

THE END AND THE BEGINNING OF ANIMUS

Andrew T. Hayashi

INTRODUCTION	698
I. THE PLAYERS	703
II. A SINGLE INTERACTION	706
III. THE REPEATED GAME	711
CONCLUSION	717
APPENDIX	718

THE END AND THE BEGINNING OF ANIMUS

Andrew T. Hayashi*

If a court determines that a law with disparate impacts had an impermissible purpose, should that influence the review of subsequent similar legislation? I argue that the answer is yes, not because it is necessary to get the review of the subsequent laws right—although it may be—but because it will deter legislatures from acting on impermissible purposes in the first place. I analyze the strategic interaction between a legislature and a court, where the legislature’s motives for passing laws with disparate impacts influence the court’s judgment about whether those laws should be upheld. The court prefers to uphold an unequal law if there is a legitimate purpose for the inequality but not if the legislature is motivated by animus. If this “game” between the legislature and court is played once, then the court cannot deter laws motivated by animus without also deterring laws with legitimate purposes. But if the game is repeated, the court can selectively deter only those laws motivated by animus if the taint is inherited by subsequent legal enactments for some length of time. By appropriately choosing the length of that time, a court can choose the “end of animus” to prevent it from ever beginning.

INTRODUCTION

During the Progressive Era, some of the economists who supported minimum wage laws argued that such laws had desirable eugenic effects.¹ Sidney Webb, the prominent British socialist economist and founder of the London School of Economics, argued that such laws would drive out of the labor market the “unemployable” and discourage the employment of “boy labor, girl labor, married women’s labor, the labor of old men, of the feeble-minded, of the decrepit and broken-down invalids and all the other alternatives to the engagement of competent male adult workers.”² He thought this result was desirable because these workers were “parasitic on other classes of the community.”³

Webb’s view of how the minimum wage created unemployment dovetailed with a racist theory of wage setting, whereby racial inferiors—specifically Chinese and African American workers—were only employable because of their tolerance for lower standards of living.⁴ Thus, Webb and other like-minded economists argued that the minimum wage could operate in a social Darwinian fashion by pushing “unemployables” into involuntary unemployment where they could be segregated or sterilized and by discouraging members of undesirable races from immigrating.⁵

* Professor of Law, University of Virginia School of Law. Thanks to the organizers and participants of the University of Alabama School of Law’s symposium on *The End of Animus: The Lifespan of Impermissible Purposes*.

1. See Thomas C. Leonard, *Retrospectives: Eugenics and Economics in the Progressive Era*, 19 J. ECON. PERSPS. 207, 212–13 (2005); THOMAS C. LEONARD, *ILLIBERAL REFORMERS* 159–60 (2016).

2. Sidney Webb, *The Economic Theory of a Legal Minimum Wage*, 20 J. POL. ECON. 973, 986 (1912).

3. *Id.*

4. Leonard, *supra* note 1, at 212–13.

5. *Id.* at 213.

Nobody advocates for minimum wage laws—or, thank goodness, most any laws—on eugenics grounds anymore.⁶ Instead, there is now a range of justifications having to do with increasing worker productivity, reducing worker turnover and absenteeism, and boosting the economy.⁷ And these proffered justifications are generally considered sincere even if there is disagreement about whether they are persuasive.⁸ No one alleges that minimum wage advocates are secretly motivated by racial animus toward Blacks or Asians or want to discourage married women from accepting paid employment. And no one thinks that the fact that minimum wage laws may have once been enacted with the taint of such motives is relevant to whether they are democratically legitimate, constitutional, or simply good policy today.

By contrast, consider the history of local zoning ordinances that require a minimum lot size for properties in specified neighborhoods or that permit only single-family residences to be built on the lots.⁹ Since larger lots are more expensive, such ordinances tend to create residential segregation by income.¹⁰ And because income is correlated with race,¹¹ they will also tend to generate racial segregation.¹² By placing the neighborhood beyond the reach of low-

6. Recent research shows that the expansion of minimum wage laws in the 1960s and 1970s was an important driver of the reduction in the white–Black earnings gap and that the laws did not have significant displacement effects on Blacks. Ellora Derenoncourt & Claire Montialoux, *Minimum Wages and Racial Inequality*, 136 Q.J. ECON. 169, 169 (2021).

7. See, e.g., *Gradually Raising the Minimum Wage to \$15: Good for Workers, Good for Businesses, and Good for the Economy: Hearing Before the H. Comm. on Educ. and Lab.*, 116th Cong. 2 (2019) (statement of Rep. Robert C. “Bobby” Scott, Chairman, H. Comm. on Educ. and Lab.).

8. The economic effects of minimum wage laws continue to be disputed. David Neumark & Peter Shirley, *Myth or Measurement: What Does the New Minimum Wage Research Say About Minimum Wages and Job Loss in the United States?*, 61 INDUS. RELS.: J. ECON. & SOC’Y 384, 385 (2022) (“[I]here are also disagreements among researchers (and others) over how to interpret the overall body of evidence on the effects of minimum wages on employment . . .”).

9. These kinds of ordinances are ubiquitous in the United States and tend to force homeowners to occupy larger lots than they would prefer. See M. NOLAN GRAY & SALIM FURTH, MERCATUS CTR., DO MINIMUM-LOT-SIZE REGULATIONS LIMIT HOUSING SUPPLY IN TEXAS? 2–3, 14 (2019).

10. Jonathan T. Rothwell & Douglas S. Massey, *Density Zoning and Class Segregation in U.S. Metropolitan Areas*, 91 SOC. SCI. Q. 1123, 1130–39 (2010) (studying metro areas where suburbs have density restrictions and have greater income segregation); Michael C. Lens & Paavo Monkkonen, *Do Strict Land Use Regulations Make Metropolitan Areas More Segregated by Income?*, 82 J. AM. PLAN. ASS’N 6, 9 (2016) (“We can conclude that density restrictions lead to increased income and racial segregation . . .”). Lens and Monkkonen find that density restrictions specifically lead to the concentration of affluent households. *Id.* at 7–9.

11. See, e.g., Raj Chetty et al., *Race and Economic Opportunity in the United States: An Intergenerational Perspective*, 135 Q.J. ECON. 711, 712 (2020) (“Differences in economic outcomes by race have persisted for centuries in the United States and continue up to the present day. For example, in 2016, the median household income of black Americans was \$39,500, compared with \$65,000 for non-Hispanic white Americans.” (citations omitted)).

12. Evidence here is more mixed and suffers from causal identification issues. Keith R. Ihlanfeldt, *Exclusionary Land-Use Regulations Within Suburban Communities: A Review of the Evidence and Policy Prescriptions*, 41 URB. STUD. 261, 279 (2004); see, e.g., Rolf Pendall, *Local Land Use Regulation and the Chain of Exclusion*, 66 J. AM. PLAN. ASS’N 125, 126 (2000) (discussing low-density zoning associated with reduced number of Hispanic and Black residents); John M. Quigley et al., *Local Land-Use Controls and Demographic Outcomes in a Booming Economy*, 41 URB. STUD. 389, 414 (2004) (noting that “local land-use policy significantly impacts the path and composition of population growth” with low-density residential development being associated with

income and minority homebuyers and thereby geographically restricting their homeownership opportunities, they will also drive up the cost of homeownership for those buyers. And in fact, at least for some of its proponents, racial segregation was once a goal of density regulation.¹³

Given this history, what should we make of local zoning decisions made today that favor lower density? Opponents of “upzoning” may express concerns about the ability of public infrastructure and services to keep up with the population growth, including the effects on traffic and school crowding. These justifications, although legitimate if taken at face value, are more apt to be taken as pretextual than the contemporary reasons that are offered for minimum wage legislation. And the conclusion that they are a fig leaf for bias is even more likely if the density opponents express concerns about the effects of greater density on neighborhood “character” or “flavor,” which may be interpreted by a suspicious observer as code for racial or economic exclusivity.¹⁴

These two examples raise a common question: does the genealogy of a law or policy matter? As Professor W. Kerrel Murray has observed, this “is an old problem, but its resolution hardly seems impending.”¹⁵ If a proposed law violated some norm—of constitutionality, say, or simply prudent policymaking—in the past, what relevance does that have as to whether a substantially similar or even identical law violates the same norm now?

