2009

Optimal Patent Jurisprudence

Scott Baker
Washington University in St. Louis Law School

Claudio Mezzetti
University of North Carolina (UNC) at Chapel Hill - Department of Economics

TILBURG UNIVERSITY

TILEC

TILEC Discussion Paper

DP 2009-006

Optimal Patent Jurisprudence

This paper results from the TILEC IIPC grant competition

By

Scott Baker & Claudio Mezzetti

February 2009

ISSN 1572-4042

http://ssrn.com/abstract=1355863

Optimal Patent Jurisprudence

Scott Baker
University of North Carolina
School of Law
Chapel Hill, NC 27599-3380
sbaker@email.unc.edu

Claudio Mezzetti
University of Warwick
Department of Economics
Coventry, CV4 7AL
c.mezzetti@warwick.ac.uk

Conference Draft[1]

February 4, 2009

[1] An earlier version of this paper was presented at the "Workshop on Innovation, Intellectual Property, and Competition Policy" held at Tilburg University. We thank the participants at the conference and especially our commentator, Jan Boone, for helpful suggestions. This research was generously supported by a grant from the Tilburg Law and Economics Center.

Abstract

We model judicial learning about optimal patent policy. The court is infinitely lived; the plaintiff and defendant are short lived. Litigated cases provide the court with information about the optimal rule. Different cases provide different sorts of information. Opinions influence the stream of future cases likely to be litigated and, as a result, change the flow of information to the court. In structuring opinions, courts make decisions whether to learn fast or slow. We have three main results. First, patent law will stabilize even if the court places zero value on the "predictability" of legal rules. Second, path dependence of law is a rare outcome. It occurs only when the court stops learning and decides that the error costs (the losses from some cases going the wrong way) are lower than the decision costs. Finally, the law can be optimally incoherent in the short run. The court will pay lip service to prior holdings, while dramatically altering the legal landscape. Patent opinion incoherence, which is often the subject of much scholarly critique, makes sense because it facilitates future learning from a population of cases most important to the court for policy-making.

1 Introduction

Courts make most patent law. Although patent reform is perpetually discussed in congressional committees and subcommittees, Congress seldom legislates in the area. That leaves a big chunk of the United States innovation policy in the hands of the judges. These judges are routinely criticized by scholars and the patent bar. As to the Supreme Court, the charge is that justices lack the necessary information and technology background to make informed patent policy. As a result, Supreme Court patent doctrine is haphazard, shoddy, difficult to understand, and inconsistent. With respect to the federal circuit, the argument is that the Patent and Trademark Office (PTO) has difficulty implementing any doctrine that court lays down. A related charge is that opinions by different panels of the federal circuit contradict each other; that is, the federal circuit rarely says the same thing twice in a row.

The fact that judges have trouble fashioning patent policy is not surprising. They don't have a lot of information, mostly what can be derived from litigated cases. And the issues are complicated. Even if the judges and justices agreed on the goal of patent law, how to get there would present a thorny problem. Yet all is not lost. Through the opinions they write, judges influence the next cases they see and, as a by-product, the information flow upon which to base future articulations of law. It is this ability, the ability to issue opinions, that enables the court to learn over time. How judicial learning takes place and what impact learning has on patent case law form the basis of our inquiry.

When should a court in construing patent law issue a narrow decision? A broad decision? Refuse to decide the case altogether, dismissing, for example, on grounds that it lacks jurisdiction? A narrow opinion doesn't preclude future litigation. It has no value as precedent. As a result, a narrow opinion doesn't alter the stream of future cases the court is likely to see. At the same time, a narrow opinion provides little notice about the contours of the law to the PTO, inventors, or potential infringers. A broad decision does the opposite. By its terms, a broad decision seeks to govern situations with facts quite different from the facts of the decided case. While a narrow decision is uninformative, a broad decision has a greater chance of being wrong.

The paper models optimal patent jurisprudence. To do so, we relax a standard assumption in the economic, political science, and law and economics literature on judicial behavior. We assume that the judges do not know what doctrine (i.e., set of legal rules) best advances their interests. Judges have a policy leaning, an ideal point, but are unaware how to effectuate that preference. At the same time, we assume that judges share a common goal for patent policy. These assumptions, we think, accurately describe how judges approach patent law. Unlike many other areas of law, much of the conflict in patent law is about means, not ends.

The distinctive feature of the analysis is that judicial learning is entirely endogenous. Judges learn from the cases they hear. The facts of the cases are the data points, the series of cases the sample. A judge estimates the optimal rule from the sample. Each opinion alters the stream of future cases, which changes the court's estimate of the optimal policy. A court's decision today influences the cases it sees tomorrow, which dictates what it learns the following day. Recast, patent jurisprudence morphs into a dynamic programming problem.

The model provides three insights. First, patent law made via judicial decisions will settle. At some point, the court stops learning and rests on the rule or standard of its last decision. The reason is simple: There is a cost to collecting data, the cost of reviewing cases on the merits. The benefit is the precision of the legal rule, how well the doctrine advances the court's interests. When the cost of collecting more data outweighs the benefit from increased precision, the court stops learning and, as a result, the law stabilizes. In fact, the law will settle even in the most extreme circumstance: where the court doesn't care about predictability of legal rules or reliance by inventors, the PTO, or industry participants. Counter-intuitively, a court's commitment to the rule of law is not necessary for a rule of law to emerge, scarce judicial resources is enough of a force.

Second, the law can experience doctrinal "traps." The court will refuse to revisit or refine a doctrine knowing full well that doctrine is imperfect. Imperfections in the doctrine mean that some cases that, from the court's perspective, should have settled to favor the plaintiff will instead settle in favor of the defendant. In these cases, the court's overarching goal will be compromised. The court tolerates this compromise to avoid the costs of having to examine another case on the merits, a prerequisite to streamlining the doctrine. Faced with a doctrinal trap, the court stalls. It dismisses on grounds that it lacks jurisdiction; it concludes that the case is moot or unripe; it writes an unpublished decision; or, in the case of the Supreme Court, it refuses to grant certiorari. In our model, these practices are justifiable because they allow the court to stick with an imperfect doctrine when the cost of perfecting is too high.

Many of the doctrinal puzzles that dog patent scholars can be seen as traps. Scholars make much out of the weaknesses of the courts' approach to patentable subject-matter (Marvin 2007; New 2007; Thomas 2002), obviousness (Ducur 1996, add cites), and prosecution history estoppel (Wagner 2002). Our model shows that, even if the court knew these areas of law were wrong-headed, it might not fix them. And that would be efficient. Sometimes the cost of error-correction outweighs the gain from a better doctrine.

If decision costs are sufficiently small, however, the law will evolve to a legal rule that exactly coincides with the court's interest. Only when doctrinal traps arise will law be path dependent, meaning that the order the cases come before the court determines the eventual legal rule. The model thus calls into question the conventional wisdom that large players in the patent game get the law they want at the expense of small inventors (cite). Control of the order of litigation is necessary, but not sufficient for large players to accomplish this feat. In fact, in equilibrium, the court estimates the losses associated with the imperfect rule, a rule that, say, advances the interests of large players at the expense of society. The court fails to revise the doctrine only if these losses are not that big of a deal.

Finally, the model demonstrates that when a court first confronts an issue or faces a new problem, the law made through opinions might jump around a lot. The court will distinguish away prior cases that seem indistinguishable. The reason is that the court learns the most from the first few cases it hears. If it goes down the wrong path in the first decision, the next case will provide enough new information that the court will recalibrate the doctrine immediately. Charting a new path ensures the court observes cases in the future that matter the most for refining the legal rule. In the short run, the case law will be inconsistent, but this inconsistency will reflect optimal behavior over the long haul. The model thus provides an alternative take on the well-documented and often-criticized inconsistency in patent case law. Inconsistency might not reveal flaws in the courts, but virtue.

