Patient Patents

Can certain types of patent litigation be beneficially delayed?

Doug Lichtman
Professor of Law
UCLA

There was a time when patent cases would routinely end with the relevant court granting an injunction against future infringement. Today, however, injunctions are increasingly rare, and courts instead are allowing adjudged infringers to continue to use patented technologies, subject to an obligation to pay a court-determined forward-looking royalty for any future use. A vast literature already exists thinking about this change; but that literature has missed one important implication: patent litigation can now be beneficially slowed. The intuition is simple. One reason why courts used to race to the finish was because they wanted to minimize the importance of their own damages calculations. Courts had no choice but to calculate damages for infringement that already occurred; but, the faster the case, the sooner the injunction, and thus the sooner that the litigants would be the ones negotiating about the future, setting prices and establishing terms. With injunctions now increasingly off the table, however, a court's attempted quantification remains important no matter when the case ends. Before the verdict, the court's influence is relevant under the banner of backward-looking patent damages. After the verdict, the court's influence is relevant in the form of court-determined forward-looking royalties. The end of the case thus no longer represents a significant reduction in the importance of the court's economic understandings. As a result, there is less of a reason to race through patent litigation, and thus a real opportunity for courts to slow down and in various ways increase the accuracy of their important work.

Patient Patents

Can certain types of patent litigation be beneficially delayed?

It is almost impossible for a judge or jury to accurately determine how much a given patent is worth. To run the math right, that determination would require a rich understanding of the technology at issue; of how that technology compares to available next-best alternatives; and of what all that might mean for real-world products, services, prices, profits, and market share. No surprise, judges and juries rarely have that rich understanding. And, while expert testimony can in theory speak to all of these issues, accurately evaluating complicated, adverse, hired-gun technical and economic analyses is itself a daunting and precarious task.

Patent courts have long recognized this challenge, of course, and their response for many years was simple: they issued injunctions. That is, at the end of a patent case, patent courts would routinely order the adjudged infringer to stop its unlawful behavior. Yes, the court would still have to quantify whatever infringement took place before the injunction issued. But, once the injunction was in place, at least, the obligation to further quantify patent value would fall entirely on the litigating parties. The adjudged infringer would have to figure out how much to offer the patent holder in exchange for the patent holder waiving the injunction and allowing the infringer to continue to use the patented technology. Likewise, the patent holder would have to determine whether to accept that amount of money or instead simply enforce the court's injunction. The court's economic imperfections would at that moment become irrelevant. Judges and juries could be happily ignorant about prices, profits, and technology details. As long as they could accurately answer the binary question of whether the accused product or service should be subject to an injunction, the parties

from that point forward would be forced to privately engage in the difficult task of quantify patent value.

Times, however, have changed. In 2006, an influential Supreme Court decision cast doubt on this conventional practice, rejecting what lower courts had understood to be a presumption in favor of injunctive relief and emphasizing instead that even a meritorious patent plaintiff might not deserve an injunction. The statistics shifted almost immediately. Whereas in 2005 district courts issued injunctions in literally every case where a patentee won on the merits and requested an injunction, in 2007 courts awarded injunctions to meritorious patent holders in only 70% of the cases, and by 2010 that rate had dropped again to 60%. Courts had finally recognized what scholars had known for some time: as imperfect as they inevitably are, court-determined forward-looking royalties are in certain instances better measures of patent value than are the private arrangements struck after an injunction is either credibly threatened or in fact put into place.

This change in patent remedies was driven by a changing understanding as to whether private parties or courts are better situated to quantify patent value. But the change has an unintentional implication: patent courts can now beneficially slow down. Why? Back when private negotiations were the preferred approach, delaying a court verdict by (say) one year was tantamount to taking a year when prices would have been set by the market and transforming that into a year when prices would instead have to be set by the court. Courts, after all, set damages for infringement that takes place prior to the final verdict, but private parties set prices once an injunction is in place. To whatever extent courts believed that prices set by private parties were more accurate than prices set by courts—and, remember, that was the dominant view—this was a

reason to hurry. Faster decisions meant less time governed by patent law's damages rules and more time governed by assumedly better private numbers.

In cases where injunctions are not going to be available regardless, however, this trade-off disappears. In those cases, delaying a court verdict by one year simply takes a year when prices would have been set by the court under the rubric of court-ordered forward-looking royalties and transforms that into a year where prices will be set by the court under the rubric of court-ordered backward-looking damages. The court's economic imperfections drive the numbers either way. Delay is suddenly not as costly as it might once have seemed.

The benefits of delay thus begin to loom large. For example, what better way to decide (under section 103 of the Patent Act) whether a given invention was "obvious to those skilled in the art" than to wait a few years and see if a sufficiently large number of skilled practitioners independently come up with the same invention? Similarly, would not court decisions on validity be more reliable if they could be delayed long enough for the Patent Office to first run its own re-examination of any disputed patent, for example under the new Inter Partes Review procedure? Moreover, fast litigations are systematically biased in favor of patent holders because a patent holder can prepare its case long before the complaint is filed, whereas accused infringers will often not even know about the patents at issue until after the complaint is served. Fast clocks exacerbate this difference in time to prepare; slower clocks would mitigate it.

