

Endogenous Cases and the Evolution of the Common Law*

Giri Parameswaran[†]

October 12, 2017

Abstract

I develop a dynamic model of judge-made law in which the ideal legal rule is unknown but can be learned. In contrast to existing papers, the flow of cases heard by the court is affected by the court's prior decisions. The model highlights the significance of this feedback in explaining when and why the court will write broader or narrower opinions, and the long-run properties of common law. In equilibrium, the law settles endogenously, since the incentives to make legally controversial choices disappear as the law evolves. Settled law exhibits residual uncertainty and ambiguity, and potentially implements inefficient outcomes.

JEL Codes: K10, K40

Key Words: Law and Economics, Incompleteness of Law, Judge-Made Law, Evolution of Legal Rules

*I would like to thank Charles Cameron, James Fogarty, John Londregan, Stephen Morris, seminar participants at Harvard, Haverford, Princeton, UC San Diego, and the 2014 American Law and Economics Association Meetings, and anonymous referees for their helpful comments.

[†]Assistant Professor, Department of Economics, Haverford College, 370 W Lancaster Rd, Haverford PA 19041. Phone: (610) 896-2905 Email: gparames@haverford.edu

1 Introduction

In common law systems, courts play an important role in creating or clarifying the law in situations where it is silent or vague. In doing so, judges are often criticized for ‘legislating from the bench’, but their law creation prerogative differs from that of legislators and other policy makers in important ways. For example, judges cannot proactively initiate ‘policy’ changes – they are limited to resolving actual controversies that arise when the existing law is unclear. Institutional norms, such as the doctrine of *stare decisis* or respect for precedent, oblige judges to extend the law in a manner consistent with previous judicial decisions. And even if they could freely amend doctrine, courts cannot directly bring their desired outcomes to bear; at most they can create incentives for agents to behave one way rather than another.

These institutional features distinguish ‘policy making’ in the judicial setting, and affect how the common law evolves. For example, when deciding cases, courts tend not to declare sweeping legal rules, but rather make incremental changes to the existing body of law. These changes are described as ‘narrow’ if the doctrine is extended only insofar as is necessary to resolve the instant case, and ‘broad’ if it implicates other cases as well. The incremental approach is beneficial in that it avoids entrenching potentially costly mistakes into the law.¹ But in failing to adequately clarify law that is vague or has gaps, the incremental approach may also deter agents from making socially efficient choices.²

Courts thus face a trade-off between writing broader rules that reduce uncertainty about the law, and narrow opinions that defer policy making to future (potentially better informed) courts. The ‘option value’ of deferral, in turn, depends on the future stream of cases that the court will likely decide. But the court’s docket is endogenous to its rules; since agents

¹For example, in *City of Ontario v Quon* (560 U.S. 746 (2010)), Justice Kennedy opined that: “Prudence counsels caution before the facts in the instant case are used to establish far-reaching premises... A broad holding ... might have implications for future cases that cannot be predicted.”

²In his dissenting opinion in *Blakely v Washington* (542 U.S. 296 (2004)), Justice Breyer argued: “But this case affects tens of thousands of criminal prosecutions... Federal prosecutors will proceed with those prosecutions subject to the risk that all defendants in those cases will have to be sentenced, perhaps tried, anew. Given this consequence and the need for certainty, I would not proceed further piecemeal.”

alter their behavior in response to changes in the law, the court’s current and prior decisions feed back into the future controversies that arise, and thus, into the court’s ability to affect the law in the future.

To explore these trade-offs, I build a dynamic model of judicial decision-making³ that extends the framework in Baker and Mezzetti (2012) by explicitly modeling the case-arrival process. For contextual concreteness, I consider the scenario of a firm (the tortfeasor) whose production process creates a (public) nuisance. Given the established law, the court should apply a balancing test that holds the firm liable if it ‘unreasonably interferes’ with the victim’s enjoyment of property – i.e. if the “gravity of harm outweighs the utility” of the conduct.⁴ The court is uncertain about the size of the harm, but can learn about it through the cases it hears. If found liable, the firm must pay compensatory damages⁵, the predominant remedy in the common law.⁶ Since the size of the nuisance is uncertain, so is the standard for ‘unreasonable interference’ to be applied. In this context, the court’s law creation role is to clarify which activity levels are reasonable and which are not. The court’s legal rule defines the set of cases that will be decided summarily; i.e. the realm of conduct which the court has classified as either definitely unreasonable or definitely not unreasonable. The legal rule may leave some cases unclassified; if so, the rule is silent as to the tortfeasor’s liability status.

³The model is one of the judiciary insofar as it incorporates the above mentioned features that are particular to the courts. (These include the inability to proactively initiate policy changes, the requirement that changes be consistent with existing doctrine, etc.) These constraints typically do not apply to other rule-making bodies, but to the extent that they do, the model’s insights can be extended more generally.

⁴Restatement (Second) of Torts §826.a. (Am. Law Inst., 1979).

⁵When damages are awarded, the standard practice is to compensate victims for harms suffered. “The fundamental principle of damages is to restore the injured party, as nearly as possible, to the position he would have been in had it not been for the wrong of the other party” (*United States v Hatahley*, 257 F.2d 920 (1958)). (See also Laycock (2012). Feldman (1996) lists other cases stating this principle.)

⁶Courts have applied both the common law remedy of damages and the equitable remedy of injunction in different cases, and the question of which remedy to apply is as yet unsettled. The Second Restatement of Torts (§826.b.) allows for damages where the balancing test is not met (i.e. the gravity of the harm does not outweigh the utility to the actor) but nevertheless “the harm caused by the conduct is serious”. The law and economics approach similarly recommends that damages be applied when there is a social interest in the continuation of the activity and transaction costs make Coasian bargaining prohibitive (see, for example, Cooter and Ulen (2011) or Bryson and Macbeth (1972)). Courts have adopted this view in a sequence of cases, most famously in *Boomer v Atlantic Cement Co.* (26 N.Y.2d 219, 309 N.Y.S.2d 312 (1970)). (See also *Northern Indiana Public Serv. Co. v Vesey* (210 Ind. 338, 200 N.E. 620 (1936)), *Maddox v International Paper Co.* (47 F.Supp. 829 (W.D. La. 1942)) amongst others.)

When confronted with a case in the silent region, the court examines the relevant evidence, learns whether the firm's output was unreasonable or not, and updates its beliefs about the size of the nuisance. In rendering its decision, the court must write an opinion that extends the legal rule in a way that resolves the case at hand. The revised legal rule will be narrow if the new set of cases to be decided summarily is limited to those whose ideal disposition became known to the court, given the new information. By contrast, the revised rule is broad if the set of summarily decided cases includes some whose ideal disposition remains uncertain. Thus, broad decisions potentially misclassify cases, and entrench these mistakes into law. The firm makes its output decision in every period, based on the anticipated costs and benefits. Under this framework, the common law only evolves if a case arises in the silent (or ambiguous) region of the law. This, in turn, will only occur if the legal rule provides the firm with the right incentives to make a choice whose legal consequence is uncertain (i.e. to 'experiment'). Hence, the court's ability to learn and revise the law is constrained by its ability to make experimentation incentive-compatible for the firm.

The model's setting within the nuisance context is made purely to provide concreteness to the analysis. The insights can be applied to any setting in which the court holds the defendant liable (and awards damages) according to a threshold (negligence) rule. For example, when determining the permissibility of policy conduct under the fourth amendment, the court must balance the state's prosecutorial interests against individuals' privacy concerns (the extents of which may be uncertain). The plaintiff will be awarded damages only if the search was deemed to be unreasonably invasive. Similarly, a tortfeasor will be held liable for damages under the tort of negligence only if he fails to take sufficient precaution.⁷

An important ingredient in my model is the feedback between legal-rules and the flow of cases that the court hears. This feedback is operationalized by explicitly modeling the behavior of the firm, and assuming that cases arise as a consequence of rational firm choices. This

⁷The negligence context is slightly different to nuisance in that taking precaution is costly to the tortfeasor and beneficial to the victim. The formal model can be adapted to this framework by applying the transformation $T(x) = 1 - x$ to all variables.

stands in strong contrast to many existing models of law creation (including Gennaioli and Shleifer (2007), Niblett (2013) , and Baker and Mezzetti (2012)) in which agent behavior is unmodeled. Those papers typically assume that cases are random draws from a legal-rule-invariant distribution. By contrast, the case-generating process in my model is endogenous to the legal rule. Since the legal rule can only evolve when a controversial case arises, the case generating process potentially limits the scope for learning and future rule making by the court. The court cannot learn whether certain choices are appropriate or not without having first observed the consequences of those choices. But the court does not make output choices; the firm does. The court does not learn by experimenting — it learns by having *others* experiment. Its learning is limited by its ability to make experimentation incentive-compatible. The endogeneity of case flow is thus essential to understanding the court's trade-off in writing narrower or broader opinions. The opportunity cost of broader opinions (in terms of learning opportunities forgone) depends crucially on the expected future flow of cases.

This paper makes several contributions. First, it characterizes the nature of decision-making by agents facing an incomplete legal rule. Given an arbitrary legal rule, the firm will generically choose either the largest 'safe' output (guaranteed to not incur a penalty) or experiment in the ambiguous region of the law. By experimenting, the firm realizes a larger share of available (pre-penalty) profits, but potentially incurs a penalty. The safe output will be chosen if the opportunity cost of pre-penalty profits forgone is smaller than the expected penalty from experimentation. This will be true if the safe output is relatively large. Conversely, experimentation is preferred when the safe output is sufficiently low. Additionally, the firm's optimal experimental output is characterized by over-deterrence; the firm skews its choice towards lower outputs that are less likely to be found unacceptable by the court. This over-deterrence is socially inefficient and reflects the uncertainty costs of implementing ambiguous legal rules. In some cases, the deterrent effect is so large as to preclude experimentation entirely. Opportunities for learning and further legal evolution are not guaranteed.

Naturally, the firm's choice will be affected by the breadth of the legal rule. Broader permissive opinions (which expand the scope for penalty-free production), for example, will tend to make the safe-output more attractive, and thus skew the firm's choice towards the safe option. The court can potentially entice the firm to make a range of different choices by writing appropriately broad permissive opinions. By contrast, the firm's choice generically does not depend on the breadth of the restrictive opinion (which determines the region over which the firm will definitely be penalized), since the firm will tend to make choices far from the region where the penalty is guaranteed. Hence, the two opinion types are not equally efficacious in affecting firm behavior. The court has much greater scope to affect firm behavior by writing broad permissive opinions than broad restrictive ones, and should thus be more inclined to make use of the former than the latter.

Second, the model characterizes the trade-off to the court between writing narrow and broad opinions. Narrow opinions make experimentation more likely, and this facilitates learning and the evolution of law. However, since ambiguity causes over-deterrence, narrow opinions result in socially inefficient short-run outcomes. The court will write narrower opinions that induce experimentation when the size of this short-run 'bias' is small. By contrast, if the bias is large (or if experimentation is deterred entirely), the court will favor broader opinions that expand the set of risk-free choices to include the efficient output. Courts thus write broader permissive opinions when the long-run gains from experimentation and learning are outweighed by the losses stemming from short-run inefficiencies.

Third, the model characterizes the evolution of legal rules through time. I show that the size of the relative 'bias' from experimentation increases over time as learning takes place. Hence, the court will tend to favor narrower rules when the law is relatively nascent, and then switch to broader rules as the law evolves and the extent of uncertainty falls. The law 'settles' (in the sense that no future controversies arise) within a finite number of cases — courts do not continue to refine the law *ad infinitum*. Moreover, the law settles before the court is able to

perfectly learn the ideal rule; there will be residual uncertainty as to the efficacy of the law and the settled law will almost certainly be *ex post* inefficient. Furthermore, with positive probability, the settled law will be *ex ante* inefficient, as well. This will occur if the law settles abruptly (i.e. before the court would have it so), prior to the court having the opportunity to appropriately amend the legal rule to entice *ex ante* efficient behavior. Anticipating the potential for this latter outcome, courts will precautionarily write slightly broad permissive opinions, even when learning is desirable, as a hedge against a legal dynamic that settles too soon. The legal rule is, thus, characterized by preemption.

This paper contributes to a growing literature that studies the role of learning in judicial decision-making (see, in particular, Hadfield (1991), Baker and Mezzetti (2012), Fox and Vanberg (2013) and Callander and Clark (2017)). The model closest to mine is Baker and Mezzetti (2012) — indeed, the court’s information structure and policy tools are identical across both models. In that model, when a case arises whose ideal disposition is unknown, the court must decide whether to summarily dispose of the case, or to conduct a costly investigation to learn the ideal outcome. The court faces a trade-off between incurring adjudication costs and (potentially) deciding cases incorrectly. In equilibrium, the court will summarily dispose of cases that are ‘similar enough’ (i.e. it will construe existing laws at least somewhat broadly) and investigate the remainder (whereupon learning takes place and the law evolves). Naturally, the notion of ‘close enough’ is relative to the size of adjudication costs. Baker and Mezzetti show that, as long as costs are not too large, the law will evolve incrementally and eventually converge to the ideal legal rule. Important to their efficiency claim is the assumption that the case generating process is stochastic and unaffected by the legal rule. This ensures that cases will always arise in the ambiguous region, and so future opportunities to learn and revise existing rules are always guaranteed.

This paper departs from Baker and Mezzetti (2012) in that the case-generating process (governed by the rational choices of ‘agents’) is responsive to the legal rule. Hence, rather

than facing a random exogenous stream of cases, the court's docket will be endogenous to the rules it makes. The evolution of common law is now constrained by the willingness of agents to make controversial choices that require adjudication. The adjudication costs in Baker and Mezzetti (2012) are replaced with the implicit (bias) costs of experimentation being subject to the agent's incentive-compatibility constraint. In contrast to their paper, I show that the optimal breadth of opinions depends on the type of opinion being written — permissive opinions should be optimally broad, whilst restrictive opinions are optimally narrow. Further, I show that the implicit (bias) cost is itself endogenous and increases as the law evolves, creating a dynamic where the law evolves for a finite period before settling. Moreover, legal evolution stops before the court can implement the ideal legal rule.

Hadfield (1991) and Fox and Vanberg (2013) both study situations where the agent's behavior is responsive to the court's rules. Hadfield (1991) argues that the differential response by heterogeneous agents to a given legal rule will cause future courts to hear a skewed sample of cases, which biases learning, and causes legal evolution to be inefficient. This paper demonstrates that heterogeneity is not crucial to the story — the court's learning will be constrained even when it faces a single representative firm.

Fox and Vanberg (2013) present a model in which a single agent rationally responds to the court's rules. They show that, by writing broad opinions, the court can force the policy maker to experiment in a region of the policy space that facilitates more efficient learning about the ideal legal rule. This paper distinguishes itself from Fox and Vanberg (2013) in two important ways. First, theirs is a reduced-form model of agent behavior in which experimentation is always feasible. By contrast, this paper presents a structural model of preferences, and shows that broad opinions tend to inhibit (rather than promote) efficient learning, and become most effective as opportunities for future learning disappear. Second, Fox and Vanberg consider a one-sided model of law creation (the court either determines a policy to be definitely unconstitutional or not), and so their model does not display the

asymmetries in opinion efficacy that this paper highlights.

The above models all considered simple doctrinal spaces characterized by threshold rules. Callander and Clark (2017) study decision-making by a court in a complicated world, where the ideal rule is non-monotone, causing the court’s learning to be localized. Their model micro-founds the common legal practice of reasoning by analogy, and demonstrates the path-dependence of law, which similarly arises in my model. Additionally, in a hierarchical structure where the superior court can choose which lower court decision to review, Callander and Clark (2017) examine the implications of case-selection on efficient learning.

