

A Theory of Rational Jurisprudence*

Scott Baker[†] and Claudio Mezzetti[‡]

[†] School of Law, Washington University in St. Louis

[‡] Department of Economics, University of Warwick.

January 4, 2010

WORKING DRAFT

Abstract

This paper uses a dynamic model to study the development of judge-made law. Judges learn about the proper scope of the legal rule from the facts in the litigated cases. They design opinions to "fish" or search for information in future disputes that will have the most value for tailoring the law. We show that an optimal approach to judge-made law results in opinions that balance two kinds of costs: (1) decision costs, the cost of hearing future cases and (2) error costs, the costs of opinions which resolve future cases incorrectly. The model generates three results. First, opinions will tend to start broad and the law will be fleshed with a series of narrow refinements. Second, a court that learns can, under certain circumstances, optimally and rationally issue inconsistent decisions. Third, if decision costs are sufficiently high, judge-made law will converge to a doctrine that admits errors in some fact scenerios. Following the formal results, doctrinal examples of each are provided.

Keywords: Search, Judge-Made Law, Evolution of Legal Rules

JEL Classification Numbers:

*For helpful comments, we thank Jan Boone, Adrienne Davis, Kim Krawiec, Gerrit De Geest, Mitu Gulati, Gillian Hadfield, Dan Klerman, Shmuel Leshem, Anup Malani, Adam Rosenzweig, Richard Posner, and workshop participants at the University of Chicago, University of Southern California, the Triangle Workshop on Law and Economics, and Tilburg University.

1 Introduction

Judge-made law regulates much of economic and private life. The political science literature and much of the empirical legal studies literature has focused on why judges or groups of judges vote the way they do (Segal & Spaeth 2002; Miles & Sunstein 2008, Tiller & Cross 1999; Sunstein et al. 2002). But judges do more than vote, they write opinions defining the law. The production of judge-made law has received far less formal attention, especially as compared to the production of legislation.

We view the creation of judge-made law as a rational dynamic choice problem.¹ Prior case law, precedent, provides information about where the legal doctrine has worked and where it hasn't. The judicial opinion reveals to future judges and litigants what a judge has learned from applying the doctrine to a specific set of facts.

In practice, judicial opinions differ in scope. Based on the investigation of the interaction between the doctrine and a set of facts, some opinions set the rules for a wide range of future activities. For example, in *Miranda v. Arizona*,² the Supreme Court held that Fifth Amendment privilege against self-incrimination required the accused be adequately and effectively apprised of his rights before an in-custody interrogation. *Miranda* laid ground rules for police in lots of situations (all those involving in-custody interrogation).

Other opinions set the rules for few future activities. *Cox Broadcasting Corp. v. Cohn*³ provides an example. There, the Court considered whether a newspaper could be liable for revealing the name of a rape victim. The newspaper argued for a broad immunity under the First Amendment. The Court refused, closely tying its decision to the facts of the case: whether a newspaper could be liable for publishing a rape victim's name obtained through public records.

In the model, the judge decides how "broadly" to write the opinion by trading off two kinds of costs: (1) error costs – the costs due to an opinion deciding future cases incorrectly, and (2) decision costs – the costs of hearing future cases. Judges learn about the proper scope of the legal rule by bombarding the legal doctrine with facts from litigated cases. The facts allow the judge to fill in the law. Each time a judge hears a

¹We use the term judge-made law to mean common law, constitutional law, and the interpretation of broad or general statutory phrases, like "reasonable restraint of trade" under the Sherman Act.

²384 U.S. 436 (1966).

³429 U.S. 469 (1975).

case he can refine the legal rule to account for any new information embedded in those facts.

The critical part of the opinion is what it doesn't say – the issues it leaves open. By leaving certain issues open, the judge invites litigation on those matters. The gaps in the opinion can be viewed as sign-posts telling litigants where to focus their attention. In so doing, judicial opinions determine the future flow of litigation and, with that, the future flow of information and facts upon which to make judge-made law. In this setup, the ability to guide the course of future litigation justifies the publication of opinions. It also distinguishes judge-made law from legislation. Finally, by focusing on opinions, the model builds off the intuition many lawyers have that the language of the opinion, including any dicta, is important. What the opinion says – and doesn't say – is crucial because it determines the kinds of facts the court will see in the future, which, in turn, determines what the court will be able to learn.

The model offers several predictions about judge-made law that are consistent with actual decisions and doctrine.

First, opinions will narrow over time. The first opinion in a doctrinal series will tend to be broad, followed by a string of more narrow refinements.

Second, judge-made law will differ in scope and breath across areas of law. The reason: Decision costs and error costs are not identical in every area of judge-made law. Consider error costs. Take a contract law opinion that incorrectly decides that certain promises are unenforceable when, in fact, a majority of contracting parties would prefer those promises be enforced. The resulting error is costly, but not dramatically so. The judicial error increases transaction costs because now contracting parties have to draft around the incorrect judicial rule.

By contrast, suppose the court incorrectly demarcates the line between a coerced confession and a non-coerced confession. The error cannot be fixed by private parties or alternative police practices. In the law and economics jargon, the coerced confession rule is an immutable rule. The contract rule is a default rule. Because errors "stick" with immutable rules, the model predicts that the court will move more slowly in areas of law that involve such rules.

Third, legal doctrine can be inconsistent over time. An opinion in time one that states activity X is illegal will be trumped by an opinion at time 2 that states activity X is legal. Such inconsistency is usually chalked up to judicial miscues or differing pref-

erences among judges. The model suggests an alternative justification for inconsistent jurisprudence: judicial learning. The judges update in light of new information and, as a result, reconsider prior decisions. Judges benefit by cloaking doctrinal inconsistency in language that distinguishes cases that seem indistinguishable. In that way, a judge can use what he has learned to improve the law while maintaining the appearance of a commitment to stability.

Finally, judge-made law will eventually settle, meaning the court will define the law for all possible activities.⁴ If decision costs are higher than twice the error costs, the settled law contains errors. The court understands that the legal doctrine works imperfectly for some scenarios, but fixing and further refining the doctrine is not worth the effort. The threshold involves double error costs because the court can make two kinds of errors. The court might find "liability" for activities where "no liability" is the correct approach or the court might find "no liability" for activities where "liability" is the correct approach.

The conventional wisdom is that the courts prefer a settled but imperfect law because of an interest in stability. Here, the law settles because the benefits of further error correction are not worth the cost of additional judicial effort.

The model makes several assumptions. We assume a single court that lives forever. We assume the court consists of judges with identical policy preferences. We assume the court can anticipate and therefore manage the flow of future litigation. Despite these assumptions, the predictions from the model are consistent with the development of a number of important strands of legal doctrine and case law.

The paper unfolds as follows. The next sub-section reviews the related literature. Section 2 develops the initial building block: an economic model of legal reasoning. Section 3 specifies the model. Here we characterize the features of rational jurisprudence, meaning the most efficient or cost-effective way for judges to learn about the appropriate scope of legal rules. The three propositions in that section contain the main results. The first proposition sets out the relationship between opinion scope and decision costs, error costs, the weight placed on the future, and the amount of information the judge feels is left to be learned. The second proposition demonstrates the conditions under which legal doctrine can be optimally and rationally inconsistent. The third proposition

⁴This notion of settled law is analogous to Lewis Kornhauser's idea of the extended rule or legal regime. Kornhauser 1992(a) and Kornhauser 1992(b).

shows that legal doctrine can converge to a set of imperfectly-tailored rules. Following each proposition are real-world examples of the formal results. Section 4 suggests some possible extensions to the model. Section 5 concludes.

1.1 Literature Review

This paper straddles the literatures on judicial behavior and the evolution of legal rules. Scholars of judicial behavior model courts or judges as rational actors seeking to maximize an objective function. The constraints in the models vary and depend on the question being asked. Sometimes the judge makes his decision anticipating the likely position of Congress or the executive (Eskridge & Ferejohn 1992). Other times, the judge feels constrained by other judges sitting on the panel (Spitzer & Talley 2009) or the likely position of the higher court (Songer et al. 1994; McNollgast 1995). Still other models view judges as interacting repeatedly with each other over time (O’Hara 1993).

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 concern since Judge Richard Posner asserted that the common law was efficient (Posner 1973). In the wake of Posner’s assertion, scholars sought to explain the mechanism by which the common law could become efficient. The early 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 sharply contested (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). There has been historical work, pointing out reasons why the common law would tend to be pro-plaintiff (Klerman 2007). For our purposes, the important feature many of these models lack is a forward-looking judge. Case selection drives the law, with judges playing little role. Here, the opposite occurs. Through opinions – broad or narrow –

judges dictate what future cases get litigated.

