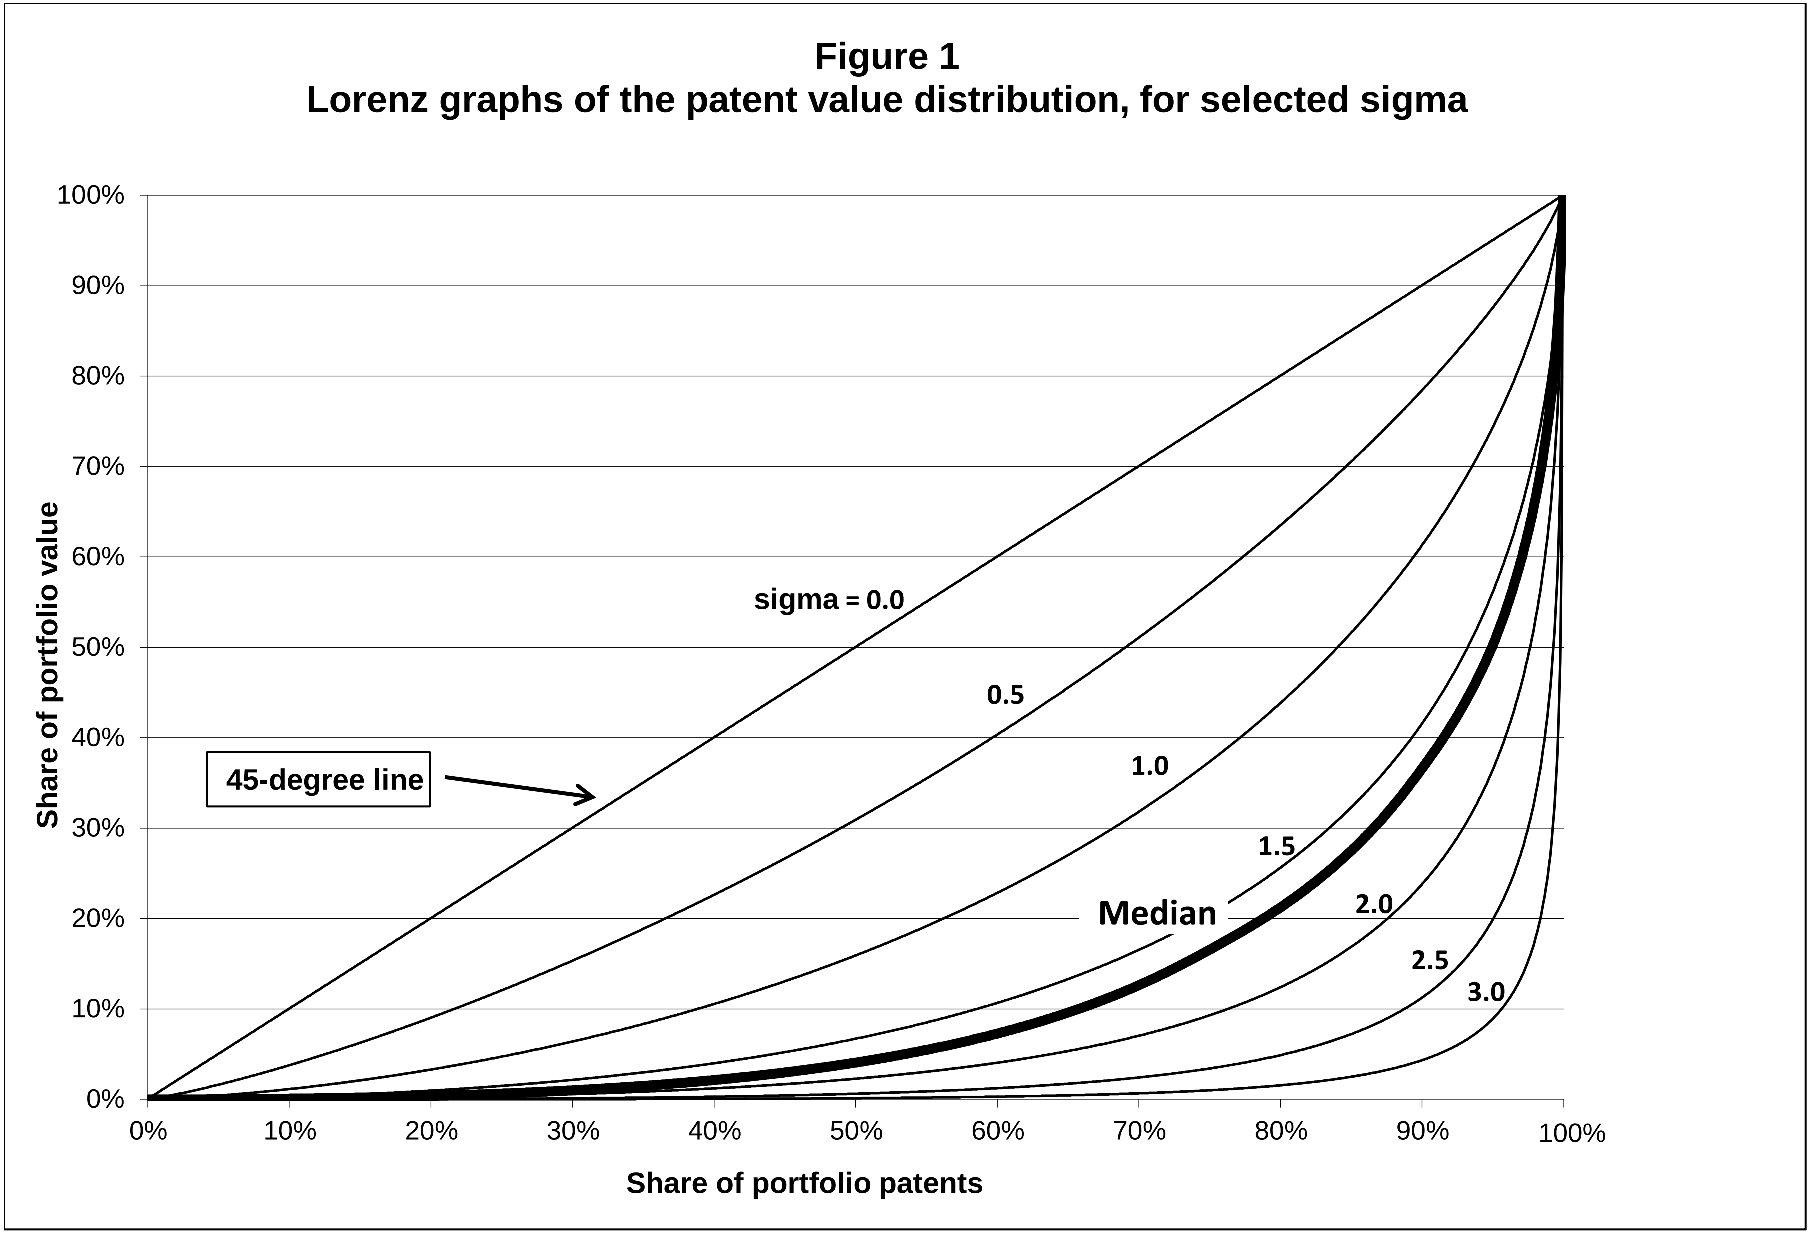
Figure 1
Lorenz graphs of the patent value distribution, for selected sigma
100%
90%
80%
70%
60%
50%