Indeed, the court will optimally make incredulous statements like, "case A, which appears to every lawyer in the country to stand for proposition X, really doesn't stand for proposition X, instead it stands for proposition Y." Scholars and legal commentators find such statements maddening. That is to say, the judge or justice prefers a certain outcome in the case, prior decisions are a road-block, so the judge re-characterizes previous decisions in a self-serving way. We show that these doctrinal moves could have a less sinister explanation. The ability to pull stunts like this allows the court to learn more quickly.

More generally, legal scholars and political scientists regard inconsistency in doctrine as inevitable, because courts have multiple judges voting and each one has a different policy preference (Easterbrook 1987). In our model, a court with a unified policy preference – i.e., a preference with which all judges agree – will nonetheless issue inconsistent doctrine. The upshot is that one cannot just look at inconsistency in the doctrine and conclude the court is acting badly.

The paper proceeds as follows. Section 2 briefly reviews the literature, mostly from political scientists and law and economics scholars, about judicial behavior. Section 3 sets out the model and Section 4 contains the main results. Section 5 concludes.

2 Literature Review

This paper sits at the intersection of two literatures: judicial behavior and the evolution of legal rules. The judicial behavior literature is well-developed by political scientists. These scholars model courts or judges as rational actors seeking to maximize some objective function. The constraints in the models vary and depend on the question being asking. Sometimes the judge makes his decision anticipating the likely position of the Congress or the executive (add cites). The judge wants to push the law just far enough that the concurrent branches won't overrule the decision. Other times, the judge feels constrained by other judges sitting on the panel (add cites) or the likely position of the higher court (add cites). Still other models view judges as interacting repeatedly with each other over time (O'Hara 1993). The repeated interaction, then, compels each judge to follow the opinions issued by his peers. In so doing, he ensures they will follow his opinions.

In all these models, the judge's objective function is given. The judge knows what he wants – the issue is whether he can actually do it given the other actors in the system . Our model drops this assumption and asks different questions: If a judge is unsure about the best course of action how might he structure opinions? And what does judicial learning mean for the evolution of law?

In the law and economics literature, the evolution of legal rules has been a topic of much research. In 1973, Judge Richard Posner famously asserted that the common law was efficient. After canvassing all of common law, he concluded that more often than not legal doctrine bore the hallmark of efficiency (Posner 1973). In the wake of Posner's assertion, scholars sought to explain the mechanism by which the common law could become efficient. The models relied on case selection by litigants (Priest 1977, Rubin 1977). Litigants present inefficient rules to the court, while settling cases involving efficient rules. Since the court only sees and agrees to reexamine inefficient rules, the law dove-tails toward efficiency. The validity of these early models has been questioned (Hadfield 1992; Hirshleifer 1982; Bailey and Rubin 2004). The literature has blossomed with many models of factors that point toward or against efficiency (see, for example, Zywicki 2003; Fon and Parisi 2008). For our purposes, the important feature these models lack is a forward-looking judge. Case selection drives the law, with judges playing no role. Here the exact opposite occurs. Through opinions, judges drive the future case selection. The judicially-constructed case selection, then, determines what the judge learns and, as a result, the law the judge eventually issues.

Finally, economists have searched for a theory to explain why a common law legal origin positively correlates with development. To understand this fact, law and finance scholars needed a model of judicial behavior. They faced the same problem as Posner did. If, in fact, the common law is efficient, why is that so? Nicola Gennaioli and Andrea Shleifer (2007) show the importance of distinguishing cases for the efficiency of legal rules. Judges distinguish prior cases to reach preferred outcomes. To distinguish, the judge must consider another dimension to the legal problem. The introduction of this extra dimensionality embeds new information into the law. Like in the political science models, judges don't learn in Gennaioli and Shleifer's model. Judges know the policy they want to reach, the issue is can they get there.

3 The Model

We model a court of last resort, the U.S. Supreme Court or, where the Supreme Court has not spoken to an issue, the federal circuit. We are not concerned with a court submitting to the opinions or precedent of a superior court. If it wants, the court can disregard its own previous opinions. Litigants on the other hand are bound by the decisions.

For concreteness, the model is presented in terms of the following question: what is patentable subject matter? This is a vexing question in patent law that courts continually struggle with (see, for example, In re Bilski, _F.3d_ (Fed. Cir. 2008)(en banc)). To receive a patent, an inventor must first show that his invention belongs to the class of patentable subject matter. Section 101 of the patent statute defines patentable subject matter as:

> Whoever invests or discovers any new and useful process, machine, manufacture, or composition of matter, or any new and useful improvement thereof, may obtain a patent therefor, subject to the conditions and requirements of this title.

In each case, the court decides whether the given innovation falls within section 101. Inventions are represented by a real number, x. One can think of x as the "facts" of the case, in our example, the characteristics of the innovation. If x_1 and x_2 are close together, the two innovations share many common features. As the distance between x_1 and x_2 grows, the two innovations share fewer and fewer commonalities, the facts of the cases become much different.

Inventions below a threshold θ are not patentable subject matter, while inventions above the threshold are patentable subject matter. At the outset, the court doesn't know θ, the threshold of patentable subject matter that best advances the interest of the patent statute. The court estimates this threshold by observing a series of cases.

The prior distribution of θ is $F(\theta)$ with positive density $f(\theta)$ over $[\underline{\theta}, \overline{\theta}]$ (with $\underline{\theta} \geq -\infty$, $\overline{\theta} \leq \infty$). If a case is brought to court and examined, the court determines whether the innovation is valid; it learns whether the innovation is above or below θ. Upon seeing a case, the court can issue an opinion.

An opinion consists of a lower bound, a, and an upper bound, b. What do these bounds mean? In an opinion, the court says that going forward innovations below a are invalid and innovations above b are valid. The bounds are the opinion's precedent, what the opinion stands for. When a and b are far apart, the opinion is narrow. It doesn't demarcate many innovations as either clearly patentable subject matter or clearly not patentable subject matter. Such an opinion provides little information to future litigants. In contrast, the narrower the bands, the broader the opinion. An opinion with narrow bands specifies many innovations as either "inside" or "outside" the patentable-subject matter box.

The law remains uncertain only for innovations inside the bounds of the opinion. For these innovations, the opinion has not stated whether the innovation is or is not patentable subject matter.

The choice of precedent bounds matters to the court. If the court sets the lower bound a too high, too many inventions will qualify as not patentable subject matter, hampering innovation. If the court sets the higher bound b too low, too many inventions will be patentable subject matter. And such an expansive subject-matter formulation might provide monopoly rents unnecessarily. Alternatively, the expansive formulation might create bottlenecks, whereby an inventor could claim as a patentable subject matter a broad principle and use the broad patent as leverage to stifle follow-on innovation.

Such concerns, for example, are found in the debate about patents that claim the "gene" involved in diagnostic testing for medical research (Merz et al. 2002). The argument is that such a patent might preclude a large swath of research into, say, a cure for breast-cancer. That is to say, a researcher wanting to conduct research on breast cancer would need a license from the holder of the diagnostic gene patent to even get started. The need to obtain such a license might then discourage innovation into a cure.

3.1 Players and Timing

Plaintiffs and defendants live for one period. The court lives an infinite number of periods. The timing of the game follows:

- At the start of period t, the plaintiff draws a case, x_t, from the distribution $G(x)$ with positive density $g(x)$ over $[\underline{x}, \overline{x}]$ with $\underline{x} \geq \underline{\theta}$, $\overline{x} \leq \overline{\theta}$.
- The plaintiff makes a settlement offer to the defendant. The settlement offer depends on the on the "law" available at period t, that is, the court's most recent opinion, $\{a_t, b_t\}$.
- The defendant accepts or rejects the settlement offer.
- If the defendant rejects the offer, the case goes to trial and up on appeal.
- If appealed, the court decides the case, updates its beliefs about θ, and issues an opinion, $\{a_{t+1}, b_{t+1}\}$.