There are a few exceptions to this pattern. Cooter et. al (1977) develop a model where the court can learn and asks whether tort law will converge to efficiency. Unlike here, the judges in their model do not write opinions to influence future litigation. All opinions are implicitly the same breath and scope. Similarly, Dari-Mattacci et al. (2010) develop a dynamic model where the litigants bring information to the courts and the courts issue decisions. The number of decisions is the "precedent" in their model, but the judges do not choose the contours of the opinions. Finally, Hadfield (2009) and Hadfield (2008) explore the institutional features that allow a judiciary to adapt rules to new circumstances. Hadfield's work sets out a theoretical framework and then applies that framework by engaging in a comparative analysis.

Turning to the economists, they have also been searching for a theory to explain why a common law legal origin positively correlates with economic development. To understand this empirical fact, law and finance scholars needed a model of judicial behavior. They faced the same problem as Judge Posner. If, in fact, the common law is efficient, why? 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 considers 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 or outcome they want to reach, the issue is can they get there.

2 Legal Reasoning

Law schools train students to "think like a lawyer." What is usually meant by this phrase is the ability to (1) compare two cases; (2) identify differences between the cases; and (3) argue why those differences are important in light of the purpose of the legal rule. If the differences are important, the good lawyer can distinguish the two cases and suggest a reason why case 1 shouldn't be followed in case 2. The prior case law provides information, which can be plumbed for clues about how to decide future cases. In the end, legal doctrine results from showering the legal rule – which will often be phrased in broad and vague terms – with new fact patterns. If the new fact patterns indicate that,

as stated, the judge-made rule no longer serves its function, the common law judge is free to reformulate the rule or create an exception.

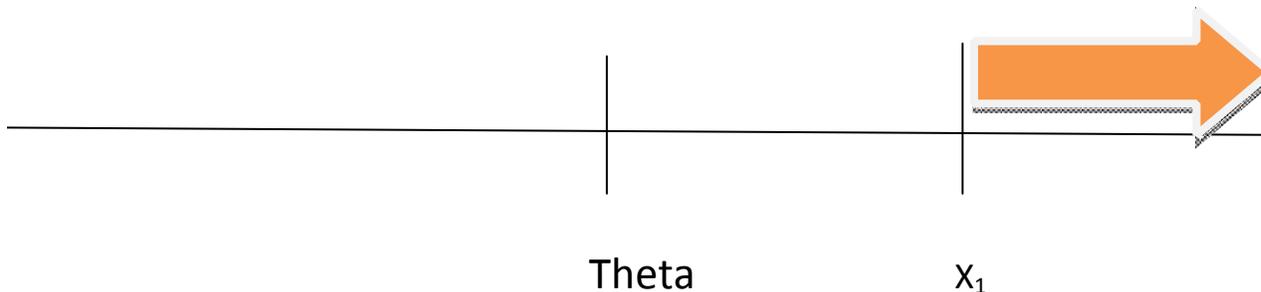
To capture this process, suppose that the judge-made law is attempting to regulate some activity, $x \in [-\infty, \infty]$. The activity could be anything. It might be the effort expended in driving safely; the actions of police officers in conducting searches; the amount of legal services afforded indigent defendants; or the amount of care owed by producers to end users of their products. Each activity carries a cost, $C(x)$, and a benefit, $B(x)$. The value of activity x is the difference between benefit and cost, $V(x) = B(x) - C(x)$. Value decreases in the activity. At some point, θ , value is zero, $V(\theta) = 0$, the benefit of the activity equal the cost. The threshold point is unknown. It is a random variable distributed according to $F(\theta)$ with positive density $f(\theta)$ over $[\underline{\theta}, \bar{\theta}]$ (with $\underline{\theta} \geq -\infty, \bar{\theta} \leq \infty$). Legal doctrine assesses the costs and benefits of some activity in a series of factual situations.⁵

Suppose that a case with facts x_1 comes to court. Upon examining the merits of the case and incurring a decision cost, D , the judge learns whether x_1 is greater than or less than θ . In the event that $x_1 > \theta$, the court learns that $C(x_1) > B(x_1)$, or $V(x_1) < 0$. Accordingly, the activity should be prohibited and, in a civil case, the defendant held liable. On the other hand, if $x_1 < \theta$, then $C(x_1) < B(x_1)$, or $V(x_1) > 0$. The activity should be permitted and, in a civil case, the defendant held not liable.

In addition, close examination of the facts teaches the court something about the world and, in particular, the random variable θ . Suppose that the court learns that $x_1 > \theta$. From that, the court can infer that cases greater than x_1 , say x_2 , also have costs greater than benefits. This knowledge allows the court to update its beliefs about θ . After observing x_1 , the court learns that the random variable must lie below the upper bound, x_1 . In this way, the prior case law provides information about the proper scope of the legal rule. Figure 1 illustrates this idea. The orange arrow represents the set of activities that the court learns have costs that exceed benefits after observing $x_1 > \theta$.

Take an example. A driver is liable for an accident if he acted without reasonable care.

⁵When speaking of costs and benefits, we take no normative position. The court could measure costs and benefits in terms of efficiency. Alternatively, the court could be searching for the most just legal rule, where the benefit and costs are measured in terms of how much the legal rule advances, say, corrective or distributive justice. The model applies equally to searching for the best legal rule plus exceptions or fleshing out the appropriate standard with a series of fact situations.



Consider first a driver who was using his cell phone before a crash. In our model, the court looking at the facts could determine whether the expected loss from the accident outweighed the cost of precaution (waiting to use the cell phone). Suppose it did and the court holds the driver liable. Up to this point, legal reasoning and prior case law have played no role in the analysis.

Subsequently, another case comes to court. In that case, a driver is using his blackberry prior to a crash. The two cases have an important difference: one driver used a cell phone, the other a blackberry. So, the first case would not control. Nonetheless, a court reading the case about the cell phone might infer that because the driver was likely more distracted with a blackberry than a cell phone, the expected loss was higher in the second case, while the burden of precaution was roughly the same. As a result, the driver in the second case should be held liable. In this sense, the case law provides information about the costs and benefits of activities "close to" the activities described in the prior cases. In our model, the court can impute from the cell phone crash case that any activity more distracting than a cell phone will have more costs than benefits. The court doesn't learn, however, how much less distracting than using a cell phone an activity must be to avoid liability – i.e., the exact value of θ .

At a higher level of abstraction, consider the following issue: whether an automobile manufacturer should be liable for negligence to the end consumer when a sale transpires through a retail dealer.⁶ List all activities x potentially subject to the negligence rule.

⁶This was the issue in *MacPherson v. Buick Motor Co.*, 217 N.Y. 382 (1916). On the relationship

The cut-off, θ , is the knife-edge activity where the imposition of negligence liability has equal benefits and costs. In deciding the automobile manufacturer case (call that case x_2), the court would scan the prior cases for similar activities, like, say, whether drug manufacturers were held liable for negligence to end consumers (call that case x_1).

From canvassing the case law, the court learns the costs and benefits from imposing a negligence regime on the drug manufacturer. Suppose $C(x_1)$ is bigger than $B(x_1)$; that is, $V(x_1) < 0$. As a result, the prior court held that the drug manufacturer was liable for negligence, $x_1 > \theta$. In our model, the court also learns whether $x_2 > x_1$. If so, $V(x_2) < V(x_1) < 0$ and the court knows that the car manufacturer should also be held liable. If, however, $x_2 < x_1$ the court only learns that $V(x_2) > V(x_1)$. The court might have a hunch that, in the car manufacturer case, the benefits from imposing liability for negligence outweigh the costs. But it isn't sure. That is to say, the knife-edge activity – where benefits equal cost – could be with car manufacturers, or a higher activity. The court updates its beliefs about the scope of the rule from the prior cases, but the threshold remains unknown.

Formally, let $F(\theta|H_t)$ be the court's estimate of the posterior distribution over θ , which depends on the case 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, declaring the activity invalid, when the case is x_t is $(1 - F(x_t|H_t))$. 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 activity that the court found to have positive value and declared valid; (2) the smallest past activity that the court found invalid, or having negative value. These two endpoints squeeze the court's estimate of the posterior distribution. The court knows that θ must lie somewhere between the highest past valid activity and the lowest past invalid activity.

Denote the highest past activity the court declared valid, W_t and the lowest past activity the court declared invalid, R_t . (We will formally define them later.) For $x \in$

of this case to prior precedent, see Posner (1990) and Levi (1948).