It may be helpful to think about cases where the past has no relevance. For example, suppose we are concerned with whether the economic benefits of a law exceed the costs, and we can determine that the benefits exceed the costs based on current information. In that case, the fact that an earlier law failed the cost–benefit test does not help us answer the question of whether the law violates the relevant norm. Or suppose that we are concerned with whether a law’s disparate treatment of two groups stigmatizes one of the groups or sends

population growth of non-Hispanic whites—and to a more modest degree, African-Americans—and declines of Asians and Hispanics).

13. See RICHARD ROTHSTEIN, *THE COLOR OF LAW: A FORGOTTEN HISTORY OF HOW OUR GOVERNMENT SEGREGATED AMERICA* 48 (2017); cf. David S. Schoenbrod, *Large Lot Zoning*, 78 *YALE L.J.* 1418, 1420–21 (1969) (“Many homeowners also want to exclude families of an economic status lower than their own, perhaps on the assumption that poorer- or darker-skinned neighbors will hinder the education of their children, commit more crime, or make living in the suburbs less prestigious.”); Paul Boudreaux, *Lotting Large: The Phenomenon of Minimum Lot Size Laws*, 68 *ME. L. REV.* 1, 11 (2016); Jenny Schuetz, *Guarding the Tonn Walls: Mechanisms and Motives for Restricting Multifamily Housing in Massachusetts*, 36 *REAL EST. ECON.* 555, 556 (2008) (“One hypothesis is that residents prefer neighbors of the same social class or race, so that affluent or largely white suburbs will use restrictive zoning to exclude lower income households and people of color.”).

14. The District Court for the Eastern District of New York in *MHANY Management, Inc. v. Incorporated Village of Garden City* found that the “sequence of events” that led to the adoption of a relatively low-density development plan with a disparate impact on racial minorities “g[a]ve rise to an inference of race-based animus.” *MHANY Mgmt., Inc. v. Incorporated Vill. of Garden City*, 985 F. Supp. 2d 390, 416 (E.D.N.Y. 2013), *aff’d sub nom.* *MHANY Mgmt., Inc. v. Cnty. of Nassau*, 819 F.3d 581 (2d Cir. 2016). The court appears to have viewed a single expression by a Garden City resident of concerns about “character” and “flavor” with particular suspicion, mentioning it twice in the opinion. *Id.* at 403, 417.

15. W. Kerrel Murray, *Discriminatory Taint*, 135 *HARV. L. REV.* 1190, 1192 (2022).

a message that it is somehow not entitled to equal respect and dignity, and we can determine based on prevailing social facts whether the law does in fact have that stigmatic or expressive meaning today. Again, whether a similar law would have had the same social meaning at some earlier date may seem irrelevant. When the relevant norm does not explicitly incorporate a regard for the past, and when all relevant facts about whether the law complies with the norm are known current facts, then the past does not matter.

Viewed in this way, we can think about what happens when not all relevant facts about whether a law complies with a norm are known. For example, suppose we have only imperfect information about whether the current economic benefits of a proposed minimum wage law exceed the costs. Knowing that the costs of a substantially identical prior law exceeded the benefits is relevant—it can tell us something about whether the currently proposed law complies with the norm—as long as there is good reason to think that there is some continuity in the economics of the minimum wage so that the costs and benefits of such laws are correlated over time. If we have perfect knowledge of the current costs and benefits, or if there is no continuity in the relevant facts over time, then the past is not informative about whether the law complies with the norm.

Now consider the case of animus. When a court rules that a law is unconstitutional because its enactment was motivated by legislative animus toward some person or group,¹⁶ how should that ruling affect the evaluation of subsequent enactments of substantially similar or even identical laws? The discussion so far suggests that the prior enactment may have evidentiary value to a current law under review if the relevant facts—the presence of legislative animus—can be only imperfectly observed and we have reason to think that there is sufficient continuity in animus over time to make the past informative of the present.

The motives and intentions of legislators are difficult to discern, particularly if they have an incentive to obscure them, so the first condition is satisfied. The question, then, is whether we have reason to think that there is sufficient continuity in a legislature's intentions between the enactments for the earlier presence of animus to be informative. Certain facts would disrupt that continuity and weaken the inference from past legislative intentions to current ones, including dramatic changes to the legislature's composition—especially if

16. On the role of animus in constitutional doctrine, see WILLIAM D. ARAIZA, *ANIMUS: A SHORT INTRODUCTION TO BIAS IN THE LAW* 159–60 (2017). The significance of legislative animus in constitutional doctrine remains unclear, with uncertainty about how animus is defined, the kinds of evidence from which it can be inferred, and its legal effect once it has been identified. For one analysis of these questions, see Susannah W. Pollvogt, Windsor, *Animus*, and the Future of Marriage Equality, 113 COLUM. L. REV. SIDEBAR 204 (2013). And yet, “animus is alive and well and is poised to increase in importance in the pantheon of equal protection arguments.” *Id.* at 210.

the new legislators were drawn from different political parties or demographic groups—or a long time between enactments.¹⁷

If the animus in a prior enactment has only evidentiary relevance to subsequent enactments, then the question of how long the “taint” should persist depends on how effectively we can ferret out the presence of animus today and how much continuity there is between past and present legislatures. But maybe the past has more than evidentiary relevance. Maybe our norms against governmental animus should directly incorporate a regard for the past because of our memories. A current law that unequally burdens different racial groups—even if motivated entirely by legitimate purposes today—may serve as a harmful reminder of a history of animus and be callous to the ways that people experience those reminders.

Thus, there are cases where one might want to take notice of earlier legal enactments motivated by animus when evaluating a current, similar enactment. That earlier animus could be evidence of current animus, or that earlier animus could provide an independent reason not to enact a similar law now. The analysis in either case takes for granted the earlier presence of animus and asks how the proper consideration of that animus should influence our judgment about the current law’s constitutionality. In this Essay, I invert this question and ask how our choice about whether the taint of animus carries over to subsequent enactments can prevent legislatures from acting out of animus in the first place.

This is a standard law and economics approach to the question of unconstitutional animus, focusing on the *ex ante* incentives created for legislatures when a finding of animus creates a taint that prevents reenactment.¹⁸ The approach I take in this Essay uses game theory to model the strategic interaction between a legislature, which may have legitimate or illegitimate motives for passing an unequal law, and a court that wants to strike down laws with disparate impacts when they are motivated by animus but not when there is a legitimate purpose for the disparity. I show that the court can do better than striking down or upholding all unequal laws by allowing a finding of animus to taint subsequent reenactments for a specified length of time. The analysis highlights the key tradeoffs faced by courts and legislatures and answers the question of how long courts should require the taint of animus to persist if they never want it to appear in the first place. By choosing when animus ends, they can stop it from beginning.

In Part I, I describe the setup of the game-theoretic model, focusing on the objectives and choices faced by the court and the legislature. Part II explains

17. There may also be actions that the current legislature can take to credibly signal that the political body no longer harbors the animus it once did. See Andrew T. Hayashi, *The Law and Economics of Animus*, 89 U. CHI. L. REV. 581, 638–40, 642 (2022).