3.2 The Litigation Game

The plaintiff sues to enjoin the defendant's infringing activity. As a defense, the defendant asserts that the plaintiff's patent is invalid because it lacks patentable subject matter. Since the plaintiff's patent is invalid, the defendant argues, there can be no infringement. For exposition purposes, think of the plaintiff as an incumbent firm and the defendant as a potential entrant. The patent is the gateway into the plaintiff's market. If the plaintiff loses the lawsuit, it will face entry and competition. If the defendant loses, it cannot enter and is left with its outside option. It is convenient to normalize the payoffs from losing the case to zero for both parties. The plaintiff's (net) benefit from obtaining injunction relief and preventing entry is B_P. The defendant's (net) benefit from winning the case, avoiding the injunction, and being able to enter the incumbent's market is B_D.

The plaintiff and the defendant incur costs from using the court system. The plaintiff's cost is C_P; the defendant's cost is C_D. After going to court, the plaintiff and defendant show the facts of the case to the court; that is, they show the value of x_t. The court then decides whether to investigate the case (at a cost C), or to declare the patent as valid or invalid without further investigation. As we shall see, it will always be optimal for the court to declare the patent invalid if $x_t < a_t$ and to declare it valid if $x_t > b_t$. It simply follows that if the plaintiff's case, x_t, falls below a_t, the plaintiff knows for sure that the court will declare its patent invalid. Thus, the plaintiff drops its case to avoid paying C_P. If the plaintiff's case lies above b_t, the plaintiff knows for sure that the court will declare its patent valid. Likewise, the defendant knows it will for sure lose at trial. As a result, the defendant is willing to pay C_D to avoid the court costs and will accept the defendant's offer to settle and not enter the defendant market.

Now suppose that the plaintiff has a case within the precedent bounds, $\{a_t, b_t\}$. Here, both parties must form expectations about what the court is likely to do.

To make the analysis as simple as possible, when the law is uncertain assume that the plaintiff can make one of two offers. First, it can either allow the defendant to enter and, in effect, concede its patent is invalid. Second, the plaintiff can assert that its patent is valid and push to forbid entry as part of any settlement. If the plaintiff caves and makes a settlement offer that allows entry, the defendant will, of course, accept. If the plaintiff's settlement offer forbids entry, the defendant may or may not accept the offer depending on what it thinks will happen on appeal.

Let $F(\theta|H_t)$ be the court's estimate of the posterior distribution over θ, which depends on the history H_t up to time t. The history is just the series of cases brought before time t. The probability that the court will rule in favor of the plaintiff when the case is x_t is $1 - F(x_t|H_t)$. Assume for concreteness that when indifferent parties prefer not to go to court. The defendant thus rejects the no-entry offer if

$$0 < B_D F(x_t|H_t) - C_D$$

or

$$F(x_t|H_t) > \frac{C_D}{B_D}$$

Understanding that the defendant will always accept the "entry" offer, the plaintiff will make such an offer only if the expected payoff from litigating is less than zero; that is

$$[1 - F(x_t|H_t)]B_P - C_P < 0$$

Rearranging gives

$$F(x_t|H_t) > \frac{B_P - C_P}{B_P}$$

Otherwise, the plaintiff will make the "no-entry" offer. The case goes to court if the plaintiff makes the no-entry offer and the defendant rejects that offer. Such a turn of events happens if

$$\frac{C_D}{B_D} < F(x_t|H_t) < \frac{B_P - C_P}{B_P}$$

Assume that C_D and C_P are very close to zero. Under these assumptions, the case goes to court for sure when $x_t \in (a_t, b_t)$. The upshot of this simple litigation game is that cases settle when the law is clearly defined and do not settle when the law is uncertain.[1] The clarity of the law rests in the hands of the court.

[1] This formulation of the settlement game precludes monetary transfers between the litigants when the law is uncertain. They can't split the difference, an obviously unrealistic simplification. The focus of this paper is judicial learning, not the decision between litigation and settlement. That decision is subject to a very large literature (for an overview, see Hay and Spier (1998). All we are trying to capture here is that cases where the law is uncertain are less likely to settle. For a model where the clarity of the law affects settlement behavior, see Easterbrook, Landes, and Posner (1981).

3.3 Judicial Learning and the Creation of Law

A. Learning Process

Assume that at date 1 there has been no precedent, $a_1 = \underline{\theta}$, $b_1 = \overline{\theta}$, and hence a case is brought to court; the court observes the case and issues an opinion. Over time, the court gradually learns about the optimal threshold value, θ. Suppose that the court hears a case with facts such that $x = 5$. After examination, the court learns that $\theta > 5$. Given this fact, the court also learns that the optimal threshold is bigger than 4 and 3 and so on. The court doesn't learn how much bigger than 5 the optimal threshold is.

More generally, as of time t, the court has observed a history or series of cases, H_t. The court wants to use information from that history to update its beliefs about the posterior distribution of the optimal threshold, $F(\theta|H_t)$.

Given the learning process described above, the only relevant history includes (1) the highest past innovation where the court declared the patent invalid and (2) the smallest past innovation where the court declared the patent valid. These two endpoints squeeze the court's estimate of the posterior distribution. The court knows that θ must lie somewhere between the highest past innovation found invalid and the lowest past case found valid.

Denote the highest past innovation the court declared invalid, W_t and the lowest past innovation the court declared valid, R_t. Formally,

$$\begin{aligned} W_t &= \max_{\tau<t}\{x_\tau : x_\tau \in [a_\tau, b_\tau], x_\tau < \theta\} \\ R_t &= \min_{\tau<t}\{x_\tau : x_\tau \in [a_\tau, b_\tau], x_\tau \geq \theta\} \end{aligned}$$

For $x \in [W_t, R_t]$, the court's updated distribution is

$$F(x|H_t) = F(x|W_t, R_t) = \frac{F(x) - F(W_t)}{F(R_t) - F(W_t)}$$

while

$$\begin{aligned} F(x|W_t, R_t) &= 0 \text{ if } x \leq W_t \\ F(x|W_t, R_t) &= 1 \text{ if } x \geq R_t \end{aligned}$$

The posterior distribution captures the well-known lawyer view on precedent. Relying on past decisions provides information, saving the court time. Assuming the prior judgments were correct, the court can take those rulings as given and focus on "new" issues.[2] The posterior formally summarizes Judge Benjamin Cardozo's intuition that "the labor of judges would be increased almost to the breaking point if every past decision could be reopened in every case, and one could not lay one's own course of brick on the secure foundation of the courses laid by others who had gone before him." (Cardozo (1921), p. 249).

B. The Court's Objective Function

The court values two things. First, the court cares about time and effort spent evaluating cases, denoted by C. Second, the court cares about errors. Errors happen when a case settles out of court in the "wrong" direction, given the court's estimate of the posterior distribution, $F(\theta|H_t)$. An error occurs if a case that should have gone in favor of the plaintiff is instead dropped, or if a case that the defendant should have won is settled in favor of the plaintiff. It may seem natural to assume that $L > C$, the loss of an error is greater than the cost of going to court; however, in the formal analysis we do not need to make such an assumption.