Other hierarchical judicial models explore the implications of political constraints on decision-making. Lax (2012) characterizes optimal rule-making when the superior court can only imperfectly monitor subsequent decisions by lower courts. Bueno de Mesquita and Stephenson (2002) micro-found the norm of *stare decisis* in a related scenario where superior courts can only imperfectly communicate their ideal rule to lower courts. Staton and Vanberg (2008) motivate vague or incomplete rule-making in the context of separation-of-powers, when compliance by other political actors is not guaranteed. My model differs in that the relevant constraint derives, not from political actors, but from the behavior of underlying agents whose behavior generates future cases.

Other dynamic models of judicial decision-making investigate how the common law evolves, and the implications for efficiency, when judges (or courts) have heterogeneous preferences (see Gennaioli and Shleifer (2007), Ponzetto and Fernandez (2008), and Niblett (2013)). Gennaioli and Shleifer (2007), for example, provide foundations for the “Cardozo Theorem”, which states that the individual biases of judges tend to wash out as law is created piece-meal, and that legal evolution is, on average, efficiency enhancing. Ponzetto and Fernandez (2008) similarly show that the common law converges towards more efficient rules, making it more effective in the long-run than statutory rule making. These models typically do not involve any uncertainty about the ideal legal rule — although the game is dynamic, there is

no role for learning.

This paper contributes to the literature on the efficiency of common law more broadly. In a seminal paper, Posner (1973) argued that the decisions of efficiency-minded judges would tend to cause the common law to produce efficient rules. Subsequent papers have examined why the common law may tend to efficiency, even if judges are imperfect in their motives, information or execution. Priest (1977) and Rubin (1977) provide a ‘demand-side’ argument, that inefficient legal rules are more likely to be litigated and subsequently overturned, than efficient ones.⁸ Cooter, Kornhauser and Lane (1979) show that incremental rule-making by courts can converge upon efficient rules, even if individual courts are imperfectly informed about the ideal rule. Hadfield (2011) asks whether legal rules will dynamically adapt to new or local conditions. By and large, the primary force that drives the law, in these papers, is case selection (i.e. which cases are litigated). The opposite effect, of the law driving the sorts of cases that arise, is typically left unexplored (although, see Png (1987)). My model considers both effects — indeed, modeling the cyclicity of these effects is an important innovation of this paper.

The remainder of this paper is organized as follows: Section 2 presents the formal model. Section 3 characterizes optimal decision-making by the firm, given an arbitrary legal rule. Section 4 studies optimal decision-making by the court, taking the firm’s anticipated response as given, and characterizes the long-run properties of the common law. Section 5 presents several extensions.

⁸Cooter and Kornhauser (1980) provide a formalization of these arguments. By contrast, Hylton (2006) shows that information asymmetries between litigants can cause biased rules, which favor the more informed party, to evolve. Zywicki (2002) stresses the importance of ‘supply-side’ factors, such as institutional rules and norms, in determining the efficiency of common law.

2 Model

There is a profit-maximizing firm that must choose a quantity of output to produce. The firm's gross profit from producing output q is $\pi = q - \frac{1}{2}q^2$.⁹ The profit-maximizing output is $q_{max} = 1$. The firm's production creates a nuisance which harms a victim. The size of the harm is θq , where $\theta \in (0, 1)$ is the constant marginal harm. The socially efficient output level is $q^{eff} = 1 - \theta < q_{max}$, which implies that an unregulated firm will unreasonably interfere with the victim's enjoyment of property.

The size of the marginal harm θ is unknown, and parties share common beliefs about its value. For simplicity, I assume that beliefs at the beginning of period t are uniformly distributed on the interval $[l_t, u_t]$ where $0 < l_t \leq u_t < 1$.¹⁰ For technical convenience, I assume that $u_0 < \hat{u}_0$, where $\hat{u}_0 > 0.9$ is a threshold which is formally defined in the proof of Lemma 2. This restriction implies that the harm is not so large an injunctive remedy would be appropriate.

There is an efficiency-minded court that seeks to maximize social welfare. The victim may bring a case before the court if it experiences harm, and will do so whenever there is a positive probability of being compensated. If held liable, the firm must compensate the victim for its expected harm, which is the standard common law remedy.

A legal rule is a pair of thresholds (λ_t, μ_t) (with $\lambda_t \leq \mu_t$) that prescribe how various cases ought to be decided. According to the rule, firms producing $q \leq \lambda$ should not be held liable, whilst firms producing $q > \mu$ should.¹¹ The legal rule is silent as to the firm's liability from producing output $q \in (\lambda, \mu]$. It is, in effect, an incomplete negligence rule, and the thresholds

⁹The functional form is chosen for simplicity. The qualitative results can be shown to generalize for any strictly concave π .

¹⁰As will become clear, the uniform distribution is the conjugate prior to the likelihood function implied by the learning technology. Hence, uniform initial priors ensure uniform beliefs throughout the dynamic game.

¹¹The legal rule is implicitly monotone — if the court would ideally penalize outcome q , then it should naturally penalize any outcome $q' > q$, since these impose even greater harms on the victim. Given the model framework, such threshold rules are consistent with rational rule making by the court.

λ and μ are permissive and restrictive thresholds, respectively, which determine regions of *per se immunity*, *ambiguity* and *strict liability*. I assume, without loss of generality, that $1 - u_t \leq \lambda_t \leq \mu_t \leq 1 - l_t$.¹²

Since courts are bound by prior decisions, cases arising in regions where the legal rule is not silent will be decided mechanically. When choosing outputs in these regions, the firm has certainty about the legal implications of its choices. By contrast, a firm that produces in the *ambiguous* region of the case space, $q \in (\lambda, \mu)$ faces uncertainty about whether it will be penalized or not. I say that a firm making such a choice *experiments*¹³, and it is in this region that there is scope for learning and law creation. When such a case arises, since it cannot be summarily disposed, the court fully investigates the case — it hears expert testimony, consults amici, etc. — and perfectly learns whether the chosen output level was above or below the socially efficient level. Given the assumptions of the model, output q will be found to be unreasonable interference if $q > q^{eff} = 1 - \theta$.¹⁴

Upon receiving this new information, the court does two things. First, it (and the firm) updates its beliefs about the size of the harm according to Bayes' Rule. For example, suppose at time t , the court's prior beliefs were $\theta^{prior} \sim U[l_t, u_t]$ and the firm's output q_t was found to be unreasonable. The court learns that $\theta > 1 - q_t$; the learning mechanism truncates the support of the belief distribution. Posterior beliefs remain uniform, and so $\theta^{post} \sim U[1 - q_t, u_t]$. Since the court carries these beliefs forward into the next period, $l_{t+1} = 1 - q_t$ and $u_{t+1} = u_t$. Similarly, if q_t was found to not be unreasonable, then the court

¹²Since $\theta_t \leq u_t$, it is commonly known that any $q \leq 1 - u_t$ cannot be unreasonable. Since an efficiency-minded court would never want to punish such behavior, all agents can intuit that any $q < 1 - u_t$ must be in the region of *per se immunity*, even if the court hasn't explicitly said so. Similarly, all agents can intuit that any $q > 1 - l_t$ will be definitely punished.

¹³A semantic note is in order. 'Experimentation' typically suggests that a risky choice is made with a view to acquire information that benefits future decision-making. As will be made clear in section 3, the firm is not forward-looking, and so its risk taking is not motivated by a desire for information. However, the court is forward-looking, and it values the information derived from the firm's risk-taking, when crafting legal rules. Hence, we can think of the court being the party that actually experiments, mediated by the firm's actions.

¹⁴The learning technology implicitly assumes that, after inspecting the evidence, the court can only observe whether the firm's output was excessive or not. It cannot observe the size of the deviation from the efficient level.

learns that $\theta < 1 - q_t$, and updates its beliefs to $\theta^{post} \sim U[l_t, 1 - q_t]$, which implies that $l_{t+1} = l_t$ and $u_{t+1} = 1 - q_t$. In the case of over-production, these updated beliefs also form the basis for assessing the appropriate penalty to impose on the firm. Since the plaintiff is entitled to compensatory damages, the court assigns a penalty equal to the size of the expected harm: $P(q_t, l_{t+1}, u_{t+1}) = q_t E_{t+1}[\theta] = \frac{1}{2}(l_{t+1} + u_{t+1})q_t = \frac{1}{2}(1 - q_t + u_t)q_t$.¹⁵

Second, the court must announce an updated legal rule $(\lambda_{t+1}, \mu_{t+1})$ that is consistent with the disposition of the case. For example, if q_t was found to be reasonable, the court must update its legal rule such that $\lambda_{t+1} \geq q_t$, which ensures that future courts will also find q_t -like cases acceptable. The court may either announce a *narrow rule* ($\lambda_{t+1} = q_t$) or a *broad rule* ($\lambda_{t+1} > q_t$).¹⁶ A narrow rule extends the settled portion of the law in the smallest possible way that is consistent with the disposition of the case, and the court's new information. By contrast, a broad rule extends the law in a way that governs additional cases whose ideal disposition is not known to the court.

To explain the decision in a case when the firm was held not liable, it suffices for the court to update its permissive threshold. In fact, there is no way that it could amend its restrictive threshold (which determines the region of *strict liability*) that would explain the case disposition. Although, in principle, the court may write an opinion that purports to amend the restrictive threshold, this portion of the opinion would necessarily be understood to be *obiter dicta*, and therefore lacking the weight of precedent. Since future courts are not bound by *dicta* (and since, in this model, *dicta* itself conveys no additional private information from the court to the litigants), I constrain the court to only revise the threshold relevant to explaining its holding in the given case. Accordingly, if it finds a given output to be unacceptably large, the court must revise its restrictive threshold, but may not revise its permissive one, and vice versa. (The alternative setup, where both thresholds may be varied, is presented as an extension.)

¹⁵For clarity, the penalty depends upon posterior beliefs at time t , which are equivalent to prior beliefs at time $t + 1$.

¹⁶This definition is consistent with Baker and Mezzetti (2012) and Fox and Vanberg (2013).

The timing of the game is as follows: at the beginning of period t , the environment is characterized by a legal rule (λ_t, μ_t) and beliefs parameterized by (l_t, u_t) . Given this environment, the firm chooses its output q_t . If $q_t \leq \lambda_t$ or $q_t > \mu_t$, the law is applied mechanically. There is no learning and the period ends. The period $t + 1$ environment remains identical to the period t environment. By contrast, if $q_t \in (\lambda_t, \mu_t]$, the court hears the case. Learning occurs, and beliefs are updated, as described above. The court updates the legal rule for the following period, and the period ends.

A strategy for the firm is a function $q(l, u, \lambda, \mu)$ that assigns a feasible output to every environment. A strategy for the court is a pair of functions $\lambda(l, u, \mu)$ and $\mu(l, u, \lambda)$ which characterize the relevant location of the threshold being updated. I focus on Markov Perfect Equilibria, which restricts players to condition their strategies only on pay-off dependent variables. The effect of this assumption is that each player must make the same choice when facing the same environment. In particular, time-dependent or history-dependent strategies are disallowed.

3 Firm's Choice

A key insight of this paper is that dynamically optimal rule making depends on the future stream of cases to be decided. This expected case flow is in turn determined by the optimal behavior of agents, given the rules they face. Accordingly, the analysis of this model is in two parts. In this section, I characterize the firm's optimal choice, in any arbitrary environment. This describes the relationship between the current rule and future controversies. In Section 4, I characterize optimal rule making by a court that anticipates the behavior (and hence, cases) that will likely follow.

The firm maximizes expected net-of-penalty profits in the current period, taking the law

and the court's expected behavior (where the law is silent), as given.¹⁷ Given the Markovian assumption, time subscripts are omitted in the analysis below.

The firm's expected net profit is $\Pi = q - \frac{1}{2}q^2 - qf(q)$, where $f(q)$ denotes the expected per-unit fine.

$$f(q) = \begin{cases} 0 & q \leq \lambda \\ \frac{u^2 - (1-q)^2}{2(u-l)} & \lambda < q \leq \mu \\ \frac{u+l}{2} & q > \mu \end{cases}$$

In the ambiguous region, the expected per-unit fine is the product of the probability of having over-produced $\frac{u-(1-q)}{u-l}$ and the (updated) expected marginal harm $\frac{u+(1-q)}{2}$. It is easily verified that $f(q)$ is weakly increasing in q , and strictly increasing in q over the ambiguous region. Hence, the expected penalty is increasing in the likelihood that the firm's choice is found to be unreasonable.

The firm's marginal profit is given by:

$$\Pi'(q) = \begin{cases} 1 - q & q < \lambda \\ 1 - q - \frac{u^2 - (1-q)^2}{2(u-l)} - \frac{1-q}{u-l}q & \lambda < q < \mu \\ 1 - q - \frac{u+l}{2} & q > \mu \end{cases}$$

These functions exhibit two important discontinuities at the thresholds λ and μ . First, the profit function Π is discontinuous whenever the legal rule is broad (i.e. there is a discontinuity at $q = \lambda$ whenever $\lambda > 1 - u$, and at $q = \mu$ whenever $\mu < 1 - l$). To see why, suppose the permissive threshold λ is broad. Then, as output increases from $1 - u$ to λ , the probability that the firm has created an unreasonable interference begins to grow, although the firm faces no penalty. However, once output exceeds λ , the firm potentially becomes liable for damages, and the probability of being penalized jumps discontinuously from 0 to its true

¹⁷In effect, I assume that the court faces a sequence of short-run firms. This assumption keeps the analysis tractable.

value. Similarly, if the upper threshold μ is construed broadly, the probability of being penalized jumps discontinuously from its true level to 1, as output increases above μ . Here we see the mechanism by which opinion breadth affects firm choice. Broad legal rules create a wedge between the probability that the firm ought to be penalized, and the probability that it actually will be. This wedge in probabilities naturally affects the incentives for the firm to choose amongst different outputs.

Second, the marginal profit function Π' is also discontinuous at the thresholds, and this is true regardless of whether the legal rules are broadly or narrowly construed. Here, the intuition is more straight-forward: the expected fine behaves differently in the different regions, and so the slopes are unlikely to coincide at the boundaries. In the region of *per se* immunity ($q \leq \lambda$), the marginal expected penalty is 0, whereas outside this region, the marginal penalty is strictly positive. In the region of strict liability, the average and marginal penalties are constant (and equal to $\frac{u+l}{2}$). By contrast, the average expected penalty is increasing in the ambiguous region, which implies that the marginal expected penalty exceeds the average penalty. Following an increase in output, the firm not only pays the average penalty on the marginal unit produced, but must also pay the higher average penalty on all infra-marginal units.

Given the above discussion, the firm's optimal choice can take one of four values, which I denote by q_λ , q_μ , q_A and q_E . The first two are associated with the aforementioned discontinuities and are akin to 'corner solutions'.¹⁸ The latter are interior solutions. q_A is the solution to the first order condition given the firm's incentives in the ambiguous region. Similarly, q_E is the solution to the first order condition given incentives in the strict liability region.¹⁹ Straight-forwardly, $q_E = 1 - \frac{u+l}{2}$, which corresponds to the *ex ante* socially efficient choice.²⁰

¹⁸The notation distinguishes the threshold from the firm's choice (which may be located at the threshold).

¹⁹Note carefully: there is no guarantee that q_A and q_E , so defined, are feasible, in the sense that they fall within their associated regions. It may be that $q_A \notin (\lambda, \mu]$ or that $q_E < \mu$. Obviously, if infeasible, such a choice cannot be optimal for the firm.