$[W_t, R_t]$, 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 a well-known view on precedent. Relying on past decisions provides information and saves judicial resources. Assuming the prior judgments were correct, the court can take those rulings as given and focus on "new" issues.⁷ 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).

3 The Model

There exists a single court of last resort that lives forever. The court is trying to learn about the policy parameter θ . In specifying the model this way, we ignore the dynamics between different judges of a court; there is no logrolling or horse-trading of judicial votes. Moreover, we assume that the judges are identical and, hence, we can use the terms "court" and "judge" interchangeably. We do not allow different judges to have different policy preferences. We also ignore the relationship between courts of different levels, like the Supreme Court and the appellate courts, or the appellate courts and the district courts. Finally, we assume there is no cost from deviating from prior precedent.

The assumption of shared common values and a single court might strike readers as odd. In many areas of law – patent, tax, antitrust – we think the assumption is a

⁷The 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 informative. Previous judicial errors would cancel out, enabling the court to extract important information relevant to updating from the prior cases.

reasonable one. In other areas of law, we might think of the model applying to a set of, say, republican or democratic-appointed judges who hold a majority on a court. Unlike the prior literature, the key ingredient is that we don't assume that the judge knows what rule or legal doctrine will accomplish their goals, whatever those goals might be. The question is: Can they learn about the impact of legal doctrine through the opinions they write and the cases they hear and the facts presented to them?

To learn about the policy threshold, the court issues opinions. An opinion consists of an interval, (a_t, b_t) , and has two components, a holding and dicta. The part of the opinion which is necessary for the result is the holding. Anything else the court says is dicta. Suppose that a case, x_1 , is not settled by the litigants. By paying D , the cost of examining the merits of the case, the court uncovers the costs and benefits of x_1 . Say x_1 is greater than θ . From this, the court knows that the activity should be declared invalid. Here, the holding of the opinion is b_t . It declares liability for all activities above b_t . More to the point, if the court wants to hold the defendant liable, they must set the bound b_t less than x_1 . The dicta is the point a_t . It says that for activities less than a_t there is no liability. This part of the opinion is not essential to the imposition of liability, but it is informative and determines which cases will be brought to court.

The court's opinion can be seen as (1) articulating a bright line legal rule with subsequent cases calling for exceptions or limitations on the applicability of the rule or (2) as filling out a legal standard. Either way, what the court does is provide a mapping for some subset of facts x into a legal consequence – liable, not liable. Generally speaking, a rule will govern more fact patterns than a standard (in the model a broader opinion will give rise to a rule), but there is no restriction that a court must stick with a rule or a standard or some combination of both over time. Indeed, the model is general enough to give the court an infinite degree of flexibility between rules and standards. Standards can become more rule-like over time. Rules can become riddled with exceptions. The key feature is that the kinds of opinions the court writes impacts what it learns in the

future.⁸

3.1 Timing

The timing of the game is as 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 $[\underline{x}, \bar{x}]$ with $\underline{x} \geq \underline{\theta}$, $\bar{x} \leq \bar{\theta}$.⁹
- 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, $\{a_{t+1}, b_{t+1}\}$.

Parties settle every case where the court has defined the law, in holding or dicta. If $x_t < a_t$ the court has previously opined that defendants should not be liable for this activity. Since the plaintiff isn't going to win anyway, he drops his case and avoids the court costs. If $x_t > b_t$ the court has previously declared the defendant should be liable. And so, the defendant immediately pays the plaintiff's damages and avoids the cost of trial. Only when $x_t \in [a_t, b_t]$ will the parties be uncertain about the court's resolution of

⁸A narrow opinion leads to more learning than a broad opinion. Other factors also might dictate the contours of an opinion, such as the need to garner a majority of the votes of the court (Cameron and Kornhauser (2008); Friedman (2009)). 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.

⁹In specifying the model in this fashion (always drawing facts from the same distribution), we abstract away from the law's impact on primary behavior. We do this to ease the analysis and focus on judicial learning. So long as parties either (1) make mistakes about the contours of the law when deciding on their primary activity or (2) face a small probability of getting caught and sued, the court should potentially see cases from the entire distribution each period. We are still working on the model when law changes the underlying distribution, that is, the primary activity levels.

their dispute. We assume that all these cases go to trial and up on appeal. All we want to capture here is that (1) appeals are more likely to occur when the law is uncertain and (2) judicial opinions control how certain the law is.

We can now formally define the highest past activity the court declared valid, W_t and the lowest past activity the court declared invalid, R_t :¹⁰

$$W_t = \max_{\tau < t} \{x_\tau : x_\tau \in [a_\tau, b_\tau], x_\tau < \theta\}$$

$$R_t = \min_{\tau < t} \{x_\tau : x_\tau \in [a_\tau, b_\tau], x_\tau \geq \theta\}$$

3.2 The Jurisprudential Trade-off

The court values two things. First, the court cares about time and effort spent evaluating cases, D . Second, the court cares about errors, L . Errors happen when a case settles out of court in the "wrong" direction, given the court's estimate of the posterior distribution, $F(\theta|H_t)$. In other words, an error occurs if, given the law either (1) a plaintiff drops a case during settlement that would have been decided in favor of the plaintiff by the court or (2) a case that the defendant should have won is settled in favor of the plaintiff. It may seem natural to assume that $L > D$, 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.

To get a better grasp on the concept of errors, an example helps. Suppose that the court sees case 5 and learns that $\theta \geq 5$. Accordingly, the court knows that activities below 5 have benefits greater than costs; the plaintiff should lose those cases. The court doesn't know what should happen with activities greater than 5. The court might extrapolate that a plaintiff with, say, case 6 should also lose. Given this belief, the court may set the lower bound $a_t = 7$. In words, the court in deciding case 5 issues an opinion that decides case 6 as well. The court feels comfortable doing this because activities 5 and 6 are quite similar, the cases are hard to distinguish. The broader opinion allows the court to hedge on decision costs. The court avoids the time and effort spent hearing the merits of case 6 because it feels confident reasoning by analogy from case 5. But the court could be wrong; θ might just equal 5, or 5.5. If so and if the next plaintiff draws

¹⁰We assume that the court would declare the threshold activity θ as invalid. This is unimportant, given that $x = 0$ is a zero probability event.

case 6, an error transpires. Given the law, the plaintiff anticipates that the court will dismiss its case based on the prior precedent. Accordingly, the plaintiff drops his case.

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 might like this one looks like a good idea. It isn't. Wider bounds mean that the chance of having to decide a case in the following period is higher because the court's opinion is less informative to future litigants. And so, the need to evaluate cases on the merits, which is costly, happens too much. The court isn't relying enough on reasoning by analogy, i.e., extrapolating costs and benefits from similar cases. Instead the court is investigating on the merits every case that comes in the courthouse door.¹¹

In setting the precedent bounds, 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. 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 learning benefit allows for course correction over the long haul.

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:

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

¹¹In our model, each judge is identical and cares about the effort expended by future judges. As a result, the judge sets the opinion bounds such that it is not cost-justified (the benefit in error reduction does not outweigh the cost of investigation of the case) for the judge in the next period to review the case on the merit if that case lies below a_t or above b_t . Litigants understand that next period's judge won't be motivated to review the case if they bring the matter to court (it is efficient for that judge to simply rely on the prior precedent). And so, these cases settle in light of the prevailing law. Note this does not reflect the court committing not to review the cases. If a case outside the bounds is brought, it is not in the future judge's interest to actually review the case and incur decision cost, D .

The first term in (1) is the cost of having to decide a case in the next period. To see why, let's say the judge sets $a_t = b_t$. Such a choice defines the law for all activities. 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 reflect the expected one-period error costs.

Consider the second term. If the judge sets $a_t \geq W_t$, there is a chance the plaintiff will draw a case x_t between W_t and a_t . Given the opinion's lower bound, the plaintiff will drop this case. The expression in square brackets is the probability the court attaches to the event that this case, x_t , should go in favor of the plaintiff instead (the activity should be declared invalid). The third term follows from a similar analysis on the upper bound of the decision; here the defendant settles but the activity should be declared valid. 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 equation (1) implies that the court will find it optimal to declare "no-liability" for that activity without investigating the case (i.e., without incurring the cost D). 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 "liability" for that activity.

With the objective function in hand, we can now describe how a court that wanted to learn would trade-off error costs and decision costs. The interval $[W_t, R_t]$ represents what is left to be learned about the policy parameter, given the case history, H_t . The tighter the interval, the more informed the court is about θ .