Suppose, for example, that the court has learned from past cases that $\theta \geq 5$. All innovation below 5 lacks patentable subject matter and so the plaintiff should lose. The court doesn't know what should happen with cases greater than 5. The court might extrapolate that a plaintiff with, say, case 6 should also lose. After all, the innovations in case 5 and case 6 share many features – the facts are similar. The court may set the lower bound, $a_t = 7$. But the court could be wrong; θ might just equal 5. There is a chance that the plaintiff will draw case 6. Given the law, the plaintiff anticipates that the court will declare its patent invalid without investigating the case (even though the court would declare the patent valid if it knew the true value of the validity threshold!). So, the plaintiff caves and drops his case, resulting in an error.

[2] The assumption that the court never makes mistakes in determining where the case x_t lies relative to the threshold θ is just a convenient simplification. Allowing for mistakes by the court wouldn't affect the main results over the long term, so long as the average decision was more right than wrong. Previous judicial errors would cancel out, enabling the court to extract important information relevant to updating from the prior cases.

Why would the court let an error occur? Shouldn't the court use all the information it has in each period, setting the precedent bounds, $\{a_t, b_t\} = \{W_t, R_t\}$? True, a "narrow" decision like this one looks like a good idea. It isn't, not always. Wider bounds mean that the chance of having to decide a case in the following period is higher. The court's opinion is less informative. And so, litigation may happen too much.

When setting the precedent bounds in each period, the court trades off the cost of having to decide a case next period against errors that accrue from an inaccurate decision. These single period trade-offs are embedded in a dynamic model, however. The dynamic model means that there is another benefit from a narrow decision: enhanced learning in the future periods. By rendering a cautious decision in period one, the court preserves a stream of future cases and the anticipated learning that case load provides. This dynamic benefit – the learning benefit – allows the court to change and correct course over the long run.

To formalize these intuitions, let δ be the discount factor and let $V(H_t)$ be the court's value function at time t, as a function of the history H_t. At the end of period $t-1$ the court chooses a_t, b_t subject to $W_t \leq a_t \leq b_t \leq R_t$ to maximize its expected payoff:

$$\begin{aligned} V(H_t) &= V(W_t, R_t) = \max_{W_t \leq \widehat{a}_t \leq \widehat{b}_t \leq R_t} \Big\{ -C \left[G(\widehat{b}_t) - G(\widehat{a}_t) \right] \\ &\quad -L \int_{W_t}^{\widehat{a}_t} \int_{W_t}^{x_t} g(x_t) \frac{f(\theta)}{F(R_t) - F(W_t)} d\theta dx_t \\ &\quad -L \int_{\widehat{b}_t}^{R_t} \int_{x_t}^{R_t} g(x_t) \frac{f(\theta)}{F(R_t) - F(W_t)} d\theta dx_t + \delta E_t V(H_{t+1}) \Big\} \end{aligned} \tag{1}$$

The first term in (1) is the cost of having to decide a case in the next period. To see why, suppose the court sets $a_t = b_t$. Such a choice defines the law for all possible innovations. As a result, all cases settle. The court observes no cases at time t and incurs no decision costs. The greater the distance between a_t and b_t, the greater the chance the plaintiff draws a case where the law is unsettled. The second and third terms in (1) reflect the expected one-period error costs. Consider the second term. If the court sets $a_t \geq W_t$, there is a chance the plaintiff will draw a case between W_t and a_t. Given the opinion's lower bound, this case will settle in favor of the defendant. The probability measure inside the integral is the court's expectation that this case should have gone in favor of the plaintiff instead.

The third term follows from a similar analysis on the upper bound of the decision. The fourth term is the expectation about the future value of the court's objective function, given its opinion choices today. This last term captures the dynamic learning considerations described above.

Note that if the plaintiff brings to court a case $x_t \in (W_t, a_t)$, then (1) implies that the court will find it optimal to declare the patent invalid without investigating the case (i.e., without incurring the cost C). Similarly, if the plaintiff brings to court a case $x_t \in (b_t, R_t)$, then the court will find it optimal to declare the patent valid without investigating the case.

4 Results

We start by looking at the evolution of law. Convergence happens when $a_t = b_t$. There are two possibilities.

The first possibility is that the law converges perfectly. The court continues to take cases until it learns the true value of θ. Alternatively, the law converges imperfectly. Here, the court defines the law for all possible innovations without uncovering the exact value of θ. More formally, the law converges imperfectly when the court sets a_t equal to b_t, but R_t, the upper bound on the estimate of θ, differs from W_t, the lower bound on the estimate. In that case, the court doesn't get the law right and it knows it. Claims by legal scholars about imperfections in the law should fall on deaf ears. Indeed, by ignoring the scholarly concerns, the court acts efficiently.

Absent from the model are all the familiar arguments why a court sticks with a less than ideal doctrine. The court bears no cost for deviating from precedent. The court doesn't care about the "predictability" of legal rules. There is no reliance on law by outsiders. We stack the deck against convergence and still over time the law becomes more and more predictable.

Whether the eventual rule is perfect or imperfect depends on a trade-off between the cost of deciding a case and the errors induced by an imperfect rule. When error costs are small relative to decision-costs, the law settles imperfectly. In fact, when error costs are small relative to decision-costs a myopic court that only looks at the current payoff would not consider any cases. It is only because of learning and potential future benefits from learning that the court may hear cases and the law evolves. Eventually, however, the benefit of learning will vanish and the law will settle on an imperfect rule, one that does not eliminate all errors.

On the other hand, when error costs are large relative to the decision costs, the law settles on a perfect rule, one that admits no errors. Even when the law settles perfectly, however, the court will not want to learn at maximum speed. Learning at maximum speed would involve setting the narrowest precedent bounds $a_t = W_t$ and $b_t = R_t$. At the margin, the cost saving impact of reducing the learning speed by a small amount is always valuable to the court. Some cases will not be brought to court, even though the court does not know how they should be decided. The following proposition formalizes these results.

Proposition 1 *The law always converges, but the court never learns at maximum speed: $W_t < R_t$ implies $a_t > W_t$ and $b_t < R_t$. There are two possible outcomes of the learning process. If $L > 2C$, the law converges to a perfect rule that admits no error. If $L < 2C$, the law converges to an imperfect rule that does not eliminate all errors.*

Proof: See the appendix.

Surprisingly, a court that places an intrinsic value on the "rule of law" is not necessary for the rule of law to happen. A court that learns will eventually articulate predictable and stable legal rules. Indeed, the "rule of law", which sits as a core value of the American legal system, can itself be seen as a fortunate by-product of scarce judicial resources. Scarcer resources means that the court promulgates settled, but imperfect, law more quickly.

With this framework, we can also identify the factors that influence the speed with which the court learns. Commentators often make disparaging remarks about the court's failure to update doctrine in light of changed circumstances (add cites). This is a common critique in patent law, where technology advances at a rapid pace (add cites). Indeed, it would be surprising if patent law could keep up. Interpret θ as the optimal subject-matter rule given the technology available in that period. In other words, let θ be a meta-rule, incorporating changes in technology.

Courts learn quickly when decision costs are relatively small. Then, the court feels comfortable articulating a narrow decision in the first period (a large value of $b_2 - a_2$). Such a decision means that a case comes before the court in the following period with a high probability. Since decision costs are small, the court doesn't mind examining this next case. It places a greater value on improving the doctrine based on the information that case brings to the court's attention. Repeating this dynamic, the court learns quickly by issuing a series of narrow opinions that effectively fish for future cases.

Likewise, high losses from errors push the court to issue narrow opinions. Given large losses from errors, the court wants to minimize their occurrence; it renders an opinion that will allow it to investigate many uncertain cases in the future. The bounds of the opinion approach the court's estimate of the posterior distribution for that period. And so, the court again sees – and learns from – a case in the next period with a high probability. The court fishes for future cases here because it fears making a ruling that results in errors in the current period.

As the court cares more about future outcomes, it issues narrower opinions. Narrow opinions increase the probability of hearing a case in the next period, hence they accelerate learning and reduce the instances of future cases heard and future errors.