None of this is to imply that delay has no costs, or that tremendously long delays are now suddenly desirable. Delay means longer periods of uncertainty for both patent holders and accused infringers. Delay also makes it more difficult for poorly capitalized patent holders to pursue even valid claims. Moreover, delay will certainly require that courts take more seriously the need to award interest to patent holders to

whom payments would now be even more overdue. That said, thanks to the change in the Supreme Court's remedies jurisprudence, these costs and benefits newly trade off in ways that favor delay in certain cases. In this Essay, I set out to make that case.

I begin in Part I by discussing more fully the difficulties that courts face when it comes to calibrating patent harms or otherwise quantifying patent value. As I hint in this Introduction, the truth is that courts are incapable of doing this work accurately, and the modern trend away from injunctive relief will unavoidably put more pressure on this significant and difficult-to-mitigate weakness.

In Part II, I turn to injunctive relief. The Supreme Court, the Federal Trade Commission, the Department of Justice, and the European Commission have all in recent years questioned whether private parties will accurately quantify patent value when negotiating under the threat of injunction. And they are right to worry. Injunctions often allow patent holders to extract payments that have very little to do with the merits of their patented technologies, and much more to do with the disruptions an injunction would cause to already-existing manufacturing facilities, already-signed purchase contracts, and other already-important relationships. In these situations, as bad as court-determined quantifications might be, private negotiations that take place under the threat of injunction will be predictably worse.

For Part III, I turn to the question of delay, starting with a discussion of the costs that might realistically be associated with delays of one to several years. As I note in the introduction, delay does have drawbacks. It means longer periods of uncertainty for both patent holders and accused infringers. It might make it more difficult for poorly capitalized patent holders to pursue even valid claims. And it will certainly require that courts take more seriously the need to award interest to patent holders to whom payments will now be even more overdue. These are real costs to be sure; however, in

certain settings, they might pale in comparison to the gains associated with a slower, more accurate patent process.

Part IV thus turns to those possible gains. As noted above, delay can help even the playing field between patent holders and accused infringers, making up for the fact that patent holders can plan their cases for years before filing, whereas accused infringers might not even be aware of the relevant patent until the patent holder's complaint is formally served. Delay also opens the door to a wide range of other substantive and procedural reforms. My point here is not that delay is warranted in every instance, nor that delay is a panacea for the patent system's ills. My point instead is that delay can meaningfully increase accuracy in certain settings.

In Part V, I focus on a specific example around which all of the prior discussions coalesce: litigation over patents relevant to a technical standards like the 3G wireless, 4G wireless, and WiFi standards. These are high-profile, big-ticket litigations where injunctive relief is almost surely no longer available. These are litigations where typical litigants can easily weather a modest delay. And these are cases where extra time would surely lead to more accurate results. The Department of Justice and the European Commission have focused on these fights for other reasons, but for me these cases are interesting because they represent strong examples of situations where the power of delay can be beneficially brought to bear.

Part VI then concludes on a pessimistic note. If court-determined valuations and injunction-inspired negotiations are both deeply imperfect, the patent system seems dangerously likely to cause harm no matter how a given case ends. The only real solution is to reign in the beast, allowing the reality of imperfect patent remedies to temper society's recent and troubling penchant to issue patents on even the smallest of technical achievements. Bluntly, given how hard it is to enforce patents efficiently,

patent protection should be reserved for only society's most major accomplishments, not its relatively minor ones.

I. Measuring Patent Damages

Under current law, a patent holder whose patent has been infringed is entitled to compensation measured in one of two ways: measured by lost profits, which is to say the profit the patent holder would have made had there been no infringement; or measured by a court-determined reasonable royalty, which is often defined as the royalty the parties would have chosen had they actually negotiated just prior to the first act of infringement. Patent holders are also entitled to be paid interest on whatever damages are owed, and certain court costs, including any fees paid to the court, but not including the costs associated with attorney time.

That might all sound sensible enough – and, at a certain level, it is -- but the real challenge in patent law comes not in articulating these high-level concepts, but in applying them to actual cases. One problem is that patent law today gives juries too much information. The main case on the aforementioned "reasonable royalty" calculation, for example, lists no fewer than fifteen factors that a jury is allowed to consider in establishing the royalty. In isolation, each of those entries makes intuitive sense. One factor, for instance, welcomes information about the "commercial relationship between the licensor and licensee, such as whether they are competitors" while another sweepingly invites testimony on the "effect of selling the patented specialty in promoting sales of other products of the licensee." In the aggregate, however, a fifteen-factor list is an engraved invitation to mischief. A good lawyer or clever expert can massage almost any story such that it seems to fit into one of those

fifteen buckets. And jurors, overwhelmed by so many arguments, details, and numbers, likely respond by judging on simpler, less accurate, emotional grounds.