Our first proposition suggests when we should observe a court issue a broad opinion versus a narrow opinion versus not investigate the merits of the case at all. The difference $b_t - a_t$ measures the set of cases left uncertain by the opinion; the smaller $b_t - a_t$, the broader is the opinion. Thus, it is appropriate to define the "jurisprudential breadth" at time t as $J_t = 1/[b_t - a_t]$; the smaller the breadth, the narrower the opinion at time t . The optimal "jurisprudential breadth" depends on four factors: (1) the decision cost D , (2) the error cost L , (3) how much the court thinks it knows about the optimal rule, and (4) future's value (as measured by the discount factor δ). We have the following result.

Proposition 1 (i) *A narrower opinion is optimal when decision costs are lower, error costs are higher, the future is more valuable, and the court knows relatively more about the optimal doctrine; (ii) A broader opinion is optimal when decision costs are higher, error costs are lower, the future is less valuable, and the court knows relatively little about the optimal doctrine. That is, formally,*

$$\frac{\partial J_t}{\partial D} > 0, \frac{\partial J_t}{\partial L} < 0, \frac{\partial J_t}{\partial W_t} > 0; \frac{\partial J_t}{\partial R_t} < 0; \frac{\partial J_t}{\partial \delta} < 0.$$

Proof:

See the appendix.



If decision costs are low, it doesn't cost much to learn about the benefits and costs of a particular activity by investigating the merits of the case. As such, there is no need for the court in an opinion to extrapolate costs and benefits from one case to another. The court can simply consider each activity independently. To do that, the court sets its opinion bounds each period very close to what it knows for sure, $\{a_t, b_t\}$ is very close to $\{W_t, R_t\}$. This move results in opinions limited closely to the facts of the case. If error costs are low, the cost of extrapolating from one case to another is small. Realizing this, the court can aggressively hedge and place the opinion bounds well in the interior of what it knows. An aggressive hedge saves on decision costs because the court sees fewer cases in the future. Since error costs from issuing a precedent that incorrectly decides future cases are small, this cost savings is worthwhile and pushes the court toward broader opinions.

As the court knows more and more about the optimal policy, it will tend to make smaller and smaller refinements to the legal doctrine. When the court first considers an issue, for example, past opinion are totally uninformative or, the same idea, maximally narrow. In that case the first jurisprudential step will tend to be big. More generally, we expect the jurisprudential steps to be inversely correlated with the current breadth of the law. The more fact situations covered by the current law, the smaller are subsequent refinements to that law.

There is a debate raging between judicial minimalists like Cass Sunstein, who advocate that in most circumstances the court be cautious, rendering a series of narrow decisions (Sunstein 1999) and 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 some light on this debate. Narrow opinions are not always called for or prudent, but rather depend on the circumstances and what the court has learned so far. The alternative 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 rationally optimally will sometimes be a minimalist and other times take broad steps. This is true both over time and across areas of law. And, indeed, when examining actual doctrine we see much of this jurisprudential heterogeneity.

3.2.1 Examples: Doctrinal Trajectory – Broad to Narrow

Proposition one suggests that opinions in a doctrinal series will start broad – when the court doesn't know very much about the contours of the legal rule – and be followed by a series of narrow refinements. In addition, the opinion language (a_t and b_t in the model) influences what cases are subsequently litigated. What follows are some examples of this phenomenon.

In *Erie v. Tompkins*, the Supreme Court ruled that federal courts hearing cases based on diversity jurisdiction must apply state substantive law, no matter whether the law comes from the state court, the state legislature, or the state constitution. The opinion is dramatic in the number of cases it touched upon. It has been called "one of the most important cases at law in American legal history." (Black 1942). After *Erie*, in every diversity case, federal courts had to apply the substantive law of the state.

The Court in *Erie* left a number of matters undecided. Specifically, in his concurrence, Justice Reed noted that the *Erie* doctrine must be applied only to state substantive law, not state procedural rules. In other words, a federal court was free to apply federal procedural rules in a case based on diversity jurisdiction. The language in the concurrence set the question for future courts to decide: What is a procedural rule and what is a substantive rule? The follow-on cases to *Erie* have, for the most part, addressed this question. Over time, litigants have attacked various rules as primarily substantive or primarily procedural. The Court has placed some rules in the substantive bin and others in the procedure bin. In *Hanna v. Plumer*,¹² the Court refined the inquiry to differentiate substantive from procedural rules, but this refinement applies to

¹²380 U.S. 460 (1965)

far fewer cases than the initial decision in *Erie* itself. In short, the trajectory of *Erie* and its prodigy maps onto – and is consistent with – an infinitely-lived court attempting to learn about the proper contours of a jurisdictional rule.

A second example concerns a state court’s power to exercise personal jurisdiction over a non-resident defendant. In *International Shoe Company v. State of Washington*, the Court articulated the test for when a state court has such jurisdiction. The Court stated that jurisdiction attached if “[the defendants] have certain minimum contacts with [the state] such that the maintenance of suit does not offend ‘traditional notions of fair play and substantial justice.’” *International Shoe* is a broad decision, greatly expanding the reach of the state courts. The prior case law required a defendant to be physically present in the state before jurisdiction attached. The Court in *International Shoe* allowed suit against defendants whose activities merely touched the state in some fashion.

The phrases “minimum contact” and “traditional notions of fair play and substantial justice” were not expressly defined by the Court. And so, subsequent litigants brought facts pushing the envelope of what might be considered a minimum contact and what counts as fair play and substantial justice. As in the model, the opinion language in *International Shoe* guided the subsequent litigation. It told the litigants what to focus on. More important, in the last fifty years, the Court has filled in those phrases, determining their content in a series of opinions, each one narrower than *International Shoe* itself.¹³ Again, we see here the evolutionary pattern – broad to narrow – consistent with the court attempting to learn about proper scope of the doctrine.

This pattern is not confined to jurisdictional matters. In *Miranda v. Arizona*, the Court ruled that the Fifth Amendment requires the police to provide protective devices to ensure that any statement made during “custodial interrogation” was truly the product of the accused’s free choice. *Miranda* led to police reading criminal defendants’ their “rights” in every police station house in the United States. As such, the

¹³See, for example, *Helicopteros Nacionales de Columbia S.A. v. Hall*, 466 U.S. 408 (1984) (holding that general jurisdiction could not be asserted by Texas court against Peruvian consortium); *McGee v. International Life Insurance Company*, 355 U.S. 220 (1957) (holding that sporadic contact with a state was enough to establish personal jurisdiction if that contact led to the cause of action); *Travelers Health Association v. Commonwealth of Virginia* (allowing jurisdiction against a non-resident insurer who had used mail to solicit business from state residents).

decision was broad. Like *Erie* however, *Miranda* left some matters unresolved. *Miranda* stated that the warnings are triggered only when a defendant is in “custody” and subject to “interrogation” – two terms left undefined. In *Miranda*’s wake, litigation fomented about the meaning of these words. The language in the *Miranda* opinion drove the future litigation. And, consistent with proposition one, the later decisions tended to be narrower, applicable to far fewer fact situations than *Miranda* itself.¹⁴

Finally, consider an example from the common law. In tort law, employers are vicariously liable for the torts committed by their employees during the course of employment. By contrast, the general rule for independent contractors is that employers are not responsible for the acts of independent contractors.¹⁵ The initial rule is broad, applying to every contractual relationship that does not rise to the level of employment no matter the tort committed. The independent contractor rule has been modified with a series of narrow exceptions. The rule does not apply when the employer hires the independent contractor to engage in an activity he knows is inherently dangerous.¹⁶ The rule does not apply when employer hires an independent contractor to do construction in a public place, like a highway.¹⁷ The rule does not apply to independent contractors employed by landlords to provide upkeep for tenants.¹⁸ Each of these exceptions is narrower than the initial rule, which is consistent with the court learning about the optimal contours of the vicarious liability rules for employers using independent contractors.

¹⁴See, for example, *Mathis v. United States*, 391 U.S. 1 (1968) (holding that *Miranda* does not apply when the defendant is in custody on a matter unrelated to the criminal investigation); *Berkemer v. McCarty*, 468 U.S. 420 (1984) (holding that whether a defendant is in “custody” turns on whether a reasonable person would feel free to leave the situation); *Oregon v. Mathiason*, 429 U.S. 492 (1977) (holding that when a defendant accepts an invitation by the police to come to the police station, he is not in custody); *Rhode Island v. Innis*, 446 U.S. 291 (1980) (holding that *Miranda* safeguards come into play whenever “a person in custody is subject to either express questioning or its functional equivalent”).