The next proposition captures these three ideas.

Proposition 2 *(A) As cases become less costly, learning speeds up and precedent bounds get closer to what the court estimates are the law's true bounds (i.e., a_t is increasing in C and b_t is decreasing in C). (B) As losses from bad decisions increase, learning becomes faster, and precedent bounds get closer to what the court estimates are the law's true bounds (i.e., a_t is decreasing in L and b_t is increasing in L). (C) As the court cares more about the future, learning speeds up and precedent bounds become broader (i.e., a_t is decreasing in δ and b_t is increasing in δ).*

Proof: See Appendix.

The next question is whether the order in which cases appear before the court influences the eventual legal rule. The doctrine of stare decisis means that previous decisions bind the court. Stated differently, the court faces a cost from deviating from case law. The precedent constraint locks a court in along one trajectory or path. So, the conventional argument goes, case order matters (Hathaway 2001). And then, yes, control of case order can be important. By targeting cases to come before the court in a certain order, litigants can achieve the legal rule they most prefer (Stern 1995; Galanter 1974). The power of the repeat litigant is analogous to the power of the agenda setter articulated in models of legislation (Romer and Rosenthal 1978, for an overview of this literature, see Krehbiel 2004).

In our model, prior decisions aren't precedent in the normal sense of the word; the court pays no price for deviating from its own case law. The court learns from previous decisions because it believes those cases were rightly decided. Without a cost of deviating from prior case law, can law still be path dependent? Yes. This is the first novel twist on path dependence of our analysis. A cost to deviating from precedent is not necessary for path dependence. If the cost of case examination is sufficiently high, a law's resting point will always depend on case order.

A second twist involves the repeat litigant. Court learning dampens the agenda setting power of the repeat litigant. From Proposition 1, the repeat litigant's power can only be unleashed when the court expects small losses from errors relative to the cost of investigating a case.[3] If these error losses are large relative to the investigating costs, the court will step up and refine the doctrine. More formally, the following is true.

Proposition 3 *Law displays limited path dependence in the long run. The only time long-run path dependence arises is when the losses associated with the "wrong" settlement outside the precedent bounds are small as compared to the decision costs,* $L < 2C$.

Proof: See Appendix.

Having considered convergence, path dependence, and the factors which influence the speed of learning, we now turn to what happens on the way. The path to a stable outcome can take two forms: coherent and incoherent.

On a coherent path, opinion bounds narrow and never expand ($a_t > a_{t-1}$ and $b_t < b_{t-1}$ for all t). An example is the Supreme Court's approach to patent prosecution estoppel.

During prosecution, applicants often narrow the patent's claim at the request of the examiner. The interaction between this narrowing and the so-called doctrine of equivalents is complicated. Under the doctrine of equivalents, the patent monopoly extends beyond inventions that literally infringe the patent claim to include inventions that are "equivalent" to literally infringing inventions.

A series of cases presented the following question: Can a patent-holder assert, under the doctrine of equivalents, a claim he jettisoned during the

[3] The investigating cost C may also include the cost of changing the precedent bounds; when the court investigates a case it will change the precedent bounds with probability one.

prosecution process? In Warner Jenkinson Co., Inc. v. Hilton Davis Chemical Co.[4], Justice Thomas said no, but added a caveat. The jettisoned claim can not be asserted so long as the modification relates to "patentability," – a term left undefined. The Court filled in the blanks in Festo Corp. v. Shoketsu Kinzoku Kogyo Kabushiki Co, Ltd.[5] There, Justice Kennedy stated:

> A patentee's decision to narrow his claims through amendment may be presumed to be a general disclaimer of the territory between the original claim and the amended claim. . . . There are some cases, however, where the amendment cannot reasonably be viewed as surrendering a particular equivalent. The equivalent may have been unforeseeable at the time of the application; the rationale underlying the amendment may bear no more than a tangential relation to the equivalent in question; or there may be some other reason suggesting that the patentee could not reasonably be expected to have described the insubstantial substitute in question. In those cases the patentee can overcome the presumption that prosecution history estoppel bars a finding of equivalence. Id. at 740.

Festo puts gloss on the meaning of patentablity. In terms of the model, the court defined the kind of modifications made during prosecution that then bar arguments based on the doctrine of equivalents. And more narrowing is likely to come, as the Court defines what it means for an equivalent to be "unforeseeable" and what criteria will be used to decide whether an amendment bears "tangential relation" to an equivalent.

More interesting, a court might issue incoherent doctrine. On this path, the opinion bounds bounce around. Sometimes $a_t > a_{t-1}$ and $b_t < b_{t-1}$; other times $a_t < a_{t-1}$ and $b_t > b_{t-1}$. The doctrine doesn't follow a clear pattern. Despite this fact, the court is taking the right jurisprudential approach. An abstract example demonstrates why this might happen:

Suppose that $[W_t, R_t]$ is $[-100, +100]$ and the court sets $[a_t, b_t] = [-50, 50]$, splitting the difference on its estimate of the posterior to save decision-costs. The next case is brought to court, say that case is $x = 49$. The court finds out that case is not valid. The updated posterior is $[W_{t+1}, R_{t+1}] = [49, 100]$. Given these beliefs, the court now faces the opinion choice. To narrow the

[4] 520 U.S. 17 (1997).
[5] 535 U.S. 722 (2001).

choice on the upper bound would involve setting $b_{t+1} = 49$. The court would have to define the law for all innovations. Such a bold move makes little sense, given the error rate and the corresponding losses with cases $[50, 100]$. So, instead the court issues an incoherent opinion, setting b_{t+1} at, say, 75. What looks like inconsistent decision-making is really the court optimally gathering and then using information from the cases.

Examples of inconsistent doctrine are sprinkled throughout patent law. Parker v. Flook[6] and Diamond v. Diehr[7] are illustrative. Both cases addressed whether computer software counts as patentable subject matter. In Parker v. Flook, the court held no, reasoning that software was akin to a mathematical formula and thus not patentable subject matter. In Diamond v. Diehr, the court answered yes. In so doing, the court did not overrule but instead distinguished Flook. The dissent pointed out that the Court's distinguishing efforts were strained, stating:

> The essence of the claimed discovery in both [Flook and Diehr] was an algorithm that could be programmed on a digital computer. . . . In Flook, the algorithm made use of multiple process variables; in this case, it makes use of only one. In Flook, the algorithm was expressed in a newly developed mathematical formula; in this case, the algorithm makes use of a well-known mathematical formula. Manifestly, neither of these differences can explain today's holding. Id. at 210

Between 1978 and 1981 – the time of the Flook and Diehr decisions – the law with respect to the patentability of computer software was incoherent, as reflected by the variety of interpretations of Flook given by the lower courts. It is now settled. The initial inconsistency could have reflected learning by the court. The computer software industry was just starting when these cases came up. As a result, the Court confronted the issue for the first time. In the model, doctrinal incoherence is most likely to happen under these conditions. The last proposition states this result formally.

Proposition 4 *Given judicial learning, the law can display optimal incoherence. Such incoherence is much more likely to occur when the court first confronts an issue.*

---

[6] 437 U.S. 584 (1978).
[7] 450 U.S. 175 (1981).

Proof: See Appendix.

Luck plays a significant role in the model. If the first case happens to be close to the optimal threshold, the court learns quickly. After the first case, the court can summarily rule lots of innovations as either patentable or not patentable subject matter. If, on the other hand, the first case happens to be far away from the optimal threshold, the court may take a long time to reach the stable outcome.

Since luck partly determines how fast the law converges, different doctrines will converge at different rates, depending on the amount of information conveyed by the first few cases considered. Some doctrines bounce around, while other doctrines reach a stable position quickly. The unpredictable doctrine doesn't necessarily involve difficult legal issues or more controversial judgments. Instead, what is being observed is the lingering effect of randomness in the initial case selection.