18. See *id.* at 582–87, 632.

how the interaction between the two might be expected to play out if it only happened once. Part III describes how the interaction can change when the legislature and court are engaged in a repeated game, a complexity that allows me to analyze how the taint of animus should persist. The mathematical details of the model are included in the Appendix.

I. THE PLAYERS

Imagine that every so often, the legislature can pass a law allocating some new good or entitlement between two groups of people. Call these two groups the “Ins” and the “Outs.” The legislature can choose whether to split the good evenly between the two groups or give it all to the Ins. The allocation can be made explicitly conditional on In-group membership or simply done in a way that has that effect.¹⁹ For a specific example of the former case, consider the facts of *U.S. Department of Agriculture v. Moreno*.²⁰ *Moreno* involved Congress’s decision to extend food stamp benefits to individuals living in households whose members are all related (the Ins) and deny them to individuals living in households with unrelated members (the Outs).²¹

The legislature’s decision is influenced by how much they value the interests of the two groups. To keep things simple, suppose that the legislature may either give equal weight to the interests of the Ins and the Outs or completely disregard the interests of the Outs. Let us say that this first kind of legislature is unbiased while the second kind of legislature is biased or harbors animus for the Outs. I will use the terms “bias” and “animus” interchangeably to refer to the legislature’s desire to harm the Outs relative to the Ins. An unbiased legislature will—absent other considerations—prefer to allocate the good equally to the two groups while a biased legislature will generally prefer to give it all to the Ins.

But there is a complication. Suppose that there may be, in certain cases, some legitimate benefit to allocating the good only to the Ins and denying it to the Outs. For example, suppose that the Ins and the Outs are geographically segregated and delivering the entitlement to each group requires incurring some fixed cost, such as establishing an office for providing the benefits where each group lives. Or perhaps a separate process needs to be established to evaluate

19. Courts give heightened scrutiny to facially neutral laws only if the law has a discriminatory effect and that precise effect was intended and motivated the legislature. *Washington v. Davis*, 426 U.S. 229, 239–43 (1976). In that case, the law is subject to the same scrutiny as if the law were facially discriminatory. *Id.* at 240–43.

20. *U.S. Dep’t of Agric. v. Moreno*, 413 U.S. 528, 529–33 (1973). *Moreno* is important for its statement of the constitutional significance of animus: “[I]f the constitutional conception of ‘equal protection of the laws’ means anything, it must at the very least mean that a bare congressional desire to harm a politically unpopular group cannot constitute a *legitimate* governmental interest.” *Id.* at 534.

21. *Id.* at 529.

the eligibility of members from each group. Administrative benefits such as these give even unbiased legislatures a reason to treat groups unequally.

As a result, if the legislature gives everything to the Ins, it could be because of its animus toward the Outs, a legitimate administrative benefit to doing so, or both. The problem for a court trying to evaluate the legislative motives behind a law that excludes the Outs is that it cannot observe the presence of either animus or the existence of any legitimate purpose. The legislature will make representations about its motives, and it may be possible to collect evidence that helps verify or falsify those representations, but courts and other observers cannot ultimately know the legislature's motives.

Since there are two factors (the possible presence of bias and the possible presence of legitimate benefits from inequality) affecting the legislature's decisions that are hidden from the court, it will be convenient to speak as though there are four "types" of legislatures, with the court not knowing which type enacted the law under review. Each type is characterized by being motivated or not by animus and by being motivated or not by a legitimate purpose. The four types are biased and unbiased legislatures deriving legitimate benefits from excluding the Outs, and biased and unbiased legislatures not deriving legitimate benefits from excluding the Outs. If we let B(U) indicate the presence (absence) of bias, and we let $\beta(0)$ indicate the presence (absence) of legitimate benefits from inequality, then for the sake of parsimony, we can refer to these four types of legislatures using the labels: $L_{B\beta}$, $L_{U\beta}$, L_{B0} , L_{U0} . The four types are summarized in the table below.

Legitimate Purpose for Inequality		
	Yes	No
Biased	$L_{B\beta}$	L_{B0}
Unbiased	$L_{U\beta}$	L_{U0}

Biased legislatures will always prefer to exclude the Outs, regardless of whether there is also a legitimate purpose for doing so. An unbiased legislature will prefer to include the Outs unless there is a legitimate reason to exclude them. The only uncertainty surrounds the preferences of the unbiased legislature when there are legitimate benefits from excluding the Outs. If the legitimate benefits from excluding the Outs are very small, then the legislature will prefer to include them, but if the legitimate benefits are large, then they will prefer to exclude them. In the following analysis, I assume that the legitimate benefits of excluding the Outs are large enough that an unbiased legislature would choose to exclude them. If we do not make this assumption, then the

problem of distinguishing between good and bad motives behind an unequal law would never arise because only biased legislatures would ever want to exclude the Outs.

The other player in this drama is the court. The court has its own preferences: it is unbiased between the Ins and the Outs, valuing the interests of each group equally. The court also values any legitimate benefits that might arise from excluding the Outs but only if an unbiased legislature made the decision to exclude the Outs. This means that the court would prefer to strike down an unequal law passed by a biased legislature even if there is also a legitimate purpose for the law. That is, legislative animus is not redeemed in the court's judgment by the presence of a legitimate purpose. The presence of animus is fatal, a "silver bullet."²²

A word about "preferences" in this context: I am using this term in the technical sense it is used in economics, which is simply to describe how an actor orders a set of outcomes according to their choice-worthiness.²³ To prefer *A* to *B* is not necessarily to derive greater pleasure from *A* than from *B* but simply to say that one will act in such a way as to bring about *A* instead of *B*. And the process that leads to the choice of *A* over *B* can be the internal deliberations of a single actor—in the case of an individual's preferences—or it can be the result of negotiations and bargaining or some kind of voting procedure among the members of a corporate body, such as a corporation or legislature.

Thus, when I describe the court as valuing any legitimate benefits that arise from an unequal law only when the legislature is unbiased, this is best understood as describing how the court applies animus doctrine rather than a description of what judges desire. The economic approach sidesteps the question (much-discussed in the scholarly literature) of whether intent-based tests make any sense as applied to the decisions of a group of people or whether only individuals can have intentions.²⁴ In the economic approach, the court is concerned with whether the legislature's behavior is consistent with giving inadequate weight to the interests of the Outs. On this view, the animus inquiry is not about what is in the minds of individual legislators but about how the legislature acts as a body. In this sense, the economic approach resembles what

22. This is how Professor Pollvogt argues that animus operates, doctrinally. Susannah W. Pollvogt, *Unconstitutional Animus*, 81 *FORDHAM L. REV.* 887, 889, 930 (2012).

23. See ANDREU MAS-COLELL ET AL., *MICROECONOMIC THEORY* 5 (1995).

24. There are difficult issues with intent-based inquiries as applied to legislation. Professors Kendrick and Schwartzman discuss the ontological, epistemic, and relevance objections to intent-based tests as well as the question of the persistence of "taint," which is the subject of this symposium. Leslie Kendrick & Micah Schwartzman, *The Etiquette of Animus*, 132 *HARV. L. REV.* 133, 146–54 (2018). On the role of intentions in constitutional doctrine, see, for example, Richard H. Fallon, Jr., *Constitutionally Forbidden Legislative Intent*, 130 *HARV. L. REV.* 523 (2016); Calvin Massey, *The Role of Governmental Purpose in Constitutional Judicial Review*, 59 *S.C. L. REV.* 1 (2007); Ashutosh Bhagwat, *Purpose Scrutiny in Constitutional Analysis*, 85 *CALIF. L. REV.* 297 (1997); Caleb Nelson, *Judicial Review of Legislative Purpose*, 83 *N.Y.U. L. REV.* 1784 (2008).