²⁰I distinguish q_E and q^{eff} . The former is the *ex ante* socially efficient output from the perspective of an imperfectly informed court. The latter is the true socially optimal output, given marginal harm θ .

Amongst these choices, q_λ is risk and penalty free, whereas the remaining options entail a positive probability of being penalized.

With this discussion in mind, I characterize the firm's optimal output as a function of the legal thresholds (λ, μ) and belief parameters (l, u) :

Proposition 1. *The firm's optimal choice is characterized as follows: There exists $\bar{\mu}(l, u)$ with $\bar{\mu} < q_A < q_E$, such that:*

1. *If $\mu \geq q_I$, then there exists $\bar{\lambda}(l, u) \in [1 - u, q_I)$, such that:*

$$q^* = \begin{cases} \lambda & \lambda \geq \bar{\lambda}(l, u) \\ q_A & \lambda < \bar{\lambda}(l, u) \end{cases}$$

2. *If $\bar{\mu}(l, u) < \mu < q_I$, then there exists $\tilde{\lambda}(l, u) \in [1 - u, \bar{\lambda}(l, u)]$ with $\tilde{\lambda}(l, u) \leq \bar{\lambda}(l, u)$, such that:*

$$q^* = \begin{cases} \lambda & \lambda \geq \tilde{\lambda}(l, u) \\ \mu & \lambda < \tilde{\lambda}(l, u) \end{cases}$$

3. *If $\mu \leq \bar{\mu}(l, u)$, then there exists $\tilde{\tilde{\lambda}}(l, u) = \min \left\{ 1 - u, 1 - \sqrt{1 - \left(1 - \frac{u+l}{2}\right)^2} \right\}$, which satisfies $\tilde{\tilde{\lambda}} \in [1 - u, \tilde{\lambda}]$, such that:*

$$q^* = \begin{cases} \lambda & \lambda \geq \tilde{\tilde{\lambda}}(l, u) \\ q_E & \lambda < \tilde{\tilde{\lambda}}(l, u) \end{cases}$$

Proposition 1 tells us three things. First, the firm will choose the safe output q_λ whenever λ is sufficiently large, and one of the non-penalty-free (hereafter 'risky') outputs, otherwise. Second, when choosing a risky output, the firm will choose q_A (the interior solution implied by incentives in the ambiguous region) whenever it is feasible (i.e. whenever $q_A \in (\lambda, \mu]$).

It turns out that this condition is guaranteed to be satisfied when the restrictive opinion is narrow (i.e. $\mu = 1 - l$). Third, if it is optimal for the firm to make a risky choice, but q_A is infeasible, it will choose q_E if the permissive threshold is sufficiently broad ($\mu < \bar{\mu}$), and q_μ otherwise.

Let us build intuition for these results. First, note that, in choosing the safe output q_λ , the firm faces a trade-off between avoiding the penalty on the one hand, and accepting lower pre-penalty profits on the other. The smaller is the safe output, the larger is the opportunity cost of these forgone profits. Hence, the firm will make the safe choice provided that λ is large enough. The thresholds $\bar{\lambda}, \tilde{\lambda}$ and $\tilde{\tilde{\lambda}}$ indicate how large the safe output must be before the safe choice is made, for each possible risky alternative.

Second, to see that q_A is most preferred amongst the risky choices, compare q_A and q_μ . Whenever feasible, q_A is the optimal choice in the ambiguous region. Since, by construction, q_μ is contained in the ambiguous region, it can only be preferred to q_A if q_A is infeasible. Similarly, compare q_A and q_E . Suppose the penalty from the ambiguous region was applied to q_E . By the previous argument, since q_A is optimal given this penalty scheme, it must be preferred to q_E . But, in fact, to be feasible, q_E must lie in the strict-liability region and incur the penalty for sure, which makes it even less preferred.

Several features of this equilibrium are worth noting. First, in most environments, the firm will either choose the safe output q_λ or the experimental output q_A . The remaining options will only be chosen if the restrictive threshold is construed very broadly. In the next section, I will show that this is never optimal, and should never arise along the equilibrium path.

Second, the experimental output q_A is socially inefficient (i.e. $q_A < q_E$). This is consistent with Calfee and Craswell (1984), who show that negligence rules can cause over-deterrence under conditions of uncertainty. Specific to this model, the firm more than internalizes the marginal harm since, in the ambiguous region, the marginal penalty exceeds the average penalty, and since compensatory damages imply an average penalty equal to the expected

harm.²¹ This observation demonstrates one of the tensions inherent in common law rule making — that ambiguity about the law causes agents to make socially inefficient choices. By reducing the extent of ambiguity, broader rules can potentially reduce the size of this distortion, *ceteris paribus*.

Additionally, with uniform beliefs, it is easy to see that a firm producing the socially efficient output will be equally likely to be penalized as not. Since the firm actually under-produces, it is less likely to be penalized rather than not. The firm’s experimentation is biased towards receiving ‘good news’ (i.e. not being penalized) than ‘bad news’.

Third, broad opinions are asymmetric in their effect on the firm’s output. The court can entice the firm to choose most outputs simply by writing a broad permissive opinion locating λ at that output level. (Formally, the court can induce any output $q \in [\bar{\lambda}, \mu]$ by locating λ at the desired output.) This follows since, absent penalties, the firm will want to produce more than is socially desirable and broad permissive opinions extend the scope for penalty-free production.²² Hence, the court has a lot of flexibility to guide the firm’s choices by writing appropriately targeted broad permissive opinions. By contrast, the court is much more limited in its ability to target the firm’s output by manipulating the restrictive threshold μ . Indeed, for any $\mu \in [q_A, 1 - l]$ (an interval which comprises more than half the region of uncertainty $[1 - u, 1 - l]$), the upper threshold has no impact on the firm’s choice. The restrictive threshold determines the region in which the firm will be penalized for sure. But the firm typically has little incentive to produce at this level, especially if there are sufficient opportunities for production at lower output levels that are associated with smaller (or zero) penalty. Unlike broad permissive opinions, where the firm typically has every incentive to

²¹In fact, as I show in section 5.1, this insight is not limited to instances in which courts apply compensatory damages. For any positive level of damages, there is a threshold extent of uncertainty below which the penalty causes over-deterrence.

²²A note of caution — it is not true that the firm’s output is monotonic in the breadth of the permissive threshold. For example, if $\lambda_1 < \bar{\lambda} < \lambda_2 < q_A$, then increasing the lower threshold from λ_1 to λ_2 causes output to fall from q_A to λ_2 . Writing a slightly broader opinion can entice the firm to switch from a mid-range output in the ambiguous region, to the smaller safe output. Monotonicity does hold whenever $\bar{\lambda} \leq \lambda_1 < \lambda_2$, in which case the higher threshold is associated with a larger output.

take advantage of the greater permissivity, broad restrictive opinions may simply make less appealing output choices that the firm would not have chosen anyway. Hence, to the extent that they write broad opinions at all, we should expect courts to make much greater use of broad permissive opinions than broad restrictive ones. For example, when ruling on the legality of police action, we may expect the court to write broader opinions in cases where it finds the police action to be appropriate, but narrower opinions when it seeks to limit police power.

Fourth, the breadth of the permissive opinion affects whether the firm will experiment, but not where it will do so. Whilst writing narrower or broader opinions affects the set of choices whose consequences are ambiguous, it does not affect the consequences of making those choices. However, broad permissive opinions do change the value of choosing the safe option, and so may make this choice more or less desirable relative to experimentation. This stands in contrast to Fox and Vanberg (2013), where an important purpose of opinion breadth is to guide the agent to experiment in more favorable regions of the case-space. Instead, this paper posits two different roles for broad opinions: (i) to expand the scope for risk-free behavior, and (ii) to affect the firm's trade-off between making safe and risky choices.

Fifth, as indicated by part 3 of Proposition 1, under certain conditions (i.e. when $\mu \in (\lambda, \bar{\mu})$), a very broad restrictive opinion *can* induce the firm to choose the *ex ante* socially efficient output q_E . However, this scenario is largely pathological, and will only trivially arise along the equilibrium path. To understand why, suppose the safe output q_λ is very low and that the court writes an extremely broad restrictive opinion locating μ very close to λ . Then, the *per se immunity* and ambiguous regions of the legal rule only apply at very low levels of output, and for all intents and purposes, the firm faces a *strict liability* regime. Obviously, a rule constructed in this way works against the spirit of a negligence-type rule in which the firm is only held liable if its conduct is deemed to be unreasonable. Moreover, the firm may

still prefer to choose the safe output level q_λ , which it will do if q_λ is not too distant from q_E . The assumption that $u_0 < \hat{u}_0$ ensures q_λ is large enough to prevent the pathological case from occurring along the equilibrium path.

Of these five features, the middle three in particular, highlight the main tensions between narrow and broad rules that animated the discussion in the introduction. When the law is ambiguous, agents tend to make socially inefficient choices. Broad opinions, by narrowing the region of ambiguity, can entice agents to make better choices. However, broad decision risk misclassifying cases, and with the norm of *stare decisis*, commit the court to implement inefficient rules in the long-run. Furthermore, by narrowing the scope for experimentation and making safe choices more desirable, broad opinions reduce the likelihood of learning through experimentation, which hinders the further evolution of law. Section 4 focuses on how the court manages this trade-off.

The remainder of this section is devoted to the comparative statics of the firm's choice with respect to the environment (l, u, λ, μ) . Proposition 1 showed that a sufficiently broad permissive rule ($\lambda \geq \bar{\lambda}(l, u)$) may preclude the possibility of experimentation. As the following Corollary demonstrates, this outcome may obtain even when the court implements the narrowest possible legal rule.

Corollary 1. $\bar{\lambda}(l, u) = 1 - u$ whenever $u - l < 1 - u$, and $\bar{\lambda}(l, u) > 1 - u$ otherwise.

Corollary 1 states that when the extent of uncertainty is small, the firm will choose the safe output, even if the legal rule is perfectly narrow. The intuition is that when $u - l$ is small, a given increment in output causes the probability of being penalized to increase by more than would the same increment when $u - l$ is large. Hence, the marginal penalty is larger when uncertainty is low, than when it is high. In high uncertainty environments, the firm can push into the ambiguous region with greater confidence of not being penalized, making such experimentation worthwhile. By contrast, when uncertainty is low enough,

the risks associated with experimentation become prohibitive, and the firm simply prefers the safe choice. Of course, how much uncertainty is ‘small enough’ is relative to the safe output level. When this is large, the firm will be less tolerant of introducing uncertainty into its profits since it is already realizing a large fraction of total available pre-penalty profits. Hence ‘small enough’ depends on the size of the safe level $1 - u$.

An immediate implication of Corollary 1 is that, in certain environments, experimentation, and thus opportunities for learning, will not be possible, regardless of the legal rule. Additionally, even when the firm can be induced to experiment, it may not wish to do so near the socially efficient level, or in the region of the policy space most advantageous to learning by the court. To formalize this, let $b(l, u) = \frac{q_E(l, u) - q_A(l, u)}{u - l}$ be an index capturing the ‘relative bias’ of the experimental output q_A , with respect to the *ex ante* efficient level $q_E = 1 - \frac{u+l}{2}$. The bias is normalized with respect to the size of the uncertain region. If $b = 0$, then the experimental and efficient outputs coincide. By contrast, if $b = \frac{1}{2}$, then $q_A = 1 - u$, and so the experimental output is at the boundary of the uncertain region. Since $1 - u < q_A < q_E$, it follows that $b \in (0, \frac{1}{2})$ whenever q_A is feasible. Clearly, as b increases, so does the size of the ‘distortion’ in the firm’s chosen output.

Corollary 2. *The firm’s experimental output q_A (whenever feasible) satisfies:*

1. $q_A(l, u)$ is decreasing in both its arguments.
2. $b(l, u)$ is decreasing in u , and there exists $\eta(l) > \frac{8+l}{10}$ such that b is increasing in l whenever $u < \eta(l)$.

The first part of Corollary 2 states that increases in either bound of the support of beliefs will cause the experimental output q_A to fall. Since the per-unit expected penalty is increasing in both l and u , an increase in either will cause the firm to reduce output at the margin. The second part of the Corollary shows that the size of the relative bias increases both when u

decreases and l increases (in the latter case, provided that u is not too large, as is assumed). Since, u can never increase and l can never decrease along the path of play, the Corollary implies that the relative bias in the firm's choice will increase over time as the law evolves. As uncertainty decreases, the firm will choose outputs that hew more closely to the safe output level than to the efficient (but uncertain) output.

I conclude this section, by briefly summarizing the main results. Given the environment it faces, the firm will typically locate its output choice either at either the safe level q_λ (where it will not be penalized) or at the higher experimental output q_A (at which it incurs the penalty with positive probability). The experimental output q_A is inefficiently low, reflecting the firm's preference to receive good news (that its choice was not unreasonable) rather than bad news. Moreover, the size of this 'bias' in the firm's output increases as the extent of uncertainty decreases. Hence, as the court learns, the firm's output will drift relatively farther from where the court would have it ideally produce. Finally, writing broad opinions alters the firm's incentives by creating a wedge between the probability that it ought to be penalized, and the probability that it actually will be. By changing these incentives, the court may entice the firm to choose more desirable outputs than could otherwise be sustained. However, the two types of opinions are not equally efficacious in directing firm behavior — the firm is more likely to respond to broad permissive opinions than to broad restrictive ones.

4 Court's Decision

This section focuses on decision-making by the court. I begin by specifying the court's preferences, and identifying the first-best policy. I then characterize the court's optimal (second-best) policy when its decision-making is constrained by the future choices of rational firms.

4.1 Court's Objective and the First-Best

The court is efficiency-minded and so seeks to maximize the discounted stream of social surpluses. The period t social surplus is $S_t = q_t + \frac{1}{2}q_t^2 - \theta q_t$, which is simply the sum of utilities across the firm and the victim.²³ If the true marginal harm were known, the efficient output level implied by the court's balancing test would be $q^{eff} = 1 - \theta$, and this output level implies a (maximized) social surplus of $\frac{1}{2}(1 - \theta)^2$. Then, for generic output q , the associated dead-weight loss is:

$$\frac{1}{2}(1 - \theta)^2 - \left(q - \frac{1}{2}q^2 - \theta q \right) = \frac{1}{2}(q - (1 - \theta))^2$$

Since θ is unknown, the court can at best target the expected dead-weight loss:

$$E \left[\frac{1}{2}(q - (1 - \theta))^2 \right] = \frac{1}{2}(q - [1 - E(\theta)])^2 + \frac{1}{2}Var(\theta)$$

The per-period expected dead-weight loss can thus be decomposed into two terms: a bias component and an uncertainty component. The bias component is related to the deviation of the chosen output from the *ex ante* efficient output. (Recall, this is $q_E = 1 - \frac{u+l}{2}$.) The uncertainty component captures the extent of uncertainty about the true marginal harm (as measured by the variance in beliefs about θ). The larger is this uncertainty, the greater the chance that the *ex ante* ideal output q_E itself deviates from the true efficient output q^{eff} . Hence, the court wishes to minimize both deviations of actual behavior from the *ex ante* efficient output, and uncertainty about the true socially optimal output. Within a given period, the uncertainty component is unaffected by the firm's choice, and so the court does best in the stage game by minimizing the bias component of the dead-weight loss.

²³This does not imply that the court equally weights the welfare of the agents. Since utilities are quasi-linear, the Pareto planner will seek to maximize the total unweighted surplus even if she assigns different values to the welfare of different agents. With this maximal surplus, the planner can then implement any desired distribution of utilities by choosing transfers as appropriate.