5 Conclusions

This paper has identified how judges can use opinion-writing to learn and what learning means for the creation of judge-made law. The path of judicial learning maps onto the way doctrines, specifically patent doctrine, actually evolve. As with any model, we have made simplifying assumptions here. First, we model a court, not individual judges. Judges on a court might have different overarching policy goals and hence learn different things from the cases considered. While we find this unlikely in patent law, it is certainly true for hot button issues, like abortion, affirmative action, privacy, and the constitutionality of the death penalty. Second, the model views opinion writing through the lens of learning. A narrow opinion generates leads to more learning than a broad opinion. Other factors also dictate the contours of an opinion, such as the need to garner a majority of the votes of the court (Cameron and Kornhauser 2008; Carrubba et al. 2008). That said, the model could be easily interpreted as the behavior of the median justice or judge, the judge whose vote controls the outcome. This swing judge can use her opinions to learn about her optimal policy. This assumes, of course, that the swing justice or judge anticipates holding the pivot position in the future.

Finally, a debate rages between judicial minimalists like Cass Sunstein, who advocate that the court makes law, especially constitutional law, in small incremental steps (Sunstein 1999) and a subset of originalist scholars and judges who believe the court owes fidelity to the plain meaning of the constitutional text, even if that reading requires a dramatic change in the law. Our model sheds formal light on this debate. Small steps are not always called for or prudent, especially when a court first considers an issue. On the other hand, the originalist's approach – big swings in doctrine followed by broad rulings – can also be sub-optimal. The reason: Such an approach blocks the court from hearing and learning from a big swath of future cases. In fact, a court acting optimally will sometimes be a minimalist and other times issue broad rulings. Even on the same issue, a judge's jurisprudential approach should change over time. As they learn more and more, smaller and smaller refinements, a more minimalist approach, becomes attractive. Furthermore, we shouldn't expect a judge to have the same approach to jurisprudence or judging independent of the area of law. The correct approach to the creation of law depends on the relationship between the cost of decisionmaking and the loss from errors. There is no reason to suspect this relationship will be similar across pockets of law.

Appendix

Proof of Proposition 1. First, we need to write the formula for the expectation $E_tV(H_{t+1})$:

$$\begin{aligned} E_tV(H_{t+1}) &= V(W_t,R_t)\left[1-G(b_t)+G(a_t)\right] \\ &+\int_{a_t}^{b_t}\int_{x_t}^{R_t}V(x_t,R_t)g(x_t)\frac{f(\theta)}{F(R_t)-F(W_t)}d\theta dx_t \\ &+\int_{a_t}^{b_t}\int_{W_t}^{x_t}V(W_t,x_t)g(x_t)\frac{f(\theta)}{F(R_t)-F(W_t)}d\theta dx_t \end{aligned}$$

The following first order conditions are obtained by differentiating (1) with respect to a_t and b_t, with multiplier λ_t associated to the constraint $a_t \leq b_t$, and multipliers μ_t^a, μ_t^b associated to constraints $W_t \leq a_t$ and $b_t \leq R_t$:

$$\begin{aligned} Cg(a_t) - L\int_{W_t}^{a_t} g(a_t)\frac{f(\theta)}{F(R_t)-F(W_t)}d\theta + \delta V(W_t,R_t)g(a_t) \\ -\delta\int_{a_t}^{R_t}V(a_t,R_t)g(a_t)\frac{f(\theta)}{F(R_t)-F(W_t)}d\theta \\ -\delta\int_{W_t}^{a_t}V(W_t,a_t)g(a_t)\frac{f(\theta)}{F(R_t)-F(W_t)}d\theta - \lambda_t + \mu_t^a = 0 \end{aligned}$$

$$\begin{aligned} -Cg(b_t) + L\int_{b_t}^{R_t} g(b_t)\frac{f(\theta)}{F(R_t)-F(W_t)}d\theta - \delta V(W_t,R_t)g(b_t) \\ +\delta\int_{b_t}^{R_t}V(b_t,R_t)g(b_t)\frac{f(\theta)}{F(R_t)-F(W_t)}d\theta \\ +\delta\int_{W_t}^{b_t}V(W_t,b_t)g(b_t)\frac{f(\theta)}{F(R_t)-F(W_t)}d\theta + \lambda_t - \mu_t^b = 0. \end{aligned}$$

They can be rewritten as

$$\begin{aligned} 0 &= C - L\frac{F(a_t)-F(W_t)}{F(R_t)-F(W_t)} + \delta V(W_t,R_t) \qquad (2) \\ &-\delta V(a_t,R_t)\frac{F(R_t)-F(a_t)}{F(R_t)-F(W_t)} - \delta V(W_t,a_t)\frac{F(a_t)-F(W_t)}{F(R_t)-F(W_t)} - \frac{(\lambda_t-\mu_t^a)}{g(a_t)} \end{aligned}$$

$$0 = -C + L\frac{F(R_t) - F(b_t)}{F(R_t) - F(W_t)} - \delta V(W_t, R_t) \quad (3)$$
$$+\delta V(b_t, R_t)\frac{F(R_t) - F(b_t)}{F(R_t) - F(W_t)} + \delta V(W_t, b_t)\frac{F(b_t) - F(W_t)}{F(R_t) - F(W_t)} + \frac{\left(\lambda_t - \mu_t^b\right)}{g(b_t)}$$

Suppose $a_t = W_t < R_t$. Then (2) becomes $C = \frac{(\lambda_t - \mu_t^a)}{g(a_t)}$, which can only be satisfied if $b_t = a_t$. Equation (3) becomes $C - L = \frac{\left(\lambda_t - \mu_t^b\right)}{g(b_t)}$, which can only be satisfied if $b_t = R_t$. This is contradiction. Hence it must be $a_t > W_t$ whenever $W_t < R_t$.

Similarly, suppose $b_t = R_t > W_t$. Then (3) becomes $C = \frac{\left(\lambda_t - \mu_t^b\right)}{g(a_t)}$, which requires $a_t = b_t$, while (2) becomes $C - L = \frac{(\lambda_t - \mu_t^a)}{g(a_t)}$, which requires $a_t = W_t$, a contradiction. Hence it must be $b_t < R_t$ whenever $R_t > W_t$. This concludes the proof of the first part of the proposition.

Now consider whether there will be full learning. For full learning to take place in the limit (i.e., $a_t \to \theta$, $b_t \to \theta$) it must be $a_t \neq b_t$ whenever $W_t \neq R_t$. Assume $a_t = b_t$ and $W_t \neq R_t$. Since, as we have just shown, it is $\mu_t^a = \mu_t^b = 0$, adding up (2) and (3) we obtain

$$F(a_t) = F(b_t) = \frac{F(R_t) + F(W_t)}{2} \quad (4)$$

Replacing such value, and $\mu_t^a = \mu_t^b = 0$, into (2) and (3) yields

$$2C - L + 2\delta V(W_t, R_t) - \delta V(a_t, R_t) - \delta V(W_t, a_t) - \frac{2\lambda_t}{g(a_t)} = 0 \quad (5)$$

$$-2C + L - 2\delta V(W_t, R_t) + \delta V(b_t, R_t) + \delta V(W_t, b_t) + \frac{2\lambda_t}{g(b_t)} = 0. \quad (6)$$

Since $V(W_t, R_t) < V(a_t, R_t)$ and $V(W_t, R_t) < V(W_t, a_t)$, the conditions above cannot be satisfied if $2C < L$. It follows that it cannot be $a_t = b_t$, and hence learning will never stop if $2C < L$.