Professor Fallon calls the “[o]bjective [c]onceptions of [l]egislative [i]ntent,” which use features of the law to attribute intentions to the legislature.²⁵

II. A SINGLE INTERACTION

Ultimately, we are interested in analyzing the relationship between the legislature and the court as a repeated game in which the legislature periodically passes legislation that is reviewed by the court. It is only with repeated interactions that we can consider the possibility for a finding of animus in one case to taint subsequent legal enactments. In that dynamic setting, the court’s decisions at any point in time can depend on earlier judgments it has made about the legislature’s motives. That dynamism is necessary to analyze the “end of animus.” But the interaction between the legislature and court with respect to a single law is the building block of the dynamic analysis, so we start by analyzing the strategic behavior of the legislature and the court in a single period.

Suppose that the legislature can enact either a law that equally, or a law that unequally, allocates some good or entitlement. If the law distributes the good equally between the Ins and the Outs, then the court will uphold the law. If the law distributes the good unequally, the court may either uphold the law or strike down the law and mandate an equal allocation. Striking down the law is a relatively bad outcome for everyone. The legislature would prefer to have enacted an equal law in the first place than to pass an unequal law that is struck down and replaced with an equal law. And the court would prefer that too. In the real world, courts sometimes strike down unequal laws and sometimes do not. From the legislature’s perspective then, it is a gamble whether to enact an unequal law. It will only want to try and pass an unequal law if the benefits of that unequal law to the legislature—derived either from animus or a legitimate purpose—are big enough and there is a high enough probability of it being upheld. Suppose the court is certain to uphold an unequal law. In that case, the legislature will enact an unequal law except in the one case that it is unbiased and has no legitimate purpose for the inequality. On the other hand, if the court is certain to strike down an unequal law, then the legislature will never try to enact one.

The legislature’s decision about what kind of law to enact is necessarily made in anticipation of what the court will do when reviewing an unequal law. By increasing the probability with which it strikes down unequal laws, the court can discourage unequal laws from being enacted in the first place because it makes the gamble on inequality less attractive to the legislature. A legislature of

25. Fallon, *supra* note 24, at 541–53. Former Professor and now Justice Elena Kagan’s analysis of speech regulation uses this approach. See Elena Kagan, *Private Speech, Public Purpose: The Role of Governmental Motive in First Amendment Doctrine*, 63 U. CHI. L. REV. 413, 505–14 (1996).

type $L_{B\beta}$ —which is both biased *and* derives legitimate benefits from inequality—is the type of legislature most willing to “roll the dice” with the court and pass an unequal law. The next most likely legislature to pass an unequal law will be either $L_{U\beta}$ or L_{B0} , depending on how big the inequality’s legitimate benefits are relative to the amount of legislative animus in the particular law at issue. Thus, there are two possible ways that the types of legislatures might be ordered in terms of how much they want to enact an unequal law:

$$L_{B0} < L_{U\beta} < L_{B\beta} \quad (1)$$

$$L_{U\beta} < L_{B0} < L_{B\beta} \quad (2)$$

In the first case, the legitimate benefits of inequality are a more powerful motive for inequality than animus, and in the second case, the opposite is true. Of course, the legislature L_{U0} will never prefer to enact an unequal law even if it is certain to be upheld.

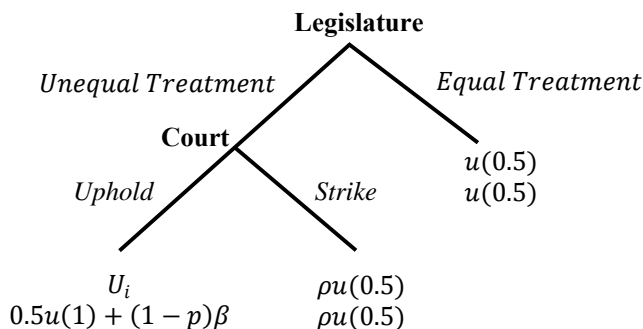
The game-theoretic analysis of this interaction assumes that the legislature and court will both choose a strategy designed to best achieve their objectives given the strategy that the other one has adopted. An *equilibrium* of this game is a situation where the strategy chosen by the legislature is its best choice given the strategy being used by the court, and the strategy being used by the court is its best choice given the strategy being used by the legislature.²⁶

The figure below is a representation of the strategic interaction in the form of a game tree, which illustrates the sequence of decisions moving from top to bottom. The tree looks the same regardless of the legislature’s type, except for the “payoffs” at the end of the branches. These payoffs are numerical representations of the preferences of the court and the legislature over the different outcomes of the game with higher numbers indicating more preferred outcomes.

First, the legislature enacts a law creating an equal or an unequal allocation of the good between the Ins and the Outs. We can think of the total amount of the good as being 1. If the allocation is equal—allocating 0.5 of the entitlement to each group—then the game ends because the court does not review that law. And in that case, the legislature and the court both get the payoff associated with an equal allocation: $u(0.5)$. If the allocation is unequal—allocating all 1 of the good to the Ins—then the court strikes down the law with some positive

26. See Stephen W. Salant & Theodore S. Sims, *Game Theory and the Law: Ready for Prime Time?*, 94 MICH. L. REV. 1839, 1855 (1996) (“More generally, a (pure strategy) Nash equilibrium is a strategy profile with the property that, *given* the strategies of all the other players, no player can do strictly better by unilaterally choosing a different strategy than he has.”).

probability, and the legislature is faced with a risky gamble that could result in the law being upheld or in its being struck down.



What does it mean that the court strikes down the law only probabilistically? The best way of understanding the court's strategy in this case is not that they actually decide whether to strike down the law randomly, such as by flipping a coin or pulling a number out of a hat or some such thing. Instead, it is that its decision appears at least partially uncertain from others' perspectives—specifically the legislature's perspective—because it is based on factors that others cannot observe. The decision may be well-reasoned internally from the court's perspective but seem somewhat random to outsiders.²⁷

If the court does strike down the law, then the court and the legislature both derive the payoff $\rho u(0.5)$, where ρ is some fraction less than one. This means, consistent with what we have assumed above, that the payoffs to both the court and the legislature are lower if an unequal law is struck down and replaced with an equal law than if an equal law was enacted in the first place.

In the bottom left corner are the payoffs if the legislature passes an unequal law that is upheld. The first term U_i is the legislature's payoff, and it depends on the type of legislature that passed the law (U stands for "utility"—a numerical index of how desirable this outcome is—and i is the type of the legislature). This payoff will be largest in the case of a biased legislature that also has a legitimate purpose for inequality. The term below U_i is the payoff that the court gets from upholding an unequal law. This also depends, in part, on the type of legislature that passed the unequal law.

The first part of the payoff $0.5u(1)$ is the payoff associated with an unequal allocation of the good, given that the court values the interests of the Ins and the Outs equally. Viewed in isolation, this outcome would be worse for the court than if the law had provided equal treatment; however, the unequal

27. See John C. Harsanyi, *Games with Randomly Disturbed Payoffs: A New Rationale for Mixed-Strategy Equilibrium Points*, 2 INT'L. J. GAME THEORY 1, 1 (1973); Robert J. Aumann, *Correlated Equilibrium as an Expression of Bayesian Rationality*, 55 ECONOMETRICA 1, 2-3 (1987).

law may have a legitimate purpose yielding social benefit β , and that purpose will incline the court to uphold the law as long as the legislature that enacted the law is not motivated by animus. The variable p is the probability that the legislature enacting the unequal law is biased. And so, one can see that as the court becomes more confident that the legislature is biased—as p increases—then the weight given to any legitimate social benefit from the unequal law declines, and it becomes more attractive for the court to strike down the law. The two extreme cases are $p = 1$, where the court knows that the legislature enacting the unequal law is biased, and $p = 0$, where the court knows that the legislature enacting the unequal law is unbiased. In general, the court will not know the legislature's type, so p will depend on which types of legislatures are willing to gamble and enact an unequal law given the likelihood of it being struck down.