In the dynamic game, the firm's choice when it experiments has the additional effect of inducing learning, which in turn affects the size of the uncertainty component of dead-weight loss in future periods. (Of course, if the firm's output falls outside the ambiguous region, there is no learning, and so there is no change in the environment. There will simply be repetition of the stage game.) Suppose the firm chooses output $q \in (\lambda, \mu]$ in the ambiguous region. This output will be held unreasonable with probability $\frac{u-(1-q)}{u-l}$, and if so, the new beliefs will be $\theta \sim U[1-q, u]$. By contrast, this output level will not be found unreasonable with probability $\frac{(1-q)-l}{u-l}$, and if so, the new beliefs satisfy $\theta \sim U[l, 1-q]$. Then, the expected uncertainty component of the next-period dead-weight loss is:

$$\begin{aligned} E_t [Var_{t+1}[\theta] | q_t] &= \frac{u - (1 - q)}{u - l} \left[\frac{1}{12} (u - (1 - q))^2 \right] + \frac{(1 - q) - l}{u - l} \left[\frac{1}{12} ((1 - q) - l)^2 \right] \\ &= \frac{1}{4} \left(q - \left(1 - \frac{u + l}{2} \right) \right)^2 + \frac{1}{48} (u - l)^2 \end{aligned}$$

The expected uncertainty component of next-period deadweight loss has two important features that are easily verified. First, the arrival of information (i.e. learning) always reduces the uncertainty component of dead-weight loss. The intuition is obvious: learning truncates the support of the distribution of θ and this decreases the variance, our measure of the uncertainty cost. Second, the gains from learning are maximized by experimenting at the socially efficient output q_E . Since the dead-weight loss is convex in the extent of uncertainty $(u - l)$, an optimizing court will seek to smooth the extent of future uncertainty across both realizations of the case disposition. With a uniform prior, this is achieved by experimenting at the socially efficient output. It also ensures that, under efficient experimentation, the associated output should be equally likely to be found unreasonable or not. By contrast, experimenting anywhere else in the distribution causes the variance to be larger or smaller depending on the outcome of the case, and the situation with larger variance will be more probable.

Imagine a world in which the court could direct the firm to choose any output it desires,

whether this is incentive compatible for the firm or not. I refer to this as the first-best world, since in it, the court can implement its unconstrained ideal policy. Let θ denote the true marginal harm and assume $\theta \in [l_0, u_0]$.

Lemma 1. *In the first-best equilibrium, the court induces the ex ante socially efficient output $q_E(l_t, u_t) = 1 - \frac{l_t + u_t}{2}$ in every period t . Furthermore, in the long-run, the court perfectly learns the truth (i.e. $[l_t, u_t] \rightarrow [\theta, \theta]$ as $t \rightarrow \infty$) and output converges to the true socially efficient output $q^{eff} = 1 - \theta$.*

The first part of the Lemma shows that the court would ideally induce the firm to produce the *ex ante* socially efficient output in every period. The court's dual goals of minimizing period t bias and minimizing the uncertainty component of future dead-weight loss are complementary; both objectives are achieved by locating output at q_E . Absent feasibility constraints (such as the firm's incentive-compatibility constraint), this is the court's optimal policy. The second part of the Lemma shows that in the long-run, the court's beliefs converge to the truth, and the induced choice converges to the true socially efficient output. Since the court can always induce experimentation, the learning process continues indefinitely, each time reducing the extent of uncertainty by a factor of $\frac{3}{4}$. This ensures that uncertainty disappears in the long-run. Similar to part (2) of Proposition 2 in Baker and Mezzetti (2012), Lemma 1 demonstrates the efficiency of the common law. Absent significant constraints on the court's ability to direct firm behavior, the common law is able to implement the *ex ante* efficient outcome when there is uncertainty about the optimal policy, and to learn and implement the true optimal policy in the long-run.

4.2 Equilibrium in the Second-Best

The long-run efficiency result in Lemma 1 depended on the Court always being able to entice the firm to experiment and thereby learn. If the firm cannot be induced to experiment, the

the common law will not evolve any further. Additionally, short-run efficiency of the common law requires that the firm produce the *ex ante* socially efficient output in each period.

However, as was shown in Section 3, such choices are not consistent with the optimizing behavior of the firm. Corollary 1 showed that the court cannot always entice the firm to experiment, and Corollary 2 showed that even when sustainable, the experimental output might involve a significant bias (i.e. deviation from the *ex ante* socially efficient level). Over-deterrence is a consequence of legal ambiguity. Proposition 1 also showed that the firm could implement a much broader set of outputs by writing broad permissive opinions — however, in so doing, it would halt the learning process, since the firm would now choose the safe output level which is governed by precedent. Hence, to the extent that experimentation is feasible, the court faces a trade-off between writing narrow(er) opinions that induce experimentation and writing targeted broad opinions. In the former case, learning continues, enabling the court to reduce the uncertainty component of future dead-weight losses, at the cost of accepting a higher bias component of current dead-weight loss. In the latter case, the court can avoid the bias costs altogether, but makes no gains on the uncertainty front.

Suppose the court hears a case in which the firm's output was held reasonable, and so the court must update its permissive threshold λ . Let $V(\mu, l, u)$ be the expected discounted stream of dead-weight losses that follow, assuming the court chooses λ optimally, and taking as given the (next period) environment (μ, l, u) and the future path of play. Similarly, suppose the court hears a case in which the firm's choice was found to be unreasonable, and so the court must update its restrictive threshold. Analogously, define $W(\lambda, l, u)$ as the expected discounted stream of dead-weight losses that follow, assuming the court optimally updates its restrictive rule μ , and taking as given the environment (λ, l, u) and the future path of play. Let $\delta \in (0, 1)$ be the court's discount factor.

Proposition 2. *There exists a Markov Perfect Equilibrium characterized by a unique pair of value functions (V, W) satisfying:*

$$\begin{aligned}
\bullet V(\mu, l, u) &= \min_{\lambda \in [1-u, \mu]} \begin{cases} \frac{1}{1-\delta} s(q^*(\lambda, \mu, l, u), l, u) & q^* \leq \lambda \text{ or } q^* > \mu \\ s(q^*(\lambda, \mu, l, u), l, u) + \\ + \delta \left[\frac{u-(1-q^*)}{u-l} W(\lambda, 1-q^*, u) + \frac{(1-q^*)-l}{u-l} V(\mu, l, 1-q^*) \right] & q^* \in (\lambda, \mu] \end{cases} \\
\bullet W(\lambda, l, u) &= \min_{\mu \in [\lambda, 1-l]} \begin{cases} \frac{1}{1-\delta} s(q^*(\lambda, \mu, l, u), l, u) & q^* \leq \lambda \text{ or } q^* > \mu \\ s(q^*(\lambda, \mu, l, u), l, u) + \\ + \delta \left[\frac{u-(1-q^*)}{u-l} W(\lambda, 1-q^*, u) + \frac{(1-q^*)-l}{u-l} V(\mu, l, 1-q^*) \right] & q^* \in (\lambda, \mu] \end{cases}
\end{aligned}$$

where $s(q, l, u) = \frac{1}{24}(u-l)^2 + \frac{1}{2}\left(q - \left(1 - \frac{u+l}{2}\right)\right)^2$ is the per-period dead-weight loss, and $q^*(\lambda, \mu, l, u)$ is the firm's optimal choice, as defined in Proposition 1.

Proposition 2 verifies that a Markov-Perfect equilibrium exists, and is characterized by a unique pair of value functions, one for each type of case disposition. The value functions must satisfy the Bellman Equations, which require that, after any history, the court write the opinion that minimizes its expected discounted stream of dead-weight losses, assuming optimal decision-making in all future periods. Importantly, the court anticipates the firm's future response to its decisions.

I now turn attention to the properties of the court's optimal policy.

Lemma 2. *In any Markovian Equilibrium, following a history in which the court is required to update its restrictive threshold μ , the court will optimally choose any μ^* satisfying $\mu^* \in \left[\frac{3}{4}(1-l), (1-l)\right]$, where $l = 1 - q$ is the updated lower bound on beliefs. In particular, it is always optimal to write a narrow opinion $\mu^* = 1 - l$.*

Lemma 2 echos the previously discussed intuition that the restrictive opinion is generally ineffective at incentivizing the firm. Given that the firm's behavior is characterized by underprovision (over deterrence), a broad restrictive opinion that increases the likelihood of being penalized will never be able to induce the firm to *increase* its output. This follows from

parts 1 and 2 of Proposition 1 which show that if the restrictive opinion affects the firm's choice at all, it will be to entice a choice below the experimental output.²⁴ Since this would be uniformly worse for the court, both in increasing policy bias and making learning less efficient, there is never a strict incentive for the court to write a broad restrictive opinion.

Nevertheless, Lemma 2 shows that there are a multiplicity of restrictive opinions that are (payoff equivalent) consistent with equilibrium. These broad opinions are equilibrium consistent only insofar as there is no subsequent path of play in which such an opinion adversely affects social welfare. Thus, a broad restrictive opinion μ is permissible, if after every future history, it neither precludes future experimentation nor prevents an optimally broad permissive opinion from being written. Intuitively, the set of 'insignificant' broad opinions, are those that are 'minimally broad', and hence unlikely to inefficiently constrain future courts.

Next, I consider optimal behavior by the court following a case in which it must revise its permissive threshold λ .

Lemma 3. *The optimal permissive rule is always broad ($\lambda^* > 1 - u$). Furthermore, there exists some threshold $\phi(l, u)$ with $\phi < \frac{1}{2}$ such that:*

- $\lambda^* < \bar{\lambda}(l, u)$ if $q_A \in (1 - u, \mu]$ and $b(l, u) < \phi(l, u)$, and
- $\lambda^* = \min \left\{ 1 - \frac{u+l}{2}, \mu \right\}$ otherwise.

The Lemma states that the court will either write a slightly broad permissive opinion that is consistent with the firm continuing to experiment, or write a broad opinion targeted at the socially efficient output q_E . (Recall, by Proposition 1, the firm will experiment provided that $\lambda < \bar{\lambda}(l, u)$.) In the former case, the court continues to learn in subsequent period(s), whilst in the latter case, the breadth of the opinion precludes future learning, and the law

²⁴The assumption that $u_0 < \hat{u}_0$ ensures that the perverse case encapsulated in part 3 of Proposition 1, in which the legal rule essentially collapses to a strict-liability rule, does not arise.

becomes settled. That the court should target the socially efficient output (or the nearest feasible output if μ is prohibitively broad) in the latter case is intuitive — absent a learning motive, the best the court can do is to minimize bias in each period.

Two features of Lemma 3 are worth highlighting. First, when $b \in (\phi, \frac{1}{2})$, the court may choose to stop learning even though further experimentation is possible. The reason is that learning is costly; its gains come at the cost of inducing inefficient choices. As the bias increases, the gains from learning decrease and the inefficiency costs increase. Hence, if the bias is large enough, the court may do better to forgo learning and simply write a broad permissive opinion targeting the *ex ante* socially efficient output, thereby avoiding the bias cost altogether. Hence, not only have I shown (in Corollary 1) that the firm's incentives cause it to eventually stop experimenting, but additionally, the court's incentives cause it to *induce* the firm to stop experimenting even sooner. The value of the new information that arrives along the equilibrium path gradually decreases. As this happens, the court turns to broad opinions to incentivize the firm to make production choices closer to the socially efficient level.

Second, even when the court induces experimentation, it nevertheless writes broad permissive opinions — albeit opinions that are not so broad as to preclude learning. In such cases, since the firm experiments, the location of the permissive opinion does not affect welfare in the following period (or in any subsequent period in which the firm experiments). However, the location of the permissive opinion may still matter if, in the continuation play, the court does not have the opportunity to revise λ before experimentation stops. (Such a history will arise, whenever all subsequent outputs chosen before experimentation ceases, are found to be unreasonable.) Since, once experimentation stops, the firm will simply choose q_λ in all future periods, the court would prefer that λ were closer to the socially efficient output q_E rather than not. Hence, the court has an incentive to *preemptively* write slightly broad permissive opinions, just in case the opportunity to revise this opinion does not arise soon

enough into the future. Preemption is a form of hedging by a court that cannot perfectly control the flow of cases it hears, and the opportunities to thereby refine legal rules.

The analogous result in Baker and Mezzetti (2012) (where courts face explicit adjudication costs and an exogenous stream of cases) is worth highlighting. In their paper, the purpose of broad opinions is to mitigate costs; expanding the set of cases to be disposed summarily reduces the likelihood of needing to adjudicate a controversial case. The court does not incur adjudication costs in this paper. Instead, the court faces the implicit bias costs which arise because it is subject to the firm's incentive-compatibility constraint and cannot induce first-best experimentation. Baker and Mezzetti show that both permissive and restrictive opinions should be optimally broad and that optimal breadth is increasing in adjudication costs. If costs are sufficiently large, opinions should be maximally broad such that no ambiguity remains. Similarly, this paper shows that, if bias costs are large enough, the court should write a 'maximally broad' permissive opinion that precludes further experimentation. However, this paper also shows that restrictive opinions can be optimally narrow, and should never be too broad. Furthermore, the optimal breadth of permissive opinions is non-monotone in bias costs.

Taken together, the above lemmata enable the characterization of the long-run behavior of the common law. Lemma 2 shows that the court will induce experimentation whenever feasible, following a 'bad-news' case (i.e. when it can update μ). Lemma 3 shows similarly that, following 'good news', the court will induce experimentation provided it is feasible and the bias is not too large. Corollary 1 shows that as the extent of uncertainty ($u-l$) decreases, so do opportunities for experimentation. Moreover, Corollary 2 shows that the relative bias in the experimental output increases as the extent of uncertainty decreases. All told, as learning occurs and uncertainty decreases, so does the court's ability and desire to learn further. Once the court has learned sufficiently much, learning will stop (either voluntarily, or because it is no longer feasible). This implies the following result:

Proposition 3. *In the second-best equilibrium, the Court induces a finite period of experimentation (at socially inefficient outputs) before the law becomes settled. Settled law is characterized by residual uncertainty, and the long-run outcome is ex ante socially inefficient with positive probability. (Formally, there exists $T < \infty$ such that for all $t > T$, beliefs are parameterized by (l_T, u_T) with $l_T < u_T$, and the legal rule satisfies $\lambda_T \leq 1 - \frac{l_T + u_T}{2}$ and $\mu_T \geq \frac{3}{4}(1 - l_T)$, with probability 1. The law will be ex ante inefficient, in the long-run, if $\lambda_T < 1 - \frac{l_T + u_T}{2}$.)*

Proposition 3 describes the long-run dynamics of the common law. Several features are worth noting. First, in contrast to other models where the court continues to refine the law *ad infinitum*, legal evolution in this model proceeds over a finite number of cases, but eventually stops. The reason is that, in low uncertainty environments, the benefits to either the firm or court (or both) of continued experimentation are outweighed by the costs (of being penalized, in the case of firms, and of tolerating inefficient outputs, in the case of the court). After this point, the firm never makes choices that warrant review by the court — the law becomes settled.

Second, this settled law is characterized by residual uncertainty; the court does not perfectly learn the true social cost θ . The behavior induced by the settled law will deviate from the true socially efficient choice, almost surely. Hence, whilst the common law provides a mechanism for institutional learning, the mechanism is imperfect. Even in the long-run, the common law cannot guarantee convergence to optimal outcomes.