¹⁵The English decision articulating this point is *Reedie v. London & N.W. Ry.*, 4 Ex. 244 (1849); an early American decision is *Blake v. Ferris*, 5 N.Y. 48 (1851).

¹⁶See *Mavrikidis v. Petullo*, 153 N.J. 117 (1998); *Benesh v. New Era, Inc.*, 207 Ill.App.3d 1049 (1991); *Miller v. Westcor Ltd. Partnership*, 171 Ariz. 387 (1991).

¹⁷See, for example, *Rose Mellon v. Shriners Hospitals for Children et al.*, 22 Mass L. R3p. 544 (2007)

¹⁸See *Otero v. Jordon Restaurant Enters.*, 119 N.M. 721 (Ct.App.1995).

3.2.2 Examples: Differential Error Costs and Decision Costs

Proposition 1 suggests that opinions will be narrower as the cost of errors goes up. In other words, as losses from errors increase, the court will be reluctant to extrapolate from one set of facts to create exceptions to bright line rules for other fact situations or to modify the prevailing standard. Instead, the court will wait for the facts of a case to tee-up the issue squarely. In this intuitive result, we see an explanation of a number judge-made doctrines.

Courts construe statutes to avoid constitutional questions.¹⁹ Since constitutional rulings can only be changed by the Supreme Court, the Court avoids them if possible. In the error-cost framework, the Supreme Court wants to avoid making – and thereby potentially making an error in – constitutional decisions because such decisions are stickier than other rulings, which can be modified by legislation.

Recent abortion jurisprudence provides another prominent example of high error costs, this time because the underlying issue – when does life begin? – is subject to significant debate. After the broad ruling in *Roe v. Wade*,²⁰ the Supreme Court set forth in *Planned Parenthood v. Casey*,²¹ the test for whether a restriction on an abortion right survives constitutional scrutiny. The test – whether the restriction places an undue burden on the abortion right – has remained in place for a number of years. Since *Casey*, the abortion cases have been narrow and fact specific, as predicted by Proposition 1. The Supreme Court has held that the requirement of spousal notification is an undue burden;²² that a ban on partial birth abortion is not an undue burden;²³ that the 24 hours waiting period prior to an abortion is not an undue burden; and that regulation requiring parental notification did not impose an undue burden if it allowed for a judicial bypass.²⁴

An example of low error costs comes from contract law. The mailbox rule states

¹⁹ See *Ashwander v. TVA*, 297 U.S. 288 (1936) ("When the validity of an act of the Congress is drawn in question, and even a serious doubt of constitutionality is raised, it is the cardinal principle that this Court will first ascertain whether a construction of the statute is fairly possible by which the question can be avoided.") (Brandeis, J. concurring).

²⁰ 410 U.S. 113 (1973)

²¹ 505 U.S. 833 (1992)

²² *Casey*, supra note __.

²³ *Gonzales v. Carhardt*, 550 U.S. 124 (2007).

²⁴ See *Ohio v. Akron Center for Reproductive Health*, 497 U.S. 502 (19910).

that acceptance of a contract occurs upon dispatch.²⁵ The courts have refined this rule rarely. There is one primary exception. In the case of option contracts, acceptance occurs upon receipt.²⁶ Why have the courts been stingy with exceptions to the rule, not allowing modifications for other circumstances (where, say, the offeree tried to retract the acceptance by phone, after he had dispatched the acceptance by mail)? Our model suggests that the cost of applying the mailbox rule to circumstances where it doesn't quite fit is small. Contracting parties, after all, can just draft around the rule. And so, it makes sense that the mailbox rule has broad applicability and is rarely subject to judicially-created exceptions.

3.3 Optimal Inconsistency in Legal Doctrine

Having considered optimal jurisprudential breadth we can now further investigate the evolution of legal doctrine. One aspect of judicial opinions that frustrates legal scholars and lower court judges is inconsistency. The court will 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." The usual justification is that 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. In our model, a court might issue inconsistent doctrine and that inconsistency will reflect learning by the court. This result is expressed in the next proposition. Formally, we say that the jurisprudence is inconsistent if either $a_t < a_{t-1}$ or $b_t > b_{t-1}$. In the former case, activity x , $a_t < x < a_{t-1}$ was deemed valid at time $t - 1$, but uncertain or invalid in the opinion at time t . In the latter case, activity x , $b_t > x > b_{t-1}$ was declared invalid at time $t - 1$, but uncertain or valid at time t .

Proposition 2 *An optimal, rational jurisprudence can be inconsistent.*

Proof:

See the appendix.



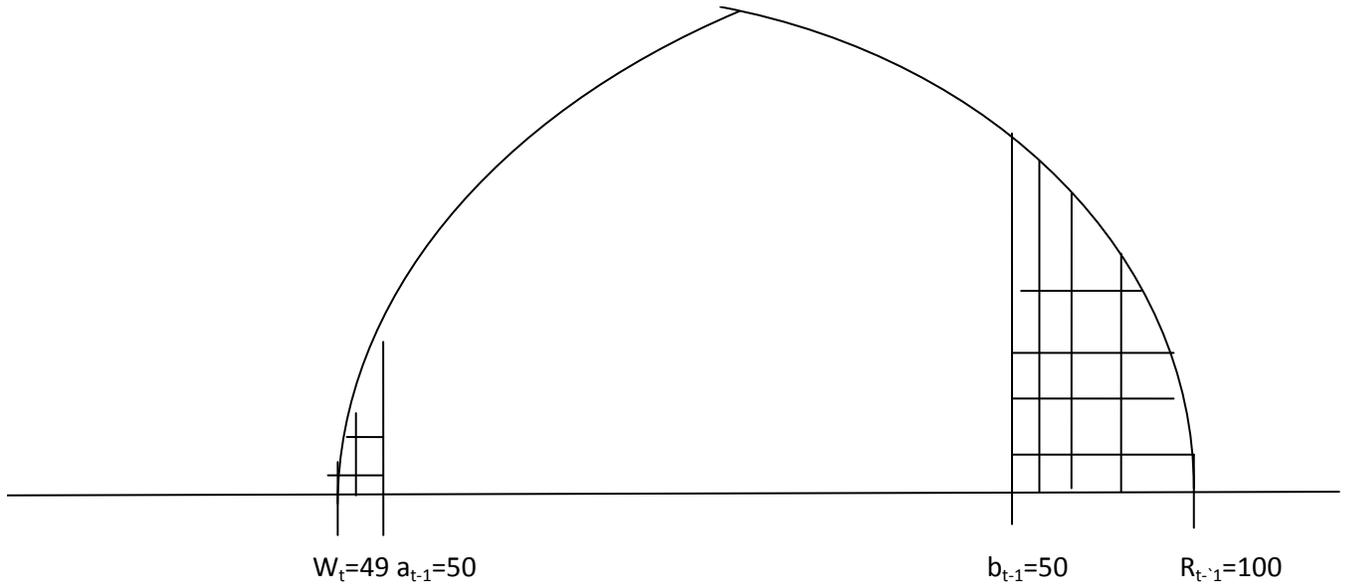
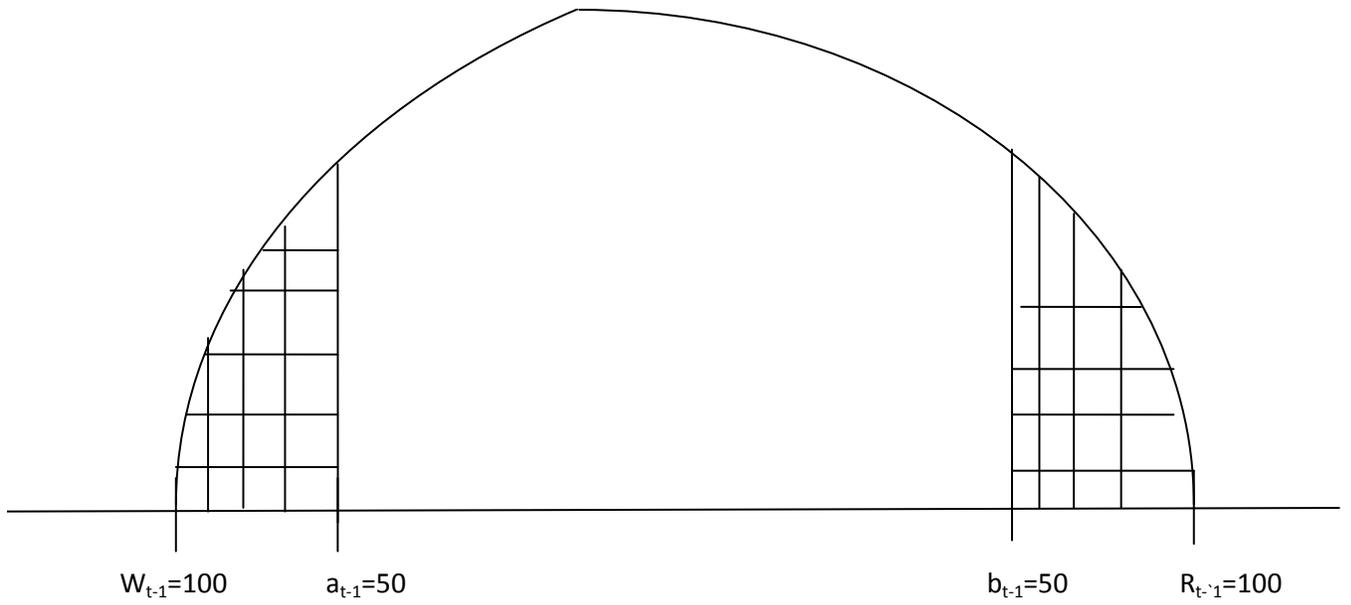
²⁵See *Adams v. Lindsell*, 106 Eng. Rep. 250 (1818); Restatement (Second) of Contracts 63 (a).