There are several possible outcomes or “equilibria” to this game. In the first equilibrium, the court strikes down any unequal law it reviews, and all legislatures, regardless of their type, choose to enact only equal laws. It is easy to check that this is an equilibrium. Given that the court will strike down any law that is passed, the legislature will always prefer to pass an equal law than to pass an unequal law that will be struck down. Given that the legislature is going to pass only equal laws, it really does not matter what the court might have done if the legislature had (counterfactually) enacted an unequal law.

In a second possible equilibrium, the court upholds all laws it reviews, and the legislature will pass an unequal law unless it is unbiased and there is no legitimate purpose for inequality. This is an equilibrium because the other types of legislatures prefer to enact an unequal law that is upheld rather than pass an equal one. It is true that the court reviewing an unequal law in this equilibrium does not know whether the legislature that enacted it is biased, but as long as the probability p of legislative animus is low enough, the court will prefer to uphold the unequal law, tolerating the relatively low risk of animus. We are more likely to see this outcome in the world if the cost to the court of striking down an unequal law is very high, the legitimate social benefits to the inequality are large, or the probability of the legislature being biased is low.

But these two equilibrium outcomes do not correspond to what we observe in the world. Sometimes legislatures do enact unequal laws—which is not predicted by the first equilibrium—and sometimes courts strike down unequal laws that have been enacted—which is not predicted by the second equilibrium. There is, however, a third possible outcome to this game that is more realistic.

The court may be indifferent between striking down an unequal law and not, given the particular probability that the law was motivated by animus. When that is the case, it can—and will be willing to—set the probability of striking down the law at a range of values without destroying the legislature's incentive to pass the unequal law. In this outcome of the game, unequal laws

will be enacted—some motivated by animus and some not—and the court will sometimes strike down an unequal law.

We can see this by thinking first about the equilibrium where the court never strikes down an unequal law and all types of legislatures (other than L_{U0}) enact unequal laws. Let us call s the probability that the court strikes down an unequal law. If the court is indifferent about striking down an unequal law—because there is some chance the legislature passing the unequal law was motivated by animus and some chance it was not—then it can increase s a little above 0, introducing just a small chance of striking down the law, and it will not deter any type of legislature from passing an unequal law, but now unequal laws will sometimes be struck down.

As the court continues to increase the probability of striking down the law, eventually some types of legislatures will no longer be willing to take their chances enacting an unequal law. For example, suppose that the potential legitimate social benefits from inequality are greater than the animus we are concerned about. This is case (1) above. Then, the court can sustain an outcome where all unequal laws have a legitimate purpose, although sometimes they are enacted by a biased legislature and sometimes by an unbiased legislature. To achieve this, the probability of striking down the law s must be low enough to ensure that an unbiased legislature with a legitimate purpose wants to pass an unequal law but not so low that a biased legislature without a good purpose (type L_{B0}) will also want to do so.²⁸ This is not a terrible outcome for the court. In this outcome, unequal laws will still be tainted in the sense that some of the laws that are upheld will have been motivated by a biased legislature. But on the positive side, the court can at least deter the biased legislature that does not also have a legitimate purpose by setting s sufficiently high.

So, in a single interaction between the legislature and the court, we can see how their strategic behaviors could result in the outcome whereby legislatures enact laws with disparate impacts and the court—unable to observe the legislature's motives—might sometimes (but not always) strike down unequal laws. This is an unsatisfactory state of affairs. The very best outcome from the court's perspective would be to uphold an unequal law only if enacted by an unbiased legislature with a legitimate purpose. But if the court wants to allow for the possibility of upholding at least some unequal laws—and it well might, if there can be important legitimate purposes for unequal treatment—then the possibility that the legislature is also expressing animus for the Outs is unavoidable. The court cannot purge the specter of animus from unequal laws because its only tool for discouraging unequal laws—increasing the probability

28. In theory, there are also equilibria where only biased legislatures choose inequality and the court does not strike down the law despite knowing this because the costs of doing so are so great. We ignore these equilibria because our focus is on the outcome where the court strikes down laws not knowing for sure whether the legislature is biased or not.

of striking down the law—does not discriminate between biased and unbiased legislatures, and no legislature values an unequal law more than a legislature motivated by both animus *and* a legitimate purpose.

III. THE REPEATED GAME

Viewed in isolation, any one-time interaction where the legislature enacts an unequal law will always leave open the possibility that animus was a motivation. Although a vigorous and thorough factfinding process may smoke out evidence of improper motives among individual legislators, it is easy enough for canny legislators to obscure their intentions by being careful about what they say. Moreover, it is conceptually unclear how the motives of individual legislators might aggregate to form intent at the level of the legislature as a body. For these reasons, our approach has been to focus on the objective actions of the legislature itself and what those actions reveal about the values that motivate its actions. As discussed at the end of Part II, the court's inability to strike down only those laws tainted by animus is due to limitations on the tools available to it. Striking down laws more frequently will not do the trick. What is needed is a punishment that deters legislatures acting out of animus *more* than unbiased legislatures with legitimate purposes.

The court faces what is known as a “screening problem” in the economics literature.²⁹ This is a situation in which one party to a transaction or strategic interaction has information relevant to that interaction that only it knows.³⁰ For example, some workers are more productive than other workers. A prospective employer would like to hire the most productive workers, but only the workers know how productive they are. Some consumers value a product more than others, and a seller would like to charge a higher price to those who value it more. But consumers know how much they value the product while the seller does not. Civil litigants have private information about the merits of the claim and about their willingness to settle, which would affect the strategy of the other party if it were common knowledge. The screening problem is how the party who is ignorant in each of these situations can cause the other party to reveal her private information.

Consider the case of an employer and her potential employees. Although she cannot directly observe how productive they would be in the position, the employer may be able to distinguish between the different kinds of workers by giving them the opportunity to undertake some task that would credibly reveal (i.e., beyond mere assertions) just how productive they will be. The key is that the task must be easier for more productive workers to perform than for less

29. For a summary of this literature, see John G. Riley, *Silver Signals: Twenty-Five Years of Screening and Signaling*, 39 J. ECON. LITERATURE 432 (2001).

30. See *id.* at 433–34.

productive workers. Education is the canonical example of this in the economics literature.³¹ If higher education is easier for more productive workers to obtain than for less productive workers, then more productive workers will have an incentive to acquire education to reveal that they are, in fact, more productive workers.

To be clear, this incentive exists even if the education itself does not affect their productivity but only because education signals to employers that they are more productive. As a result, employers can offer employment or higher wages to workers with higher education levels knowing they are more productive. As long as it is sufficiently easier for high-productivity workers to obtain education than for lower productivity workers, less productive workers will be unwilling to acquire additional education, even knowing they would be paid more. This is known as a “separating equilibrium” because workers with different productivity levels acquire different amounts of education and receive different wages—they separate according to their productivity “types.”³²

The interaction between the legislature—which has private information about its motives in enacting unequal laws—and the court—which wants to evaluate those laws based on the legislature’s motives—can be analyzed as a screening game. The problem is that the court does not have something like education in the employment context to discriminate between biased and unbiased legislatures. But things could be different. In other work, I have argued that courts could try to screen out unequal laws passed by legislatures with bad intentions by requiring the legislature to compensate the Outs whenever they enact an unequal law.³³ The compensation requirement acts as a useful screen in this context because a legislature that values the interests of the Outs less than the interests of the Ins will be much less willing to transfer resources to that group than an unbiased legislature.³⁴ As a result, biased legislatures will be discouraged from enacting unequal laws, whereas unbiased legislatures will not.