Third, given this residual uncertainty, the best long-run rule that the (informationally constrained) court could potentially implement is the one that implements the *ex ante* efficient output. However, this is not guaranteed. Moreover the efficiency properties of the common law are path dependent. It matters whether the firm’s most recent experimental choice was held to be acceptable or not. In the former case, the court may update its permissive threshold and write a broad opinion targeting the *ex ante* efficient output. By contrast, in the

latter case, the court may only update its restrictive threshold, and as I have shown, this tool is inefficacious at incentivizing the firm. The firm will simply choose the safe output implied by the previously chosen permissive opinion, which must be below the *ex ante* efficient level. Hence, if the law settles following a ‘bad news’ case — a positive probability event — the firm’s steady-state output will be *ex ante* inefficient; there will be over-deterrence even in the long-run.

In this latter case, I say that the law settles ‘abruptly’; the court would ideally wish to resolve another case, in order to amend its permissive threshold. Indeed, it is the inefficiency stemming from this constraint on the court’s decision-making that generates the motive for preemption. Anticipating that it may not have the opportunity to amend its permissive threshold in the future, the court will write somewhat broad permissive opinions whenever it has the opportunity, even if it ultimately seeks to induce experimentation. The opinion is the broadest one that does not preclude experimentation in either the current period, or any future period when experimentation would be optimal. Preemption ensures that, if the law settles abruptly, the induced steady-state output is less inefficient, in the sense of being closer to the *ex ante* efficient output than would be the case if opinions were perfectly narrow.

The analogous result in Baker and Mezzetti (2012) is again worth highlighting. In that paper, adjudication costs are fixed and exogenous. If costs are large, then opinions will be maximally broad and the law will settle immediately. Alternatively, if costs are small, opinions will be narrower, and an ambiguous region of law will always exist. Stochastic arrival of cases guarantees that opportunities for learning persist, and so the law continues to evolve indefinitely. By contrast, in this paper, implicit bias costs are endogenous and become larger over time as the extent of uncertainty decreases. Hence, in this paper, the law evolves for some finite period of time, and then settles, when implicit costs become sufficiently large. Descriptively, this makes for a more compelling account of the evolutionary dynamics

of legal doctrine.

The notion of settled law also differs subtly between the two models. Functionally, in both models, the law settles when it ceases to change any further. In Baker and Mezzetti (2012), this requires that the ambiguous region be empty, and so in their model, settled law is unambiguous. By contrast, in this paper, settled law may still remain ambiguous in part, provided that the firm would never make choices in the ambiguous region. Hence, even as the common law settles, the legal status of certain behaviors may never be resolved.

Note, importantly, that the mechanism that induces short-run inefficiency is conceptually distinct from the one that potentially causes the long-run inefficiency. The common law induces inefficient choices in the short-run because legal ambiguity causes over-deterrence. By contrast, the potential long-run inefficiency arises both because experimentation does not continue long enough for the court to learn the truth, and because constraints on rule making prevent the court from amending the law *ex post* to correct these defects. As the example in the following footnote makes clear, if experimentation were always guaranteed, then the common law could converge to the efficient rule, even if it always produced biased outcomes along the equilibrium path.²⁵

I conclude this section by briefly demonstrating the robustness of the model to variant specifications in two contexts. First, the results in this paper do not hinge crucially on the assumption that legal rules are irreversible. The inefficiency does not arise because the court had previously written a broad opinion, that proved to be inefficient *ex post*, after the arrival of new information. Instead, the inefficiencies stem from the endogeneity of the case-generating process, and the firm's reluctance to experiment in low uncertainty environments. Indeed, as Corollary 1 demonstrates, experimentation would eventually cease,

²⁵Suppose the court always writes narrow rules and consider a firm whose output satisfies: $\hat{q}(l, u) = 1 - \frac{3u+l}{4} < q_E$. Such a firm always experiments, and this experimental output is always inefficiently low. Nevertheless, the court always learns, and in the long-run beliefs converge to the truth (albeit at a slower rate than would be the case if experimentation were efficient). Hence, the firm's output converges to the efficient level in the long-run, even though every short-run choice is *ex ante* inefficient.

even if the court were constrained to only write narrow opinions (which it would never have an incentive to undo). The mechanism that causes learning to stop operates at a deep level that is independent of the court's strategy.

Second, the assumption of a representative firm is not crucial, and the model can readily accommodate heterogeneous firms. To see this, suppose there are n firms which differ in their production technology, and that firm i 's profit function is $\alpha_i q - \frac{1}{2}q^2$. Suppose $\frac{1}{n} \sum_i \alpha_i = 1$, so that the representative firm has the 'average' profit function. Firm i 's optimal choice is qualitatively unchanged, and is given by $\hat{q}(\lambda, \mu, l, u; \alpha_i) = \alpha q^*(\frac{\lambda}{\alpha}, \frac{\mu}{\alpha}, \frac{l}{\alpha}, \frac{u}{\alpha})$, where q^* is optimal choice defined in Proposition 1.²⁶ Since different firms will make different choices under the same rule, the case-generating process will now be stochastic. Moreover, firms with larger α will be more likely to experiment, and so the court will see a biased stream of cases, as in Hadfield (1991). The court's per-period deadweight loss function becomes²⁷:

$$\sigma = \frac{1}{2} (\bar{q} - (1 - E[\theta]))^2 + \frac{1}{2} Var(\theta) + Var(\alpha_i - q_i)$$

which is analogous to the deadweight loss in the representative-firm case, except for the introduction of the third component. With heterogeneous production technologies, the efficient output varies across firms. But this ideal dispersion of outputs may not be consistent with the legal rule, which might, for example, induce all firms to make the same safe choice. The new component of deadweight loss captures the inefficiency that arises when the actual dispersion of firm outputs disagrees with the ideal dispersion.

²⁶This follows by showing that firms' expected net-profit is homogeneous of degree 2.

²⁷ $\bar{q} = \frac{1}{n} \sum_i q_i$ is the average output, and $Var(\alpha_i - q_i) = \frac{1}{n} \sum_i [(\alpha_i - q_i) - (1 - \bar{q})]^2$.

5 Extensions

5.1 Optimal Penalty

In the baseline model, compensatory damages were the assumed penalty for unreasonable interference. This is the standard remedy awarded by common law courts, and reflects the principle that the purpose of damages is to make the victim whole (see footnote 5). The over-deterrence induced by this penalty was crucial to generating both the short- and long-run inefficiency results. Whilst the ubiquity of its use recommends the modeling choice, it is instructive to explore the implications of the penalty function being otherwise — if only to understand the implications of this institutional norm.

Consider the following variant model, which is identical to the baseline model except in the structure of the penalty function, in the ambiguous region. The per-unit penalty is unchanged in the regions of *per se immunity* and *strict liability*. In the ambiguous region, the court may now freely set the penalty. Let the per-unit penalty be $\varphi(u, l) \cdot \frac{u+l}{2}$, where the scale factor φ describes the fraction of the expected harm that are paid in damages.²⁸ In the variant model, the law is not silent in the ambiguous region; although the liability status of different choices remains unspecified, the court now announces what the penalty will be for choice that are found unreasonable, *ex post*. By contrast, in the baseline model, the law was entirely silent in the ambiguous region, about both liability and penalty. Instead, the firm formed rational beliefs about the probability of being held liable and the expected penalty, assuming that compensatory damages would be applied in the event that it is held liable.

Lemma 4. *With a flexible penalty, the court can implement the first-best outcome by announcing the per-unit penalty implied by the scale factor $\varphi(u, l) = \frac{u-l}{1-l} < 1$. Doing so induces*

²⁸In principle, we can have $\varphi \leq 1$. If $\varphi > 1$, the firm over-compensates the victim, which is consistent with the payment of punitive damages. If $\varphi < 1$, we can think of the court applying a different standard to firms that knowingly overproduce (by choosing $q > \mu$), and those whose conduct was legally unclear, and only found to be unreasonable, *ex post*.

the firm to choose the ex ante socially efficient output in every period , and enables the court to discover the true socially efficient rule in the long-run.

This result is easily understood by recalling that compensatory damages caused over-deterrence, whilst zero damages would not deter the firm at all. Intuitively, there must be some intermediate level of damages – less than full damages – that exactly incentivizes the firm to make the socially efficient choice. Moreover, Lemma 1 then entails that, by implementing this penalty, the court will learn efficiently and will discover the true efficient policy q^{eff} in the long-run.

Lemma 4 shows that the court can achieve the first-best outcome by imposing ‘discounted damages’. Corollary 3 shows that the size of the discount must increase as the law evolves.

Corollary 3. *The sequence of scale factors $\{\varphi_t\}$ that are applied along the equilibrium path is monotonically decreasing and converges to 0.*

The intuition is a familiar one. As the law evolves and the extent of uncertainty decreases, the marginal probability of being penalized increases (for a given increment in output). At the margin, this causes the firm to decrease output below the socially efficient level. The court must then reduce the penalty to entice the firm back to the socially efficient choice. In the limit, the optimal penalty must fall to zero.

Taken together, these results demonstrate that over-deterrence is not a feature of compensatory damages, *per se*. A similar dynamic as in the baseline analysis would arise whenever the court committed to a rule that at least partially compensated the victim. Whatever the level of compensation, it will eventually be too high relative to the efficient penalty, and over-deterrence will result.

The results also point to the role of broad opinions in a second-best world²⁹, where the court’s institutional commitment to compensatory damages prevents first-best outcomes.

²⁹The notion of second-best in this section differs from that in Section 4. In that section, second-best

Broad opinions become increasingly valuable as the disparity between compensatory and first-best damages increases. This accords with the result in Lemma 3, which states that the court will write broad opinions when the relative bias in the firm's output becomes sufficiently large. As Lemma 4 demonstrated, that bias is a consequence of the penalty being too large.

5.2 *Dicta* is Persuasive

In the baseline model, when deciding cases, I assumed that the court could only update the threshold relevant to explaining the case disposition. I motivated this modeling choice by noting that any attempt to update the other threshold would be plainly understood to be *obiter dicta*, and thus lacking precedential value. Additionally, the court does not have any private information that it could credibly communicate to the litigants. The baseline model, in effect, assumes that courts will never issue *dicta* because they assume it would be rationally ignored.

However, it may be argued that *dicta* can be persuasive, foreshadowing to the community the future decisions that the court intends to enshrine. In this sub-section, I consider the alternative formulation, in which (as is the case in Baker and Mezzetti (2012)), the court may update both thresholds when deciding a case in the ambiguous region. Doing so leaves the firm's decision unaffected, since the firm makes its choice taking the current legal rule as given and without reference to the future path of the law. Accordingly, there is no strict incentive for the court to write broad restrictive opinions (although it remains equilibrium consistent for the court to write 'minimally broad' opinions).

outcomes arose because the court was subject to the firm's incentive-compatibility constraint. The relevant constraint, here, is the court's ability to choose the level of damages; in the first-best world it can do so freely, whilst in the second-best world it must impose compensatory damages. (As the proof of Lemma 4 shows, optimal penalty choice by the court in the first-best world causes the firm's incentive-compatibility constraint to be effectively non-binding.)

The change does implicate the court’s optimal policy when writing permissive opinions. In the baseline model, the court would preemptively write somewhat broad opinions when it wanted to induce experimentation, as a hedge against being later unable to revise its permissive threshold before the law settles. Now that the court can always adapt either threshold, this precautionary motive no longer exists. The court *could* write purely narrow permissive opinions whenever experimentation remained optimal, and only write broad opinions when the law finally settled. This motivates the following adaptation of Lemma 3 to the alternative framework:

Lemma 5. *Suppose the court is able to change both thresholds. There exists $\hat{\phi}(l, u) > \phi(l, u)$ such that the following policy is equilibrium consistent:*

- $\lambda^* = 1 - u$ if $q_I \in (\lambda, \mu]$ and $b(l, u) < \hat{\phi}(l, u)$, and
- $\lambda^* = \min \left\{ 1 - \frac{u+l}{2}, \mu \right\}$ otherwise.

The court will write narrow permissive opinions whenever experimentation is feasible and not too biased, and will write a broad permissive opinion targeting the *ex ante* efficient output if experimentation is either no longer feasible or the bias is sufficiently large. Importantly, the bias that the court is willing to tolerate is larger than was the case in the baseline model. The court is now less constrained in its ability to use its policy tools. This makes experimentation less costly in the event of ‘bad news’ that causes the common law to settle at an inefficient output. Given the higher option value of experimentation, the court will be willing to tolerate a larger short-run bias in order to facilitate learning. Along the equilibrium path, we should expect the law to evolve for a longer period of time before settling. Furthermore, although the settled law will continue to exhibit residual uncertainty, it should never settle at an *ex ante* inefficient output.

These results suggest that the institutional norm distinguishing precedent from *dicta* negatively impacts the efficiency of common law. Understand the value of this distinction (if it

exists at all) may require richer models of judicial decision-making in which, for example, the court has private information, or judges have heterogeneous preferences. This suggests fruitful directions for future research.

5.3 Non-Efficiency-Minded Courts

Judges in the baseline model were assumed to be efficiency-minded. This assumption was useful in highlighting the paper's key insights, that structural factors, such as endogenous case flow, can limit the efficacy of the common law, even when judges are efficiency-minded. However, the assumption itself was not crucial to the model's insights, and can be relaxed.

For example, suppose the judges instead sought to maximize $q - \frac{1}{2}q^2 - (1 + \alpha)\theta q$, where $\alpha \neq 0$. If $\alpha > 0$, the court over weights the harm to the victim when applying the balancing-test, and vice versa. The court's ideal policy is now $\tilde{q} = 1 - (1 + \alpha)\theta$, but this policy is no longer Pareto optimal (see footnote 23). As before, when hearing a case, the court observes whether the chosen output was above or below its ideal, \tilde{q} , and updates its beliefs accordingly.

Consider the first-best world in which the court is unconstrained by incentive-compatibility, and can implement its *ex ante* ideal policy $\tilde{q}_E = 1 - (1 + \alpha)\frac{u+l}{2}$. Although the choice is biased (relative to the socially efficient output $q_E = 1 - \frac{u+l}{2}$), it is consistent with efficient learning.³⁰ In the long-run, the court will discover the true marginal harm θ , and implement its ideal policy \tilde{q} . The first-best results are largely unchanged.

The incentives for the firm, however, do change, since changing the standard for unreasonable interference also changes the likelihood of being penalized. Let \tilde{q}_A denote the new experimental output. An increase in α causes both \tilde{q}_A and \tilde{q}_E to decrease, but \tilde{q}_A decreases

³⁰The proof is directly analogous to the proof of Lemma 1.

by less.³¹ There are now two distinct notions of bias. The ‘court bias’ $\frac{\tilde{q}_E - \tilde{q}_A}{(1+\alpha)(u-l)}$ measures the deviation between the firm’s choice and the court’s ideal, whereas the ‘social bias’ measures the deviation from the socially efficient output q_E . The comparative statics on \tilde{q}_A and \tilde{q}_E imply that the court bias is decreasing in α , whilst the social bias is increasing. The latter result follows intuitively by noting that as α increases, the court applies a more stringent standard for reasonable conduct on the firm, which magnifies the over-deterrence.

Since the opportunity cost of experimentation is increasing in the court bias, the court should be more willing to experiment, and should learn more efficiently, as α increases. There is a tension between efficient learning and efficient stage game outcomes. A court that over-weights the harm to the victim ($\alpha > 0$) will experiment for longer and learn more efficiently, but along the way, will introduce larger true bias costs. By contrast, a court that under-weights the harm will over-deter the firm by less, but will also experiment less, and accept larger uncertainty costs in the long-run.