Better jury instructions could mitigate the above concerns, of course. More detailed verdict forms could also help, specifically by making the jury's math more vulnerable to after-the-fact judicial review. District court judges could in theory act as gatekeepers for damages analysis, excluding theories that are permissible but implausible, or in other ways allowing only the most important information to be presented in court. Moreover, judges could make more use of court-appointed damages experts, asking them to weigh the relevant economic evidence impartially but then exposing them to vigorous cross-examination in order to increase the reliability of their work. But experience teaches that district court judges are reluctant to use any of these tools to discipline the damages conversation. After all, district court judges have long had substantial power to craft jury instructions, to require detailed verdict forms, to appoint unbiased experts, and to toss expert testimony to the extent it does not live up to acceptable scientific standards. Judges, however, use these powers only lightly, and patent holders then exacerbate the problem by strategically choosing to file their cases in jurisdictions where the judges are most lax.

Jurors receiving too much information is one problem; jurors receiving too little information is another. For example, a patent holder will typically focus its case on the couple patents it owns, not say anything about other patents, owned by other patent holders, that are nevertheless relevant to the infringing product. However, this can be an omission of enormous importance. The typical cell phone, for example, implicates thousands of patents held by dozens of firms. What confidence can we have that a jury looking at one or two of those patents in isolation will come up with a remotely plausible estimate for its relative value? An analogous interaction would be to ask a

layperson to value a car's rearview mirror without telling that layperson anything about the car's engine, its doors, or its windshield. One can imagine a lawyer in that context delivering grandiose arguments that are completely true—"no one would have bought that car without a rearview mirror!"—but it is also easy to see how those arguments would be misleading in the absence of other information.

Accused infringers cannot plausibly solve this problem. One challenge is that some of the patents will be missing from the analysis because even the accused infringer will not know about them. After all, patent applications can languish for years at the Patent Office, completely out of view, only to then issue and become relevant to already existing products and services. Moreover, even issued patents are often hidden as a practical matter. The Patent Office issues over 400,000 patents every year; they are often written to be intentionally vague; and patent drafters are allowed to coin and then use their own idiosyncratic words and phrases. Moreover, the Patent Act ironically discourages the act of reading a patent, because under modern patent law the very act of reading is the predicate for enhanced damages and also for certain theories of third-party liability. Given all that, accused infringers often do not know the full list of patents implicated by their own products and services.

Even where an accused infringer does know that full list, the accused infringer will be reluctant to mention those patents for fear of alienating the jury by honestly reporting that the product at issue actually infringes not only the patents at issue in the case but also (say) another dozen patents that have yet to be litigated. And even if the accused infringer knew about all of the relevant patents and was willing to catalogue them for the jury, what then? No court is going to be willing to construe and evaluate dozens of patents that are not directly at issue in the case just to provide the jury with necessary context about a given product's patent footprint.

The dynamics of a patent trial impose yet another constraint on a jury's ability to value patents: accused infringers cannot afford to spend too much trial time or credibility arguing about damages. Lawyers, of course, are enormously comfortable with the idea that an accused infringer can simultaneously argue that the patent is invalid, that the patent is not infringed, and that, if it is valid and infringed, the damages ought to be very small. But juries are not similarly at ease with these types of "That is not my dog; if that is my dog, he did not bite you; and if that is my dog and he did bite you, you deserved it" patterns. Thus, as a strategic matter, defense counsel in a patent case will often have no choice but to downplay arguments about why damages should be low and focus instead on validity and infringement. That of course presents a problem in cases where those first two moves fail. The jury simply has not been given what it needs to do the damage numbers right, and the jury's momentum might at that point significantly favor the patent holder regardless.

Need more? A conventional approach to "reasonable royalty" analysis is to establish some percentage – say, 2% -- and then to calculate damages by applying that percentage to the retail price of the infringing item. Patent lawyers speak of this as setting a royalty rate and applying it to a royalty base. In most patent cases, there is not an obviously correct choice for the royalty base. In a case where the patent covers a camera lens, for instance, plausible arguments can be made in favor of using the price of the camera itself as the royalty base, and plausible arguments can be made in favor of using the price that the camera company pays for the lens when it acquires the lens from the relevant lens manufacturer. From an economic perspective, this choice does not matter. If the higher royalty base is used, the jury should see that and scale the royalty rate down proportionally. If the lower royalty base is used, again the jury should compensate by scaling up. Yet patent lawyers fight intensely over this issue in most cases, clearly believing that jurors are unable to make these sorts of intuitive

adjustments. (One theory is that jurors are not willing to think about percentages much below 0.25%, and so the theoretical "scaling down" response has a floor that makes large royalty bases strategically attractive. Another theory is that a large royalty base will set up a David-versus-Goliath dynamic in the case, with the jury thinking that the infringer can easily afford to pay the patent holder in the case at hand.)

Thus far, I have focused on damages measures that are designed simply to compensate patent holders for the monies they would have earned had their relevant patents not been infringed. The mathematics become even more intractable when damages attempt to do something more. Consider, for instance, deterrence. If damages in a patent case simply required the accused infringer to pay exactly what he would have paid had he negotiated ahead of time, accused infringers would have no incentive to negotiate. Instead, an infringer's dominant strategy would be to infringe, hope to avoid detection, and then, in any case where the infringement is detected, pay. By doing so, the infringer would benefit from the possibility of not being detected, and the infringer at worst would have to pay exactly what he would have paid had he announced himself from the start. Of course, knowing this, patent law damages cannot merely be calculated to be the amount of money the infringer would have paid had he paid in the first place. Damages must instead be higher – although good luck quantifying exactly how much higher given that the correct answer there turns on unknowable information such as the infringer's own estimate as to the probability of being caught.