But that analysis, like the one in Part II, focused on a one-time interaction. In fact, the legislature and the courts will repeat this interaction many times over the years. The strategies available to the court and the legislature multiply significantly in a repeated game. Let us call each play of the game—each enactment of a law—a “period” of the repeated game. The court and legislature could choose their actions in each period independently of their previous interactions. In this case, the outcome in each period of the repeated game will be one of the equilibria of the one-shot game described in Part II. But the court

31. See Michael Spence, *Job Market Signaling*, 87 Q.J. ECON. 355, 358 (1973); Andrew Weiss, *Human Capital vs. Signalling Explanations of Wages*, 9 J. ECON. PERSPS. 133 (1995).

32. See, e.g., Jiwoong Lee et al., *Separating Equilibrium in Quasi-Linear Signaling Games*, 48 INT’L J. GAME THEORY 1033, 1034 (2019).

33. Hayashi, *supra* note 17, at 633–47.

34. *Id.*

and legislature could also make their actions contingent on the actions taken in the previous periods.

This dynamic framework allows us to explore the question at the heart of this symposium: How long should the taint of animus last? That is, when a law is judged to have been motivated by animus, should our evaluation of subsequent—perhaps very similar—legal enactments be influenced by that prior determination? One perspective on this question focuses on the subsequent enactments and on ensuring that the taint of the earlier enactment is purged.³⁵ The economic approach inverts that question, making the persistence of that taint part of the court's strategy for deterring the legislature from acting out of bias in the first instance.

This approach asks whether the persistence of animus's taint can be a useful screen by which the court can sort good legislative motives from bad motives. I show that if the court plays a "trigger" strategy, whereby a finding of animus in one period is followed by n periods in which the court strikes down all unequal laws, then it may be able to ensure that unequal laws are only passed with legitimate motives—at least temporarily. This is an improvement on the outcome of the single-shot game and shows how making the taint of animus persist can be instrumentally useful in purging unequal laws of the specter of legislative bias.

We need to make some assumptions to formalize this repeated interaction for analysis. Suppose that the legislature's type, as either biased or unbiased, is determined at the beginning of the game and is fixed for the duration. Then, in each period of the game, the legislature confronts a new issue involving allocating some good or entitlement. Sometimes there will be a legitimate purpose for excluding the Outs from the allocation, and sometimes there will not. Only the legislature knows. From this point, play in each period proceeds just as it does in the single-shot game, with the legislature choosing an allocation and the court then deciding whether to strike it down—perhaps only with some probability—if it is unequal.

An equilibrium in the repeated game is defined in the same way as an equilibrium in the single-shot game: it is a pair of strategies for the court and for the legislature that specifies what each will do in each period of the game, such that the strategy played by each yields the best outcome, given the strategy that the other one is using. The dynamic nature of the game means that there can be outcomes in a single period of the repeated game that differ from the outcomes that could occur if the interaction occurred only once. Recall that in the one-shot version of the game there are two equilibria.³⁶ In the first equilibrium, the legislature passes only equal laws regardless of its type because

35. See, e.g., Murray, *supra* note 15.

36. We shall ignore the equilibrium in which the court simply upholds all unequal laws because this is a special case of the inequality equilibrium.

the court would strike down any unequal law. Call this the “equality equilibrium.” In the second equilibrium, the legislature will pass an unequal law unless it is both unbiased and there is no legitimate purpose for unequal treatment, and the court strikes down those laws with some probability. Call this the “inequality equilibrium.” In fact, there are actually a whole set of inequality equilibria in which both biased and unbiased legislatures will pass an unequal law, and each equilibrium involves a different probability that the unequal law is struck down. At one extreme, there is an inequality equilibrium where the court *never* strikes down an unequal law, but the court can choose any probability for striking down the law as long as it is not too high because beyond that point the legislature will no longer want to take its chances by passing an unequal law.

In the repeated game, it is possible to ensure that—at least sometimes—unequal laws are only passed by unbiased legislatures so unequal laws are purged of animus. This is the best outcome of all. The key to achieving this outcome is the fact that there are two equilibria of the single-shot game. Crucially, although both the biased legislature and the unbiased legislature with a legitimate purpose prefer the inequality equilibrium to the equality equilibrium, the biased legislature prefers the inequality equilibrium *more*.

Consider three possible ways that the repeated game might unfold over $n + 1$ periods.³⁷ The first possibility is that the court would strike down any unequal law, so the equality equilibrium is played each period. The second possibility is that the inequality equilibrium is played in each period. These are represented by the following sequences of outcomes (the subscript indicates the time period):

$$\textit{Equality}_1, \textit{Equality}_2, \textit{Equality}_3, \dots, \textit{Equality}_{n+1} \quad (\text{E})$$

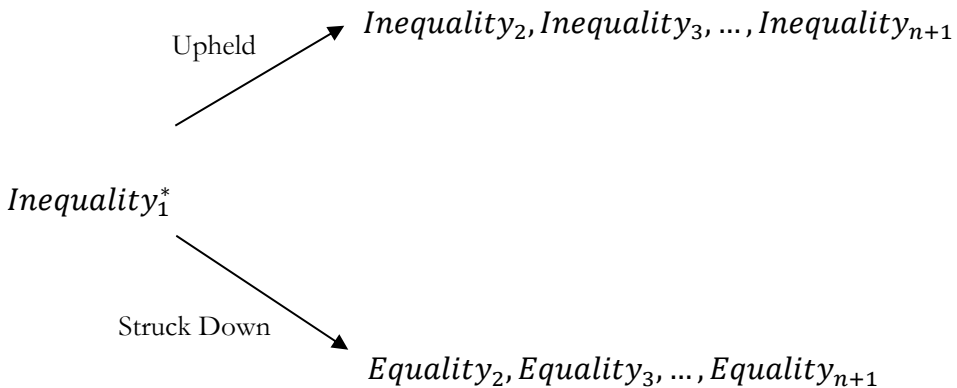
$$\textit{Inequality}_1, \textit{Inequality}_2, \textit{Inequality}_3, \dots, \textit{Inequality}_{n+1} \quad (\text{I})$$

Viewed from the beginning of the game, both the biased and the unbiased legislatures will prefer that their interactions with the court looks like (I) rather than (E). For the biased legislature, this is easy to see. The biased legislature will always prefer the inequality equilibrium to the equality equilibrium, regardless of whether or not it turns out that they happen to also have a legitimate purpose for enacting an unequal law in a given period. But the unbiased legislature will also prefer (I) to (E) because in any period in which they have a legitimate

37. There are other equilibria of the repeated game, but my task here is simply to prove that the trigger strategy can yield a better outcome than repeated play of the mixed-strategy or pure-strategy equilibria of the period game.

purpose for an unequal law, they will prefer to have the option to enact that law. In any period in which it turns out that there is no legitimate purpose for inequality, they will simply enact an equal law. But the biased legislature prefers (I) to (E) *more* than the unbiased legislature does. This is because although both may benefit from being able to pass an unequal law if a legitimate purpose for one arises, the biased legislature will always benefit from enacting an unequal law just because it has animus for the Outs.