6 Conclusion

This paper presented a dynamic model of adjudication to study the evolution of common law. The model provided foundations to explain the incremental nature of the evolution of legal rules, and to explain when and why courts might write narrower or broader rules. An important feature of the model, that distinguishes it from existing papers, is that the court’s future docket (i.e. the case-generating process) is endogenous to the current legal rule adopted by the court. Hence, there is a feedback between the court’s current decisions,

³¹To see this, note that the firm’s profit in the ambiguous region is given by $\tilde{\pi} = q - \frac{1}{2}q^2 - \frac{q}{2(u-l)} \left[u^2 - \left(\frac{1-q}{1+\alpha} \right)^2 \right]$. It follows that \tilde{q}_A is the solution to $\left[\frac{3}{2} + \alpha \right] \chi^2 - [1 - (1 + \alpha)(u - l)] \chi - \frac{1}{2}u^2 = 0$, where $\chi = \frac{1-q}{1+\alpha}$. (It is easily verified that taking $\alpha = 0$ gives the first order condition defining q_A in the baseline analysis.) Then, by the implicit function theorem, $\frac{\partial \chi}{\partial \alpha} = \frac{\chi^2 + (u-l)\chi}{[1 - (1+\alpha)(u-l)] - (3+2\alpha)\chi} < 0$. By contrast $\frac{\partial}{\partial \alpha} \left(\frac{1-\tilde{q}_E}{1+\alpha} \right) = \frac{\partial}{\partial \alpha} \left(\frac{u+l}{2} \right) = 0$.

and its ability to learn and affect the law through the future controversies that may arise. To operationalize this feedback, I explicitly modeled the choices of the agent whose behavior the court sought to regulate.

The paper showed that, as the legal rule evolves, agents will be less likely to make controversial choices that require further adjudication, such that the flow of cases that the court hears will eventually stop. Thus, in contrast to other models of the common law, where the law typically evolves *ad infinitum*, this paper shows that legal evolution is characterized by an endogenous stopping point. Moreover, the fact that the law settles in finite time has several implications for the efficiency of common law. First, settled law will exhibit residual uncertainty — the common law is an imperfect institutional mechanism to acquire information. Second, the settled law will almost certainly not implement the true socially optimal outcome, and with positive probability will even implement *ex ante* inefficient outcomes. As the paper documents, these failures of the common law arise precisely because the case-generating process is endogenous. This demonstrates the importance of taking seriously the feedback between rule-making and future agent behavior.

The model shares several commonalities with Baker and Mezzetti (2012), and some of the results are analogous (although others are importantly different). However, the mechanisms underlying these papers are quite different. In that paper, the court trades off the benefits of learning against the explicit costs of adjudicating controversial cases; their paper is essentially a model of work-load management. By contrast, the pertinent costs in this paper are the implicit costs of being subject to the firm’s rational decision-making process. Ambiguities in the law cause the firm to make socially inefficient choices, and so the court must trade-off the cost of short-run inefficiencies against the benefits from learning. The model thus captures in a strong sense the tensions that courts face in determining how broadly to construe their decisions. Whilst narrow opinions maximize the scope for future learning and retains for the court the option for more informed future decision-making, they also leave significant

ambiguities about the law that can deter agents from making desirable choices. By reducing ambiguity, broad opinions can induce agents to make choices that are more socially efficient.

7 Appendix

Proof of Proposition 1 and Corollary 1. Let $\Pi(q)$ denote the expected net-profit function. $\Pi(q)$ is continuous and differentiable everywhere except (possibly) at $q = \lambda$ and $q = \mu$. Moreover, it is piece-wise strictly concave in the relevant regions. Hence, to find the optimal policy, it suffices to find the optimal policy in each region, and then choose the maximizer amongst the these. Recall $0 \leq l \leq u \leq 1$, and $0 < 1 - u \leq \lambda \leq \mu \leq 1 - l < 1$.

Consider the region of *per-se immunity*: $q \leq \lambda$. The net profit is: $\Pi(q) = q - \frac{1}{2}q^2$ and so $\Pi'(q) = 1 - q > 0$. The optimizing output in this region is therefore at the endpoint: $q = \lambda$, and the maximal profit is $\Pi^\lambda(\lambda) = \lambda - \frac{1}{2}\lambda^2$. Next, consider the region of *strict liability*: $q > \mu$. The net profit is : $\Pi(q) = q - \frac{1}{2}q^2 - \left(\frac{u+l}{2}\right)q$, and so $\Pi'(q) = \left[1 - \frac{u+l}{2}\right] - q$. If $\mu > 1 - \frac{u+l}{2}$, then $\Pi'(q) < 0$ for all $q > \mu$, and so the optimizing output is again at the endpoint $q = \mu$. If, instead, $\mu < 1 - \frac{u+l}{2}$, then $q^* = q_E = 1 - \frac{u+l}{2}$. Denote $\Pi^E = q_E - \frac{1}{2}q_E^2 - \left(\frac{u+l}{2}\right)q_E = \frac{1}{2}\left(1 - \frac{u+l}{2}\right)^2$.

Finally, suppose $q \in (\lambda, \mu]$. Then $\Pi(q) = q - \frac{1}{2}q^2 - \frac{1}{2}\frac{u^2 - (1-q)^2}{u-l}q$, and $\Pi'(q) = 1 - q - \frac{u^2 - (1-q)^2}{2(u-l)} - \frac{1-q}{u-l}q$. (For notational simplicity, denote these by $\tilde{\Pi}(q)$ and $\tilde{\Pi}'(q)$.) Since $\tilde{\Pi}'(q)$ is quadratic in q , it admits two roots. Furthermore, since $\tilde{\Pi}''(q) = \frac{3q-2}{u-l} - 1$ is increasing in q , the smaller of these roots is a maximizer, whilst the larger is a minimizer. Let q_A be the smaller solution to $\tilde{\Pi}(q) = 0$. By the quadratic formula, $q_A = 1 - \frac{(1-(u-l)) + \sqrt{(1-(u-l))^2 + 3u^2}}{3}$. We can show that $q_A < 1 - \frac{1+l}{2} < 1 - \frac{u+l}{2}$. (To see this, show that in the ambiguous region: $\tilde{\Pi}'\left(1 - \frac{1+l}{2}\right) = -\frac{\left(u - \frac{1}{2}(1+l)\right)^2}{2(u-l)} < 0$. Then since $\tilde{\Pi}'(q)$ is decreasing at q_A , it must be that $q_A < 1 - \frac{1+l}{2}$.) Finally, note that if $q_A \geq 1 - u$, it must be that $\tilde{\Pi}'(1 - u) = u - \frac{u(1-u)}{u-l} \geq 0$, and so $1 - u \leq u - l$. Hence, if $1 - u > u - l$, then $q_A < 1 - u$, which obviously falls outside the ambiguous region. There are three possibilities: (i) $q_A < \lambda$, in which case $\tilde{\Pi}(q) < 0$ for all

$q \in (\lambda, \mu]$, and so the optimum is at the end point $q = \lambda$; (ii) $q_A > \mu$, in which case $\tilde{\Pi}(q) > 0$ for all $q \in (\lambda, \mu]$, and so the optimum is at the end point $q = \mu$, and denote $\Pi_\mu(\mu) = \tilde{\Pi}(\mu)$; and (iii) $q_A \in (\lambda, \mu]$, in which case q_A is optimal. Denote $\Pi^A = \tilde{\Pi}(q_A)$.

Note — the above expressions: $\Pi^\lambda, \Pi^\mu, \Pi^A$ and Π^E are technically defined for all (u, l, λ, μ) , even if these are not always feasible/equilibrium-consistent (although the conditions under which feasibility failed were discussed above). Hence, to find the optimal output choice in any given environment, one need only compare the expected net profits, from amongst the feasible choices. Of course, that $q = \lambda$ is always feasible. First, I show that $\Pi^\mu < \Pi^E < \Pi^A$ (whenever these are feasible). To see this, note that:

$$\Pi^E = \Pi^A + \int_{q_A}^{q_E} \tilde{\Pi}'(q) dq - \left[\frac{u+l}{2} - \tilde{f}(q_E) \right] q_E$$

where the second term is the additional profit gained by increasing output from q_A to q_E assuming the fine structure in the ambiguous environment, and the third term is an adjustment for the fact that the average expected fine at q_E is assessed under the strict liability region. Noting that $\tilde{\Pi}(q) < 0$ whenever $q > q_A$, and $\tilde{f}(q_E) = \frac{u^2 - (\frac{u+l}{2})^2}{2(u-l)} = \frac{3u+l}{8} < \frac{u+l}{2}$, gives $\Pi_E < \Pi_A$. Furthermore: $\Pi^\mu = \max_{q \leq \mu < q_A} \tilde{\Pi}(q) < \max_q \tilde{\Pi}(q) = \Pi^A$, since the constraint $\mu < q_A$ binds strictly. Hence, the firm will never choose q_E (or q_μ) when q_A is feasible.

Next, analyze the choice between q_λ and q_A , whenever q_A is feasible. Note that Π^A is a constant function of λ , whilst $\Pi^\lambda(\lambda)$ is continuous and strictly increasing in λ for $\lambda < 1$. Moreover, $\Pi^\lambda > \Pi^A$ when $\lambda = q_A$, since the firm produces the same output, but is not held liable. There are two cases to consider: First, from above, we know that $q_A \leq 1 - u$ when $1 - u > u - l$. If so, then $\Pi^\lambda(1 - u) \geq \Pi^A$, and then since Π^λ is increasing in λ , $\Pi^\lambda(\lambda) \geq \Pi^A$ for all $\lambda \in [1 - u, 1]$ (with equality only potentially at the left-hand end-point). Second, suppose $q_A > 1 - u$. Then $\Pi^\lambda(1 - u) < \Pi^A$. (This will be shown, below.) Then, by the intermediate value theorem, there exists $\bar{\lambda} \in (1 - u, q_I)$ s.t. $\Pi^\lambda(\lambda) \geq \Pi^A$ whenever $\lambda \geq \bar{\lambda}$.

We can see this clearly by noting that for $\lambda < q_A$:

$$\Pi^A = \Pi^\lambda(\lambda) + \int_\lambda^{q_A} \tilde{\Pi}'(q) dq - \frac{u - (1 - \lambda)}{u - l} \lambda$$

where the third term is an adjustment for the discontinuous jump in the penalty when moving from the per-se immunity to the ambiguous region. By construction, $\tilde{\Pi}'(q) > 0$ for $q < q_I$, and so the second term is positive. The third term is generically positive (although subtracted), but is zero when $\lambda = 1 - u$. In this case, the second term dominates the third term, and so $\Pi^A > \Pi^\lambda$. As λ increases, the second term decreases, and the third term increases, such that at some threshold, $\Pi^\lambda \geq \Pi^A$.

It remains to examine the firm's choice when q_I is not feasible (i.e. when $\mu < q_I$, so that q_I falls into the strict liability region). In lieu of choosing q_I , the firm has two choices — either to choose the largest output in the ambiguous region (i.e. $q = \mu$) or the best output in the strict-liability region ($q = q_E$). Now, from above, $\Pi^\mu(\mu)$ is continuous and strictly decreasing in μ , and $\Pi^\mu \leq \Pi^A$, with strict inequality whenever $\mu < q_I$, and $\Pi^\mu(0) = 0 < \Pi^E$. Moreover, Π^E is constant in μ and $\Pi^E < \Pi^A$. Hence, by the intermediate value theorem, there exists some $\bar{\mu} \in (0, q_I)$ s.t. $\Pi^\mu \geq \Pi^E$ whenever $\mu \geq \bar{\mu}$ (noting that possibly $\bar{\mu} < \lambda$). Now, if $\mu \geq \bar{\mu}$, then the firm's ultimate choice is between q_λ and q_μ . Using the logic in the previous paragraph, there exists some threshold $\tilde{\lambda}$ s.t. $\Pi^\lambda \geq \Pi^\mu > \Pi^E$ whenever $\lambda \geq \tilde{\lambda}$. Moreover, since $\Pi^A > \Pi^\mu$, it must be that $\tilde{\lambda} < \bar{\lambda}$. Similarly, if $\mu < \bar{\mu}$, then there exists some threshold $\tilde{\tilde{\lambda}}$ s.t. $\Pi^\lambda \geq \Pi^E$ whenever $\lambda \geq \tilde{\tilde{\lambda}}$. Moreover, it must be that $\tilde{\tilde{\lambda}} \leq \tilde{\lambda}$. Finally, we can find a closed form expression for $\tilde{\tilde{\lambda}}$, by solving:

$$\begin{aligned} \lambda - \frac{1}{2}\lambda^2 &= \frac{1}{2} \left(1 - \frac{u+l}{2}\right)^2 \\ \lambda &= 1 - \sqrt{1 - \left(1 - \frac{u+l}{2}\right)^2} = \tilde{\tilde{\lambda}}(l, u) \end{aligned}$$

where we naturally take the negative root (or the bound would be larger than 1). It is

easily verified that $\tilde{\lambda} < 1 - \frac{u+l}{2}$. (To see this, denote $\xi(\lambda) = \frac{1}{2}\lambda^2 - \lambda + \frac{1}{2}\left(1 - \frac{u+l}{2}\right)^2$, and note that $\xi(\tilde{\lambda}) = 0$ and $\xi'(\lambda) < 0$ for $\lambda \in (0, 1)$. Then $\xi\left(1 - \frac{u+l}{2}\right) = \left(1 - \frac{u+l}{2}\right)^2 - \left(1 - \frac{u+l}{2}\right) = -\left(1 - \frac{u+l}{2}\right) \frac{u+l}{2} < 0$, which implies $\tilde{\lambda} < 1 - \frac{u+l}{2}$.) Moreover, $\tilde{\lambda} \leq 1 - u$ provided $u \leq \frac{2-l+2\sqrt{1+4l-l^2}}{5} = \bar{u}(l)$. Note that $\bar{u}(l)$ is increasing in l and $\bar{u}(0) = \frac{4}{5}$, and so $\bar{u}(l) \geq \frac{4}{5}$, which implies that $\bar{\lambda}_2 = 1 - u$ whenever $u < \frac{4}{5}$. \square

Proof of Corollary 2. First, I demonstrate the comparative statics on q_A . Since $q_A = 1 - \frac{(1-(u-l)) + \sqrt{(1-(u-l))^2 + 3u^2}}{3}$, we have:

$$\begin{aligned}\frac{\partial q_A}{\partial l} &= -\frac{1}{3} - \frac{1}{3} \frac{(1-(u-l))}{\sqrt{(1-(u-l))^2 + 3u^2}} \\ \frac{\partial q_A}{\partial u} &= \frac{1}{3} + \frac{1}{3} \frac{(1-(u-l)) - 3u}{\sqrt{(1-(u-l))^2 + 3u^2}}\end{aligned}$$

Clearly, $\frac{\partial q_A}{\partial l} < -\frac{1}{3} < 0$. To show that $\frac{\partial q_A}{\partial u} < 0$, it suffices to show that $\frac{(1-(u-l)) - 3u}{\sqrt{(1-(u-l))^2 + 3u^2}} < -1$. This is true iff $[(1-(u-l)) - 3u]^2 > (1-(u-l))^2 + 3u^2$ iff $1 - u < u - l$. But this is precisely the condition under which q_A is feasible. Hence, $\frac{\partial q_A}{\partial u} < 0$ whenever q_A is feasible. By a similar method, we can verify that $\frac{\partial q_A}{\partial u} > -\frac{1}{2}$.