An even harder adjustment along these lines is the adjustment that must be made to account for the fact that, prior to litigation, there is almost always some uncertainty as to whether the patent at issue is valid, infringed, or both. That is, the accused infringer might plausibly believe that the patent should never have been

issued, and the accused infringer might similarly believe that its technology is not covered by the patent's claims. Litigation resolves that uncertainty. Thus, when a patent holder prevails, the damages awarded naturally must be higher than the royalties the parties would have negotiated prior to verdict. But how much higher?

Think of it this way: if prior to litigation a patent holder and a would-be licensee both agree that there is a 50% chance that the asserted patent is invalid, their private deal would reflect those doubts. The licensee would demand a discount as compared to a sure-thing royalty, and the patent holder would accept that discount in order to avoid the risk of a bad outcome. If that patent holder ends up successfully litigating the issue, however, the resulting court-ordered royalty should no longer reflect that 50% discount. Had the patent holder lost the case, he would have earned nothing. Given that he won, he should earn the undiscounted award. Again here, however, saying that in theory is a far cry from accurately implementing that adjustment in a real case. After all, any unrelated real-world deals made by the patent holder likely understate this target because those deals would have been subject to the very uncertainties that this litigation in the end removes.

My list could go on, but the point by now is obvious: no matter whether the goal of the patent system is compensation, deterrence, punishment or some combination thereof, actually choosing a sensible cash response to patent infringement is enormously difficult. Court-appointed damages experts might help. Bifurcation of the liability and damages phases might help too. But, even with those adjustments, quantification will always be a deeply imperfect science. The key question, then, is whether there is a plausible, superior alternative.

II. Injunctions

For decades, patent courts thought that there was a superior option: injunctions. The idea was simple. As soon as the court knew that a given patent was both valid and infringed, the court could order the relevant infringer to stop its unlawful activities. At that point, the court's own inability to value the patent would become irrelevant. Either the infringer would stop using the technology, and hence there would be no need to further quantify patent value; or the infringer would negotiate a deal directly with the patent holder, and hence the private parties would quantify patent value untainted by the court's imperfect understanding.

That all sounds good in theory, but in practice injunctions cause their own substantial distortions. Consider a typical scenario where a company happens to independently invent some technology that later turns out to also be subject to someone else's patent. Not knowing about the patent, the company might make substantial investments that are specifically tied to the technology. The company might build manufacturing facilities that are specifically tailored to include the feature or component at issue. The company might sign long-term contracts to procure inputs specially suited for that now-vulnerable use. The company might also make other long-term commitments, such as promising to deliver the patented functionality to particular customers and downstream partners.

These are normal commitments for a company to make. But think about what happens if a patent holder later emerges, sues, and wins an injunction. Remember, the purpose of an injunction in this context is to force the infringer to negotiate with the patent holder and ideally come to an agreement that accurately reflects the value of the patented technology. But, on these facts, the private negotiation will sound nothing like that. The patent holder will begin by pointing out that, by virtue of the injunction, the

infringer can be forced to idle that manufacturing facility, break those supply contracts, and disappoint those customers and partners. The patent holder will then offer to allow the infringer to avoid those disruptions, but only if the infringer is willing to pay a price that reflects those potential costs. That is, the price set in the shadow of the injunction will have almost nothing to do with the merits of the patented technology as compared to some next-best alternative. Instead, the price will be driven by questions about how hard it would be for the infringer to abandon the technology now, given whatever commitments the infringer made before it even knew the patent existed. Put differently, injunctions in these situations allow patent holders to hold hostage infringers' technology-specific investments.

This hostage-taking dynamic would not be an issue if companies could reliably identify relevant patents before building manufacturing facilities, signing contracts, and otherwise making commitments. And some commentators believe that companies can in fact do that. As they point out, the quid pro quo of the patent system is that inventors disclose their inventions to the public, and in exchange the government grants those inventors exclusive rights to make, use, or sell the disclosed technologies. That would seem to suggest, they say, that there is an accessible public record of patented technologies, and that firms vulnerable to injunctive relief could simply flip through that record and identify potential obstacles to their work. In practice, however, that approach simply does not work.

Trouble begins with the fact that every patent is written in its own vocabulary. Two patents might therefore describe the exact same technology, but the descriptions would look nothing alike, and might similarly bear no resemblance to how the potential infringer talks or thinks about its own products and services. To make matters worse, patent language is subject to hopelessly nuanced rules of interpretation. Indeed, there

are actually cases where the Federal Circuit has struggled to decide what it describes as "plausible disagreements" as to the meanings of seemingly innocuous words like "to,""on," "about," and "through." In a world with that much hairsplitting—let alone the large number of patents in force—identifying and interpreting every relevant patent is just implausible. Of course, this is not to imply that no patents can be identified by means of a careful search. Often even an amateur eye can spot at least a few relevant patents in short order. In practice, however, a firm cannot hope to reliably identify all patents relevant to a given product or service, and identifying even a subset of those patents is likely an expensive, time-consuming, and deeply flawed process.