40%
30%
20%
10%
0%
Share of portfolio value
sigma = 0.0
0.5
1.0
1.5
Median
2.0
2.5
3.0
45-degree line
0%
10%
20%
30%
40%
50%
60%
70%
80%
90%
100%
Share of portfolio patents

Figure 2
Selected statistics for the median patent value Lorenz graph

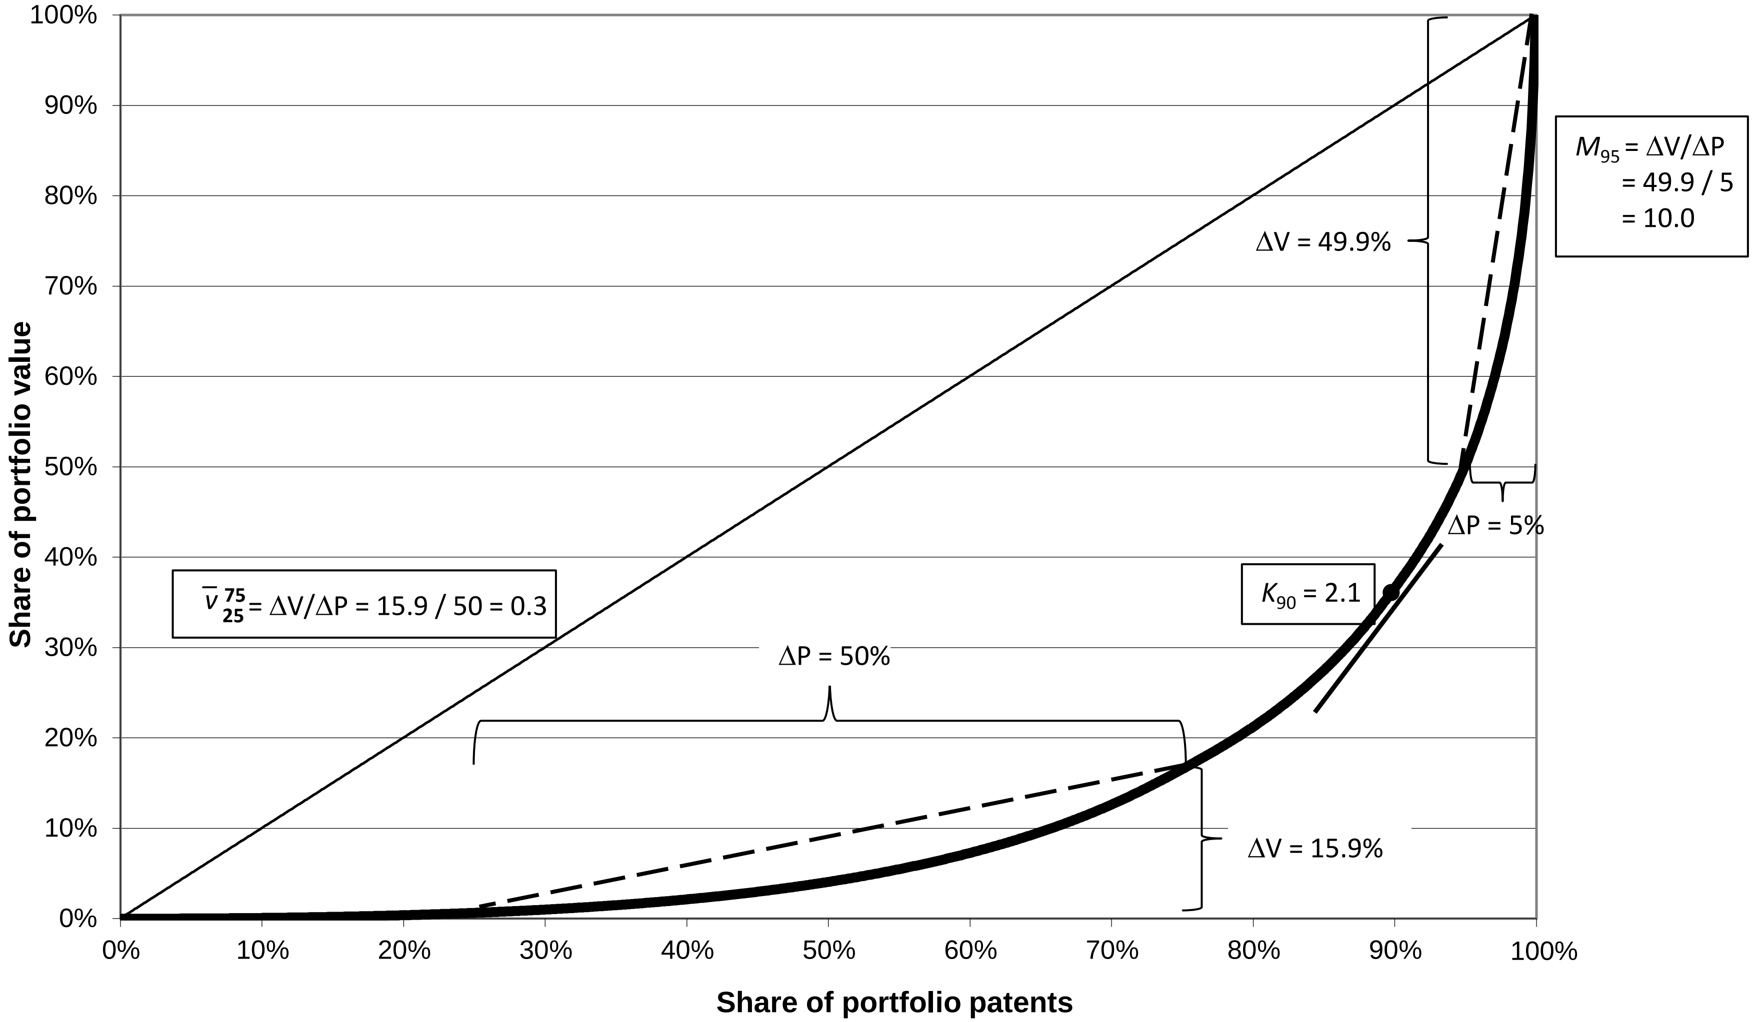