Suppose, however, that instead of always striking down unequal laws (E) or upholding all unequal laws with some probability (I), the court implemented the following legal rule: uphold an unequal law with some probability, and then, if the law is in fact struck down, strike down all unequal laws with certainty for n periods. And if the law is not in fact struck down, then uphold unequal laws with the same probability for n periods. Another way of describing this is to say that the court plays the inequality equilibrium in Period 1, and then plays the equality equilibrium for the next n periods if the law is struck down, and plays the inequality equilibrium for the next n periods if the law is upheld in the first period.



This strategy discourages both biased and unbiased legislatures from passing an unequal law in Period 1 because both value the option of being able to pass an unequal law in the future. The crucial thing to observe is that the threat of animus's taint—the consequence of an unequal law being struck down in Period 1—matters *more* to biased legislatures than unbiased legislatures. An unbiased legislature worries about the possibility that there will be significant legitimate benefits from allocating some future entitlement unequally. The biased legislature worries about this too but because biased legislatures *also* prefer inequality in the absence of a legitimate purpose, they are less willing to risk the taint of animus and therefore less inclined to take their chances enacting an unequal law in Period 1. The threat of imposing the equality equilibrium acts as a screen, with a greater deterrent effect for biased legislatures.

In the Appendix, I describe the formal conditions under which animus's taint makes it possible to deter only biased legislatures from passing unequal laws. In general, it is easier to deter biased legislatures from passing unequal laws when they do not also have a legitimate purpose, although in some cases it is possible to screen out all legislative animus even if the legislature also has a legitimate purpose. And it is easier to deter legislatures from acting out of animus the longer the taint of animus lasts and the more the legislature cares about its ability to enact its preferred legislation in the future. To illustrate with an extreme example, consider a legislature about to pass the final bill in a lame-duck session. A court cannot deter the legislature from acting out of animus by threatening that the animus would taint subsequent legislation because the legislature (as currently constituted) will not be around to suffer the consequences.

By using persistent taint to deter legislative animus, courts can achieve a new outcome in Period 1 that is better than both the equality and inequality equilibria of the single-shot game; it is an outcome where courts—and the public—can be more confident that a law providing differential treatment or merely disparate impact was not motivated by animus because legislatures motivated by animus have been selectively deterred. It is an inequality equilibrium, but one that is better from the court's perspective than the best inequality equilibrium that can be achieved without persistent taint. Call this new Period 1 outcome *Inequality*₁^{*}, and call the equilibrium of this repeated game, whereby the first period is followed by *n* occurrences of either the equality or the inequality equilibrium—depending on whether the unequal law was struck down in Period 1—the “taint” or (I) equilibrium.

By requiring the taint of animus to persist, selectively deterring legislative animus in Period 1, the court can generate the (I) outcome. But will it want to? Compare (I) with (I) and with (E). From the court's perspective, (I) is clearly better than (I)—playing a single-shot inequality equilibrium each period—both because of the benefits of screening out legislative animus in Period 1 (*Inequality*₁^{*} is preferred to *Inequality*₁) and because the court prefers to mandate equal laws for the remaining *n* periods than to play the inequality equilibrium (*Equality* is preferred to *Inequality* in each subsequent period). Thus, whether an unequal law is struck down in the first period or not, the court prefers the sequences of outcomes to those in (I).

The trickier question is whether the court also prefers (I) to (E), i.e., to simply mandating equality in each period by striking down all unequal laws. Allowing inequality when there is a legitimate purpose and no animus is the best outcome for the court, so in the first period, *Inequality*₁^{*} is preferred to *Equality*₁. However, by committing to the (I) strategy, the court runs the risk that it will have to play the *Inequality* equilibrium and therefore tolerate the possible existence of legislative animus for the next *n* periods (*Equality* is preferred to *Inequality* in each subsequent period). This is the price that the

court must pay for deterring legislative animus in Period 1; it must commit to playing the inequality equilibrium that the legislature—and especially the biased legislature—prefers for the next n periods. Whether that is worth it or not for the court depends on how long it must commit to tolerating inequality in the future to screen out biased legislatures and the size of the potential *legitimate* benefits that might justify unequal treatment. I summarize the conditions under which the court will prefer (I) to (E) in the Appendix.

CONCLUSION

A judicial determination that a law with disparate impacts was motivated by animus should influence the review of subsequent unequal laws. From one perspective, this is because taking notice of the earlier law is necessary to get the review of the later law right. Intuitively, evidence of legislative animus at some earlier date may be evidence of legislative animus at a later date. Or, as Professor Murray has argued, perhaps evidence of animus at an earlier date should trigger heightened review and a heavier justificatory burden for subsequent unequal laws even if they are known not to have impermissible purposes.³⁸

But from the economic perspective I adopt in this Essay, the question is inverted: Can taking notice of legislative animus from an earlier law when reviewing subsequent laws give the legislature the incentive to get the *earlier law* right? The answer to this question is yes, at least under certain circumstances. Allowing the taint of animus to persist can be a useful tool for helping to purge legal enactments with a disparate impact from the specter of legislative animus.

38. Murray, *supra* note 15.

APPENDIX

Single Period Game

A legislature (L) must allocate 1 unit of a good between two groups, each of which values the good according to the concave utility function u . There are two features of L's preferences which are the private information of L. First, θ is the relative weight L assigns to the first group's utility. With probability p , the legislature is biased and cares only about the first group ($\theta = 1$), and with probability $(1 - p)$, the legislature is unbiased and cares equally about both groups ($\theta = 0.5$). Second, with probability q , the legislature derives legitimate benefit β from allocating the entire good to the first group. An example of β might be the cost savings if there are fixed administrative costs to allocating the good to each group. Assume that L can either split the good equally between the two groups or give all of it to the first group. L's utility from the allocation is:

$$\theta u(x) + (1 - \theta)u(1 - x) + \beta \cdot 1\{x = 1\}$$

We assume that $\beta \geq u(0.5) - 0.5u(1)$ which implies that the legitimate benefits are large enough to cause an unbiased legislature to prefer $x = 1$. This is necessary to avoid trivializing the problem. If it does not hold, then unbiased legislatures will never want to choose an unequal allocation. The assumption holds if the benefits are large or if utility is less concave.

If L chooses an equal allocation, the court does nothing. If the allocation is unequal, the court strikes down the law with probability s and implements an equal allocation, in which case L receives payoff $\rho u(0.5)$; $\rho < 1$. The reason for ρ is to capture the idea that L would rather choose an equal allocation than have the court impose equality after striking down an unequal allocation.

The court is unbiased, valuing the utility of the two groups equally. The court's payoff from striking down a law is the same as the legislature's: $\rho u(0.5)$. The court also values any administrative benefits from inequality but only if the inequality was chosen by an unbiased legislature. This assumes that a legitimate benefit does not purge the taint of bias. The court's utility is:

$$0.5u(x) + 0.5u(1 - x) + (1 - \tilde{p})\beta \cdot 1\{x = 1\}$$

where \tilde{p} is the probability that the legislature is biased conditional on choosing $x = 1$.

We begin with L's choice about how to allocate the good. There are four types of legislatures, characterized by whether they are biased (B) or unbiased (U) and whether there is (β) or is not (0) a legitimate benefit from choosing an unequal allocation. The legislature's expected utility from choosing $x = 1$

depends on its type and the court's choice s . The legislature's strategies are as follows:

- $L_{B\beta}$ chooses $x = 1$ if $U_{B\beta}(s) = (1 - s)(u(1) + \beta) + \rho u(0.5) \geq u(0.5)$
- $L_{U\beta}$ chooses inequality if $U_{U\beta}(s) = (1 - s)(0.5u(1) + \beta) + \rho u(0.5) \geq u(0.5)$
- L_{B0} chooses $x = 1$ if $U_{B0}(s) = (1 - s)u(1) + \rho u(0.5) \geq u(0.5)$
- L_{U0} never chooses $x = 1$

Note that $L_{B\beta}$ values inequality the most and L_{U0} never chooses inequality. $L_{U\beta}$ values inequality more than L_{B0} if $\beta > 0.5u(1)$.