All that is the analysis as it applies to issued patents. Patents that have yet to issue pose even more significant problems. Patent applications can be kept from public view for at least eighteen months after filing, and a strategic applicant can maintain secrecy even longer by (for example) certifying to the Patent Office that the relevant application has not been filed in any country that requires publication. Moreover, even years after an original patent filing, a patent applicant can return to the Patent Office and file new claims based on his old submission. Those new claims can be broader than were the original claims, and yet the patent system will still treat the application as if those claims were filed back when the original application was submitted. Thus a technology might be patent-free when first evaluated, but years later that technology might be subject to perfectly valid but at-the-time-undetectable patent protection.

Add to this the concern that the Patent Office is not particularly reliable when it comes to evaluating proposed inventions and weeding out those that cover already known achievements. This means that, no matter how careful a company might be, there will always remain the real risk that some later patent applicant will claim to have invented a relevant technology, and, despite the fact that the invention was already well

known and also in use, nevertheless convince the Patent Office to issue the patent. Against those sorts of mistakes even careful attempts at search are no answer.

The hostage-taking dynamic is therefore inevitable. Firms will regularly create patented technology and make investments related to that technology. Firms will regularly do so without knowing that a patent already was, or later will be, relevant. And if those firms are vulnerable to injunctive relief, the resulting negotiation between patent holder and infringer will systematically overvalue the patent. The patent holder will be able to demand a price that reflects the value of the patented technology as compared to the next-best option; but the patent holder will also be able to extract an additional amount that reflects the disruptions the infringer would suffer if it were forced to make that switch now, after having already committed to what turned out to be the patented option. Injunctions, then, are not a promising mechanism by which to achieve accurate patent quantification. Like court-determined valuations, in many instances, injunctions are highly unlikely to ultimately get the numbers remotely right.

III. Delay: The Costs

In patent law, it is widely assumed that faster is better. The International Trade Commission races to decide its cases in under eighteen months. Patent courts in Virginia and Texas pride themselves as being "rocket docket" jurisdictions. The recently enacted America Invents Act gives the Patent Office only eighteen months to reexamine a patent that it previously issued but later doubts.

Why the hurry? At the start of this Essay, I argued that one reason for all this haste is that the courts and Congress want to minimize the importance of court-

determined quantifications. Put another way, patent policy-makers are well aware of how difficult it is for judges and juries to value patent harms, and so patent law was built to reach injunctive relief with reasonable speed and thereby move valuation questions out of the courtroom and into the boardroom. As I also pointed out, however, with injunctions now increasingly off the table, that explanation no longer works. At the end of a modern patent case, the court will often keep the valuation question for itself and impose a forward-looking court-determined royalty. Fast decisions therefore no longer help from a valuation perspective. For infringement that takes place prior to the court's final decision, the court must value the patent under the framework of traditional damages analysis; and for infringement that takes place after the court's final decision, the court must still value the patent, albeit it under the newly important framework of court-determined forward-looking royalties. Difficulties in valuation no longer argue in favor of fast decision-making.

So what does? One reason for courts to hurry is that a dollar paid today is not the same as a dollar paid next year. This concept is often referred to as the "time value of money" and the intuition is likely familiar. If an infringer owes one dollar to some patent holder, the patent holder is obviously better off if that dollar is paid immediately rather than paid months later. If the patent holder had possession of that dollar, after all, he could have used it to support his business, he could have used it to purchase some consumption item for himself or his family, or he could have even put it in the bank and earned interest. Delay thus threatens to underpay patent holders. If an infringer is determined to owe a particular royalty, but that determination is delayed for several years, the value of the payment to the patent holder is lower than it would have been had the payment been prompt.

This, of course, is not a reason to rush; it is instead a reason to award prejudgment interest. That is, patent courts can fully compensate for this harm by accounting for the time value of money when issuing payment instructions. Section 284 of the Patent Act provides the necessary legal authority: "Upon finding for the claimant, the court shall award the claimant damages adequate to compensate for the infringement, . . . together with interest . . . as fixed by the court." And the Federal Circuit has made clear that the purpose of this language is "to compensate for the delay a patentee experiences in obtaining money he would have received sooner if no infringement had occurred." The time value of money, then, is a reason to pay attention to delay, but not a reason to avoid it.

A more serious reason to be nervous about delay is the concern that some patent holders might not have adequate liquidity to survive while waiting for their patent verdict. Yes, some patent litigants have tremendous resources (hello, Nokia) and can patiently wait for a court to determine what monies are due. But some patents are held by smaller entities for whom their patent case might be their primary or indeed only asset. The problem here is not that longer litigation is more expensive. It might be, for instance if lawyer time is continually wasted coming up to speed on the case in order to deal with some issue, only to then have the case sit idle again until the next big event. And it might not be, for instance if a slower pace allows the legal team to sequence the work in a more efficient manner or to run the case with a smaller staff. The problem instead is that small patent holders might not have the capital they need to pay key staffers and literally keep the lights on. Bluntly, a patent holder cannot cover his gas bill by telling the gas company about his expected, one-year-away, highly-probable $100 million patent verdict.