Now suppose $a_t \neq b_t$ for all values of $W_t < R_t$, and as a consequnce,

$\lambda_t = 0$. Using this and $\mu_t^a = \mu_t^b = 0$, by subtracting (3) from (2) we obtain

$$\begin{aligned} 0 &= 2C - L\left(1 - \frac{F(b_t) - F(a_t)}{F(R_t) - F(W_t)}\right) + 2\delta V(W_t, R_t) \\ &\quad -\delta V(a_t, R_t)\frac{F(R_t) - F(a_t)}{F(R_t) - F(W_t)} - \delta V(b_t, R_t)\frac{F(R_t) - F(b_t)}{F(R_t) - F(W_t)} \\ &\quad -\delta V(W_t, a_t)\frac{F(a_t) - F(W_t)}{F(R_t) - F(W_t)} - \delta V(W_t, b_t)\frac{F(b_t) - F(W_t)}{F(R_t) - F(W_t)} \end{aligned} \tag{7}$$

Note that for any $\varepsilon > 0$, there exists W_t and R_t sufficiently close to each other, so that the left hand side of (7) is greater than $2C - L + \varepsilon$. It follows that if $2C > L$, then (7) cannot hold for such values of W_t and R_t; hence it cannot be the case that $a_t \neq b_t$ for all values of $W_t < R_t$. Learning will eventually stop if $2C > L$. ■

Proof of Proposition 2. At an interior solution, the first order condition (2) and (3) can be rewritten as

$$\begin{aligned} \Phi_a \equiv C - L\frac{F(a_t) - F(W_t)}{F(R_t) - F(W_t)} + \delta V(W_t, R_t) - \delta V(a_t, R_t)\frac{F(R_t) - F(a_t)}{F(R_t) - F(W_t)} \\ - \delta V(W_t, a_t)\frac{F(a_t) - F(W_t)}{F(R_t) - F(W_t)} = 0 \end{aligned}$$

$$\begin{aligned} \Phi_b \equiv -C + L\frac{F(R_t) - F(b_t)}{F(R_t) - F(W_t)} - \delta V(W_t, R_t) + \delta V(b_t, R_t)\frac{F(R_t) - F(b_t)}{F(R_t) - F(W_t)} \\ + \delta V(W_t, b_t)\frac{F(b_t) - F(W_t)}{F(R_t) - F(W_t)} = 0 \end{aligned}$$

Totally differentiating with respect to a_t, b_t and the parameter z, $z \in \{C, L, \delta\}$, yields

$$\begin{aligned} \frac{\partial \Phi_a}{\partial a_t}\partial a_t + \frac{\partial \Phi_a}{\partial b_t}\partial b_t + \frac{\partial \Phi_a}{\partial z}\partial z &= 0 \\ \frac{\partial \Phi_b}{\partial a_t}\partial a_t + \frac{\partial \Phi_b}{\partial b_t}\partial b_t + \frac{\partial \Phi_b}{\partial z}\partial z &= 0 \end{aligned}$$

Using Cramer's rule and the second order conditions we obtain:

$$sign\left(\frac{\partial a_t}{\partial z}\right) = sign\left(-\frac{\partial \Phi_a}{\partial z}\frac{\partial \Phi_b}{\partial b_t} + \frac{\partial \Phi_b}{\partial z}\frac{\partial \Phi_a}{\partial b_t}\right)$$

$$sign\left(\frac{\partial b_t}{\partial C}\right) = sign\left(-\frac{\partial \Phi_b}{\partial z}\frac{\partial \Phi_a}{\partial a_t} + \frac{\partial \Phi_a}{\partial z}\frac{\partial \Phi_b}{\partial a_t}\right)$$

First note that $\frac{\partial \Phi_a}{\partial b_t} = \frac{\partial \Phi_b}{\partial a_t} = 0$. Then use the second order conditions $\frac{\partial \Phi_a}{\partial a_t} < 0$ and $\frac{\partial \Phi_b}{\partial b_t} < 0$ to write the comparative statics signs as:

$$\begin{aligned} sign\left(\frac{\partial a_t}{\partial z}\right) &= sign\left(\frac{\partial \Phi_a}{\partial z}\right) \\ sign\left(\frac{\partial b_t}{\partial z}\right) &= sign\left(\frac{\partial \Phi_b}{\partial z}\right) \end{aligned}$$

The proposition follows from the following conditions:

$$\begin{aligned} \frac{\partial \Phi_a}{\partial C} &> 0; \qquad \frac{\partial \Phi_b}{\partial C} < 0 \\ \frac{\partial \Phi_a}{\partial L} &< 0; \qquad \frac{\partial \Phi_b}{\partial L} > 0 \\ \frac{\partial \Phi_a}{\partial \delta} &< 0; \qquad \frac{\partial \Phi_b}{\partial \delta} > 0 \end{aligned}$$

■

Proof of Proposition 3: By Proposition 1, there cannot be any path dependence when $2C < L$; in such a case full learning takes place and $\lim_{t\to\infty} a_t = \lim_{t\to\infty} b_t = \theta$. Suppose instead $L < 2C$. To see that path dependence may occur, assume that θ and x are both drawn from the uniform distribution over the interval $[-M, M]$ (all values of θ and x are equally likely). Assume the true value of θ is $\theta = 0$.

From Proposition 1, we know that learning will stop. Suppose $a_t = b_t$ and $W_t \neq R_t$, so that a steady state is reached when no more learning takes place. Then, it is $V(W_t, R_t) = V(H_{t+1})$ and

$$
\begin{aligned}
V(W_t, R_t) &= -\frac{L}{1-\delta} \max_{W_t \le \widehat{a}_t \le \widehat{b}_t \le R_t} \left\{ \int_{W_t}^{\widehat{a}_t} \frac{1}{2M} \frac{x_t - W_t}{R_t - W_t} dx_t \right. \\
&\quad \left. + \int_{\widehat{b}_t}^{R_t} \frac{1}{2M} \frac{R_t - x_t}{R_t - W_t} dx_t \right\} \\
&= -\frac{L}{1-\delta} \left\{ \frac{(a_t^2 - W_t^2) - 2W_t (a_t - W_t)}{4M (R_t - W_t)} \right. \\
&\quad \left. + \frac{2R_t (R_t - b_t) - (R_t^2 - b_t^2)}{4M (R_t - W_t)} \right\} \\
&= -\frac{L}{1-\delta} \left\{ \frac{(a_t - W_t)^2 + (R_t - b_t)^2}{4M (R_t - W_t)} \right\}
\end{aligned}
$$

Recalling that adding up the foc's gives (4), we have

$$
F(a_t) = F(b_t) = \frac{a_t + M}{2M} = \frac{W_t + R_t + 2M}{4M}
$$

or

$$
a_t = b_t = \frac{W_t + R_t}{2}
$$

which gives

$$
V(W_t, R_t) = -\frac{L (R_t - W_t)}{8 (1-\delta) M}
$$

Replacing such value into (5) and (6) yields the first order conditions

$$
2C - L - \frac{\delta L (R_t - W_t)}{4 (1-\delta) M} + \frac{\delta L (R_t - W_t)}{8 (1-\delta) M} - \frac{2\lambda_t}{g(a_t)} = 0
$$

$$
-2C + L + \frac{\delta L (R_t - W_t)}{4 (1-\delta) M} - \frac{\delta L (R_t - W_t)}{8 (1-\delta) M} + \frac{2\lambda_t}{g(b_t)} = 0.
$$