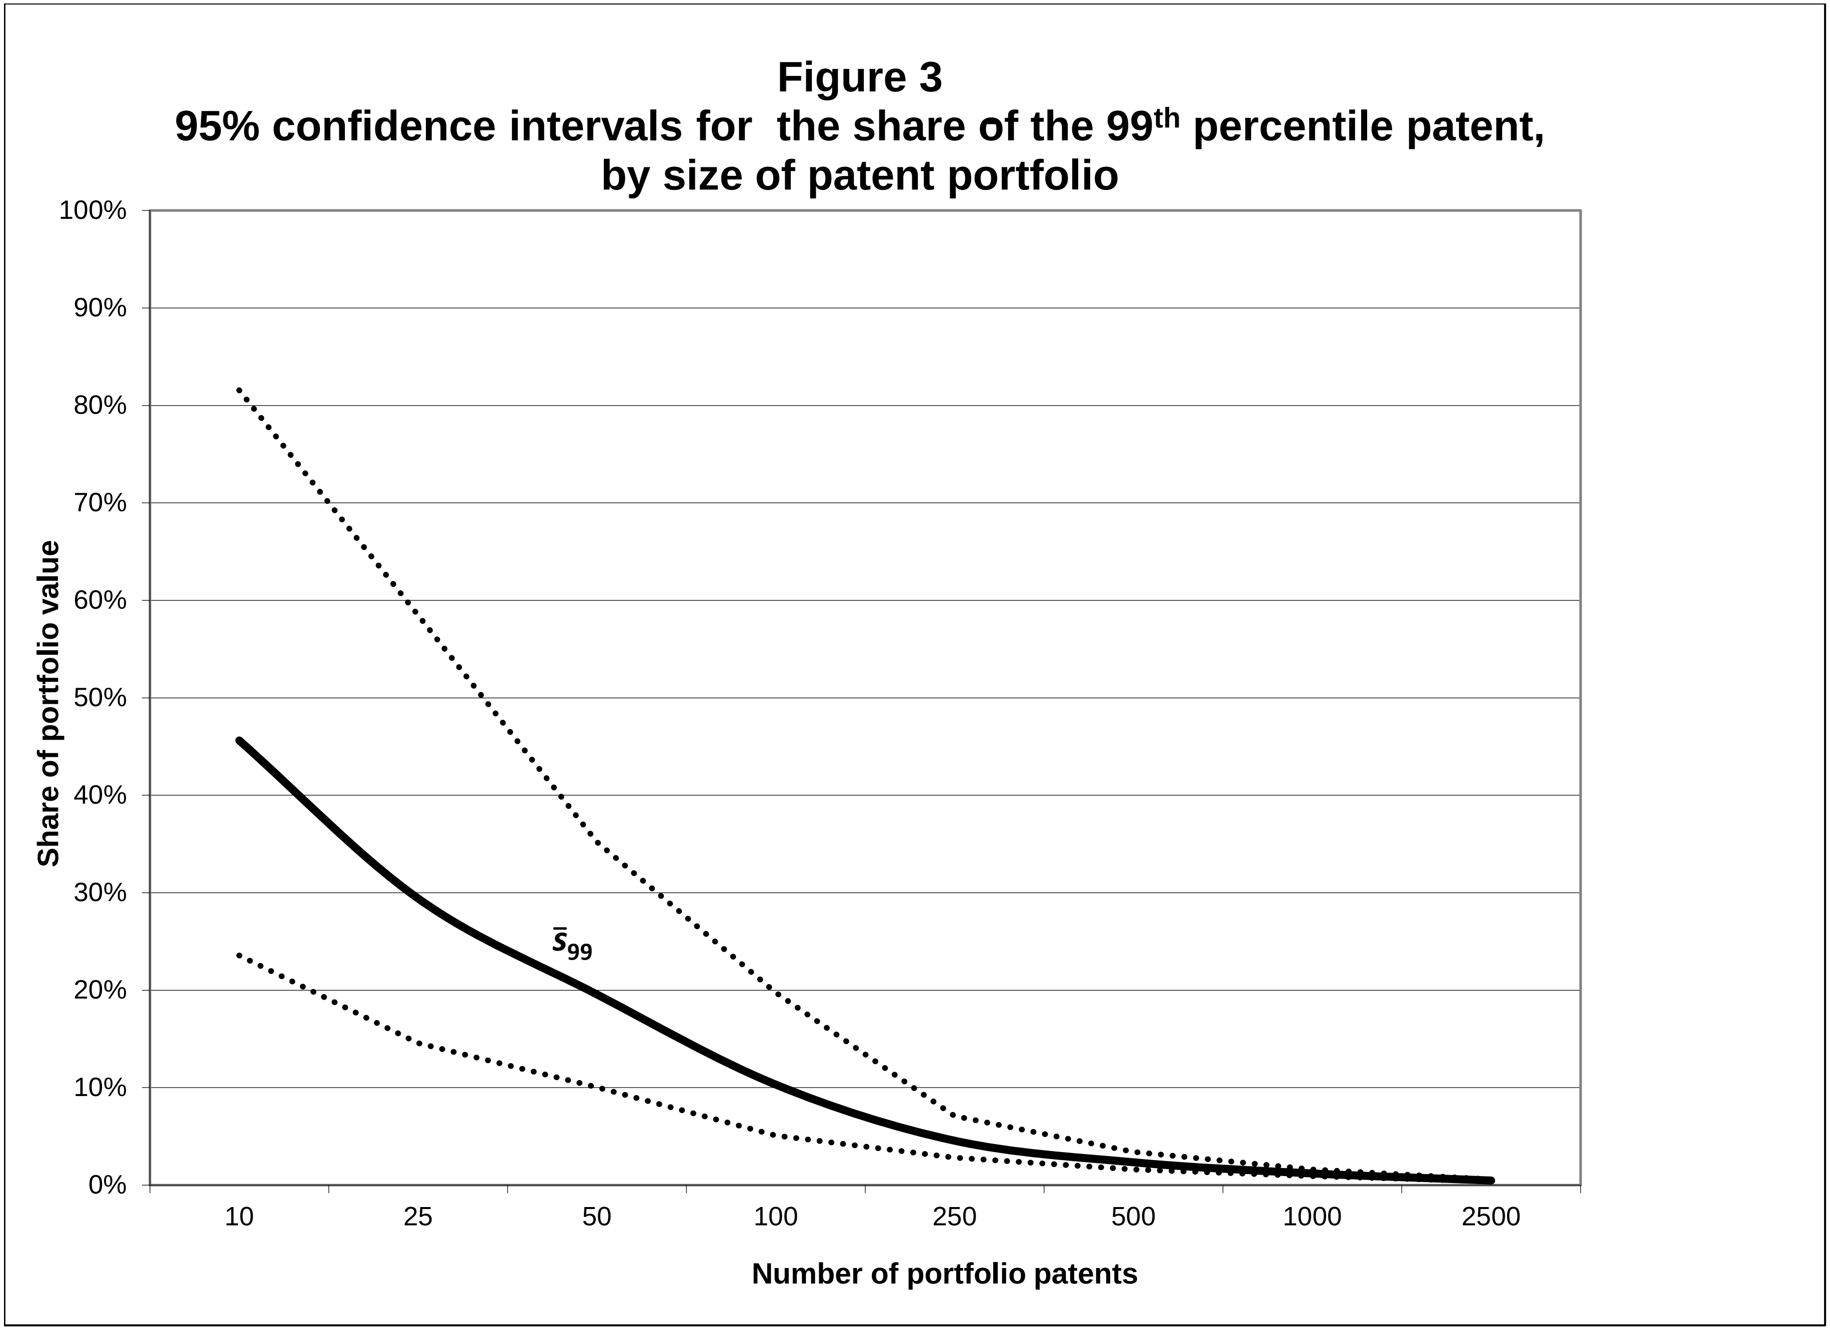
Figure 3
95% confidence intervals for the share of the 99th percentile patent, by size of patent portfolio
Share of portfolio value
100%
90%
80%
70%
60%
50%
40%
30%
20%
10%
0%
$\bar{s}_{99}$
10
25
50
100
250
500
1000
2500
Number of portfolio patents