This concern has more purchase than the concern about the time value of money. There presumably are patent holders who might not be able to tolerate an extra year or two or three of delay, even if the costs of litigation stay the same and even if the court promises to ultimately award not only the monies owed but also an appropriate interest payment. That said, this concern has been substantially dampened by the rise of third-party litigation finance. Significant capital is today available to patent holders who might need help paying for litigation and covering other expenses while they wait for a final court ruling. True, that capital comes at a cost; an investor who funds a patent litigant for several years will demand some sort of return on his investment in the event the case ultimately goes well. But courts could account for that cost when calculating damages, reimbursing for this loss in much the same way that courts reimburse for the lost time value of money. How all this balances out will admittedly vary by case, by plaintiff, and over time; but the core point for now is that while liquidity is a cost of delay, it is a cost that can be managed. Besides, this same point applies today, when patent cases routinely drag on for five, six or seven years – and thus the real question relevant to this Essay is the degree to which an extra year or two would, at the margin, significantly change the overall value of a given patent right.

That takes us to the third and primary reason to avoid delay: delay increases the duration of patent uncertainty. Until a court definitively rules on a question of patent infringement, neither the patent holder nor the accused infringer know for sure whether the patent is valid, whether the patent is infringed, and what costs will be imposed for past as well as future use. And the range of plausible outcomes can be large. In many patent cases, there is a realistic chance that the court will decide the patent is invalid or not infringed, leading to a zero outcome. In those same cases, however, patent holders are often credibly seeking tens to hundreds of millions of dollars. Uncertainty between

those extremes obviously matters, and thus maintaining uncertainty for a longer time is clearly a real cost that must be considered in any conversation about delay.

That said, be careful not to overemphasize certainty in this context given how little weight certainty is accorded almost everywhere else in patent practice. Consider, for example, the rules that govern when a court determination regarding patent validity binds later litigants. A patent holder who successfully defends patent validity in the context of a first infringement suit must start afresh when he sues a second infringer. Again, the patent holder must rebuff arguments that the patent was improvidently granted. Again, the patent holder must establish his desired claim constructions. A patent holder whose patent is found invalid in some first case, by contrast, is barred from ever again enforcing that patent. If there is some randomness in litigation, the result here is to shift significant uncertainty onto patent holders. A lucky draw has implications only for the specific litigation at hand. An unlucky one has implications for every future interaction.

The interpretive rules under which patent claims are analyzed similarly undermine patent certainty, not because of their substance but because they are constantly in flux. One minute the PTO is approving claim language where some new apparatus is described in part by articulating how the apparatus should be used; the next, the Federal Circuit retroactively declares all such claims to be so unclear as to be invalid. One minute the practice of altering claim language during patent prosecution is understood to be a natural part of the give-and-take between applicant and examiner; the next, the Federal Circuit and the Supreme Court combine to announce that almost every such language alteration will be construed as a concession that limits patent scope, and that the new rule will apply retroactively.

And this is just the tip of the iceberg. The Federal Circuit regularly reverses lower court claim construction decisions. The Supreme Court recently clarified the rules that govern what types of inventions are eligible subject matter for patent law (Mayo), only to basically ignore their own clarification twelve months later (Myriad). A patent can be held invalid because someone uncovers "secret" prior art—art that was not public at the time of invention, but that is nevertheless admissible in court under one of several special exceptions. And on and on and on.

I provide these examples not to question whether certainty has value – of course it does – nor even to criticize these specific rules and decisions. I provide these examples instead to point out how disingenuous it would be to put certainty on an untouchably high pedestal in just this one context. The lesson from patent law more generally seems to be that stability is desirable, but the patent system is often willing to pay only a modest price to achieve it.

One reason that patent law is so willing to sacrifice stability is that, in practice, legal uncertainty is only one among many types of uncertainty in play. Pharmaceutical companies, for instance, admittedly worry about the strength of their patent portfolios. But a little less certainty there is unlikely to radically alter behavior given that success in the pharmaceutical industry critically depends on other, unavoidable uncertainties such as the uncertainty associated with FDA review and the very real risk that, because of some unexpected side effect, a blockbuster drug will suddenly become a source of devastating legal liability. Similarly, small firms and start-ups confront enormous risks above and beyond the risks associated with patent validity. Indeed, every venture capitalist in the country can list dozens of innovative start-ups that today hold issued, presumptively valid patents but have yet to generate a penny of revenue. Again, patent

uncertainty is important; litigation delay increases the length of time during which uncertainty will prevail; but the importance of uncertainty ought not be overstated.

Note, too, that prolonged uncertainty in this context would disproportionately hurt weak claims. A patent that is clearly valid, clearly infringed, and clearly valuable does not suffer much harm if it takes the courts an extra year or two to finalize the exact monies owed. Even while waiting, that patent holder can be reasonably confident that the patent will survive court challenge. Even while waiting, that patent holder can be reasonably confident that the patent will ultimately generate substantial income. A patent holder relying on a suspect patent, by contrast, will feel uncertainty more sharply. Uncertainty in this light might be a feature of delay, not a bug.

Moreover, remember that uncertainty is felt by both patent holders and accused infringers. Given that, it might be that increased uncertainty would simply pressure both sides to settle their dispute earlier in time. Such a settlement would reflect each side's expectation as to what the court would do if the case were allowed to progress to final judgment at an appropriately slow pace. But it would be attractive because it would spare each side the costs of litigation, and the costs of continued uncertainty. Those savings could be shared by the parties, and in some instances that might close the gap between what the patent holder is willing to accept and what the accused infringer is willing to pay.