Next, let $b = \frac{(1-q_A) - \frac{u+l}{2}}{u-l}$. Then:

$$\begin{aligned}\frac{\partial b}{\partial u} &= -\frac{1}{u-l} \left[\frac{\partial q_A}{\partial u} + \frac{1}{2} + b \right] \\ \frac{\partial b}{\partial l} &= -\frac{1}{u-l} \left[\frac{\partial q_A}{\partial l} + \frac{1}{2} - b \right]\end{aligned}$$

Since $\frac{\partial q_A}{\partial u} > -\frac{1}{2}$ and $b \geq 0$, it follows that $\frac{\partial b}{\partial u} < 0$. Suppose $\frac{\partial b}{\partial l} < 0$. Then $\frac{\partial q_A}{\partial l} + \left(\frac{1}{2} - b\right) > 0$ and so $b < \frac{1}{2} + \frac{\partial q_A}{\partial l} < \frac{1}{6}$ (since $\frac{\partial q_A}{\partial l} < -\frac{1}{3}$). Fix some $l \in [0, 1]$. Feasibility implies $u > \frac{1+l}{2}$. We know that if $u = \frac{1+l}{2}$, then $q_A(l, u) = 1 - u$, and so $b = \frac{1}{2} \not< \frac{1}{6}$. Moreover, we just showed that $\frac{\partial b}{\partial u} < 0$. Hence $b = \frac{1}{2} + \int_{\frac{1+l}{2}}^u \frac{\partial R(l, z)}{\partial u} dz$. To ensure $b < \frac{1}{2} + \frac{\partial q_A}{\partial l}$, we need u to be sufficiently large. It can be verified (by direct computation) that $b < \frac{1}{2} + \frac{\partial q_A}{\partial l}$ when $u = 1$ for any l .

Hence, by the intermediate value theorem, there exists $\eta(l) \in (\frac{1+l}{2}, 1)$ such that $\frac{\partial R}{\partial l} < 0$ only if $u > \eta(l)$. Moreover, since $\frac{\partial q_A}{\partial l} < -\frac{1}{3}$, we know $\eta(l) > \frac{(2-l) \pm \sqrt{4+6l+6l^2}}{5} > \frac{8+l}{10}$. (To see this, note that $b < \frac{1}{6}$ implies $5u^2 - [4 - 2l]u - (l^2 + 2l) > 0$. The result follows by solving for u .) Taking the contrapositive gives $\frac{\partial b}{\partial l} > 0$ whenever $u < \eta(l)$. \square

Proof of Lemma 1. Let $U(l, u)$ be the court's value function which captures the expected discounted stream of dead-weight losses given beliefs (l, u) . (It is easily shown, using an analogous argument to the that in the proof of Proposition 2, that such a value function exists, and is unique and convex.) For notational convenience, let $s(q; l, u) = \frac{1}{24}(u-l)^2 + \frac{1}{2}(q - (1 - \frac{u+l}{2}))^2$ denote be the per-period dead-weight cost associated with output q and beliefs (l, u) . The value function must satisfy the Bellman Equation:

$$U(l, u) = \min_{q \in [1-u, 1-l]} \left\{ s(q; l, u) + \delta \left[\frac{u - (1-q)}{u-l} U(1-q, u) + \frac{(1-q) - l}{u-l} U(l, 1-q) \right] \right\}$$

It is easily verified (by substitution) that $U(l, u) = \frac{1}{6(4-\delta)}(u-l)^2$. Hence, the optimal output q minimizes:

$$\frac{1}{24}(u-l)^2 + \frac{1}{2} \left((1-q) - \frac{u+l}{2} \right)^2 + \frac{\delta}{6(4-\delta)} \left[\frac{[u - (1-q)]^3}{u-l} + \frac{[(1-q) - l]^3}{u-l} \right]$$

and it is easy to verify that the minimizer is $q = 1 - \frac{u+l}{2}$.

Let (l_t, u_t) be the beliefs at time t . Since $q_t = 1 - \frac{u_t+l_t}{2}$, period $t+1$ beliefs must either be $(l_{t+1}, u_{t+1}) = (\frac{l_t+u_t}{2}, u_t)$ or $(l_{t+1}, u_{t+1}) = (l_t, \frac{l_t+u_t}{2})$. Note that $u_t \geq u_{t+1}$ and $l_t \leq l_{t+1}$ for each t . Let $\{u_t\}$ and $\{l_t\}$ be the equilibrium sequence of beliefs. Since $\{u_t\}$ is a bounded, monotonically decreasing sequence and $\{l_t\}$ is a bounded monotonically increasing sequence, it follows that $u_t \rightarrow u_\infty$ and $l_t \rightarrow l_\infty$. Moreover, since $u_{t+1} - l_{t+1} = \frac{1}{2}(u_t - l_t)$ for every t , and $u_t - l_t \geq 0$, it follows that $\{(u_t - l_t)\}$ is a bounded and (strictly) monotonically decreasing sequence. Hence $(u_t - l_t) \rightarrow 0$, which implies $u_\infty = l_\infty$. Now, the learning technology ensures that $\theta \in [l_t, u_t]$ for every t , and so $\theta \in [l_\infty, u_\infty]$. Hence $l_\infty = \theta = u_\infty$. \square

Proof of Proposition 2. Let F be the set of bounded functions on $[0, 1]^4$. Let $\mathbf{1}[q^*(\lambda, \mu, l, u)]$ be the indicator function which takes value 1 if $q^* \in (\lambda, \mu]$. Let $T : F^2 \rightarrow F^2$ be an operator $T = (T^v, T^w)$, defined as follows: for all $(v, w) \in F^2$:

$$\begin{aligned} T^v[v, w](\lambda, \mu, l, u) &= \min_{\mathcal{L} \in [1-u, \mu]} \mathbf{1}[q^*] \left\{ s(q^*, l, u) + \delta \left[\frac{u - (1 - q^*)}{u - l} v(\mathcal{L}, \mu, 1 - q^*, u) + \frac{(1 - q^*) - l}{u - l} w(\mathcal{L}, \mu, l, 1 - q^*) \right] \right\} \\ &\quad + (1 - \mathbf{1}[q^*]) \frac{1}{1 - \delta} s(q^*, l, u) \end{aligned}$$

where $q^* = q^*(\mathcal{L}, \mu, l, u)$, and

$$\begin{aligned} T^w[v, w](\lambda, \mu, l, u) &= \min_{\mathcal{M} \in [\lambda, 1-l]} \mathbf{1}[q^*] \left\{ s(q^*, l, u) + \delta \left[\frac{u - (1 - q^*)}{u - l} v(\lambda, \mathcal{M}, 1 - q^*, u) + \frac{(1 - q^*) - l}{u - l} w(\lambda, \mathcal{M}, l, 1 - q^*) \right] \right\} \\ &\quad + (1 - \mathbf{1}[q^*]) \frac{1}{1 - \delta} s(q^*, l, u) \end{aligned}$$

where $q^* = q^*(\lambda, \mathcal{M}, l, u)$.

By construction, the value function pair (V, W) are a fixed point of this operator. To prove existence, it suffices to show that T is a contraction mapping. To do so, it suffices to show that Blackwell's conditions are satisfied. Let $(v^A, w^A), (v^B, w^B) \in F^2$ and suppose $(v^A, w^A) \leq (v^B, w^B)$. Let $(\mathcal{L}^i, \mathcal{M}^i)$ be the optimal policies when continuation utilities are given by (v^i, w^i) , and let $q^i = q^*(\mathcal{L}^i, \mu, l, u)$ or $q^i = q^*(\lambda, \mathcal{M}^i, l, u)$ (as appropriate), for $i \in \{A, B\}$. Then:

$$\begin{aligned} T^v[v^A, w^A] &= \mathbf{1}[q^A] \left\{ s(q^A, l, u) + \delta \left[\frac{u - (1 - q^A)}{u - l} v^A(\mathcal{L}^A, \mu, 1 - q^A, u) + \frac{(1 - q^A) - l}{u - l} w^A(\mathcal{L}^A, \mu, l, 1 - q^A) \right] \right\} \\ &\quad + (1 - \mathbf{1}[q^A]) \frac{1}{1 - \delta} s(q^A, l, u) \\ &\leq \mathbf{1}[q^B] \left\{ s(q^B, l, u) + \delta \left[\frac{u - (1 - q^B)}{u - l} v^A(\mathcal{L}^B, \mu, 1 - q^B, u) + \frac{(1 - q^B) - l}{u - l} w^A(\mathcal{L}^B, \mu, l, 1 - q^B) \right] \right\} \\ &\quad + (1 - \mathbf{1}[q^B]) \frac{1}{1 - \delta} s(q^B, l, u) \\ &\leq \mathbf{1}[q^B] \left\{ s(q^B, l, u) + \delta \left[\frac{u - (1 - q^B)}{u - l} v^B(\mathcal{L}^B, \mu, 1 - q^B, u) + \frac{(1 - q^B) - l}{u - l} w^B(\mathcal{L}^B, \mu, l, 1 - q^B) \right] \right\} \\ &\quad + (1 - \mathbf{1}[q^B]) \frac{1}{1 - \delta} s(q^B, l, u) \\ &= T^v[v^B, w^B] \end{aligned}$$

where the first inequality follows from the fact that \mathcal{L}^A is optimal given (v^A, w^A) whilst \mathcal{L}^B is merely feasible, and the second inequality follows from the monotonicity assumption.

By a similar argument, we can show that $T^w [v^A, w^A] \leq T^w [v^B, w^B]$. Hence $T [v^A, w^A] \leq T [v^B, w^B]$ whenever $(v^A, w^A) \leq (v^B, w^B)$, which verifies the monotonicity property. Next, let $(v^A, w^A) = (v^B + c, w^B + c)$. Then:

$$\begin{aligned}
T^v [v^A, w^A] &= \mathbf{1} [q^A] \left\{ s(q^A, l, u) + \delta \left[\frac{u - (1 - q^A)}{u - l} v^A (\mathcal{L}^A, \mu, 1 - q^A, u) + \frac{(1 - q^A) - l}{u - l} w^A (\mathcal{L}^A, \mu, l, 1 - q^A) \right] \right\} \\
&\quad + (1 - \mathbf{1} [q^A]) \frac{1}{1 - \delta} s(q^A, l, u) \\
&= \mathbf{1} [q^A] \left\{ s(q^A, l, u) + \delta \left[\frac{u - (1 - q^A)}{u - l} v^B (\mathcal{L}^A, \mu, 1 - q^A, u) + \frac{(1 - q^A) - l}{u - l} w^B (\mathcal{L}^A, \mu, l, 1 - q^A) + c \right] \right\} \\
&\quad + (1 - \mathbf{1} [q^A]) \frac{1}{1 - \delta} s(q^A, l, u) \\
&= T^v [v^B, w^B] + \mathbf{1} [q^A] \delta c \\
&\leq T^v [v^B, w^B] + \delta c
\end{aligned}$$

where the third line follows from the recognition that the maximization problems in the two cases differ only by a constant, and so the maximizers will coincide. A similar argument holds for T^w , and so the operator satisfies discounting. Hence, T is a contraction mapping, and it admits a unique fixed point (V, W) . \square

Proof of Lemma 2 . Consider the court's decision after a history in which output q was found to be inefficiently large. The new beliefs are (l, u) , with $l = 1 - q$ and λ is unchanged from the previous period. It is left to characterize the optimal restrictive opinion μ . Recall — the court can induce one of four outcomes $q \in \{q_\lambda, q_\mu, q_A, q_E\}$. First note that, assuming $u < \bar{u}(1 - q)$ (which I verify below), it follows that $\tilde{\lambda}(l, u) = 1 - u$. Then, Proposition 1 entails that the court cannot write a broad restrictive opinion that will entice the firm to choose q_E . Next, suppose $\lambda \geq \bar{\lambda}(l, u)$, which implies that the firm will definitely choose q_λ in the current and all future periods. Then the choice of μ is inconsequential. Hence, the choice of μ only matters if $\lambda < \bar{\lambda}(l, u)$. Again, by Proposition 1, it follows that q_I is chosen if $\mu \geq q_A$, and either q_λ or q_μ is chosen if $\mu < q_A$. In the latter case, $\lambda \leq \mu < q_A$, which guarantees an outcome even more biased than q_A . Clearly the court would prefer to induce q_A when it is available. Hence the optimal opinion must satisfy $\mu \geq q_A$.

In fact, we can find a tighter bound on μ . We have just shown that broad restrictive opinions do not provide positive benefits to the court. Nevertheless, a broad opinion can be equilibrium consistent, as long as there is no sub-history in which such an opinion harms the court. With positive probability, following this decision, the court will decide a sequence of cases in which the firm's output is always found to be acceptable. In each case, the court will update its beliefs in such a way that causes u to fall (leaving l unchanged), and will be able to revise λ , but not μ . To be precise, define $q_A^1(l, u) = q_A(l, u)$ and $u^1(l, u) = 1 - q_A^1(l, u)$, and then for each $k = 2, 3, \dots$ define $q_A^k(l, u) = q_A(l, u^{k-1}(l, u))$. Similarly, define $u^k(l, u) = 1 - q_A^k(l, u)$. The sequence $\{q_A^k\}$ is the set of outputs that will follow if the court always induces experimentation whenever possible, and these outputs are always found to be acceptable. The sequence $\{u^k\}$ specifies the new upper bound on beliefs. Clearly, u^k is monotonically decreasing in k , and then by Corollary 2, q_A^k is monotonically increasing. By Corollary 1, for a fixed l , experimentation is feasible whenever $u > \frac{1+l}{2}$. Moreover, since $q_A(l, \frac{1+l}{2}) = 1 - \frac{1+l}{2} = \frac{1-l}{2}$, and q_A^k is decreasing in its second argument, it follows that each q_A^k along this history satisfies $q_A^k \leq \frac{1-l}{2}$. Finally, μ should not be broad enough as to preclude the court writing a broad opinion λ located at the *ex ante* efficiently level, after the k^{th} -iteration. Hence, we need $\mu \geq 1 - \left(\frac{u^k(l, u) + l}{2}\right)$ for each k , and so, in the limit: $\mu \geq 1 - \left(\frac{\frac{1+l}{2} + l}{2}\right) = \frac{3}{4}(1 - l)$.

It remains to show that $u < \bar{u}(1 - q)$. Recall that $\bar{u}(l) = \frac{2-l+2\sqrt{1+4l-l^2}}{5}$ is increasing in l , and that $\frac{d\bar{u}}{dl} = -\frac{1}{5} + \frac{4(2-l)}{5\sqrt{5-(2-l)^2}} \in \left(\frac{1}{5}, \frac{7}{5}\right)$. Let $\kappa(l, u) = \bar{u}(1 - q_A(l, u)) - u$. We know, from Corollary 2, that $\frac{\partial q_I}{\partial u} \in \left(-\frac{1}{2}, 0\right)$. Then: $\frac{\partial \kappa}{\partial u} = -\bar{u}'(1 - q_I) \frac{\partial q_I}{\partial u} - 1 < -\frac{3}{10}$. Hence κ is monotone in u . Moreover, $\kappa(l, 1) < 0$ (since $q_I(l, 1) > 0$ for every $l < 1$) and $\kappa(l, 0.7) > 0$ since $\bar{u} \geq 0.8$ for all l . Hence, by the intermediate value theorem, there exists a unique $\tilde{u}(l)$ s.t. $\kappa(l, \tilde{u}(l)) = 0$. Moreover, by the implicit function theorem, $\frac{d\tilde{u}}{dl} = -\frac{\partial \kappa / \partial l}{\partial \kappa / \partial u} > 0$, since $\frac{\partial q_A}{\partial l} < 0$. Now, define $\hat{u}_0 = \min\{\eta(0), \tilde{u}(0)\}$. Then, since $0 \leq l_t \leq u_t \leq u_0$ for every t , and $u_0 < \hat{u}_0$, it follows that $\tilde{u}(l_t) \geq \tilde{u}(l_0) > u_0 \geq u_t$. But then, $\bar{u}(1 - q_A(l_t, u_t)) - u_t > 0$, which completes the proof. \square

Proof of Lemma 3 . Consider the court's decision after a history in which output q was found to be acceptable. The new beliefs are (l, u) , with $u = 1 - q$ and μ is unchanged from the previous period. It is left to characterize the optimal restrictive opinion λ . For the second part of the lemma, suppose $q_A(l, u)$ is not feasible (i.e. $q_A \notin [1 - u, \mu]$). Then experimentation, and hence learning, is not possible. The best the court can do is to entice the firm to implement the best static outcome. This is the *ex ante* efficient outcome, if feasible. To guarantee this outcome, it suffices to write a broad opinion $\lambda^* = \min \left\{ 1 - \frac{u+l}{2}, \mu \right\}$. Moreover, by Proposition 1, if $1 - \frac{u+l}{2} < \mu$, then this is the unique threshold that implements this outcome.