²⁶See *Santos v. Dean*, 96 Wash.App. 849 (1999); Restatement (Second) of Contracts 63(b).

In an incoherent or inconsistent jurisprudence, the opinion bounds bounce around. Activities that were deemed invalid become valid or uncertain, and levels that were valid become uncertain or invalid. 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 benefits outweigh the costs for this activity. 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 choice on the upper bound would involve setting $b_{t+1} = 49$. The court would have to define the law for all cases. 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. Indeed the court may also reasonably decide to set $a_{t+1} = 60$. Note that activity levels like $x = 55$ were deemed invalid at time t , but become valid at time $t + 1$; levels like $x = 70$ were invalid and become uncertain. What looks like inconsistent decision-making is really the court optimally gathering and using information from the case 49. Basically, if the court sticks to the old doctrine, the probability of error – the case settling for the plaintiff when it shouldn't – is much higher. The expected error cost (the probability of an error times the loss) is bigger, so the court recalibrates the opinion bounds to reflect the learning contained in case 49. In so doing, it goes back on what it has previously said.

Figures 2 and 3 illustrate this example.



The shaded regions in Figure 1 are the probability of an error when the court sets the bounds at $[-50, 50]$, knowing that the knowledge bounds derived from the precedent are $[-100, +100]$. The shaded regions in Figure 2 incorporate the new information from

case 49 into the calculus, setting the lower knowledge bound from the precedent at 49. The error, then, of remaining with an upper opinion bound of 50 is much larger in the next period. So, the court issues an inconsistent decision.

Adding a cost of deviating from precedent doesn't change the analysis. The court will still want to have inconsistent doctrine so long as what is saved in error costs outweighs the cost of not following precedent. And, in fact, if the court can minimize the "perception" that it is not following precedent, it gets the best of both worlds: learning and reducing errors, while at the same time avoiding any reputational cost from failing to articulate a stable doctrine.

To sum up, legal scholars and political scientists regard inconsistency in doctrine as inevitable, because either (1) courts have multiple judges voting and each one has a different preference over outcomes (Easterbrook 1987) or (2) courts make mistakes. In our model, we observe inconsistent doctrine under stark assumptions: where judges share a common policy preference and they learn about the effects of doctrine. Inconsistent doctrine, then, can be seen as a result of rational judging by judges committed to the same normative values. It is not necessarily the result of Arrow vote aggregation problems or judicial miscues.

3.3.1 Examples of Inconsistent Doctrine

Together, *Parker v. Flook*²⁷ and *Diamond v. Diehr*²⁸ provide our first example of inconsistent doctrine. Both cases addressed whether computer software counts as patentability subject matter. In *Parker v. Flook*, the Supreme 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 majority'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

²⁷437 U.S. 584 (1978).

²⁸450 U.S. 175 (1981).

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.

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, then, settled. The initial confusion is consistent with learning by the Supreme 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 path taken by majority, distinguishing cases that seem indistinguishable, allowed the court to improve the law based on new information, while saving face by paying lipservice to the prior case law.

The second example comes from the Supreme Court's standing jurisprudence. In *Fronthingham v. Mellon*,²⁹ the Court said that a taxpayer did not have standing to sue to enjoin expenditures under the Federal Majoritary Act of 1921, an act providing subsidies to reduce infant and maternal mortality. The *Fronthingham* court said that the plaintiff lacked standing because his interest was only in preserving "monies from the Treasury."³⁰ In *Flast v. Cohen*,³¹ the Supreme Court relaxed this general restriction, allowing a taxpayer suit to prevent expenditures of public funds to subsidize parochial schools. The plaintiff in *Flast* claimed the expenditures violated the establishment clause. The *Flast* Court allowed the suit because there was a sufficient nexus between the plaintiff's taxpayer status and alleged constitutional violation.

In *Valley Forge Christian College v. American United for Separation of Church and State*,³² the Court again confronted a taxpayer suit against an alleged Establishment clause violation. The plaintiff in *Valley Forge* challenged a grant of surplus government property to a Christian college. The Court did not allow standing. Scholars have suggested that the reasons given in *Valley Forge* for not following *Flast* are untenable.³³ The *Valley Forge* court said that it made a difference that the plaintiff challenged an

²⁹262 U.S. 447 (1923)

³⁰*Id.* at 487.

³¹392 U.S. 83 (1968).

³²454 U.S. 464 (1982).

³³Chemerinsky (2000 p. 93).

executive transfer of property authorized under Article IV section 3 of the Constitution, rather than a spending program authorized under Article 1, section 8. In addition, the Court said that the *Flast* facts involved a direct Congressional authorization, while *Valley Forge* involved an action by the Executive. Why either of these distinctions would make a difference is a hard to understand.

In *Heen v. Freedom from Religous Foundation, Inc.*³⁴, the Court drew another tenuous distinction. *Heen* involved a challenge to the government paying for President Bush and the Secretary of Education to attend and give speeches to faith-based organizations. The Court found that the taxpayer lacked standing. In so doing, the Court distinguished *Flast* by saying the payments weren't specifically authorized by Congress (as in *Flast*), but instead came out of the general fund allocated to the executive branch. In his concurrence, Justice Scalia said that the consistency in the taxpayer standing cases "lies in the creation of utterly meaningless distinction which separate the case at hand from the precedents that have come out differently, but which cannot possibly be (in any sane world) the reason it comes out differently."

Following Justice Scalia, we might view these three cases as simply inconsistent with one another. Perhaps changing court membership generated the results. But the result is also congruent with the court learning more and more about the impact of authorizing taxpayer suits, disguising inconsistency by creating arbitrary distinctions among the cases.

3.4 Convergence without Full Learning

The final point the model makes concerns how much time courts should spend refining doctrine. Legal academics, policy-makers, and advocates often critique the law articulated by courts. Especially in technical areas like patent, tax, antitrust and environmental law, arguments are made that the judicial approach is flawed or unwise. The model shows that imperfections in doctrine are inevitable when the cost of deciding cases is sufficiently high. Learning and refining the doctrine is costly. At some point, the benefits of further refinement — tweaking the doctrine to better advance the court's interest — are smaller than the costs. The court, then, refuses to say any more on the issue. The next proposition spells out this result. We say that the law converges in finite

³⁴551 U.S. __ (2007).

time if $a_t = b_t$ at some time T

Proposition 3 *If $D > 2L$, the law converges (with probability one) without the court fully learning about θ ; there is a T and a level x such that $a_t = b_t = x \neq \theta$ for all $t \geq T$. That is to say, the court understands the doctrine is imperfect (it admits errors in some fact situations), but chooses not to refine the doctrine. If $D < 2L$ the court eventually fully learns; $\lim_{t \rightarrow \infty} a_t = \lim_{t \rightarrow \infty} b_t = \theta$.*

Proof

See the appendix ■

If decision costs are greater than double the error costs, the doctrine stabilizes with imperfections remaining in the law. In other words, the court sets $a_t = b_t$ without knowing the exact location of θ . The opinion defines the law for all activities, but the court realizes the doctrine might not apply well in some circumstances. That said, correcting those imperfections is not cost-justified.