IV. Delay: The Benefits

My focus thus far has been to explain why delay might not be as bad as it might at first seem. My points are simple. Thanks to the decline in the tendency to award injunctive relief, delay no longer exacerbates the problem of patent valuation. The time value of money can be fully accounted for by court-ordered prejudgment interest. Any

concern about liquidity applies only to patent holders with liquidity challenges, and even for them can be dampened by third-party litigation finance and factored into the court's final monetary award. And, while delay might prolong uncertainty, that uncertainty might beneficially promote settlement, and, besides, patent law and entrepreneurial activity more generally are so rife with uncertainty that a little more might not actually matter. But none of that makes a case in favor of delay. Here, I turn to that important task.

A fast case will often be heavily biased in favor of the patent holder simply because a fast case in practice gives the patent holder more time than his adversary. Patents holders can prepare their case before it is filed, and thus much of their preparation can be done off the clock. Before filing, for instance, a patent holder can line up its infringement, validity and damages experts and prepare much of the relevant testimony. An accused infringer, by contrast, might not even know the relevant patent exists until the infringer is served with the complaint.

The faster a case hits key milestones, the more this imbalance matters. For instance, in Texas, an accused infringer has only three months to disclose to the patent holder whatever evidence it will later use to argue that the patented idea was in fact obvious at the time it was supposedly invented. Three months is not a lot of time, especially because the hunt for prior art cannot really commence until after the accused infringer has first hired counsel, studied the patent, and studied its own accused products in light of the patent language.

This is a particularly severe problem in cases with a large number of patents. When ParkerVision sued Qualcomm in 2010, for instance, it alleged infringement of eighteen patents and over two hundred patent claims. Just reading all those documents presumably consumed a significant percentage of Qualcomm's litigation bandwidth,

and yet local rules still required that Qualcomm commit to its non-infringement and invalidity theories within the first few months of litigation. ParkerVision, by contrast, had spent literally decades building its portfolio and otherwise preparing for this fight.

In addition to reducing the importance of timing imbalances, delay can improve accuracy in another way: it can open the door to greater involvement by the Patent Office. The Patent Office today offers several different procedural mechanisms through which a suspect patent can be brought back for a second look. These procedures are promising because they are run by experienced patent examiners, and because, unlike courts, patent examiners are not required to show any deference toward their colleagues' original decision to issue the patent. Many of these review processes are also adversarial, unlike the initial process of patent application which is always uncomfortably ex parte.

Courts cannot benefit from these administrative reviews, however, unless they are willing to slow their own processes. The Patent Office needs time to hear evidence, reach a decision, and then defend that decision on appeal. If the Eastern District of Virginia is going to turn patent cases start-to-finish in under two years, the Patent Office simply cannot possibly help.

Delay can also increase participation by private parties. When the patent holding company VirnetX sued Microsoft for patent infringement, for instance, other technology companies were caught flat-footed. Microsoft alone defended the case. Microsoft alone petitioned the Patent Office to re-examine the patents. Within two years, however, it became clear that VirnetX was going to assert those same patent families against dozens of technology companies and hundreds of networking products. With that, Cisco and Apple joined the fray, filing their own petitions with the Patent Office and in that process delivering over six hundred pages of evidence that

Microsoft had not itself found. The Microsoft litigation moved too quickly for Microsoft to benefit from those extra disclosures, however, because the district court moved the case from start to finish in under two years. The result as it stands today is an embarrassingly split outcome under which Microsoft has had to pay VirnetX over $200 million, Apple currently labors under a $350-plus million judgment, and Cisco has been deemed to not infringe at all. [Note: I need to find a better example than VirnetX. This example does not exactly work because, as of this writing, Cisco and Apple both failed to defeat patent validity. A better example would be one where the late-comers actually were successful in invalidating the relevant patent.]

Patent holders, too, might have more of an opportunity to join the conversation if patent cases were to move a bit more slowly. Motorola, for instance, recently sued Microsoft, alleging that Microsoft's implementation of the WiFi standard infringes a handful of Motorola patents. Motorola admittedly holds only a small percentage of all the patents relevant to WiFi, however; and, until that total patent universe is clear, the court overseeing that fight will be hard pressed to accurately value Motorola's relative contribution. But that information is just not available yet. Patent holders large and small are still themselves studying the WiFi standard, iterating with the Patent Office on the scope of their claims, and gradually figuring out which patents are essential to that standard. Delay would allow more time for that process to play out, ultimately increasing the likelihood that the court's valuation will accurately quantify Motorola's proportional achievements. [Again, I might swap this example to something else.]

Evidence of independent invention is yet another type of information that might become available if courts were to slow down long enough to allow it. Inventors are not entitled to patent protection if their invention was something that would have been obvious to someone skilled in the relevant art. One good way to objectively test for

obviousness is to wait and see if a large number of people skilled in the art in fact do independently come up with the same idea. If so, the idea might well have been obvious, and the patent applicant's only real claim to fame might be that it was first to commemorate that obvious idea in the form of a patent application. This evidence takes time to percolate, however, particularly if patent holders tend to write up obvious ideas long before any commercial entity would plausibly consider implementing it.