The court's strategy after observing $x = 1$ must be optimal given its belief \tilde{p} , which we assume is correct in equilibrium.

Pure Strategy Equilibria

The court strikes down any unequal allocation ($s = 1$), and all types choose $x = 0.5$. Both the legislature and the court get payoffs of $u(0.5)$, and no laws are struck down in equilibrium.

The court does not strike down unequal allocations ($s = 0$), and all types other than L_{U0} choose $x = 1$. The court's expected payoff is $0.5u(1) + (1 - \tilde{p})\beta$ while the legislature's expected payoff depends on its type.

Mixed Strategy Equilibria

We are interested in situations where the court strikes down some unequal allocations, which requires that the court play a mixed strategy. Mixed strategy equilibria where $s \in (0,1)$ exist if:

$$0.5u(1) + (1 - \tilde{p})\beta = \rho u(0.5)$$

If only biased types choose inequality, then $\tilde{p} = 1$. Mixed strategy equilibria could exist in this case but only because the cost of striking down an unequal law is sufficiently great. There could also be mixed strategy equilibria where only types $L_{U\beta}$ and $L_{B\beta}$ choose inequality, but this would be an equilibrium where the court knows there is a legitimate legislative purpose for the inequality. We ignore these equilibria to focus on equilibria where there is uncertainty about the existence of both β and θ , so that all of $L_{U\beta}$, $L_{B\beta}$, and L_{B0} choose

inequality and the court chooses s sufficiently low to induce all three types to choose inequality. Since the type that least prefers inequality could be either $L_{U\beta}$ or L_{B0} , this means that $s \leq \min\{s_{U\beta}^*, s_{B0}^*\}$ where:

$$s_{U\beta}^* = \frac{\beta - (u(0.5) - 0.5u(1))}{\beta - (\rho u(0.5) - 0.5u(1))}$$

$$s_{B0}^* = \frac{u(1) - u(0.5)}{u(1) - \rho u(0.5)}$$

The court's expected payoff is $0.5u(1) + (1 - \tilde{p})\beta$, where:

$$\tilde{p} = \frac{pq + p(1 - q)}{pq + p(1 - q) + (1 - p)q}$$

Multiple-Period Game

The court's expected payoff in each period in which it plays the mixed strategy equilibrium is $0.5u(1) + (1 - \tilde{p})\beta = \rho u(0.5)$, and in each period that the legislature chooses equality, the court's payoff is $u(0.5)$. In this section, we try to identify an equilibrium of the repeated game that improves on both of these outcomes for the court by using a trigger strategy: the court will play s that sustains a mixed strategy equilibrium in the current period. If it strikes down the law, it plays $s = 1$ for n periods and if it does not strike down the law, then it plays s for the next n periods. For a discount factor of $\delta \in (0, 1)$, let $D^n u = \frac{\delta(1-\delta^n)}{1-\delta} u$ be the discounted sum of u for the next n periods.

The legislature's type—biased or not—is fixed at the beginning of the game, but there is a new allocation and new draw of β made in each period. This means that a biased legislature does not know whether they will be of type L_{B0} or $L_{B\beta}$ in Period 2. A legislature of type i 's expected utility in the next period given the equilibrium in that next period is:

$$E[U_i(s)] = qU_{i\beta}(s) + (1 - q)U_{i0}(s)$$

The legislature's expected utility in the next period depends on (1) the probability that the court will strike down an unequal law, and (2) the probability that there will be legitimate benefits from unequal treatment in the next period.

If the court is playing the “taint” equilibrium trigger strategy, then L_{B0} chooses $x = 1$ in Period 1 if the expected utility from choosing inequality this period, given the courts trigger strategy, is greater than the utility from choosing equality:

$$(1 - s)u(1) + s(\rho u(0.5) + D^n(u(0.5) - E[U_B(s)])) \geq u(0.5)$$

Similarly, $L_{U\beta}$ chooses $x = 1$ in Period 1 if:

$$(1 - s)(0.5u(1) + \beta) + s(\rho u(0.5) + D^n(u(0.5) - E[U_U(s)])) \\ \geq u(0.5)$$

Finally, $L_{B\beta}$ chooses $x = 1$ in Period 1 if:

$$(1 - s)(u(1) + \beta) + s(\rho u(0.5) + D^n(u(0.5) - E[U_B(s)])) \\ \geq u(0.5)$$

Note that $E[U_i(s)] > u(0.5)$ for each type because the mixed strategy is an equilibrium, so the trigger strategy makes it less attractive to choose $x = 1$ in the current period by threatening lower payoffs in the future than the type would otherwise get. The factor D^n , which is increasing in n and δ , affects the size of that threat. Note also that $E[U_B(s)] > E[U_U(s)]$ because the payoff from being biased in an inequality equilibrium is higher than the payoff of being unbiased in an inequality equilibrium. The difference in the expected utilities of being a biased—rather than an unbiased—legislature playing the inequality equilibrium is:

$$E[U_B(s)] - E[U_U(s)] = \\ q(1 - s)0.5u(1) + (1 - q)((1 - s)u(1) - (1 - s\rho)u(0.5)) > 0$$

The derivative of this difference with respect to s is negative, so by increasing the probability of striking down an unequal law, the court reduces the difference in expected benefits received by biased and unbiased legislatures from the inequality equilibrium.

We now consider how the trigger strategy can deter biased—but not unbiased—legislatures from passing an unequal law. For L_{B0} to choose equality and $L_{U\beta}$ to choose inequality, it must be the case that:

$$(1 - s)(0.5u(1) - \beta) < sD^n(E[U_B(s)] - E[U_U(s)])$$

For $L_{B\beta}$ to choose equality and $L_{U\beta}$ to choose inequality, it must be the case that:

$$(1 - s)0.5u(1) < sD^n(E[U_B(s)] - E[U_U(s)])$$

The right-hand side of these last two inequalities is the same, but the left-hand side of the second inequality is greater than the left-hand side of the first. Thus, it is easier to deter L_{B0} than to deter $L_{B\beta}$ from passing an unequal law in Period 1.

The next question is whether the court will prefer the outcome associated with the trigger strategy to the alternatives. Instead of playing the trigger strategy, the court can get $(1 + D^n)(0.5u(1) + (1 - \tilde{p})\beta)$ by playing the mixed strategy equilibrium in each period, and it can get $(1 + D^n)u(0.5)$ by striking down all unequal laws in each period. If it is possible to drive out all bad types in Period 1, the court will choose this strategy over always striking down unequal laws if:

$$(1 - s)(0.5u(1) + \gamma) + s\rho u(0.5) - u(0.5) \\ \geq D^n u(0.5)(1 - (1 - s)\rho - s)$$

That is, it will prefer the taint equilibrium associated with its trigger strategy if the benefit in the current period from screening out the biased legislatures exceeds the cost to the court of committing to the mixed strategy to reward the legislature. Note that satisfying this condition implies that the court will also prefer this equilibrium to playing the mixed equilibrium each period because $u(0.5) > \rho u(0.5)$. Increasing D^n makes the taint equilibrium less desirable for the court because it involves tolerating more inequality in the future. The court may deter only L_{B0} from passing unequal laws by choosing a shorter taint

period.³⁹ This will yield a smaller payoff in the current period for the court, but it is easier to achieve and involves smaller subsequent costs.

39. If only $L_{B\gamma}$ and $L_{U\gamma}$ choose $x = 1$, the conditional probability of legislative bias is $\tilde{p} = (pq) / (pq + (1-p)q)$.