3.4.1 Examples

Many common law doctrines can be seen as well-settled, but imperfect. We end this section with an examples from tort and property.

A tortfeasor will be found negligent if he acts without reasonable care. The negligence standard is a knife-edge inquiry. If the defendant's is found negligent and the negligence causes the injury, the defendant is liable for all the resulting damage. Alternatively, if the defendant is found not negligent, he pays nothing. Calfee and Craswell (1984) show that, when the defendant is uncertain about the legal standard, negligence can result in too much or too little deterrence. There will be too much deterrence if the marginal investment in safety both reduces the amount the defendant can expect to pay (the expected loss) and the chance he has to pay anything at all. In that case, the private gains to an one-unit investment in, say, safety are bigger than the social gains. The courts have not fine-tuned negligence law to account for the risk of over-deterrence identified by Calfee and Craswell. On this score, negligence law is imperfect. Yet it is probably not worth the judiciary's time to tailor for these concerns (which involves estimating the probability distribution of damages the defendant thinks he will have to pay). Instead, a rough rule that works most of the time is sufficient.

Property law provides property-holders some freedom to engage in self-help in protecting their property. The logic is that the fear of self-help provides additional deterrence, deterrence not provided by uncertain criminal sanctions or civil actions against a judgment-proof tortfeasor (Epstein 1999 p. 51). The self-help remedy is narrowly tailored. In the context of a trespass to chattel or conversion, the property holder is entitled to use reasonable force to defend chattel attempting to be secured by force or fraud.³⁵ The reasonable force requirement is limited to actions taken in "hot pursuit" of the property.³⁶ Once the property-holder loses possession, he must resort to the courts to get it back. The "hot pursuit" requirement is well-settled, but likely imperfect.

Think of the cases in a jurisdiction with limited resources to provide criminal sanctions for property theft. There, in terms of optimal deterrence, it might make sense to weaken or jettison the "hot pursuit" requirement. In so doing, the court would boost the deterrence benefit from the self-help and, as a result, make up for the weak deterrence from the criminal sanction. The courts, however, do not consider the public money available to combat property crime in construing the hot pursuit doctrine. The rule is fixed, but probably not appropriate in every fact scenario (where the "facts" include the amount of deterrence coming from all possible sources).³⁷ Yet the cost of fixing the doctrine – in terms refining the requirement to account for how much deterrence is coming from public officials – is probably not worth the judicial effort.

4 Extensions

The analysis might be expanded in a number of ways. First, in the model the judge writing the decision decides on the opinion's scope. In practice, it is often the second judge in the series that decides what the initial opinion means. The future judge can, for example, decide to limit the holding of the prior case to its facts, reading the prior case law narrowly. Alternatively, the future judge could decide to read the opinion

³⁵ See *Vacanti v. Master Electronics Corp.*, 245 Neb. 586 (1994).

³⁶ See *McClean v. Colf*, 179 Cal. 237 (1918).

³⁷In terms of the model, assume x represents all cases where a "lack of public enforcement" exception for the hot pursuit requirement might be raised. As x gets bigger, the benefit for creating an exception gets smaller, while the cost of creating the exception gets bigger. The law settles on no exception whatsoever, despite the fact that in some circumstances an exception might make sense.

broadly, as applicable and informative about situations far afield. On this point, we agree. Nonetheless, the words of the initial opinion limit the range of plausible readings by the future judge. A future judge, for example, would find it difficult to read *Bush v. Gore*,³⁸ as a decision with far-reaching effects. In that decision, the Supreme Court stated: "Our consideration is limited to the present circumstances, for the problem of equal protection in election processes generally present many complexities." Similarly, in the case the article began with, *Cox Broadcasting Co. v. Cohen*, the Court explicitly stated that the decision was limited to the publication of a rape victim's name obtained through public records.

The model might be expanded to include the possibility that a future judge could select from a range of plausible readings, where the judge faced a reputational cost for deviating too far from the most sensible reading of the prior opinion. In our model, introducing this tweak wouldn't make a difference because all the judges want the same thing. There is, as a result, no reason for a judge to read a case more narrowly or more broadly than intended by the opinion's author. This result would change if we allowed the judges to have different policy preferences, which takes us to the second possible extension.

We assume that judges share the same normative commitments or values. This is obviously not true, especially in the "hot button" cases. What's surprising is that the model has descriptive power, while maintaining the assumption. We might expand the model to include judges with different preferences. The interesting twist comes from viewing opinions, as we do here, as a way of shaping the flow of future information to the court. Speculating, we might think of judges as writing opinions, not just to gloss the law with their own values, but also to shape the kinds of cases that will be litigated – and hence what can be learned – by future judges. A judge could write an opinion to settle the law in an area imperfectly (from his perspective), just to prevent the future judges from hearing cases and refining the doctrine more to their liking. This is where the range of plausible readings of an opinion comes into play. If the second judge can always limit the first judge's opinion to the facts (at zero cost), this move won't work. The reason: Litigants understand that the first opinion is not informative to the future judge with the alternative preference. As a result, anything the first judge says will not impact the future case flow. The first judge can't settle the law, no matter what he says

³⁸531 U.S. 98 (2000).

in the opinion.

Finally, the model assumes that judges get information from the case at hand and by canvassing the prior case law. Judges get information from elsewhere too. We might extend the model so that each judge receives a private signal. The informational basis of the opinion, then, is threefold: the private signal, the facts of the case at hand, and the prior case law. In this extension, two judges, facing the same case with the same prior precedent, could reach different conclusions and write different opinions. The judges would disagree because they know different things, not because they have different preferences.

5 Conclusion

Judges in our framework produce law by learning from prior precedent and the facts of the case litigated before them. Once judges can learn, the process by which that learning takes place becomes key to understanding the evolution of doctrine and the production of judge-made law. Opinions define the search, with future disputes the pool from which judges fish for information.

The model has predictive power for the kinds of opinions we observe, and the vast heterogeneity in the scope of actual judicial opinions. Although we have given a series of examples in the paper, the predictions can also be taken to the data. The central concern here is the informational value and creation of precedent. Precedent and doctrine play a limited role in the vast political science literature on judging, both theoretical and empirical. We have attempted to intergrate those concerns – on the minds of most legal scholars – into a formal framework.

Finally, our model is consistent with what judges claim to be doing, searching for the best resolution of a case in light of the prior precedent, the lawyers' briefs (the information from the case at hand), and their own efforts and hunches (expending decision costs).

References

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

2. Calfee, John E. & Craswell, Richard. 1984. "Some Effects of Uncertainty on Compliance with Legal Standards." 70 *Virginia Law Review* 965-__
3. 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
4. Cardozo, Benjamin. 1921 *The Nature of the Judicial Process*.
5. 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.
6. Chermersky, Erwin. 2000. *Constitutional Law, Principles and Policies* (2d ed.)
7. Cross, Frank B. & Tiller, Emerson H. 1998. "Judicial Partisanship and Obedience to Legal Doctrine: Whistleblowing on the Federal Courts of Appeals." 107 *Yale Law Journal* 2155-__.
8. Cooter Dari-Mattiacci, Giuseppe, Deffainis, Bruno & Lovat, Bruno, 2010 "The Dynamics of the Legal System," *Journal of Economic Behavior and Organization* (*forthcoming*).
9. Easterbrook, Frank A. 1982. "Ways of Criticizing the Court." 95 *Harvard Law Review* 802-832.
10. Eskridge, William & Ferejohn, John. 1992. "The Article I Section 7 Game." 80 *Georgetown Law Journal* 523-564
11. Gennaioli, Nicola and Shleifer, Andrei. 2007. "The Evolution of the Common Law," 115 *Journal of Political Economy* 43-68.
12. Klerman, Daniel. 2007. "Jurisdictional Competition and the Evolution of the Common Law," 74 *University of Chicago Law Review* 1179-1226.
13. Kornhauser, Lewis A. 1992(a). "Modeling Collegial Courts. I. Path Dependence," 12 *International Review of Law and Economics* 169-85.
14. Kornhauser, Lewis A. 1992(b). "Modeling Collegial Courts II. Legal Doctrine," 441 *Journal of Law, Economics, and Organization* 441-470.
15. Hadfield, Gillian K. 1992. "Bias in the Evolution of Legal Rules." 80 *Georgetown Law Journal* 583-616.