I have focused thus far on specific types of information that might be more readily available if patent courts were to slow down their work; but for innovative technologies there is often a more general understanding that only comes with time. Back when the Internet was first being developed, for instance, no one really understood the impact those protocols would have on commerce, culture and communication. Much the same, when the now-familiar 2G wireless standard was first promulgated, even that technology's strongest proponents could not have foreseen the degree to which cell phone usage would permeate both work and play. Patent decisions will often be more accurate in the long run if they can be made after the parties and the courts more fully understand how the technology at issue is going to be used, and what it is going to mean for business, culture, and life.

V. Patented Standards

Patents that relate to technical standards have been the subject of particular scrutiny over the past several years. Part of the explanation is that these patents have been dragged into several high-profile, big-ticket litigation. When Samsung uses standard-essential patents to sue Apple, or Motorola uses standard-essential patents to sue Microsoft, it is no surprise that policy-makers, regulators, and commentators take notice. Part of the explanation is that these patents have been the subject of blockbuster patent sales. Microsoft, Apple, EMC, RIM, Ericsson and Sony teamed up to buy

Nortel's standard-essential patents for $4.5 billion. A few weeks later, Google spent over $12 billion to acquire Motorola, and conventional wisdom is that the bulk of the money was justified by the value of Motorola's standards-relevant patents. And part of the explanation is that technical standards are simply of enormous importance to modern society. The fact that Verizon, AT&T, Sprint, Google, Samsung, Apple, HTC and Microsoft all were able to agree on what it means to deploy first a 3G and now a 4G wireless network means that consumers have been able to use their cell phones and computers to communicate wirelessly across these various networks and brands. The fact that a similar who's-who of technology companies agreed on what it means to encode music in MP3 format, encode motion pictures on a DVD, or encode web content using HTML5 similarly made possible industries and interactions that today seem central. That each of those standards is governed by hundreds to thousands of patents meant that the patent issues were bound to attract substantial attention as well.

For my purposes, these patents make for an informative case study, because all of the issues I raised above resonate in the context of patented standards. [Note: This is a placeholder for now, because I want to finalize the front portion of the paper before writing this section and here applying those ideas in the standard-setting context. That said, my plan here is to go through all the issues from the front of the paper and show how they play out for standards-associated patents. I will point out that essential and also non-essential patents are difficult for courts to value, among other reasons because they are so numerous; because standards-compliance is just one of dozens of features that impacts consumer demand; and because the list of patents is constantly changing, even years after a given standard is deployed. I will then talk about how injunctions would make possible not only a hostage-taking dynamic here, but also an interaction where a patent holder would be able to charge a price that reflects the value of having been chosen for the standard, even if other technologies would have been equally good.

I will make the injunction points for both essential and non-essential patents; but I will flag RAND and its implications for essential patent particularly. On delay, I will then rehearse the various points made previously, noting that entities like Samsung and Motorola clearly do not have problems with liquidity and clearly can weather additional uncertainty. The last move in this section will then be to emphasize how additional time could help get outcomes right, among other things by giving infringers time to evaluate the swamp of patents relevant to a given standards, and by making sure that all the relevant patents are known before the court sets out to assign value to each. Again, I'll write this section only after the front has coalesced, because ultimately I want this section to parallel the order and structure of the front four.]

VI. Conclusion

Patent system remedies are deeply problematic. Court-determined quantifications are problematic because judges and juries simply do not have the skills or background to do that work. Injunctions are problematic because a patent holder armed with an injunction can hold hostage whatever technology-specific investments the relevant infringer previously made. As I have argued here, the fact that patent law is increasingly relying on court-determined quantifications and increasingly rejecting market-forcing injunctions opens the door to a beneficial slowing of patent litigation timetables. But the fact that both alternatives are so deeply flawed brings forward an entirely different point: patent law should raise the bar on what is deemed to be eligible for protection.

Think of it this way. Every time the Patent Office lets issue a patent, society is put on a path toward someday having to either quantify the value of that patent in court or allow a private party to enforce an injunction against a rival. Yes, some patents will then disappear without a whimper, for instance because the patented technology turns out

not to have commercial value, or because no firm ends up deploying the patented technology without the patent holder's permission. And for some patents, enforcement by injunction is an easy call because the relevant infringer knowingly copied. But in any case where the patented technology turns out to be independently invented by some other entity--a common and likely situation--society is stuck. The parties might settle their disagreements in the shadow of the law, but the imperfections inherent in both court-determined valuations and court-ordered injunctions will still impact real world behavior and influence the allocation of resources.

To me, this argues for an incredibly cautious patent system. Is it really plausible that it takes thousands of patents to motivate the development of a standard like 3G wireless? Does anyone really believe that Facebook would not have existed as soon or as early but for the incentives created by the patent system? Society has no choice to suffer the imperfections of patent law remedies in instances where patents are necessary to accelerate the development and deployment of new technologies. But a sober look at those imperfections makes clear how reluctant society should be to start down that path in the first place. Patent protection should be reserved for major accomplishments and denied to the modest ones.