Hence any value satisfying the following inequality is consistent with a steady state:

$$
2C - L - \frac{\delta L (R_t - W_t)}{8 (1-\delta) M} \ge 0
$$

which can be written as

$$
R_t - W_t \le \frac{8 (2C - L) (1-\delta) M}{\delta L}
$$

Take, for example, $C = 1.5$, $L = 1$, $M = 9$, $\delta = .9$. Then the inequality becomes:

$$R_t - W_t \leq 16$$

It follows that if the first case drawn at $t = 1$ is $x_1 = 3$, then $W_2 = -9$, $R_2 = 3$ and $a_2 = b_2 = -3$. If, on the other hand, the first case drawn is $x_1 = -3$, then $W_2 = -3$, $R_2 = 9$ and $a_2 = b_2 = 3$. (Note: if the first case drawn is less than -7, or higher than 7, then learning will continue past $t = 2$). ■

Proof of Proposition 4. As in the proof of Proposition 3, assume that θ and x are both drawn from the uniform distribution over the interval $[-M, M]$ and that the true value of θ is $\theta = 0$. Then it is:

$$\begin{aligned}
V(H_t) &= V(W_t, R_t) = \max_{W_t \leq \widehat{a}_t \leq \widehat{b}_t \leq R_t} \Big\{ -C \left[G(\widehat{b}_t) - G(\widehat{a}_t) \right] \\
&\quad -L \int_{W_t}^{\widehat{a}_t} \int_{W_t}^{x_t} g(x_t) \frac{f(\theta)}{F(R_t) - F(W_t)} d\theta dx_t \\
&\quad -L \int_{\widehat{b}_t}^{R_t} \int_{x_t}^{R_t} g(x_t) \frac{f(\theta)}{F(R_t) - F(W_t)} d\theta dx_t + \delta E_t V(H_{t+1}) \Big\} \\
&= \max_{W_t \leq \widehat{a}_t \leq \widehat{b}_t \leq R_t} \left\{ -C \frac{\widehat{b}_t - \widehat{a}_t}{2M} - L \left\{ \frac{(\widehat{a}_t - W_t)^2 + \left(R_t - \widehat{b}_t\right)^2}{4M\,(R_t - W_t)} \right\} - L\delta E_t V(H_{t+1}) \right\}
\end{aligned}$$

Differentiating with respect to $\widehat{b}_t$ (and assuming an interior solution) yields:

$$\begin{aligned}
0 &= -\frac{C}{2M} + \frac{L\,(R_t - b_t)}{2M(R_t - W_t)} - L\delta \frac{\partial E_t V(H_{t+1})}{\partial b_t} \\
b_t &= R_t - \frac{C(R_t - W_t)}{L} - 2M(R_t - W_t)\delta \frac{\partial E_t V(H_{t+1})}{\partial b_t}
\end{aligned}$$

Since $\frac{\partial E_t V(H_{t+1})}{\partial b_t} \geq 0$, this gives

$$b_t \leq R_t - \frac{C(R_t - W_t)}{L}$$

Assume $M = 9$, $C = 1$, and $L = 5$, so that learning eventually converges, $\lim_{t\to\infty} a_t = \lim_{t\to\infty} b_t = 0$. Suppose the first case drawn is $x_1 = 1$; then $W_2 = -9$, $R_2 = 1$ and $b_2 \leq 1 - \frac{10}{5} = -1$. Eventually $b_t > b_2$, and hence the law displays optimal incoherence. ■

References

Bailey, Martin and Rubin, Paul. 1994. "A Positive Theory of Legal Change." 14 *International Review of Law and Economics* 467-77

Cameron, Charles A. and Kornhauser, Lewis A. 2008. "Modeling Collegial Courts (3): Judicial Objectives, Opinion Content, Voting and Adjudication Equilibria," *NYU Law and Economics Working Paper* 08-54

Cardozo, Benjamin. 1921 The Nature of the Judicial Process.

Carrubba, Cliff, Friedman, Barry, Martin, Andrew and Vanberg, George. 2008. "Does the Median Justice Control the Content of Supreme Court Opinions?," *mimeo* New York University School of Law.

Ducur, Philippe. 1996. "New Drug Discovery Technologies and Patents," 22 *Rutgers Computer and Technology Law Journal* 369-477.

Easterbrook, Frank A. 1982. "Ways of Criticizing the Court." 95 *Harvard Law Review* 802-832.

Easterbrook, Frank A. , Landes, William and Posner, Richard A. 1980. "Contributions Among Antitrust Defendants: A Legal and Economic Analysis." 23 *Journal of Law and Economics* 331-__.

Galanter, Marc. 1974. "Why the 'Haves' Come Out Ahead: Speculations on the Limits of Legal Change," 9 *Law & Society Review* 95-160.

Gennaioli, Nicola and Shleifer, Andrei. 2007. "The Evolution of the Common Law," 115 *Journal of Political Economy* 43-68.

Hadfield, Gillian K. 1992. "Bias in the Evolution of Legal Rules." 80 *Georgetown Law Journal* 583-616.

Hathaway, Oona A. 2001. "Path Dependence in the Law: The Course and Pattern of Legal Change in a Common Law System." 86 *Iowa Law Review* 601-65.

Hay, Bruce L. and Spier, Kathryn E. 1998. "Settlement of Litigation" in *The New Palgrave Dictionary of Economics and the Law*, P. Newman ed. (New York, NY: Stockton Press).

Hirshleifer, J. 1982. "Evolutionary Models in Economics and Law. In *Research in Law and Economics*, P. Rubin and R. Zerbe (eds).(Greenwich, CT: JAI Press).

Krehbiel, Keith. 2004. "Legislative Organization," 18 *Journal of Economic Perspectives* 113-128.

Marvin, Daniel T. 2007. "The Supreme Court's Missed Opportunity to Settle the Handiwork of Nature Exception to Patentable Subject Matter in Laboratory Corporation of America v. Metabolite Laboratories, 26 *Temple Journal of Science, Technology, and Environmental Law* 113-148.

Merz, Jon. F., Kriss, Antigonne G., Debra, G. and Cho, Mildred K. 2002. "Diagnostic Testing Fails the Test: The Pitfalls of Patents are Illustrated in the Case of Haemochromatosis," 415 *Nature* 577-579.

New, John G. 2007. "Patently Wrong: The U.S. Supreme Court Punts in the Case of Labcorp v. Metabolite," 10 *Vanderbilt Journal of Entertainment and Technology Law* 147-177.

O'Hara, Erin. 1993. "Social Constraint or Implicit Collusion? Toward a Game Theoretic Analysis of Stare Decisis," 24 *Seton Hall Law Review* 736-778/

Parisi, Francesco and Fon, Vincy. 2008. The Economics of Lawmaking (New York, NY: Oxford University Press).

Priest, George L. 1977. "The Common Law Process and the Selection of Efficient Rules." 6 *Journal of Legal Studies* 65-82.

Posner, Richard A. 1972. Economic Analysis of Law (1st ed. Boston: Little Brown).

Rosenthal, Howard and Romer, Thomas. 1978. "Political Resource Allocation, Controlled Agendas, and the Status Quo," 33 *Public Choice* 27-43.

Rubin, Paul H. 1977. "Why is the Common Law Efficient?" 6 *Journal of Legal Studies* 51-63.

Sterns, Maxwell L. 1995. "Standing Back from the Forest: Justiciabiity and Social Choice." 83 *California Law Review* 1309-1413.

Sunstein, Cass R. 1991. One Case At A Time: Judicial Minimalism on the Supreme Court (Boston: Harvard University Press)..

Thomas, John R. 2002. "Liberty and Property in the Patent Law," 39 *Houston Law Review* 569-618.

Wagner, Polk R. 2002. "Reconsidering Estoppel: Patent Adminstration and the Failure of Festo," 151 *University of Pennsylvania Law Review* 159-244.

Zywicki, Todd. 2003. "The Rise and Fall of Efficiency in the Common Law: A Supply-Side Explanation." 97 *Northwestern Law Review* 1551-163.