16. Hadfield, Gillian K. 2008. "The Levers of Legal Design: Institutional Determinants of the Quality of Law." 36 *Journal of Comparative Economics* 43-73.
17. Hadfield, Gillian K. 2009. "The Dynamic Quality of Law: The Role of Judicial Incentives and Legal Human Capital." *mimeo*, USC law school.
18. 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.
19. 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).
20. Krehbiel, Keith. 2004. "Legislative Organization," 18 *Journal of Economic Perspectives* 113-128.
21. McNollgast. 1995. "Politics and the Courts: A Positive Theory of Judicial Doctrine and the Rule of Law." 68 *Southern California Law Review* 1631-
22. O'Hara, Erin. 1993. "Social Constraint or Implicit Collusion? Toward a Game Theoretic Analysis of Stare Decisis," 24 *Seton Hall Law Review* 736-778.
23. Levi, Edward H. 1948. *An Introduction to Legal Reasoning* 9-25.
24. Parisi, Francesco and Fon, Vincy. 2008. *The Economics of Lawmaking* (New York, NY: Oxford University Press).
25. Priest, George L. 1977. "The Common Law Process and the Selection of Efficient Rules." 6 *Journal of Legal Studies* 65-82.
26. Posner, Richard A. 1972. *Economic Analysis of Law* (1st ed. Boston: Little Brown).
27. Posner, Richard A. 1990. *The Problems of Jurisprudence*.
28. Posner, Richard A. 1996. *Federal Courts: Challenge and Reform*
29. Posner, Richard A. 2006. *Economic Analysis of Law* (6th ed. Boston: Little Brown)

30. Rubin, Paul H. 1977. "Why is the Common Law Efficient?" 6 *Journal of Legal Studies* 51-63.
31. Segal, Jeffrey & Spaeth, Harold J. 2002. The Supreme Court and the Attitudinal Model Revisited)
32. Sterns, Maxwell L. 1995. "Standing Back from the Forest: Justiciability and Social Choice." 83 *California Law Review* 1309-1413.
33. Sunstein, Cass R. et al., 2006. Are Judges Political? An Empirical Analysis of the Federal Judiciary
34. Sunstein, Cass R. 1991. One Case At A Time: Judicial Minimalism on the Supreme Court (Boston: Harvard University Press).
35. Songer, Donald, Segal, Jeffrey A. & Cameron, Charles M. 1994. "The Hierarchy of Justice: Testing a Principal-Agent Model of Supreme Court-Circuit Court Interactions." 38 *American Journal of Political Science* 673-
36. Spitzer, Matthew & Talley, Eric. 2009. "Panel Effects and Judicial Auditing." mimeo, University of Southern California
37. Zywicki, Todd. 2003. "The Rise and Fall of Efficiency in the Common Law: A Supply-Side Explanation." 97 *Northwestern Law Review* 1551-163.

Appendix

Proof of Proposition 1

First, we need to write the formula for the expectation $E_t V(H_{t+1})$:

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

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 $b_t \leq R_t$:

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

They can be rewritten as

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

$$\begin{aligned} 0 &= -D + L \frac{F(R_t) - F(b_t)}{F(R_t) - F(W_t)} - \delta V(W_t, R_t) \\ &+ \delta V(b_t, R_t) \frac{F(R_t) - F(b_t)}{F(R_t) - F(W_t)} + \delta V(W_t, b_t) \frac{F(b_t) - F(W_t)}{F(R_t) - F(W_t)} + \frac{(\lambda_t - \mu_t^b)}{g(b_t)} \end{aligned} \quad (3)$$

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

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

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

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

$$\frac{\partial \Phi_a}{\partial a_t} \partial a_t + \frac{\partial \Phi_a}{\partial b_t} \partial b_t + \frac{\partial \Phi_a}{\partial z} \partial z = 0$$

$$\frac{\partial \Phi_b}{\partial a_t} \partial a_t + \frac{\partial \Phi_b}{\partial b_t} \partial b_t + \frac{\partial \Phi_b}{\partial z} \partial z = 0$$

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

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

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

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:

$$\text{sign} \left(\frac{\partial a_t}{\partial z} \right) = \text{sign} \left(\frac{\partial \Phi_a}{\partial z} \right)$$

$$\text{sign} \left(\frac{\partial b_t}{\partial z} \right) = \text{sign} \left(\frac{\partial \Phi_b}{\partial z} \right)$$

The proposition follows from the following conditions:

$$\begin{array}{ll}
\frac{\partial \Phi_a}{\partial D} > 0; & \frac{\partial \Phi_b}{\partial D} < 0 \\
\frac{\partial \Phi_a}{\partial L} < 0; & \frac{\partial \Phi_b}{\partial L} > 0 \\
\frac{\partial \Phi_a}{\partial \delta} < 0; & \frac{\partial \Phi_b}{\partial \delta} > 0 \\
\frac{\partial \Phi_a}{\partial W_t} > 0; & \frac{\partial \Phi_b}{\partial W_t} < 0 \\
\frac{\partial \Phi_a}{\partial R_t} < 0; & \frac{\partial \Phi_b}{\partial R_t} > 0
\end{array}$$

■

Proof of Proposition 2

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:

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

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

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

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

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

Assume $M = 9$, $D = 1$, and $L = 5$, so that learning eventually converges, $\lim_{t \rightarrow \infty} a_t = \lim_{t \rightarrow \infty} b_t = 0$. 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.

■

Proof of Proposition 3

The first order conditions can be rewritten as

$$0 = D - L \frac{F(a_t) - F(W_t)}{F(R_t) - F(W_t)} + \delta V(W_t, R_t) - \delta V(a_t, R_t) \frac{F(R_t) - F(a_t)}{F(R_t) - F(W_t)} - \delta V(W_t, a_t) \frac{F(a_t) - F(W_t)}{F(R_t) - F(W_t)} - \frac{(\lambda_t - \mu_t^a)}{g(a_t)} \quad (4)$$

$$0 = -D + L \frac{F(R_t) - F(b_t)}{F(R_t) - F(W_t)} - \delta V(W_t, R_t) + \delta V(b_t, R_t) \frac{F(R_t) - F(b_t)}{F(R_t) - F(W_t)} + \delta V(W_t, b_t) \frac{F(b_t) - F(W_t)}{F(R_t) - F(W_t)} + \frac{(\lambda_t - \mu_t^b)}{g(b_t)} \quad (5)$$

Suppose $a_t = W_t < R_t$. Then (2) becomes $D = \frac{(\lambda_t - \mu_t^a)}{g(a_t)}$, which can only be satisfied if $b_t = a_t$. Equation (3) becomes $D - L = \frac{(\lambda_t - \mu_t^b)}{g(b_t)}$, 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 $D = \frac{(\lambda_t - \mu_t^b)}{g(a_t)}$, which requires $a_t = b_t$, while (2) becomes $D - L = \frac{(\lambda_t - \mu_t^a)}{g(a_t)}$, 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., $a_t \rightarrow \theta$, $b_t \rightarrow \theta$) 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

$$F(a_t) = F(b_t) = \frac{F(R_t) + F(W_t)}{2} \quad (6)$$

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

$$2D - L + 2\delta V(W_t, R_t) - \delta V(a_t, R_t) - \delta V(W_t, a_t) - \frac{2\lambda_t}{g(a_t)} = 0 \quad (7)$$

$$-2D + L - 2\delta V(W_t, R_t) + \delta V(b_t, R_t) + \delta V(W_t, b_t) + \frac{2\lambda_t}{g(b_t)} = 0. \quad (8)$$

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 $2D < L$. It follows that it cannot be $a_t = b_t$, and hence learning will never stop if $2D < L$.

Now suppose $a_t \neq b_t$ for all values of $W_t < R_t$, and as a consequence, $\lambda_t = 0$. Using this and $\mu_t^a = \mu_t^b = 0$, by subtracting (3) from (2) we obtain

$$\begin{aligned}
0 &= 2D - L \left(1 - \frac{F(b_t) - F(a_t)}{F(R_t) - F(W_t)} \right) + 2\delta V(W_t, R_t) \\
&\quad - \delta V(a_t, R_t) \frac{F(R_t) - F(a_t)}{F(R_t) - F(W_t)} - \delta V(b_t, R_t) \frac{F(R_t) - F(b_t)}{F(R_t) - F(W_t)} \\
&\quad - \delta V(W_t, a_t) \frac{F(a_t) - F(W_t)}{F(R_t) - F(W_t)} - \delta V(W_t, b_t) \frac{F(b_t) - F(W_t)}{F(R_t) - F(W_t)}
\end{aligned} \tag{9}$$

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 (9) is greater than $2D - L + \varepsilon$. It follows that if $2D > L$, then (9) 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 $2D > L$.

■