For the first part of the lemma, suppose $q_A(l, u)$ is feasible. For concreteness suppose $\mu > 1 - \frac{u+l}{2}$, which (by Lemma 2) should be the case along the equilibrium path. The firm can either write a broad opinion at the *ex ante* efficient level, or write an opinion consistent with experimentation. In the former case, experimentation stops, and the outcome repeats in all future periods. This is optimal provided:

$$\begin{aligned} \frac{1}{1-\delta} \left[\frac{1}{24} (u-l)^2 \right] &\leq \left[\frac{1}{24} (u-l)^2 + \frac{1}{2} \left(q_A - \left[1 - \frac{u+l}{2} \right] \right)^2 \right] + \delta \left\{ \frac{u - (1-q_A)}{u-l} W(\lambda, 1 - q_A, u) + \frac{(1-q_A) - l}{u-l} V(\mu, l, 1 - q_A) \right\} \\ \frac{\delta}{1-\delta} \frac{1}{24} (u-l)^2 &\leq \frac{1}{2} \left(q_A - \left[1 - \frac{u+l}{2} \right] \right)^2 + \delta \left\{ \frac{u - (1-q_A)}{u-l} W(\lambda, 1 - q_A, u) + \frac{(1-q_A) - l}{u-l} V(\mu, l, 1 - q_A) \right\} \\ \left(\frac{q_A - \left[1 - \frac{u+l}{2} \right]}{u-l} \right)^2 &\geq \frac{1}{12} \frac{\delta}{1-\delta} - \frac{2\delta}{(u-l)^2} \left\{ \frac{u - (1-q_A)}{u-l} W(\lambda, 1 - q_A, u) + \frac{(1-q_A) - l}{u-l} V(\mu, l, 1 - q_A) \right\} \end{aligned}$$

Letting $\phi(l, u) = \sqrt{\frac{1}{12} \frac{\delta}{1-\delta} - \frac{2\delta}{(u-l)^2} \left\{ \frac{u - (1-q_A)}{u-l} W(\lambda, 1 - q_A, u) + \frac{(1-q_A) - l}{u-l} V(\mu, l, 1 - q_A) \right\}}$, where λ^* and q_A are implicitly functions of (l, u) , gives the desired result.

To show that $\phi(l, u) < \frac{1}{2}$, suppose not. Then we can find a pair (l, u) with $\phi(l, u) \geq \frac{1}{2}$ and $b(l, u) = \frac{1}{2}$ such that experimentation is ideal even when the bias is most extreme. We have previously shown that $q_A = 1 - u$ when $b = \frac{1}{2}$. The expression for $\phi(l, u)$ simplifies to: $\phi = \frac{1}{12} \frac{\delta}{1-\delta} - \frac{2\delta}{(u-l)^2} V(\mu, l, 1 - q_A)$. Consider the continuation game captured in V . There are two possibilities: (i) experimentation persists, or (ii) the court writes a broad opinion located at the socially efficient level. In the first case, we have $v^1(\mu, l, u) = s(1 - u, l, u) + \delta V(\mu, l, u) =$

$\frac{1}{6(1-\delta)}(u-l)^2$. In the second case we have $v^2 = \frac{1}{1-\delta}s\left(1 - \frac{u+l}{2}, l, u\right) = \frac{1}{24(1-\delta)}(u-l)^2$. Clearly $V(\mu, l, u) = \min\{v^1, v^2\} = \frac{1}{24(1-\delta)}(u-l)^2$, and so $\phi(l, u) = 0$. But this contradicts $\phi \geq \frac{1}{2}$. (In fact, we can prove a stronger claim. Since $\phi(l, u)$ is continuous, then for $b(l, u) = \frac{1}{2} - \varepsilon$, $\phi(l, u) < \delta(\varepsilon)$. Let $\bar{\varepsilon} > 0$ satisfy $\bar{\varepsilon} < \frac{1}{2} - \delta(\bar{\varepsilon})$. Then $\phi(l, u) < \frac{1}{2} - \bar{\varepsilon} = \bar{\phi}$ for all (l, u) .) \square

Proof of Proposition 3. Let θ be the true marginal cost. Let $\varepsilon_1 = (l_1, u_1, \lambda_1, \mu_1)$ parameterize the initial beliefs and legal rule. Let ε_t denote the environment at time t , and let q_t be the output chosen at time t , assuming both the firm and court choose equilibrium strategies in every period. Let $\Delta_t = u_t - l_t$, and let $b_t = b(l_t, u_t)$.

Suppose the claim is wrong. Then there exists ε_0 , such that for every t , $\varepsilon_t \neq \varepsilon_{t+1}$ (i.e. the law never settles). This requires experimentation in every period, and by Corollary 1, this requires that $\Delta_t > 1 - u_t$ for every t . Now, by the definition of bias, $(1 - q_t) = \frac{1}{2}(u_t + l_t) + b_t(u_t - l_t)$. After each experimental output, there are two possibilities: If q_I is found acceptable, then $u_{t+1} = 1 - q_t$, and $\Delta_{t+1} = \left(\frac{1}{2} + b_t\right)\Delta_t > 0$. By contrast, if q_I is found unacceptable, then $l_{t+1} = 1 - q_t$, and $\Delta_{t+1} = \left(\frac{1}{2} - b_t\right)\Delta_t > 0$. Since $b_t < \phi(l_t, u_t) < \bar{\phi} < \frac{1}{2}$, we have $\Delta_{t+1} < \left(\frac{1}{2} + \bar{\phi}\right)\Delta_t$, and $\left(\frac{1}{2} + \bar{\phi}\right) < 1$. This implies $0 < \Delta_{t+1} < \left(\frac{1}{2} + \bar{\phi}\right)^t \Delta_1$.

Now, let $\tau = \left\lceil \frac{\ln(1-u_1) - \ln(u_1-l_1)}{\ln\left(\frac{1}{2} + \bar{\phi}\right)} \right\rceil > 0$, where $\lceil x \rceil$ is the smallest integer at least as large as x . ($u_1 - l_1 > 1 - u_1$, and $\frac{1}{2} + \bar{\phi} < 1$ guarantees that $\tau > 0$.) Then for all $t > \tau$, $\Delta_t < 1 - u_1 < 1 - u_t$ (where the final inequality follows since $\{u_t\}$ is a weakly decreasing sequence). But this contradicts $\Delta_t > 1 - u_t$ for all t . Hence, there must be some stopping time T , such that $\varepsilon_t = \varepsilon_T$ for all $t > T$. Moreover, since under experimentation $\Delta_t > 0$ implies $\Delta_{t+1} > 0$, it must be that $\Delta_T > 0$. \square

Proof of Lemma 4 & Corollary . The new per-unit expected penalty function is:

$$f(q) = \begin{cases} 0 & q \leq \lambda \\ \frac{u-(1-q)}{u-l} \cdot \frac{u+l}{2} \varphi & \lambda < q \leq \mu \\ \frac{u+l}{2} & q > \mu \end{cases}$$

The firm's optimal choice as implied by the incentives in the ambiguous region is the solution to the first order condition: $1 - q - \frac{2q-(1-u)}{u-l} \frac{u+l}{2} \varphi = 0$. Solving for q gives: $q = \frac{u-l+\frac{1}{2}(1-u)(u+l)\varphi}{(u-l)+(u+l)\varphi}$. In the first-best, the court induces the firm to choose $q = 1 - \frac{u+l}{2}$. It is easily verified that this will be the case when $\varphi = \frac{u-l}{1-l}$. The optimality result is an immediate consequence of Lemma 1.

The proof of the Corollary follows immediately by noting that φ is increasing in u and decreasing in l . (In particular, $\frac{\partial \varphi}{\partial l} = -\frac{1-u}{(1-l)^2} < 0$.) Since any learning episode (whether on the equilibrium path or not) causes u to decrease or l to increase, φ must decrease. Moreover, since by Lemma 1, $(u_t, l_t) \rightarrow (\theta, \theta)$ and $\theta < 1$, then $\varphi_t \rightarrow 0$. \square

Proof of Lemma 5 . Let $\hat{V}(l, u)$ be the value function in this new environment, where both thresholds can be amended after any controversial case. We have:

$$\hat{V}(l, u) = \min_{\{(\lambda, u) | 1-u \leq \lambda \leq \mu \leq 1-l\}} \left\{ s(q^*; l, u) + \delta \left[\frac{u - (1 - q^*)}{u - l} \hat{V}(1 - q^*, u) + \frac{(1 - q^*) - l}{u - l} \hat{V}(l, 1 - q^*) \right] \right\}$$

where $s(q, l, u) = \frac{1}{24} (u - l)^2 + \frac{1}{2} (q - (1 - \frac{u+l}{2}))^2$ is the per-period dead-weight loss, and $q^*(\lambda, \mu, l, u)$ is as defined in Proposition 1. Following the same argument as in the proof of Lemma 3, the court will optimally stop experimentation and write a targeted broad opinion when $b(l, u) > \hat{\phi}(l, u)$, where

$$\hat{\phi}(l, u) = \sqrt{\frac{1}{12} \frac{\delta}{1 - \delta} - \frac{2\delta}{(u - l)^2} \left\{ \frac{u - (1 - q_A)}{u - l} \hat{V}(1 - q_A, u) + \frac{(1 - q_A) - l}{u - l} \hat{V}(l, 1 - q_A) \right\}}$$

Now, notice that the Bellman Equations in the baseline model are the same as the Bellman Equation here, except that they each include an additional constraint. (The Bellman equation defining V has the additional constraint that μ cannot be changed. Similarly, the Bellman equation defining W has the additional constraint that λ cannot be changed.) Hence $\hat{V}(l, u) \leq V(\mu, l, u)$ for any $\mu \in [1 - u, 1 - l]$ and $\hat{V}(l, u) \leq W(\lambda, l, u)$ for any $\lambda \in [1 - u, 1 - l]$. Moreover, since μ is an ineffective policy tool, the shadow cost of the constraint that it cannot be changed is zero. By contrast, since there is a history arising with positive probability along which the court would wish to change λ , but cannot, this shadow cost is positive. Hence $V(\mu, l, u) = \hat{V}(l, u)$, whilst $W(\lambda, l, u) \geq \hat{V}(l, u)$. It follows that:

$$\frac{u - (1 - q^*)}{u - l} \hat{V}(1 - q^*, u) + \frac{(1 - q^*) - l}{u - l} \hat{V}(l, 1 - q^*) < \frac{u - (1 - q^*)}{u - l} W(\lambda, 1 - q^*, u) + \frac{(1 - q^*) - l}{u - l} V(\mu, l, 1 - q^*)$$

But, this implies that $\hat{\phi}(l, u) > \phi(l, u)$, where $\phi(l, u)$ is as defined in the proof of Lemma 3. □

References

- Baker, Scott and Claudio Mezzetti. 2012. “A Theory of Rational Jurisprudence.” *Journal of Political Economy* 120(3):513–551.
- Bryson, John E and Angus Macbeth. 1972. “Public Nuisance, the Restatement (Second) of Torts, and Environmental Law.” *Ecology Law Quarterly* pp. 241–281.
- Bueno de Mesquita, Ethan and Matthew Stephenson. 2002. “Informative Precedent and Intrajudicial Communication.” *American Political Science Review* 96(04):755–766.
- Calfee, J.E. and R. Craswell. 1984. “Some effects of uncertainty on compliance with legal standards.” *Virginia Law Review* pp. 965–1003.

- Callander, Steven and Tom S Clark. 2017. "Precedent and doctrine in a complicated world." *American Political Science Review* 111(1):184–203.
- Cooter, R., L. Kornhauser and D. Lane. 1979. "Liability Rules, Limited Information, and the Role of Precedent." *Bell Journal of Economics* 10(1):366–373.
- Cooter, Robert and Lewis Kornhauser. 1980. "Can Litigation Improve the Law without the Help of Judges?" *The Journal of Legal Studies* 9(1):139–163.
- Cooter, Robert and Thomas Ulen. 2011. *Law and economics*. 6th ed. Addison Wesley.
- Feldman, Heidi Li. 1996. "Harm and Money: Against the Insurance Theory of Tort Compensation." *Tex. L. Rev.* 75:1567.
- Fox, Justin and Georg Vanberg. 2013. "Narrow versus broad judicial decisions." *Journal of Theoretical Politics* p. 0951629813502709.
- Gennaioli, N. and A. Shleifer. 2007. "The evolution of common law." *Journal of Political Economy* 115(1):43–68.
- Hadfield, Gillian K. 2011. "The dynamic quality of law: The role of judicial incentives and legal human capital in the adaptation of law." *Journal of Economic Behavior & Organization* 79(1):80–94.
- Hadfield, G.K. 1991. "Bias in the evolution of legal rules." *Geo. Lj* 80:583.
- Hylton, Keith N. 2006. "Information, litigation, and common law evolution." *American Law and Economics Review* 8(1):33–61.
- Lax, Jeffrey R. 2012. "Political constraints on legal doctrine: how hierarchy shapes the law." *The Journal of Politics* 74(3):765–781.
- Laycock, Douglas. 2012. *Modern American Remedies: Concise Edition*. Wolters Kluwer Law & Business.

- Niblett, Anthony. 2013. "Case-by-Case Adjudication and the Path of the Law." *The Journal of Legal Studies* 42(2):303–330.
- Png, Ivan PL. 1987. "Litigation, liability, and incentives for care." *Journal of Public Economics* 34(1):61–85.
- Ponzetto, G.A.M. and P.A. Fernandez. 2008. "Case law versus statute law: an evolutionary comparison." *The Journal of Legal Studies* 37(2):379–430.
- Posner, R.A. 1973. "An economic approach to legal procedure and judicial administration." *The Journal of Legal Studies* 2(2):399–458.
- Priest, George L. 1977. "The Common Law Process and the Selection of Efficient Rules." *The Journal of Legal Studies* 6(1):65–82.
- Rubin, Paul H. 1977. "Why is the Common Law Efficient?" *The Journal of Legal Studies* 6(1):51–63.
- Staton, Jeffrey K and Georg Vanberg. 2008. "The value of vagueness: delegation, defiance, and judicial opinions." *American Journal of Political Science* 52(3):504–519.
- Zywicki, Todd J. 2002. "The rise and fall of efficiency in the common law: a supply-side analysis." *Nw. UL Rev.* 97:1551.